

IBM XL C/C++ for Linux, V11.1



Getting Started with XL C/C++

Version 11.1

IBM XL C/C++ for Linux, V11.1



Getting Started with XL C/C++

Version 11.1

Note

Before using this information and the product it supports, read the information in “Notices” on page 43.

First edition

This edition applies to IBM XL C/C++ for Linux, V11.1 (Program 5724-X14) and to all subsequent releases and modifications until otherwise indicated in new editions. Make sure you are using the correct edition for the level of the product.

© Copyright IBM Corporation 1996, 2010.

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

Contents

About this document v

Conventions	v
Related information	ix
IBM XL C/C++ information	ix
Standards and specifications	x
Other IBM information	xi
Other information	xi
Technical support	xi
How to send your comments	xi

Chapter 1. Introducing XL C/C++ 1

Commonality with other IBM compilers	1
Hardware and operating system support	1
A highly configurable compiler	1
Language standards compliance	2
Compatibility with GNU	3
Source-code migration and conformance checking	3
Libraries	4
Tools and utilities	5
Program optimization	6
64-bit object capability	6
Shared memory parallelization	7
Diagnostic listings	7
Symbolic debugger support	8

Chapter 2. What's new for IBM XL C/C++ for Linux, V11.1 9

Operating system support	9
Support for POWER7 processors	9
C++0x	10
Performance and optimization	13
New diagnostic reports	15
Utilization tracking and reporting tool	17
New or changed compiler options and directives	17
Built-in functions new for this release	21

Chapter 3. Enhancements added in previous versions 23

Enhancements added in Version 10.1	23
--	----

Operating system support	23
C++0x	23
Other XL C/C++ language-related updates	25
OpenMP 3.0	25
Performance and optimization	26
New or changed compiler options and directives	27
Enhancements added in Version 9.0	28
C/C++ language-related updates	28
Architecture and processor support	28
Performance and optimization	29
Other new or changed compiler options	31

Chapter 4. Setting up and customizing XL C/C++ 33

Using custom compiler configuration files	33
Configuring compiler utilization tracking and reporting	33

Chapter 5. Developing applications with XL C/C++ 35

The compiler phases	35
Editing C/C++ source files	35
Compiling with XL C/C++	35
Invoking the compiler	36
Compiling parallelized XL C/C++ applications	36
Specifying compiler options	37
XL C/C++ input and output files	38
Linking your compiled applications with XL C/C++	38
Dynamic and static linking	39
Running your compiled application	39
XL C/C++ compiler diagnostic aids	40
Debugging compiled applications	40
Determining what level of XL C/C++ is installed	41

Notices 43

Trademarks and service marks	45
--	----

Index 47

About this document

This document contains overview and basic usage information for the IBM® XL C/C++ for Linux®, V11.1 compiler.

Who should read this document

This document is intended for C and C++ developers who are looking for introductory overview and usage information for XL C/C++. It assumes that you have some familiarity with command-line compilers, a basic knowledge of the C and C++ programming languages, and basic knowledge of operating system commands. Programmers new to XL C/C++ can use this document to find information on the capabilities and features unique to XL C/C++.

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 icons, as described in “Conventions.”

Throughout this document, the `xlc` and `xlc++` compiler invocations are used to describe the actions of the compiler. You can, however, substitute other forms of the compiler invocation command if your particular environment requires it, and compiler option usage will remain the same unless otherwise specified.

While this document covers information on configuring the compiler environment, and compiling and linking C or C++ applications using the XL C/C++ compiler, it does not include the following topics:

- Compiler installation: see the *XL C/C++ Installation Guide* for information on installing XL C/C++.
- Compiler options: see the *XL C/C++ Compiler Reference* for detailed information on the syntax and usage of compiler options.
- The C or C++ programming languages: see the *XL C/C++ Language Reference* for information on the syntax, semantics, and IBM implementation of the C or C++ programming languages.
- Programming topics: see the *XL C/C++ Optimization and Programming Guide* for detailed information on developing applications with XL C/C++, with a focus on program portability and optimization.

Conventions

Typographical conventions

The following table explains the typographical conventions used in the IBM XL C/C++ for Linux, V11.1 information.

Table 1. *Typographical conventions*

Typeface	Indicates	Example
bold	Lowercase commands, executable names, compiler options, and directives.	The compiler provides basic invocation commands, xlc and xlc (xlc++), along with several other compiler invocation commands to support various C/C++ language levels and compilation environments.
<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.
<u>underlining</u>	The default setting of a parameter of a compiler option or directive.	nomaf <u>maf</u>
monospace	Programming keywords and library functions, compiler builtins, examples of program code, command strings, or user-defined names.	To compile and optimize myprogram.c, enter: xlc myprogram.c -O3.

Qualifying elements (icons)

Most features described in this information apply to both C and C++ languages. In descriptions of language elements where a feature is exclusive to one language, or where functionality differs between languages, this information uses icons to delineate segments of text as follows:

Table 2. *Qualifying elements*







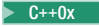






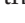
Qualifier/Icon	Meaning
C only, or C only begins   C only ends	The text describes a feature that is supported in the C language only; or describes behavior that is specific to the C language.
C++ only, or C++ only begins   C++ only ends	The text describes a feature that is supported in the C++ language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

Table 2. Qualifying elements (continued)

Qualifier/Icon	Meaning
C++0x, or C++0x begins   C++0x ends	The text describes a feature that is introduced into standard C++ as part of C++0x.

Syntax diagrams

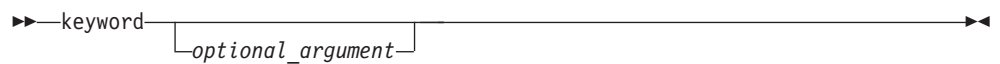
Throughout this information, diagrams illustrate XL C/C++ syntax. This section will help you to interpret and use those diagrams.

- 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.

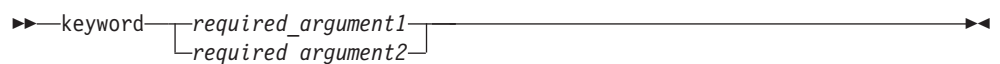
- Required items are shown on the horizontal line (the main path):



- 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:



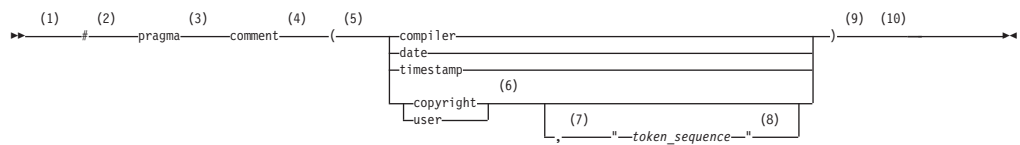
- The item that is the default is shown above the main path.



- 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 in this information

The examples in this information, 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.

The examples for installation information are labelled as either *Example* or *Basic example*. *Basic examples* are intended to document a procedure as it would be performed during a basic, or default, installation; these need little or no modification.

Related information

The following sections provide related information for XL C/C++:

IBM XL C/C++ information

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

- README files

README files contain late-breaking information, including changes and corrections to the product information. README files are located by default in the XL C/C++ directory and in the root directory of the installation CD.

- 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 *IBM XL C/C++ for Linux, V11.1 Installation Guide*.

- Information center

The information center of searchable HTML files can be launched on a network and accessed remotely or locally. Instructions for installing and accessing the online information center are provided in the *IBM XL C/C++ for Linux, V11.1 Installation Guide*.

The information center is viewable on the Web at <http://publib.boulder.ibm.com/infocenter/lnxpcmp/v111v131/index.jsp>.

- PDF documents

PDF documents are located by default in the `/opt/ibmcmp/vacpp/11.1/doc/LANG/pdf/` directory, where *LANG* is one of `en_US`, `zh_CN`, or `ja_JP`. The PDF files are also available on the Web at <http://www.ibm.com/software/awdtools/xlcpp/linux/library/>.

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

Table 3. XL C/C++ PDF files

Document title	PDF file name	Description
<i>IBM XL C/C++ for Linux, V11.1 Installation Guide, GI11-7914-00</i>	install.pdf	Contains information for installing XL C/C++ and configuring your environment for basic compilation and program execution.
<i>Getting Started with IBM XL C/C++ for Linux, V11.1, GI11-7913-00</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.

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

Document title	PDF file name	Description
<i>IBM XL C/C++ for Linux, V11.1 Compiler Reference, SC23-8606-00</i>	compiler.pdf	Contains information about the various compiler options, pragmas, macros, environment variables, and built-in functions, including those used for parallel processing.
<i>IBM XL C/C++ for Linux, V11.1 Language Reference, SC23-8607-00</i>	langref.pdf	Contains information about the C and C++ programming languages, as supported by IBM, including language extensions for portability and conformance to nonproprietary standards.
<i>IBM XL C/C++ for Linux, V11.1 Optimization and Programming Guide, SC23-8608-00</i>	proguide.pdf	Contains information on advanced programming topics, such as application porting, interlanguage calls with Fortran code, library development, application optimization and parallelization, 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 information related to XL C/C++ including IBM Redbooks® publications, white papers, tutorials, and other articles, is available on the Web at:

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

For more information about boosting performance, productivity, and portability, see the C/C++ café at <http://www-949.ibm.com/software/rational/cafe/community/ccpp>.

Standards and specifications

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 information.

- *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.
- *ANSI/IEEE Standard for Binary Floating-Point Arithmetic*, ANSI/IEEE Std 754-1985.
- *OpenMP Application Program Interface Version 3.0*, available at <http://www.openmp.org>

Other IBM information

- *ESSL for AIX V4.4 - ESSL for Linux on POWER V4.4 Guide and Reference* available at the Engineering and Scientific Subroutine Library (ESSL) and Parallel ESSL Web page.

Other information

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

Technical support

Additional technical support is available from the XL C/C++ Support page at <http://www.ibm.com/software/awdtools/xlcpp/linux/support/>. This page provides a portal with search capabilities to a large selection of Technotes and other support information.

If you cannot find what you need, you can send e-mail to compinfo@ca.ibm.com.

For the latest information about XL C/C++, visit the product information site at <http://www.ibm.com/software/awdtools/xlcpp/linux/>.

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 information or any other XL C/C++ information, send your comments by e-mail to compinfo@ca.ibm.com.

Be sure to include the name of the information, the part number of the information, 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. Introducing XL C/C++

IBM XL C/C++ for Linux, V11.1 is an advanced, high-performance compiler that can be used for developing complex, computationally intensive programs, including interlanguage calls with C and Fortran programs.

This section discusses the features of the XL C/C++ compiler at a high level. It is intended for people who are evaluating the compiler, and for new users who want to find out more about the product.

Commonality with other IBM compilers

IBM XL C/C++ for Linux, V11.1 is part of a larger family of IBM C, C++, and Fortran compilers.

XL C/C++, together with XL Fortran, comprise the family of XL compilers.

These compilers are derived from a common code base that shares compiler function and optimization technologies for a variety of platforms and programming languages. Programming environments include IBM AIX®, IBM Blue Gene®/L™, IBM Blue Gene®/P™, the Cell Broadband Engine™ architecture, IBM i, selected Linux distributions, IBM z/OS®, and IBM z/VM®. The common code base, along with compliance with international programming language standards, helps support consistent compiler performance and ease of program portability across multiple operating systems and hardware platforms.

Hardware and operating system support

IBM XL C/C++ for Linux, V11.1 supports several Linux distributions. See the README file and "Before installing XL C/C++" in the *XL C/C++ Installation Guide* for a complete list of requirements.

The compiler, its libraries, and its generated object programs will run on POWER4™, POWER5™, POWER5+™, POWER6®, POWER7™, PowerPC®, and PowerPC 970 systems with the required software and disk space.

To exploit the various supported hardware configurations, the compiler provides options to tune the performance of applications specific to the type of hardware that will be used to execute the compiled applications.

A highly configurable compiler

You can use a variety of compiler invocation commands and options to tailor the compiler to your unique compilation requirements.

Compiler invocation commands

XL C/C++ provides several different commands that you can use to invoke the compiler, for example, **xlC**, **xlC++**, and **xlC**. Each invocation command is unique in that it instructs the compiler to tailor compilation output to meet a specific language level specification. Compiler invocation commands are provided to support all standardized C/C++ language levels, and many popular language extensions as well.

The compiler also provides corresponding "_r" versions of most invocation commands, for example, `xlcr_r` and `xlC_r`. The "_r" invocations instruct the compiler to link and bind object files to thread safe components and libraries, and produce thread safe object code for compiler-created data and procedures.

For more information about XL C/C++ compiler invocation commands, see "Invoking the compiler" in the *XL C/C++ Compiler Reference*.

Compiler options

You can choose from a large selection of compiler options to control compiler behavior. Different categories of options help you to debug your applications, optimize and tune application performance, select language levels and extensions for compatibility with non-standard features and behaviors supported by other C or C++ compilers, and perform many other common tasks that would otherwise require changing the source code.

XL C/C++ lets you specify compiler options through a combination of environment variables, compiler configuration files, command line options, and compiler directive statements embedded in your program source.

For more information about XL C/C++ compiler options, see "Compiler options reference" in the *XL C/C++ Compiler Reference*.

Custom compiler configuration files

The installation process creates a default compiler configuration file containing stanzas that define compiler option default settings.

Your compilation needs may frequently call for specifying compiler option settings other than the default settings provided by XL C/C++. If so, you can use makefiles to define your compiler option settings, or alternatively, you can create custom configuration files to define your own sets of frequently used compiler option settings.

For more information about using custom compiler configuration files, see "Using custom compiler configuration files" on page 33.

Utilization tracking configuration file

The utilization tracking and reporting feature of the compiler has its own configuration file. The main compiler configuration file contains an entry that points to this file. The different installations of the compiler product can use a single utilization tracking configuration file to centrally manage the functionality of the utilization tracking and reporting feature. This utilization and reporting tool can be used to detect whether your organization's use of the compiler exceeds your license entitlements. For detailed information about the utilization tracking and reporting feature, see "Tracking and reporting compiler usage" in the *XL C/C++ Compiler Reference*.

Language standards compliance

The compiler supports the following programming language specifications for C/C++:

- ISO/IEC 9899:1999 (C99)
- ISO/IEC 9899:1990 (referred to as C89)
- ISO/IEC 14882:2003 (referred to as Standard C++)
- ISO/IEC 14882:1998, the first official specification of the language (referred to as C++98)

In addition to the standardized language levels, XL C/C++ supports language extensions, including:

- OpenMP V3.0 to support portable parallelized programming
- Language extensions to support vector programming
- A subset of GNU C and C++ language extensions
- C++0x

See C++0x in the *Getting Started with XL C/C++ for Linux* for more details.

See "Language levels and language extensions" in the *XL C/C++ Language Reference* for more information about C/C++ language specifications and extensions.

Compatibility with GNU

XL C/C++ supports a subset of the GNU compiler command options to facilitate porting applications developed with **gcc** and **g++** compilers.

This support is available when the **gxlc** or **gxlc++** invocation command is used together with select GNU compiler options. Where possible, the compiler maps GNU options to their XL C/C++ compiler option counterparts before invoking the compiler.

These invocation commands use a plain text configuration file to control GNU-to-XL C/C++ option mappings and defaults. You can customize this configuration file to better meet the needs of any unique compilation requirements you may have. See "Reusing GNU C/C++ compiler options with **gxlc** and **gxlc++**" in the *XL C/C++ Compiler Reference* for more information.

XL C/C++ uses GNU C and GNU C++ header files together with the GNU C and C++ runtime libraries to produce code that is binary-compatible with that produced by the GNU Compiler Collection (GCC). Portions of an application can be built with XL C/C++ and combined with portions built with GCC to produce an application that behaves as if it had been built solely with GCC.

To achieve binary compatibility with GCC-compiled code, a program compiled with XL C/C++ includes the same headers as those used by a GNU compiler residing on the same system. To ensure that the proper versions of headers and runtime libraries are present on the system, the prerequisite GCC compiler must be installed before installing XL C/C++.

Some additional noteworthy points about this relationship are:

- IBM built-in functions coexist with GNU C built-ins.
- Compilation of C and C++ programs uses the GNU C and GNU C++ header files.
- Compilation uses the GNU assembler for assembler input files.
- Compiled C code is linked to the GNU C runtime libraries.
- Compiled C++ code is linked to the GNU C and GNU C++ runtime libraries.
- Debugging uses the GNU debugger, **gdb**

Source-code migration and conformance checking

XL C/C++ helps protect your investment in your existing C/C++ source code by providing compiler invocation commands that instruct the compiler to compile your application code to a specific language level.

You can also use the `-qlanglvl` compiler option to specify a given language level, and the compiler will issue warnings, errors, and severe error messages if language or language extension elements in your program source do not conform to that language level.

See "qlanglvl" in the *XL C/C++ Compiler Reference* for more information.

Libraries

XL C/C++ includes a runtime environment containing a number of libraries.

Mathematical Acceleration Subsystem library

The Mathematical Acceleration Subsystem (MASS) library consists of scalar and vector mathematical intrinsic functions tuned specifically for optimum performance on supported processor architectures. You can choose a MASS library to support high-performance computing on a broad range of processors, or you can select a library tuned to support a specific processor family.

The MASS library functions support both 32-bit and 64-bit compilation modes, are thread-safe, and offer improved performance over the default `libm` math library routines. They are called automatically when you request specific levels of optimization for your application. You can also make explicit calls to MASS library functions regardless of whether optimization options are in effect or not.

See "Using the Mathematical Acceleration Subsystem" in the *XL C/C++ Optimization and Programming Guide* for more information.

Basic Linear Algebra Subprograms

The Basic Linear Algebra Subprograms (BLAS) set of high-performance algebraic functions are shipped in the `libxlopt` library. These functions let you:

- Compute the matrix-vector product for a general matrix or its transpose.
- Perform combined matrix multiplication and addition for general matrices or their transposes.

For more information about using the BLAS functions, see "Using the Basic Linear Algebra Subprograms" in the *XL C/C++ Optimization and Programming Guide*.

Other libraries

The following are also shipped with XL C/C++:

- The SMP runtime library supports both explicit and automated parallel processing. See "SMP Runtime Library" in the *XL C/C++ Optimization and Programming Guide*.
- XL C++ Runtime Library contains support routines needed by the compiler.

Support for Boost libraries

XL C/C++ for Linux, V11.1 partially supports the Boost V1.40.0 libraries. A patch file is available that modifies the Boost 1.40.0 libraries so that they can be built and used with XL C/C++ applications. The patch or modification file does not extend nor provide additional functionality to the Boost libraries.

To access the patch file for building the Boost libraries, go to this page at <http://www.ibm.com/support/docview.wss?uid=swg27006911> and follow the link in the download required Boost modification file section.

You can download the latest Boost libraries at <http://www.boost.org/>.

For more information on support for libraries, search on the XL C/C++ Compilers support page at <http://www.ibm.com/software/awdtools/xlcpp/linux/support/>.

Tools and utilities

There are many tools and utilities that are included with XL C/C++.

new_install

After you install IBM XL C/C++ for Linux, V11.1, running this utility will configure the compiler for use on your system.

vac_configure

Use this utility to create additional compiler configuration files to contain your own custom sets of compiler option default settings.

cleanpdf command

A command related to profile-directed feedback (PDF), **cleanpdf** removes all profiling information from the directory to which profile-directed feedback data is written.

mergepdf command

A command related to profile-directed feedback (PDF), **mergepdf** provides the ability to weigh the importance of two or more PDF records when combining them into a single record. The PDF records must be derived from the same executable.

resetpdf command

The current behavior of the **cleanpdf** command is the same as the **resetpdf** command, and is retained for compatibility with earlier releases on other platforms.

showpdf command

The **showpdf** command displays the call and block counts for all procedures executed in a profile-directed feedback training run (compilation under the options **-qpdf1** and **-qshowpdf**).

gxlc and gxlc++ utilities

The **gxlc** and **gxlc++** invocations translate GNU C or GNU C++ invocation commands into corresponding **xlc** or **xlc++** commands before invoking the IBM XL C/C++ for Linux, V11.1 compiler. The purpose of these utilities is to minimize the number of changes to makefiles used for existing applications built with the GNU compilers and to facilitate the transition to IBM XL C/C++ for Linux, V11.1.

Utilization reporting tool

The utilization reporting tool generates a report describing your organization's utilization of the compiler. These reports help determine whether your organization's use of the compiler matches your compiler license entitlements. You can use the **urt** command to control how the report is generated. For more information about this tool, see "Tracking and reporting compiler usage" in the *XL C/C++ Compiler Reference*.

Program optimization

XL C/C++ provides several compiler options that can help you control the optimization and performance of your programs.

With these options, you can perform the following tasks:

- Select different levels of compiler optimizations.
- Control optimizations for loops, floating point, and other types of operations.
- Optimize a program for a particular class of machines or for a very specific machine configuration, depending on where the program will run.

Optimizing transformations can give your application better overall execution performance. XL C/C++ provides a portfolio of optimizing transformations tailored to various supported hardware. These transformations offer the following benefits:

- Reducing the number of instructions executed for critical operations
- Restructuring generated object code to make optimal use of the Power Architecture[®]
- Improving the usage of the memory subsystem
- Exploiting the ability of the architecture to handle large amounts of shared memory parallelization

For more information, see these related topics:

- "Optimizing your applications" in the *XL C/C++ Optimization and Programming Guide*
- "Optimizing and tuning options" in the *XL C/C++ Compiler Reference*
- "Compiler built-in functions" in the *XL C/C++ Compiler Reference*

64-bit object capability

The XL C/C++ compiler's 64-bit object capability addresses increasing demand for larger storage requirements and greater processing power.

The Linux operating system provides an environment that allows you to develop and execute programs that exploit 64-bit processors through the use of 64-bit address spaces.

To support larger executables that can be fit within a 64-bit address space, a separate 64-bit object format is used. The linker binds these objects to create 64-bit executables. Objects that are bound together must all be of the same object format. The following scenarios are not permitted and will fail to load, execute, or both:

- A 64-bit object or executable that has references to symbols from a 32-bit library or shared library
- A 32-bit object or executable that has references to symbols from a 64-bit library or shared library
- A 64-bit executable that explicitly attempts to load a 32-bit module
- A 32-bit executable that explicitly attempts to load a 64-bit module

XL C/C++ supports 64-bit mode mainly through the use of the **-q64** and **-qarch** compiler options. This combination determines the bit mode and instruction set for the target architecture.

For more information, see "Using 32-bit and 64-bit modes" in the *XL C/C++ Optimization and Programming Guide*.

Shared memory parallelization

XL C/C++ supports application development for multiprocessor system architectures.

You can use any of the following methods to develop your parallelized applications with XL C/C++:

- Directive-based shared memory parallelization
- Instructing the compiler to automatically generate shared memory parallelization
- Message passing based shared or distributed memory parallelization (MPI)

For more information, see "Parallelizing your programs" in the *XL C/C++ Optimization and Programming Guide*.

OpenMP directives

OpenMP directives are a set of API-based commands supported by XL C/C++ and many other IBM and non-IBM C, C++, and Fortran compilers.

You can use OpenMP directives to instruct the compiler how to parallelize a particular loop. The existence of the directives in the source removes the need for the compiler to perform any parallel analysis on the parallel code. OpenMP directives require the presence of Pthread libraries to provide the necessary infrastructure for parallelization.

OpenMP directives address three important issues of parallelizing an application:

1. Clauses and directives are available for scoping variables. Frequently, variables should not be shared; that is, each processor should have its own copy of the variable.
2. Work sharing directives specify how the work contained in a parallel region of code should be distributed across the processors.
3. Directives are available to control synchronization between the processors.

As of XL C/C++ for Linux, V10.1, XL C/C++ supports the OpenMP API Version 3.0 specification. See "OpenMP 3.0" on page 25 for an overview of the support provided by this feature.

For more information about program performance optimization, see:

- "Optimizing your applications" in the *XL C/C++ Optimization and Programming Guide*
- www.openmp.org

Diagnostic listings

The compiler output listing and XML reports can provide important information to help you develop and debug your applications more efficiently.

Listing information is organized into optional sections that you can include or omit. For more information about the applicable compiler options and the listing itself, refer to "Compiler messages and listings" in the *XL C/C++ Compiler Reference*.

It is also possible to get information from the compiler in XML 1.0 format about some of the optimizations that the compiler was able to perform and also which optimization opportunities were missed. This information can be used to reduce programming effort when tuning applications, especially high-performance applications. The report is defined by an XML schema and easily consumable by tools that you can create to read and analyze the results. For detailed information about this report and how to use it, see "Using reports to diagnose optimization opportunities" in the *XL C/C++ Optimization and Programming Guide*.

Symbolic debugger support

You can instruct XL C/C++ to include debugging information in your compiled objects. The debugging information can be examined by **gdb** or any other symbolic debugger to help you debug your programs.

Chapter 2. What's new for IBM XL C/C++ for Linux, V11.1

This section describes features and enhancements added to the compiler in IBM XL C/C++ for Linux, V11.1.

Operating system support

This section contains information about the supported operating systems.

IBM XL C/C++ for Linux, V11.1 supports the following operating systems supported by IBM Power Systems™ servers:

- SUSE Linux Enterprise Server 11 Service Pack 1 (SLES 11 SP1)
- Red Hat Enterprise Linux 5.5 (RHEL 5.5)
- SUSE Linux Enterprise Server 10 Service Pack 2 (SLES 10 SP2)

Support for POWER7 processors

IBM XL C/C++ for Linux, V11.1 supports POWER7 processors.

The new features and enhancements introduced in support for the POWER7 processors, fall under the following four categories:

- vector scalar extension data types and intrinsic functions
- MASS libraries for POWER7 processors
- built-in functions for POWER7 processors
- compiler options for POWER7 processors

Vector scalar extension data types and intrinsic functions

This release of the compiler supports the Vector Scalar eXtension (VSX) instruction set in the POWER7 processors. New data types and intrinsic functions are introduced to support the VSX instructions. With the VSX intrinsic functions and the original Vector Multimedia eXtension (VMX) intrinsic functions, you can efficiently manipulate vector operations in your application.

For more information about the VSX data types and intrinsic functions, see Vector types in the *XL C/C++ Language Reference* and Vector built-in functions in the *XL C/C++ Compiler Reference*.

Mathematical Acceleration Subsystem (MASS) libraries for POWER7 processors

Vector libraries

The vector MASS library **libmassvp7.a** contains vector functions that have been tuned for the POWER7 architecture. The functions can be used in either 32-bit mode or 64-bit mode.

Functions supporting previous POWER® processors, either single-precision or double-precision, are included for POWER7 processors.

The following new functions are added, in both single-precision and double-precision function groups:

- exp2

- `exp2m1`
- `log21p`
- `log2`

For more information about the vector libraries, see Using the vector libraries in the *XL C/C++ Optimization and Programming Guide*.

SIMD libraries

The MASS SIMD library `libmass_simdp7.a` contains an accelerated set of frequently used math intrinsic functions that provide improved performance over the corresponding standard system library functions.

For more information about the SIMD libraries, see Using the SIMD library for POWER7 in the *XL C/C++ Optimization and Programming Guide*.

POWER7 hardware intrinsics

New hardware intrinsics are added to support the following POWER7 processor features:

- new POWER7 prefetch extensions and cache control
- new POWER7 hardware instructions

For more information, see “Built-in functions new for this release” on page 21.

New compiler options for POWER7 processors

New arch and tune compiler options

The `-qarch` compiler option specifies the processor architecture for which code is generated. The `-qtune` compiler option tunes instruction selection, scheduling, and other architecture-dependent performance enhancements to run best on a specific hardware architecture.

`-qarch=pwr7` produces object code containing instructions that will run on the POWER7 hardware platforms. With `-qtune=pwr7`, optimizations are tuned for the POWER7 hardware platforms.

For more information, see `-qarch` in the *XL C/C++ Compiler Reference* and `-qtune` in the *XL C/C++ Compiler Reference*.

C++0x

C++0x is the working draft of the new C++ programming language standard. Additional C++0x features are supported in this release of XL C/C++.

Note: C++0x is a new version of the C++ programming language standard. This is a draft standard and has not been officially adopted in its entirety. The implementation of C++0x is based on IBM's interpretation of the draft C++0x standard and is subject to change at any time without notice. IBM makes no attempt to maintain compatibility with earlier releases and therefore the C++0x language extension should not be relied on as a stable programming interface.

The following features are introduced in XL C/C++ V11.1:

- Auto type deduction
- C99 `long long`
- C99 preprocessor features adopted in C++0x
- `Decltype`

- Delegating constructors
- Explicit instantiation declarations
- Extended friend declarations
- Inline namespace definitions
- Static assertion
- Variadic templates

Auto type deduction

With the auto type deduction feature, you no longer need to specify a type while declaring a variable. This is because auto type deduction delegates the task of deducing the type of an auto variable to the compiler from the type of its initializer expression.


You can use the individual suboption **-qclanglvl=autotypededuction** or the group option **-qclanglvl=extended0x** to enable this feature.

For more information, see "The auto type specifier (C++0x)" in the *XL C/C++ Language Reference*.

C99 long long

The C++ compiler can use the C99 long long feature, which improves source compatibility between the C and C++ languages.

You can use the individual suboption **-qclanglvl=c99longlong** or the group option **-qclanglvl=extended0x** to enable the C99 long long feature.

 After this feature is enabled, if a decimal integer literal that does not have a suffix containing u or U cannot be represented by the long long int type, you can decide whether to use the unsigned long long int type to represent the literal or not by specifying the **-qclanglvl=[no]extendedintegersafe** option.

For more information, see "Integer literals" in the *XL C/C++ Language Reference*.

C99 preprocessor features adopted in C++0x

With several C99 preprocessor features adopted in C++0x, C and C++ compilers provide a more common preprocessor interface, which can ease porting C source files to the C++ compiler, eliminate semantic differences between the C and C++ preprocessors, and avoid preprocessor compatibility issues or diverging preprocessor behaviors.

You can use the individual suboption **-qclanglvl=c99preprocessor** or the group option **-qclanglvl=extended0x** to enable this feature.

For more information, see "C99 preprocessor features adopted in C++0x)" in the *XL C/C++ Language Reference*.

Decltype

With the decltype feature, you can get a type that is based on the resultant type of a possibly type-dependent expression.

You can use the individual suboption **-qlanglvl=decltype** or the group option **-qlanglvl=extended0x** to enable this feature.

For more information, see "The decltype(expression) type specifier (C++0x)" in the *XL C/C++ Language Reference*.

Delegating constructors

With the delegating constructors feature, you can concentrate common initializations in one constructor, which makes programs more readable and maintainable.

You can use the individual suboption **-qlanglvl=delegatingctors** or the group option **-qlanglvl=extended0x** to enable this feature.

For more information, see "Delegating constructors (C++0x)" in the *XL C/C++ Language Reference*.

Explicit instantiation declarations

With the explicit instantiation declarations feature, you can suppress the implicit instantiation of a template specialization or its members.

You can use the individual suboption **-qlanglvl=externtemplate** or the group options **-qlanglvl=extended** and **-qlanglvl=extended0x** to enable this feature.

For more information, see "Explicit instantiation (C++ only)" in the *XL C/C++ Language Reference*.

Extended friend declarations

The extended friend declarations feature relaxes the syntax rules governing friend declarations as follows:

- Template parameters, typedef names, and basic types can be declared as friends.
- The class-key in the context for friend declarations is no longer necessary in C++0x.

You can enable this feature with the individual suboption **-qlanglvl=extendedfriend** or the group option **-qlanglvl=extended0x**.

For more information, see "Friends (C++ only)" in the *XL C/C++ Language Reference*.

Inline namespace definitions

Inline namespace definitions are namespace definitions with an initial `inline` keyword. You can define or specialize the members of an inline namespace as if they belong to the enclosing namespace that contains the inline namespace.

You can enable this feature with the individual suboption **-qlanglvl=inlinenamespace** or the group option **-qlanglvl=extended0x**.

For more information, see "Inline namespace definitions (C++0x)" in the *XL C/C++ Language Reference*.

Static assertion

The static assertion feature provides you with the following benefits:

- Libraries can detect common usage errors at compile time.
- Implementations of the C++ Standard Library can detect and diagnose common usage errors, thus improving usability.

You can use a `static_assert` declaration to check important program invariants at compile time.

You can enable the static assertion feature with the individual suboption **-qlanglvl=static_assert** or the group option **-qlanglvl=extended0x**.

For more information, see "static_assert declaration (C++0x)" in the *XL C/C++ Language Reference*.

Variadic templates

With the variadic templates feature, you can define class or function templates that have any number (including zero) of parameters.

You can use the individual suboption **-qlanglvl=variadic[templates]** or the group option **-qlanglvl=extended0x** to enable this feature.

For more information, see "Variadic templates (C++0x)" in the *XL C/C++ Language Reference*.

Related information in the *XL C/C++ Compiler Reference*

 **-qlanglvl**

Performance and optimization

Additional features and enhancements assist with performance tuning and application optimization.

Enhancements to -qpdf

The use of the **-qpdf** option consists of two steps. First, compile your program with **-qpdf1** and run it with a typical set of data to generate the profiling data. Second, compile your program again with **-qpdf2** to optimize the program based on the profiling data.

In previous releases, if you modified the source file and compiled with the **-qpdf2** option, the compilation would stop with an error. As of IBM XL C/C++ for Linux, V11.1, you can use profiling data after you modify your source files. To do this, compile your application using the stale profiling data at the second stage of the PDF process.

Three new suboptions are added to the **-qpdf** option. These new suboptions allow more fine-grained control over performance improvements and extend **-qpdf** to support multiple pass profiling, cache miss profiling, block counter profiling, call counter profiling, and extended value profiling.

The three new **-qpdf** suboptions are:

level Supports multiple-pass profiling, block counter profiling, call counter

profiling, and extended value profiling. You can compile your application with `-qpdf1=level=0|1|2` to generate profiling data with different levels of optimization.

Note: Both `-qpdf1=level=0` and `-qpdf1=level=1` support single-pass profiling, whereas `-qpdf1=level=2` supports multiple-pass profiling.

exename

Generates the name of the PDF file you specify with the `-o` parameter.

defname

Reverts the PDF file to its default file name.

For detailed information about these suboptions, see `-qpdf1`, `-qpdf2` in the *XL C/C++ Compiler Reference*.

Reports about compiler optimizations

There are a number of enhancements to the listing reports to give you more information on how the compiler optimized your code. You can use this information to get further benefits from the compiler's optimization capabilities. For more details about these enhanced reports, see "New diagnostic reports" on page 15.

Performance-related compiler options and directives

The entries in the following table describe new or changed compiler options and directives.

Information presented here is a brief overview. For detailed information about these and other performance-related compiler options, see "Optimization and tuning options" in the *XL C/C++ Compiler Reference*.

Table 4. Performance-related compiler options and directives

-qinline=level=number	A new option is added to <code>-qinline</code> to provide guidance to the compiler about the relative value of inlining in relation to the default value of 5. <i>number</i> is a range of integer values between 0 and 10 that indicates the level of inlining you want to use. For details, see <code>-qinline</code> in the <i>XL C/C++ Compiler Reference</i> .
-qpdf	-qpdf provides suboptions to give you more control flexibility in controlling different PDF optimizations. For more information, see the <code>-qpdf1</code> , <code>-qpdf2</code> section in the <i>XL C/C++ Compiler Reference</i> .
-qprefetch	A new enhancement is added to -qprefetch for inserting prefetch instructions automatically where there are opportunities to improve code performance: <code>-qprefetch=assistthread</code> . -qprefetch inserts prefetch instructions automatically where there are opportunities to improve code performance. For details, see <code>-qprefetch</code> in the <i>XL C/C++ Compiler Reference</i> .

For additional information about performance tuning and program optimization, see "Optimizing your applications" in the *XL C/C++ Optimization and Programming Guide*.

New diagnostic reports

The new diagnostic reports can help you identify opportunities to improve the performance of your code.

Compiler reports in XML format

It is now possible to get information in XML format about the optimizations that the compiler was able to perform and also which optimization opportunities were missed. This information can be used to reduce programming effort for tuning applications, especially high-performance applications.

The information from the compiler is produced in XML 1.0 format. The report is defined by an XML schema and is easily consumable by tools that you can create to read and analyze the results. A stylesheet, `xlstyle.xsl`, is provided to render the report into a human readable format that can be read by anyone with a browser which supports XSLT.

In this release, the following four optimization categories are available in the report:

- Inlining
- Loop transformations
- Data reorganizations
- Profile-directed feedback information

The new **-qlistfmt** option and its associated suboptions can be used to generate the new XML 1.0 report.

For detailed information about this report and how to use it, see "Using reports to diagnose optimization opportunities" in the *XL C/C++ Optimization and Programming Guide*.

Enhancements to profiling reports

Additional sections of the listing report have been added to help you understand your programs. When using the **-qreport** option with the **-qpdf2** option, you can get the following additional sections added to your listing file in the section entitled PDF Report:

Loop iteration count

The most frequent loop iteration count and the average iteration count, for a given set of input data, is calculated for most loops in a program. This information is only available when the program is compiled at optimization level -O5.

Block and call count

This section of the report covers the call structure of the program and the respective execution count for each called function. It also includes block information for each function. For non-user defined functions, only execution count is given. The total block and call coverage, and a list of the user functions ordered by decreasing execution count are printed in the end of this report section. In addition, the block count information is printed at the beginning of each block of the pseudo-code in the listing files.

Cache miss

This section of the report is printed in a single table. It reports the number

of cache misses for certain functions, with additional information about the functions such as: cache level, cache miss ratio, line number, file name, and memory reference.

Note: You must use the **-qpdf1=level=2** option to get this report. You can also select the level of cache to profile using the **PDF_PM_EVENT** environment variable during run time.

For detailed information about the profile-directed feedback, see "Using profile-directed feedback" in the *XL C/C++ Optimization and Programming Guide*.

For additional information about the listing files, see "Compiler listings" in the *XL C/C++ Compiler Reference*

Report of data reorganization

The compiler can generate the following information in the listing files:

- Data reorganizations (a summary of how program variable data gets reorganized by the compiler)
- The location of data prefetch instructions inserted by the compiler

To generate data reorganization information, specify the optimization level **-qipa=level=2** or **-O5** together with **-qreport**. The data reorganization messages for program variable data are added to the data reorganization section of the listing file with the label DATA REORGANIZATION SECTION during the IPA link pass.

Reorganizations include:

- array splitting
- array transposing
- memory allocation merging
- array interleaving
- array coalescing

To generate information about data prefetch insertion locations, use the optimization level of **-qhot**, or any other option that implies **-qhot** together with **-qreport**. This information appears in the LOOP TRANSFORMATION SECTION of the listing file.

Additional loop analysis

A new suboption has been added to **-qhot** to add more aggressive loop analysis. **-qhot=level=2** together with **-qsmp** and **-qreport** add information about loop nests on which the aggressive loop analysis was performed to the LOOP TRANSFORMATION SECTION of the listing file. This information can also appear in the XML listing file created with the **-qlistfmt** option.

New and enhanced diagnostic options

The entries in the following table describe new or changed compiler options and directives that give you control over compiler listings.

Information presented here is a brief overview. For detailed information about these and other performance-related compiler options, see "Listings, messages and compiler information" in the *XL C/C++ Compiler Reference*.

Table 5. Listings-related compiler options and directives

Option/directive	Description
-qlistfmt	Generates a report in an XML 1.0 format containing information about optimizations performed by the compiler and missed optimization opportunities. The report contains information about inlining, loop transformations, data reorganization and profile-directed feedback.
-qreport	The listing now contains a PDF report section when used with -qpdf2 . Another new section in the listing files is a DATA REORGANIZATION section when used with -qipa=level=2 or -O5 .
-qskipsrc	Determines whether the source statements skipped by the compiler are shown in the SOURCE section of the listing file.

Utilization tracking and reporting tool

The utilization tracking and reporting feature is a lightweight and simple mechanism for tracking the compiler utilization within your organization. It is disabled by default. You can use this feature to detect whether your organization's use of the compiler exceeds your compiler license entitlements.

When utilization tracking is enabled, each invocation of the compiler is recorded in a compiler utilization file. You can run the utilization reporting tool to generate a report from one or more of these files to get a picture of the overall usage of the compiler within your organization. The **urt** command can be used to control how the report is generated. In particular, the report indicates the number of concurrent users using the compiler.

The utilization tracking and reporting feature is easy to set up and manage, and utilization tracking does not impact the usage or performance of the compiler.

For detailed information about the utilization tracking and reporting feature, see "Tracking and reporting compiler usage" in the *XL C/C++ Compiler Reference*.

New or changed compiler options and directives

New and changed compiler options and directives are described in this section.

Compiler options can be specified on the command line or through directives embedded in your application source files. See the *XL C/C++ Compiler Reference* for detailed descriptions and usage information for these and other compiler options.

Table 6. New or changed compiler options and directives

Option/directive	Description
-qarch	A new suboption has been added to -qarch , specifying -qarch=pwr7 produces object code that contains instructions that run on the POWER7 hardware platforms.
-qassert	-qassert is a new option for XL C/C++. It is used to provide information about the characteristics of the files that can help to fine-tune optimizations.

Table 6. New or changed compiler options and directives (continued)

Option/directive	Description
-qfunctrace	Traces the entry and exit points of functions in a compilation unit or only for a specific list of functions.
-qhot	<p>A new suboption has been added for -qhot. The -qhot compiler option is a powerful alternative to hand tuning that provides opportunities to optimize loops and array language.</p> <p>The -qhot=fastmath option enables the replacement of math routines with available math routines from the XLOPT library only if -qstrict=nolibrary is enabled. -qhot=nofastmath disables the replacement of math routines by the XLOPT library. -qhot=fastmath is enabled by default if -qhot is specified regardless of the hot level.</p>
-qinline	Attempts to inline functions instead of generating calls to those functions, for improved performance.
-qipa	You can generate relinkable objects while preserving IPA information by specifying -r -qipa=relink .

Table 6. New or changed compiler options and directives (continued)



Option/directive	Description
-qlanglvl	<p> C++0x New suboptions have been added to -qlanglvl:</p> <ul style="list-style-type: none"> • -qlanglvl=autotypededuction: Controls whether the auto type deduction feature is enabled. This feature can be used to delegate the task of type deduction of an auto variable to the compiler from the type of its initializer expression. • -qlanglvl=c99longlong: Controls whether the C99 long long feature is enabled. This feature improves source compatibility between the C and C++ languages. • -qlanglvl=c99preprocessor: Controls whether the C99 preprocessor features adopted in C++0x are enabled. This feature can be used to provide a more common preprocessor interface for C and C++ compilers. • -qlanglvl=decltype: Controls whether the decltype feature is enabled. This feature can be used to get a type that is based on the resultant type of a possibly type-dependent expression. • -qlanglvl=delegatingctors: Controls whether the delegating constructors feature is enabled. This feature can be used to concentrate common initializations in one constructor. • -qlanglvl=extendedfriend: Controls whether the extended friend declarations feature is enabled. This feature can be used to accept additional forms of non-function friend declarations. •  IBM -qlanglvl=extendedintegersafe: Controls whether or not unsigned long long int can be used as the type for decimal integer literals that do not have a suffix containing u or U and cannot be represented by the long long int type. This option takes effect only when the -qlanglvl=c99longlong option is specified. • -qlanglvl=externtemplate: Controls whether the explicit instantiation declarations feature is enabled. This feature can be used to suppress the implicit instantiation of a template specialization or its members. • -qlanglvl=inlinenamespace: Controls whether the inline namespace definitions feature is enabled. This feature can be used to define and specialize members of an inline namespace as if they were also members of the enclosing namespace. • -qlanglvl=static_assert: Controls whether the static assertions feature is enabled. This feature can be used to produce compile-time assertions for which a severe error message is issued on failure. • -qlanglvl=variadic[templates]: Controls whether the variadic templates feature is enabled. This feature can be used to define class or function templates that have any number (including zero) of parameters.

Table 6. New or changed compiler options and directives (continued)

Option/directive	Description
-qlibmpi	Tunes code based on the known behavior of the Message Passing Interface (MPI) functions.
-qlistfmt	Generates a report in an XML 1.0 format containing information about some optimizations performed by the compiler and some missed optimization opportunities for inlining, loop transformations, data reorganization and profile-directed feedback.
-qpdf1,-qpdf2	New options have been added to -qpdf1,-qpdf2 .
-qprefetch	A new suboption has been added to -qprefetch . When you work with applications that generate a high cache-miss rate, you can use -qprefetch=assistthread to exploit assist threads for data prefetching.
-qrestrict (C only)	You can use -qrestrict to indicate to the compiler that no other pointer can access the same memory that has been addressed by function parameter pointers.
-qsaveopt -qnosaveopt	The existing -qsaveopt option is enhanced to also include the user's configuration file name and the options specified in the configuration files.
-qstackprotect	Protects your applications against malicious code or programming errors that overwrite or corrupt the stack.
-qstaticlink	Controls how shared and nonshared runtime libraries are linked into an application.
-qstrict	A new suboption has been added to the -qstrict option to allow more control over optimizations and transformations that violate strict program semantics. -qstrict=vectorprecision disables vectorization in loops where it might produce different results in vectorized iterations than in nonvectorized ones.
-qtune	A new suboption has been added to -qtune . If you specify -qtune=pwr7 , optimizations are tuned for the POWER7 hardware platforms.

Table 7. Deprecated directives and options

Option/directive	Description
-Q	This option is deprecated and replaced with -qinline .
-qenablevmx	This option is deprecated and replaced with the -qsimd=auto option.
-qhot=simd nosimd	-qhot=simd nosimd are deprecated and might be removed in a future release. You can use -qsimd .
-qinfo=private	-qinfo=private is deprecated and replaced with -qreport .
-qinfo=reduction	-qinfo=reduction is deprecated and replaced with -qreport .
-qipa=inline noinline	-qipa=inline noinline are deprecated and might be removed in a future release. You can use -qinline .

Table 7. Deprecated directives and options (continued)

Option/directive	Description
-qipa=clonearch noclonearch	-qipa=clonearch noclonearch is no longer supported. You can use -qtune=balanced .
-qipa=clonearch noclonearch	-qipa=cloneproc nocloneproc is no longer supported. You can use -qtune=balanced .

Built-in functions new for this release

This section lists built-in functions that are new for this release.

For more information about built-in functions provided by XL C/C++, see Compiler built-in functions in the *XL C/C++ Compiler Reference*.

VSX built-in functions

Vector Scalar eXtension (VSX) is newly added for POWER7 processors.

For more information about VSX built-in functions, see Vector built-in functions.

POWER7 prefetch extensions and cache control

The POWER7 processor has cache control and stream prefetch extensions that support store stream prefetch and prefetch depth control. XL C/C++ provides the following new built-in functions to provide direct programmer access to these instructions:

- `__protected_stream_stride`
- `__transient_protected_stream_count_depth`
- `__unlimited_protected_stream_depth`
- `__transient_unlimited_protected_stream_depth`
- `__partial_dcbt`
- `__dcbtt`
- `__dcbtstt`
- `__dcbflp`

The compiler can insert the built-in functions automatically when it optimizes the code. You can disable automatic use of these instructions with **-qnoprefetch**.

For more information about the directives, see built-in functions in the *XL C/C++ Compiler Reference*.

POWER7 hardware built-in functions

New XL C/C++ built-in functions corresponding to each new POWER7 hardware instruction are added in this release. With these functions, you can directly manipulate specific hardware instructions in your code, which can improve the performance of your application.

- `__bpermd`
- `__cbcdtd`
- `__cdtbcd`
- `__load8r`

- `__store8r`
- `__divde`
- `__divdeu`
- `__cmpb`
- `__divwe`
- `__divweu`
- `__addg6s`

Conversion functions

These new functions convert between Declets and Binary Coded Decimal.

- `__cbcdtd`
- `__cdtbcd`

Comparison functions

This new function compares bytes.

- `__cmpb`

Decimal floating-point functions

This new function adds and generates sixes.

- `__addg6s`

Chapter 3. Enhancements added in previous versions

This section lists the enhancements added to the compiler in the previous releases.

Enhancements added in Version 10.1

This section describes features and enhancements added to the compiler in Version 10.1.

Operating system support

IBM XL C/C++ for Linux, V10.1 supports these operating systems supported by IBM Power Systems servers: Red Hat Enterprise Linux 5.2 (RHEL 5.2) and SUSE Linux Enterprise Server 10 SP 2 (SLES10 SP2)..

Predefined macros

There are 4 new macros:

`__ILP32` **`__ILP32__`**

Defined to 1 only when the compilation is for a target where long int, int and pointers all use 32 bits. Otherwise it is not defined.

`__LP64` **`__LP64__`**

Defined to 1 only when the compilation is for a target where long int and pointers both use 64 bits and int uses 32 bits. Otherwise it is not defined

The compiler no longer supports the `__C99_COMPLEX_HEADER__` macro.

For a complete list of the predefined macros for XL C/C++, see "Compiler predefined macros" in the *XL C/C++ Compiler Reference*.

C++0x

This release introduces support for a new version of the standard for the C++ programming language - specifically C++0x. This standard has not yet been officially adopted but we are beginning to support some of its features.

Note: C++0x is a new version of the C++ programming language standard. This is a draft standard and has not been officially adopted in its entirety. The implementation of C++0x is based on IBM's interpretation of the draft C++0x standard and is subject to change at any time without notice. IBM makes no attempt to maintain compatibility with earlier releases and therefore the C++0x language extension should not be relied on as a stable programming interface.

Specifically, in this release:

- we add a new language level
- we introduce new integer promotion rules for arithmetic conversions with long long data types
- the C++ preprocessor now supports C99 features

New language level - extended0x

The default **-qlanglvl** compiler option remains **extended** when invoking the C++ compiler.

A new suboption has been added to the **-qlanglvl** option in this release. **-qlanglvl=extended0x** is used to allow users to try out early implementations of any features of C++0x that are currently supported by XL C/C++.

C99 long long under C++

With this release of XL C/C++ for Linux, V11.1, compiler behavior changes when performing certain arithmetic operations with integral literal data types. Specifically, the integer promotion rules have changed.

Previously, in C++ (and as an extension to C89), when compiling with **-qlonglong**, an unsuffixed integral literal would be promoted to the first type in this list into which it fitted:

```
int
long int
unsigned long int
long long int
unsigned long long
```

Starting with this release and when compiling with **-qlanglvl=extended0x**, the compiler now promotes unsuffixed integral literal to the first type in this list into which it fits:

```
int
long int
long long int
unsigned long long
```

Note: Like our implementation of the C99 Standard in the C compiler, C++ will allow promotions from long long to unsigned long long if a value cannot fit into a long long type, but can fit in an unsigned long long. In this case, a message will be generated.

The macro `__C99_LLONG` has been added for compatibility with C99. This macro is defined to 1 with **-qlanglvl=extended0x** and is otherwise undefined.

For more information, see "Integral and floating-point promotions" in the *XL C/C++ Language Reference*.

Preprocessor changes

The following changes to the C++ preprocessor make it easier to port code from C to C++:

- Regular string literals can now be concatenated with wide-string literals.
- The `#line <integer>` preprocessor directive has a larger upper limit. It has been increased from 32767 to 2147483647 for C++.
- C++ now supports `_Pragma` operator.
- These macros now apply to C++ as well as C:

- `__C99_MACRO_WITH_VA_ARGS` (also available with `-qlanglvl=extended`)
- `__C99_MAX_LINE_NUMBER` (also available with `-qlanglvl=extended`)
- `__C99_PRAGMA_OPERATOR`
- `__C99_MIXED_STRING_CONCAT`

Note: Except as noted, these C++ preprocessor changes are only available when compiling with `-qlanglvl=extended0x`.

For additional information about the language standards supported by XL C/C++, see "Language levels and extensions" in the *XL C/C++ Language Reference*.

Other XL C/C++ language-related updates

Vector data types

Vector data types can now use some of the operators that can be used with base data types such as:

- unary operators
- binary operators
- relational operators

OpenMP 3.0

IBM XL C/C++ for Linux, V10.1, supports the OpenMP API Version 3.0 specification. The XL C/C++ implementation is based on IBM's interpretation of the OpenMP Application Program Interface Draft 3.0 Public Comment.

The main differences between Version 2.5 and Version 3.0 are:

- Addition of task level parallelization. The new OpenMP constructs `TASK` and `TASKWAIT` give users the ability to parallelize irregular algorithms, such as pointer chasing or recursive algorithms for which the existing OpenMP constructs were not adequate.
- for loops can now contain *var* values of unsigned int and pointer type as well as signed int.
- Stack size control. You can now control the size of the stack for threads created by the OMP runtime library using the new environment variable `OMP_STACKSIZE`.
- Users can give hints to the desired behavior of waiting threads using new environment variables `OMP_WAIT_POLICY` and `OMP_SET_POLICY`.
- Storage reuse. Some restrictions on the `PRIVATE` clause have been removed. A list item that appears in the reduction clause of a parallel construct can now also appear in a private clause on a work-sharing construct.
- Scheduling. A new `SCHEDULE` attribute, `auto`, allows the compiler and runtime system to control scheduling.
- Consecutive loop constructs with `STATIC` schedule can now use `nowait`.
- Nesting support - a `COLLAPSE` clause has been added to the `DO`, `FOR`, `PARALLEL FOR`, and `PARALLEL DO` directives to allow parallelization of perfect loop nests. This means that multiple loops in a nest can be parallelized.
- `THREADPRIVATE` directives can now apply to variables at class scope in addition to file and block scope.
- Parallelization of iterator loops of canonical form including those with random access iterators.

For more information, see:

- "Using OpenMP directives" in the *XL C/C++ Optimization and Programming Guide*
- www.openmp.org

Performance and optimization

Some features and enhancements can assist with performance tuning and optimization of your application.

Enhancements to **-qstrict**

Many suboptions have been added to the **-qstrict** option to allow more fine-grained control over optimizations and transformations that violate strict program semantics. In previous releases, the **-qstrict** option disabled all transformations that violate strict program semantics. This is still the behavior if you use **-qstrict** without suboptions. Likewise, in previous releases **-qnostrict** allowed transformations that could change program semantics. Because a higher level of optimizations might require relaxing strict program semantics, the addition of the suboptions relaxes selected rules to get specific benefits of faster code without turning off all semantic verifications.

You can use 16 new suboptions separately or use a suboption group. Here is a list of suboption groups:

all Disables all semantics-changing transformations, including those controlled by the other suboptions.

ieeefp Controls whether individual operations conform to IEEE 754 semantics.

order Controls whether individual operations can be reordered in a way that violate program language semantics.

precision Controls optimizations and transformations that can affect the precision of program results.

exceptions Controls optimizations and transformations that can affect the runtime exceptions generated by the program.

For detailed information about these suboptions, see "**-qstrict**" in the *XL C/C++ Compiler Reference*.

Performance-related compiler options and directives

The entries in the following table describe new or changed compiler options and directives.

Information presented here is a brief overview. For detailed information about these and other performance-related compiler options, see "Optimization and tuning options" in the *XL C/C++ Compiler Reference*.

Table 8. Performance-related compiler options and directives

Option/directive	Description
-qstrict	Many new suboptions have been added to give you more control over the relaxation of program semantic rules in order to gain some performance benefits.

Table 8. Performance-related compiler options and directives (continued)

Option/directive	Description
-qfloat	Some -qfloat suboptions are affected by the new suboptions for -qstrict .
-qreport	The listing now contains information about how many streams are created for each loop and which loops cannot be SIMD vectorized due to non-stride-one references. You can use this information to improve the performance of your applications.
-qsmp	When -qsmp=omp is in effect, the additional functionality of OpenMP API 3.0 is now available. For more information, see “OpenMP 3.0” on page 25.

For additional information about performance tuning and program optimization, see “Optimizing your applications” in the *XL C/C++ Optimization and Programming Guide*.

New or changed compiler options and directives

Compiler options can be specified on the command line or through directives embedded in your application source files. See the *XL C/C++ Compiler Reference* for detailed descriptions and usage information for these and other compiler options.

Table 9. New or changed compiler options and directives

Option/directive	Description
-qstrict	Many suboptions have been added to the -qstrict option to allow more control over optimizations and transformations that violate strict program semantics. See “Performance and optimization” on page 13 for more information.
-qshowmacros	When used in conjunction with the -E option, the -qshowmacros option replaces preprocessed output with macro definitions. There are suboptions provided to control the emissions of predefined and user-defined macros more precisely.
-qreport	When used together with compiler options that enable automatic parallelization or vectorization, the -qreport option now reports the number of streams in a loop and produces information when loops cannot be SIMD vectorized due to non-stride-one references.
-qsmp	When -qsmp=omp is in effect, the additional functionality of OpenMP API 3.0 is now available. For more information, see “OpenMP 3.0” on page 25.
-qtimestamps	This option can be used to remove timestamps from generated binaries.
-qtls	The thread local storage support has been enhanced to include <code>__attribute__((tls-model("string")))</code> where <i>string</i> is one of <code>local-exec</code> , <code>initial-exec</code> , <code>local-dynamic</code> , or <code>global-dynamic</code> .
-qinfo	The suboptions <code>als</code> and <code>noals</code> have been added to the qinfo option to report (or not report) possible violations of the ANSI aliasing rule.

Enhancements added in Version 9.0

This section describes features and enhancements added to the compiler in Version 9.0.

C/C++ language-related updates

The default language level for C compilations changed, and new behavior was introduced when doing arithmetic conversions with long long data types.

Default language level changed for C - extc99

The default **-qlanglvl** compiler option setting is **extc99** when invoking the C compiler with the **xl** invocation. This change allows you to use C99 features and headers without having to explicitly specify the **extc99** suboption.

You might encounter issues with the following when compiling with the new default **-qlanglvl=extc99** setting:

- Pointers can be qualified with **restrict** in C99, so **restrict** can not be used as an identifier.
- C99 treatment of long long data differs from the way long long data is handled in C89.
- C99 header files define new macros: **LLONG_MAX** in **limits.h**, and **va_copy** in **stdarg.h**.
- The value of macro **__STDC_VERSION__** changes from 199409 to 19990.

To revert to previous **xl** behavior, specify **-qlanglvl=extc89** when invoking the compiler.

Arithmetic conversions with long long data types

With XL C/C++ Version 9.0, compiler behavior changes when performing certain arithmetic operations with long long data types.

Assume an arithmetic expression where:

- One operand has type long long int or long long, and,
- The other operand has type unsigned long int, but its value cannot be represented in a long long int or long long.

Previous releases of XL C/C++ converted both operands to type long long.

The compiler now converts both operands into type unsigned long long int or unsigned long long. This new behavior is consistent with GCC compiler behavior.

For more information, see "Integral and floating-point promotions" in the *XL C/C++ Language Reference*.

Architecture and processor support

The **-qarch** and **-qtune** compiler options control the code generated by the compiler. These compiler options adjust the instructions, scheduling, and other optimizations to give the best performance for a specified target processor or range of processors.

New default setting for -qtune

The new default **-qtune** setting is:

- **-qtune=balanced**

The **-qtune=balanced** suboption is new for this release, and becomes the default **-qtune** setting when certain **-qarch** settings are specified. Using **-qtune=balanced** instructs the compiler to tune generated code for optimal performance across a range of recent processor architectures, including POWER6.

New support for POWER6 processors

XL C/C++ Version 9.0 expanded the list of **-qarch** and **-qtune** suboptions to support the newly-available POWER6 processors.

The following **-qarch** and **-qtune** options are now available:

- **-qarch=pwr6**
- **-qarch=pwr6e**
- **-qtune=pwr6**

Performance and optimization

Many enhancements were made to assist with performance tuning and program optimization.

Performance-related compiler options and directives

The entries in the following table describes new or changed compiler options and directives.

Information presented here is just a brief overview. For more information about these and other performance-related compiler options, refer to "Optimization and tuning options" in the *XL C/C++ Compiler Reference*.

Table 10. Performance-related compiler options and directives

Option/directive	Description
-qalias= global noglobal	These new -qalias suboptions enable or disable the application of language-specific aliasing rules across compilation units during link time optimization.
-qalias= restrict norestrict	These new -qalias suboptions enable or disable optimization for restrict qualified pointers. Specifying -qalias=restrict will usually improve performance for code that uses restrict qualified pointers. You can use -qalias=norestrict to preserve compatibility with code compiled with versions of the compiler previous to V9.0.
-qnofdpr -qfdpr	Specifying the -qfdpr option instructs the compiler to store optimization information in the created object file. This information is used by the Feedback Directed Program Restructuring (FDPR) performance-tuning utility.

Table 10. Performance-related compiler options and directives (continued)

Option/directive	Description
-qfloat= fenv nofenv	These new -qfloat suboptions inform the compiler if code has a dependency on the floating-point hardware environment, such as explicitly reading or writing the floating-point status and control register. Specifying -qfloat=nofenv indicates that there is no dependency on the hardware environment, allowing the compiler to perform aggressive optimizations.
-qfloat= gcclongdouble nogcclongdouble	These new -qfloat suboptions have effect only when the -qldbl128 option is in effect. They instruct the compiler to use either GCC-supplied or IBM-supplied library functions for 128-bit long double operations.
-qfloat= hscmplx nohscmplx	Specifying -qfloat=hscmplx improves optimization of operations involving complex division and complex absolute values.
-qfloat= rngchk norngchk	Specifying -qfloat=rngchk enables range checking on input arguments for software divide and inlined sqrt operations. Specifying -qfloat=norngchk instructs the compiler to skip range checking, allowing for better performance in certain circumstances. Specifying the -qnostrict compiler option sets -qfloat=norngchk .
-qipa=threads= [auto noauto number]	This new -qipa suboption lets you specify how many threads the compiler will assign to code generation during the second IPA pass.
-qnoldbl128 -qldbl128	Specifying -qldbl128 increases the size of long double types from 64 bits to 128 bits.
-qpdf	The -qpdf option can now be used to provide profile-directed feedback on specific objects. See "Object level profile-directed feedback" in the <i>XL C/C++ Optimization and Programming Guide</i> for more information.
-qsmp= threshold=n	When -qsmp=auto is in effect, this new suboption lets you specify the amount of work required in a loop before the compiler will consider it for automatic parallelization.
#pragma expected_value(param, value)	Use the #pragma expected_value directive to specify a value that a parameter passed in a function call is most likely to take at run time. The compiler can use this information to perform certain optimizations, such as function cloning and inlining.

Built-in functions in Version 9.0

Some built-in functions were added in Version 9.0.

For more information on built-in functions provided by XL C/C++, see "Compiler built-in functions" in the *XL C/C++ Compiler Reference*.

Conversion functions

These new functions convert long double data types from IBM style to GCC style.

- `long double __ibm2gccldbl (long double);`
- `_Complex long double __ibm2gccldbl_cmplx (_Complex long double);`

PowerPC cache control

The PowerPC architecture specifies the **dcbst** and **dcbf** cache copy instructions. The following new built-in functions provide direct programmer access to these instructions.

- `void __dcbst(const void* addr);` /* Data Cache Block Store */
- `void __dcbf(const void* addr);` /* Data Cache Block Flush */

POWER6 prefetch extensions and cache control

The POWER6 processor has cache control and stream prefetch extensions with support for store stream prefetch and prefetch depth control. XL C/C++ provides the following new built-in functions to provide direct programmer access to these instructions.

- `void __dcbfl(const void* addr);` /* pwr6 - Data Cache Block Flush from L1 data cache only */
- `void __protected_unlimited_stream_set(unsigned int direction, const void* addr, unsigned int ID);` /* Supported by pwr5 and pwr6 */
- `void __protected_unlimited_store_stream_set(unsigned int direction, const void* addr, unsigned int ID);` /* Supported by pwr6 */
- `void __protected_store_stream_set(unsigned int direction, const void* addr, unsigned int ID);` /* Supported by pwr6 */
- `void __protected_stream_count_depth(unsigned int unit_cnt, unsigned int prefetch_depth, unsigned int ID);` /* Supported by pwr6 */

Other new or changed compiler options

Compiler options can be specified on the command line or through directives embedded in your application source files. See the *XL C/C++ Compiler Reference* for detailed descriptions and usage information for these and other compiler options.

Table 11. Other new or changed compiler options

Option/directive	Description
-C!	Specifying the -C! compiler option removes comments from preprocessed output.
-qcommon -qnocommon	With -qcommon in effect, uninitialized global variables are allocated in the common section of the object file. When -qnocommon is in effect, uninitialized global variables are initialized to zero and allocated in the data section of the object file.
-qoptdebug -qnooptdebug	When used with optimization levels of -O3 or higher, the new -qoptdebug option instructs the compiler to produce optimized pseudocode that can be read by a symbolic debugger.
-qpack_semantic= ibm gnu	The -qpack_semantic option is a portability option that instructs the compiler to use either IBM or GCC syntax and semantics for the #pragma pack directive.

Table 11. Other new or changed compiler options (continued)

Option/directive	Description
-qreport	When used together with compiler options that enable automatic parallelization or vectorization, the -qreport option produces a pseudo-code listing showing how program loops are parallelized and vectorized. The report also provides diagnostic information if the compiler is not able to parallelize or vectorize a given loop.
-qsaveopt -qnosaveopt	In previous releases, the -qsaveopt option stored the command line options used to compile a file into the resulting object file. In Version 9.0, the information stored in the object file expanded to also include version and level information for each compiler component invoked during compilation.
-qsmp=stackcheck	This new -qsmp suboption instructs the compiler to check for stack overflow by slave threads at run time, and issue a warning if the remaining stack size is less than the number of bytes specified by the stackcheck option of the XLSMPOPTS environment variable.
-qtemplatedepth=number	-qtemplatedepth specifies the maximum number of recursively-instantiated template specializations that the compiler will process.
-qversion=verbose	The -qversion option adds a new verbose suboption. Specifying -qversion=verbose instructs the compiler to display the version and level information for each compiler component invoked during compilation.

Chapter 4. Setting up and customizing XL C/C++

For complete prerequisite and installation information for XL C/C++, refer to "Before installing" in the *XL C/C++ Installation Guide*.

Using custom compiler configuration files

You can customize compiler settings and options by modifying the default configuration file or by creating your own.

You have the following options to customize compiler settings:

- The XL C/C++ compiler installation process creates a default compiler configuration file. You can directly modify this configuration file to add default options for specific needs. However, if you later apply updates to the compiler, you must reapply all of your modifications to the newly installed configuration file.
- You can create your own custom configuration file that either overrides or complements the default configuration file. The compiler can recognize and resolve compiler settings you specify in your custom configuration files together with compiler settings specified in the default configuration file. Compiler updates that might later affect settings in the default configuration file does not affect the settings in your custom configuration files.

For more information, see "Using custom compiler configuration files" in the *XL C/C++ Compiler Reference*.

Configuring compiler utilization tracking and reporting

In addition to the compiler configuration file, there is a separate configuration file for the utilization tracking and reporting feature. Utilization tracking is disabled by default, but you can enable it by modifying an entry in this configuration file. Various other aspects of utilization tracking can also be configured using this file.

Although the compiler configuration file is separate from the utilization tracking configuration file, it contains an entry that specifies the location of the utilization tracking configuration file so that the compiler can find this file.

For more information about how to configure the utilization tracking and reporting feature, see Tracking and reporting compiler usage in the *XL C/C++ Compiler Reference*.

Chapter 5. Developing applications with XL C/C++

C/C++ application development consists of repeating cycles of editing, compiling and linking (by default a single step combined with compiling), and running.

Notes:

1. Before you can use the compiler, you must first ensure that XL C/C++ is properly installed and configured. For more information see the *XL C/C++ Installation Guide*.
2. To learn about writing C/C++ programs, refer to the *XL C/C++ Language Reference*.

The compiler phases

A typical compiler invocation executes some or all of these activities in sequence. For link time optimizations, some activities will be executed more than once during a compilation. As each program runs, the results are sent to the next step in the sequence.

1. Preprocessing of source files
2. Compilation, which may consist of the following phases, depending on what compiler options are specified:
 - a. Front-end parsing and semantic analysis
 - b. High-level optimization
 - c. Low-level optimization
 - d. Register allocation
 - e. Final assembly
3. Assemble the assembly (.s) files, and the unpreprocessed assembler (.S) files after they are preprocessed
4. Object linking to create an executable application

To see the compiler step through these phases, specify the **-v** compiler option when you compile your application. To see the amount of time the compiler spends in each phase, specify **-qphsinfo**.

Editing C/C++ source files

To create C/C++ source programs, you can use any text editor available to your system.

Source programs must be saved using a recognized file name suffix. See the “XL C/C++ input and output files” on page 38 for a list of suffixes recognized by XL C/C++.

For a C or C++ source program to be a valid program, it must conform to the language definitions specified in the *XL C/C++ Language Reference*.

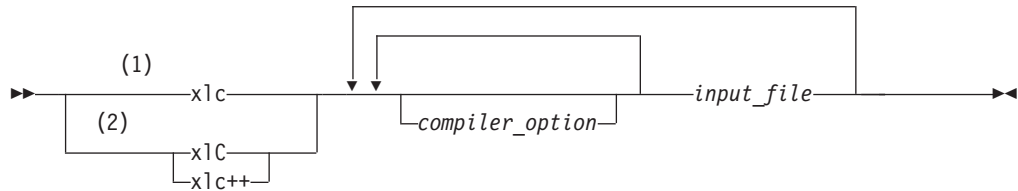
Compiling with XL C/C++

XL C/C++ is a command-line compiler. Invocation commands and options can be selected according to the needs of a particular C/C++ application.

Invoking the compiler

The compiler invocation commands perform all necessary steps to compile C or C++ source files, assemble any `.s` and `.S` files, and link the object files and libraries into an executable program.

To compile a source program, use the basic invocation syntax shown below:



Notes:

- 1 Basic invocation to compile C source code
- 2 Basic invocations to compile C++ source code

For most applications, you should compile with `xlc`, `xlc++`, or a thread safe counterpart. You can use `xlc++` to compile either C or C++ program source, but compiling C++ files with `xlc` may result in link or run time errors because libraries required for C++ code are not specified when the linker is called by the C compiler.

Additional invocation commands are available to meet specialized compilation needs, primarily to provide explicit compilation support for different levels and extensions of the C or C++ language. See "Invoking the compiler" in the *XL C/C++ Compiler Reference* for more information about compiler invocation commands available to you, including special invocations intended to assist developers migrating from a GNU compilation environment to XL C/C++.

Compiling parallelized XL C/C++ applications

XL C/C++ provides thread-safe compiler invocation commands that you can use when compiling parallelized applications for use in multiprocessor environments.

These invocations are similar to their corresponding base compiler invocations, except that they link and bind compiled objects to thread-safe components and libraries. The generic XL C/C++ thread-safe compiler invocations include:

- `xlc_r`
- `xlc++_r`
- `xlc_r`

XL C/C++ provides additional thread-safe invocations to meet specific compilation requirements. See "Invoking the compiler" in the *XL C/C++ Compiler Reference* for more information.

Note: Using any of these commands alone does not imply parallelization. For the compiler to recognize OpenMP directives and activate parallelization, you must also specify `-qsmp` compiler option. In turn, you should specify the `-qsmp` option only in conjunction with one of these thread-safe invocation commands. When you specify `-qsmp`, the driver links in the libraries specified on the `smp libraries` line in the active stanza of the configuration file.

For more information on parallelized applications see "Parallelizing your programs" in the *XL C/C++ Optimization and Programming Guide*.

Specifying compiler options

Compiler options perform a variety of functions, such as setting compiler characteristics, describing the object code to be produced, controlling the diagnostic messages emitted, and performing some preprocessor functions.

You can specify compiler options:

- On the command-line with command-line compiler options
- In your source code using directive statements
- In a makefile
- In the stanzas found in a compiler configuration file
- Or by using any combination of these techniques

It is possible for option conflicts and incompatibilities to occur when multiple compiler options are specified. To resolve these conflicts in a consistent fashion, the compiler usually applies the following general priority sequence to most options:

1. Directive statements in your source file *override* command-line settings
2. Command-line compiler option settings *override* configuration file settings
3. Configuration file settings *override* default settings

Generally, if the same compiler option is specified more than once on a command-line when invoking the compiler, the last option specified prevails.

Note: Some compiler options do not follow the priority sequence described above.

For example, the **-I** compiler option is a special case. The compiler searches any directories specified with **-I** in the `vac.cfg` file before it searches the directories specified with **-I** on the command-line. The option is cumulative rather than preemptive.

See the *XL C/C++ Compiler Reference* for more information about compiler options and their usage.

Other options with cumulative behavior are **-R** and **-l** (lowercase L).

You can also pass compiler options to the linker, assembler, and preprocessor. See "Compiler options reference" in the *XL C/C++ Compiler Reference* for more information about compiler options and how to specify them.

Reusing GNU C/C++ compiler options with `gxlc` and `gxlc++`

XL C/C++ includes various features to help you transition from GNU C/C++ compilers to XL C/C++ including `gxlc` and `gxlc++` commands.

Each of the `gxlc` and `gxlc++` utilities accepts GNU C or C++ compiler options and translates them into comparable XL C/C++ options. Both utilities use the XL C/C++ options to create an `xlc` or `xlcpp` invocation command, which is then used to invoke the compiler. These utilities are provided to help you reuse makefiles created for applications previously developed with GNU C/C++. However, to fully exploit the capabilities of XL C/C++, you should use the XL C/C++ invocation commands and their associated options.

The actions of `gxc` and `gxc++` are controlled by the configuration file `gxc.cfg`. The GNU C/C++ options that have an XL C/C++ counterpart are shown in this file. Not every GNU option has a corresponding XL C/C++ option. `gxc` and `gxc++` return warnings for input options that were not translated.

The `gxc` and `gxc++` option mappings are modifiable. For information on using the `gxc` or `gxc++` configuration file, see "Reusing GNU C/C++ compiler options with `gxc` and `gxc++`" in the *XL C/C++ Compiler Reference*.

XL C/C++ input and output files

These file types are recognized by XL C/C++.

For detailed information about these and additional file types used by the compiler, see "Types of input files" in the *XL C/C++ Compiler Reference* and "Types of output files" in the *XL C/C++ Compiler Reference*.

Table 12. Input file types

Filename extension	Description
.c	C source files
.C, .cc, .cp, .cpp, .cxx, .c++	C++ source files
.i	Preprocessed source files
.o	Object files
.s	Assembler files
.S	Unpreprocessed assembler files
.so	Shared object or library files

Table 13. Output file types

Filename extension	Description
a.out	Default name for executable file created by the compiler
.d	Make dependency file
.i	Preprocessed source files
.lst	Listing files
.o	Object files
.s	Assembler files
.so	Shared object or library files

Linking your compiled applications with XL C/C++

By default, you do not need to do anything special to link an XL C/C++ program. The compiler invocation commands automatically call the linker to produce an executable output file.

For example, running the following command:

```
xlc++ file1.C file2.o file3.C
```

compiles `file1.C` and `file3.C` to produce the object files `file1.o` and `file3.o`, then all object files (including `file2.o`) are submitted to the linker to produce one executable.

Compiling and linking in separate steps

To produce object files that can be linked later, use the `-c` option.

```
xlc++ -c file1.C          # Produce one object file (file1.o)
xlc++ -c file2.C file3.C  # Or multiple object files (file1.o, file3.o)
xlc++ file1.o file2.o file3.o # Link object files with default libraries
```

For more information about compiling and linking your programs, see:

- "Linking" in the *XL C/C++ Compiler Reference*
- "Constructing a library" in the *XL C/C++ Optimization and Programming Guide*

Dynamic and static linking

XL C/C++ allows your programs to take advantage of the operating system facilities for both dynamic and static linking.

Dynamic linking means that the code for some external routines is located and loaded when the program is first run. When you compile a program that uses shared libraries, the shared libraries are dynamically linked to your program by default. Dynamically linked programs take up less disk space and less virtual memory if more than one program uses the routines in the shared libraries. During linking, they do not require any special precautions to avoid naming conflicts with library routines. They may perform better than statically linked programs if several programs use the same shared routines at the same time. They also allow you to upgrade the routines in the shared libraries without relinking.

Because this form of linking is the default, you need no additional options to turn it on.

Static linking means that the code for all routines called by your program becomes part of the executable file.

Statically linked programs can be moved to run on systems without the XL C/C++ runtime libraries. They may perform better than dynamically linked programs if they make many calls to library routines or call many small routines. They do require some precautions in choosing names for data objects and routines in the program if you want to avoid naming conflicts with library routines. They also may not work if you compile them on one level of the operating system and run them on a different level of the operating system.

Running your compiled application

The default file name for the program executable file produced by the XL C/C++ compiler is **a.out**. You can select a different name with the `-o` compiler option.

To run a program, enter the name of the program executable file together with any run time arguments on the command line.

You should avoid giving your program executable file the same name as system or shell commands, such as `test` or `cp`, as you could accidentally execute the wrong command. If you do decide to name your program executable file with the same name as a system or shell command, you should execute your program by specifying the path name to the directory in which your executable file resides, such as `./test`.

Canceling execution

To suspend a running program, press the **Ctrl+Z** key while the program is in the foreground. Use the **fg** command to resume running.

To cancel a running program, press the **Ctrl+C** key while the program is in the foreground.

Setting runtime options

You can use environment variable settings to control certain runtime options and behaviors of applications created with the XL C/C++ compiler. Other environment variables do not control actual runtime behavior, but can have an impact on how your applications will run.

For more information on environment variables and how they can affect your applications at run time, see the *XL C/C++ Installation Guide*.

Running compiled applications on other systems

If you want to run an application developed with the XL C/C++ compiler on another system that does not have the compiler installed, you will need to install a runtime environment on that system.

You can obtain the latest XL C/C++ Runtime Environment PTF images, together with licensing and usage information, from the XL C/C++ Support page at:

www.ibm.com/software/awdtools/xlcpp/support

XL C/C++ compiler diagnostic aids

XL C/C++ issues diagnostic messages when it encounters problems compiling your application. You can use these messages and other information provided in compiler output listings to help identify and correct such problems.

For more information about listing, diagnostics, and related compiler options that can help you resolve problems with your application, see the following topics in the *XL C/C++ Compiler Reference*:

- "Compiler messages and listings"
- "Error checking and debugging options"
- "Listings, messages, and compiler information options"

Debugging compiled applications

You can use a symbolic debugger to debug applications compiled with XL C/C++.

Specifying the **-g** or **-qlinedebug** compiler options at compile time instructs the XL C/C++ compiler to include debugging information in compiled output. For more information debugging options, see "Error checking and debugging" in the *XL C/C++ Compiler Reference*.

You can then use **gdb** or any other symbolic debugger to step through and inspect the behavior of your compiled application.

Optimized applications pose special challenges when debugging. When debugging highly optimized applications, you should consider using the **-qoptdebug** compiler

option. For more information about optimizing your code, see "Optimizing your applications" in the *XL C/C++ Optimization and Programming Guide*.

Determining what level of XL C/C++ is installed

When contacting software support for assistance, you will need to know what level of XL C/C++ is installed on a particular machine.

To display the version and release level of the compiler you have installed on your system, invoke the compiler with the **-qversion** compiler option.

For example, to obtain detailed version information, enter the following at the command line:

```
xlc++ -qversion=verbose
```

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, 2010. All rights reserved.

Trademarks and service marks

IBM, the IBM logo, and `ibm.com` are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at “Copyright and trademark information” at <http://www.ibm.com/legal/copytrade.shtml>.

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.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft and Windows are trademarks of Microsoft Corporation in the United States, other countries, or both.

Cell Broadband Engine is a trademark of Sony Computer Entertainment, Inc. in the United States, other countries, or both and is used under license therefrom.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Other company, product, or service names may be trademarks or service marks of others.

Index

Special characters

- .a files 38
- .c and .C files 38
- .i files 38
- .lst files 38
- .mod files 38
- .o files 38
- .s files 38
- .S files 38
- .so files 38

Numerics

- 64-bit environment 6

A

- archive files 38
- assembler
 - source (.s) files 38
 - source (.S) files 38

B

- basic example, described ix
- built-in functions 21, 30

C

- C++0x
 - auto type deduction 10
 - C99 long long 10
 - C99 preprocessor features adopted in C++0x 10
 - decltype 10
 - delegating constructors 10
 - explicit instantiation declarations 10
 - extended friend declarations 10
 - inline namespace definitions 10
 - static assertion 10
 - variadic templates 10
- code optimization 6
- compilation
 - sequence of activities 35
- compiler
 - controlling behavior of 37
 - invoking 36
 - running 36
- compiler directives
 - new or changed 17
- compiler options
 - conflicts and incompatibilities 37
 - new or changed 17
 - specification methods 37
- compiling
 - SMP programs 36
- customization
 - for compatibility with GNU 3

D

- debugger support 40
 - output listings 40
 - symbolic 8
- debugging 40
- debugging compiled applications 40
- debugging information, generating 40
- dynamic linking 39

E

- editing source files 35
- executable files 38
- executing a program 39
- executing the linker 39

F

- files
 - editing source 35
 - input 38
 - output 38

G

- GNU
 - compatibility with 3

I

- input files 38
- invocation commands 36
- invoking a program 39
- invoking the compiler 36

L

- language standards 2
- language support 2
- level of XL C/C++, determining 41
- libraries 38
- linking
 - dynamic 39
 - static 39
- linking process 38
- listings 38

M

- migration
 - source code 37
- mod files 38
- multiprocessor systems 7, 25

O

- object files 38
 - creating 39
 - linking 39
- OMP directives 25
- OpenMP 7
- optimization
 - programs 6
- output files 38

P

- parallelization 7, 25
- performance
 - optimizing transformations 6
- problem determination 40
- programs
 - running 39

R

- running the compiler 36
- runtime
 - libraries 38
- runtime environment 40
- runtime options 40

S

- shared memory parallelization 7, 25
- shared object files 38
- SMP
 - programs, compiling 36
- SMP programs 7
- source files 38
- source-level debugging support 8
- static linking 39
- symbolic debugger support 8

T

- tools 5
 - cleanpdf utility 5
 - configuration file utility 5
 - gxlc and gxlc++ utilities 5
 - mergepdf utility 5
 - new install configuration utility 5
 - new_install utility 5
 - resetpdf utility 5
 - showpdf utility 5
 - xlc_configure 5

U

- utilities 5
 - cleanpdf 5
 - gxlc and gxlc++ 5
 - mergepdf 5

utilities (*continued*)
 new_install 5
 resetpdf 5
 showpdf 5
 xlc_configure 5

V

vac.cfg file 37

X

xlc_configure 5



Program Number: 5724-X14

Printed in USA

GI11-7913-00

