

IBM XL C for AIX, V12.1



Getting Started with XL C

Version 12.1

IBM XL C for AIX, V12.1



Getting Started with XL C

Version 12.1

Note

Before using this information and the product it supports, read the information in "Notices" on page 39.

First edition

This edition applies to IBM XL C for AIX, V12.1 (Program 5765-J01; 5725-C71) 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, 2012.

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	viii
IBM XL C information	viii
Standards and specifications	ix
Other IBM information	x
Other information	x
Technical support.	x
How to send your comments	xi

Chapter 1. Introducing XL C	1
Commonality with other IBM compilers	1
Operating system support	1
A highly configurable compiler	2
Language standard compliance	3
Compatibility with GNU	3
Source-code migration and conformance checking	3
Libraries.	3
Tools, utilities, and commands	4
Program optimization	6
64-bit object capability	6
Shared memory parallelization	7
Diagnostic listings	8
Symbolic debugger support	8

Chapter 2. What's new for IBM XL C for AIX, V12.1	9
C1X features	9
OpenMP 3.1	10
Performance and optimization	10
New diagnostic reports	11
New built-in functions.	12
New or changed compiler options and pragma directives	12

Chapter 3. Migrating from earlier versions	15
Enhancements added in Version 11.1	15
Operating system support	15
Support for POWER7 processors	16
Performance and optimization	17

New diagnostic reports	19
Utilization tracking and reporting tool	21
New or changed compiler options and directives	21
Built-in functions new for this release	24
Compatibility of redistributable library libxlopt.a	25
Enhancements added in Version 10.1	25
Operating system support	25
XL C language-related updates	25
OpenMP 3.0	26
Performance and optimization	26
New or changed compiler options and pragma directives	28

Chapter 4. Setting up and customizing XL C	29
Using custom compiler configuration files	29
Configuring compiler utilization tracking and reporting	29

Chapter 5. Developing applications with XL C	31
The compiler phases	31
Editing C source files	31
Compiling with XL C	31
Invoking the compiler	32
Compiling parallelized XL C applications	32
Specifying compiler options	32
XL C input and output files	33
Linking your compiled applications with XL C	34
Relinking an existing executable file	34
Dynamic and static linking	35
Running your compiled application	35
XL C compiler diagnostic aids	36
Debugging compiled applications	36
Determining what level of XL C is installed	37

Notices	39
Trademarks and service marks	41

Index	43
------------------------	-----------

About this document

This document contains overview and basic usage information for the IBM® XL C for AIX®, V12.1 compiler.

Who should read this document

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

How to use this document

Throughout this document, the `xlc` compiler invocation is 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 applications using the XL C compiler, it does not include the following topics:

- Compiler installation: see the *XL C Installation Guide* for information on installing XL C.
- Compiler options: see the *XL C Compiler Reference* for detailed information on the syntax and usage of compiler options.
- The C programming language: see the *XL C Language Reference* for information on the syntax, semantics, and IBM implementation of the C programming language.
- Programming topics: see the *XL C Optimization and Programming Guide* for detailed information on developing applications with XL 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 for AIX, V12.1 information.

Table 1. *Typographical conventions*

Typeface	Indicates	Example
bold	Lowercase commands, executable names, compiler options, and directives.	The compiler provides basic invocation commands, <code>xlc</code> , along with several other compiler invocation commands to support various C language levels and compilation environments.

Table 1. *Typographical conventions (continued)*

Typeface	Indicates	Example
<i>italics</i>	Parameters or variables whose actual names or values are to be supplied by the user. Italics are also used to introduce new terms.	Make sure that you update the <i>size</i> parameter if you return more than the <i>size</i> requested.
<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)

In descriptions of language elements where a feature is exclusive to the C1X standard, or where a feature is an IBM extension of the C standard, this information uses icons to delineate segments of text as follows:

Table 2. *Qualifying elements*

Qualifier/Icon	Meaning
IBM extension begins   IBM extension ends	The text describes a feature that is an IBM extension to the standard language specifications.
C1X, or C1X begins   C1X ends	The text describes a feature that is introduced into standard C as part of C1X.

Syntax diagrams

Throughout this information, diagrams illustrate XL 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):



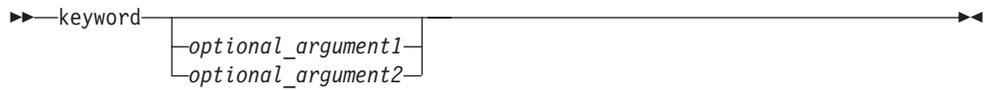
- Optional items are shown below the main path:



- 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:



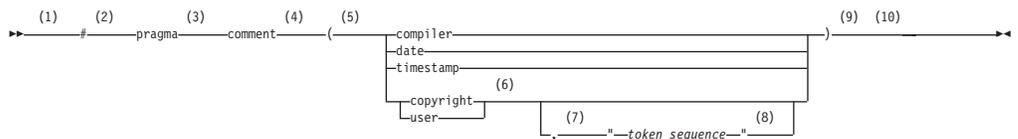
- The item that is the default is shown above the main path.



- 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:

IBM XL C information

XL 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 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 for AIX, V12.1 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 for AIX, V12.1 Installation Guide*.

The information center is viewable on the web at <http://publib.boulder.ibm.com/infocenter/comphelp/v121v141/index.jsp>.

- PDF documents

PDF documents are located by default in the `/usr/vac/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/xlc/aix/library/>.

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

Table 3. XL C PDF files

Document title	PDF file name	Description
<i>IBM XL C for AIX, V12.1 Installation Guide, GC14-7324-00</i>	install.pdf	Contains information for installing XL C and configuring your environment for basic compilation and program execution.
<i>Getting Started with IBM XL C for AIX, V12.1, SC14-7323-00</i>	getstart.pdf	Contains an introduction to the XL C product, with information on setting up and configuring your environment, compiling and linking programs, and troubleshooting compilation errors.
<i>IBM XL C for AIX, V12.1 Compiler Reference, SC14-7325-00</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 for AIX, V12.1 Language Reference, SC14-7326-00</i>	langref.pdf	Contains information about the C programming languages, as supported by IBM, including language extensions for portability and conformance to nonproprietary standards.
<i>IBM XL C for AIX, V12.1 Optimization and Programming Guide, SC14-7327-00</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 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 website at <http://www.adobe.com>.

More information related to XL C including IBM Redbooks® publications, white papers, tutorials, and other articles, is available on the web at:

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

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

Standards and specifications

XL 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 - 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>.
- *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.
- *Information Technology - Programming Languages - Extension for the programming language C to support decimal floating-point arithmetic, ISO/IEC WDTR 24732*. This draft technical report has been submitted to the C standards committee, and is available at <http://www.open-std.org/JTC1/SC22/WG14/www/docs/n1176.pdf>.
- *ANSI/IEEE Standard for Binary Floating-Point Arithmetic, ANSI/IEEE Std 754-1985*.
- *OpenMP Application Program Interface Version 3.1*, available at <http://www.openmp.org>

Other IBM information

- *Parallel Environment for AIX: Operation and Use*
- The IBM Systems Information Center, at <http://publib.boulder.ibm.com/infocenter/systems/index.jsp?topic=/com.ibm.aix.doc/doc/base/aixparent.htm> is a resource for AIX information.
You can find the following books for your specific AIX system:
 - *AIX Commands Reference, Volumes 1 - 6*
 - *Technical Reference: Base Operating System and Extensions, Volumes 1 & 2*
 - *AIX National Language Support Guide and Reference*
 - *AIX General Programming Concepts: Writing and Debugging Programs*
 - *AIX Assembler Language Reference*
- *ESSL for AIX V5.1/ESSL for Linux on POWER® V5.1 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 Support page at <http://www.ibm.com/software/awdtools/xlc/aix/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 email to compinfo@ca.ibm.com.

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

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 information, send your comments by email 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, 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

IBM XL C for AIX, V12.1 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 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 for AIX, V12.1 is part of a larger family of IBM C, C++, and Fortran compilers.

XL C, together with XL C++ and 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[®]/P, 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.

Operating system support

This section describes the operating systems that IBM XL C for AIX, V12.1 supports.

IBM XL C for AIX, V12.1 supports the following operating systems:

- AIX
 - AIX V5.3 TL 5300-07 or later
 - AIX V6.1
 - AIX V7.1
- PASE
 - IBM i V6.1 PASE V6.1 with PTF SI30636 or later
 - IBM i V7.1 PASE V7.1

See the README file and "Before installing XL C" in the *XL C Installation Guide* for a complete list of requirements.

The compiler, its libraries, and its generated object programs will run on 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 provides several different commands that you can use to invoke the compiler, for example, `xlC`, `c99`, and `c89`. 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 language levels, and many popular language extensions as well.

The compiler also provides corresponding "`_r`" versions of most invocation commands, for example, `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 compiler invocation commands, see "Invoking the compiler" in the *XL 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 compilers, and perform many other common tasks that would otherwise require changing the source code.

XL 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 compiler options, see "Compiler options reference" in the *XL C Compiler Reference*.

Custom compiler configuration files

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

Your compilation needs may frequently involve specifying compiler option settings other than the default settings provided by XL 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 29.

Utilization tracking configuration file

The utilization and reporting tool can be used to detect whether your organization's use of the compiler exceeds your license entitlements.

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.

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

Language standard compliance

This section provides language standard compliance information for IBM XL C for AIX, V12.1.

The compiler supports the following programming language specifications for C:

- ISO/IEC 9899:1999 (C99)
- ISO/IEC 9899:1990 (referred to as C89)

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

- OpenMP Application Program Interface V3.1
- Language extensions to support vector programming
- A subset of GNU C language extensions
- A subset of C1X features

See “C1X features” on page 9 for more details.

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

Compatibility with GNU

XL C supports a subset of the GNU compiler command options to facilitate porting applications developed with `gcc` compilers.

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

These invocation command uses a plain text configuration file to control GNU-to-XL 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 compiler options with `gxlc` " in the *XL C Compiler Reference* for more information.

Source-code migration and conformance checking

XL C helps protect your investment in your existing 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 Compiler Reference* for more information.

Libraries

XL 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 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 Optimization and Programming Guide*.

Other libraries

The following are also shipped with XL C:

- The SMP runtime library supports both explicit and automated parallel processing. See "SMP Runtime Library" in the *XL C Optimization and Programming Guide*.
- The memory debug runtime library is used for diagnosing memory leaks. See "Using memory heaps" in the *XL C Optimization and Programming Guide*.

For more information on support for libraries, search on the XL C Compiler support page at <http://www.ibm.com/software/awdtools/xlc/aix/support/>.

Tools, utilities, and commands

This topic introduces the main tools, utilities, and commands that are included with XL C. It does not contain all compiler tools, utilities, and commands.

Tools

IBM Debugger for AIX

The IBM Debugger for AIX can help you detect and diagnose errors in programs that are running locally or remotely. You can control the execution of your programs by setting compiled language-specific breakpoints, suspending execution, stepping through your code, and examining and changing the contents of variables.

The debugger contains views and functionality specific to a given programming language. With the compiled language views, you can monitor variables, expressions, registers, memory, and application modules of the application you are debugging.

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. The **urt** command contains options that can be used to customize the report. For more information, see Tracking and reporting compiler usage in the *XL C Compiler Reference*.

Utilities

CreateExportList utility

The **CreateExportList** utility creates a file that contains a list of all the global symbols found in a given set of object files. For more information, see Exporting symbols with the CreateExportList utility in the *XL C Optimization and Programming Guide*.

gxc utility

The **gxc** utility translates a GNU C invocation command into a corresponding **xlc** command and invokes the XL C compiler. The purpose of the utility is to minimize the number of changes to makefiles used for existing applications built with the GNU compiler and to facilitate the transition to the XL C compiler. For more information, see Reusing GNU C compiler options with **gxc** in the *XL C Compiler Reference*.

Commands

genhtml command

The **genhtml** command converts an existing XML diagnostic report produced by the **-qlistfmt** option. You can choose to produce XML or HTML diagnostic reports by using the **-qlistfmt** option. The report can help with finding optimization opportunities. For more information about how to use this command, see **genhtml** command in the *XL C Compiler Reference*.

Profile-directed feedback (PDF) related commands

cleanpdf command

The **cleanpdf** command removes all profiling information from the directory to which profile-directed feedback data is written.

mergepdf command

The **mergepdf** command 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 following types of profiling information for all the procedures executed in a PDF run (compilation under the **-qpdf1** option):

- Block-counter profiling
- Call-counter profiling
- Value profiling
- Cache-miss profiling, if you specified the **-qpdf1=level=2** option during the **-qpdf1** phase.

You can view the first two types of profiling information in either text or XML format. However, you can view value profiling and cache-miss profiling information only in XML format.

For more information, see `-qpdf1`, `-qpdf2` in the *XL C Compiler Reference*.

vacndi

The **vacndi** script installs XL C to a non-default directory location. For more information, see Updating an advanced installation using **vacndi** in the *XL C Installation Guide*.

Program optimization

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

64-bit object capability

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

The AIX 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 binder 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
- Attempts to run 64-bit applications on 32-bit platforms

On both 64-bit and 32-bit platforms, 32-bit executables will continue to run as they currently do on a 32-bit platform.

XL 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 Optimization and Programming Guide*.

Shared memory parallelization

XL C supports application development for multiprocessor system architectures.

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

- Directive-based shared memory parallelization (OpenMP, SMP)
- Instructing the compiler to automatically generate shared memory parallelization
- Message passing based shared or distributed memory parallelization (MPI)
- POSIX threads (Pthreads) parallelization
- Low-level UNIX parallelization using `fork()` and `exec()`

The parallel programming facilities of the AIX operating system are based on the concept of threads. Parallel programming exploits the advantages of multiprocessor systems, while maintaining a full binary compatibility with existing uniprocessor systems. This means that a multithreaded program that works on a uniprocessor system can take advantage of a multiprocessor system without recompiling.

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

OpenMP directives

OpenMP directives are a set of API-based commands supported by XL 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 for AIX, V12.1, XL C supports the OpenMP API Version 3.1 specification. See "OpenMP 3.1" on page 10 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 Optimization and Programming Guide*
- www.openmp.org

Diagnostic listings

The compiler output listings and the XML or HTML reports 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, see "Compiler messages and listings" in the *XL C Compiler Reference*.

It is also possible to get information from the compiler in XML or HTML 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 is 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 Optimization and Programming Guide*.

Symbolic debugger support

You can instruct XL C to include debugging information in your compiled objects by using the `-g` compiler option.

The debugging information can be examined by **dbx**, the IBM Debugger for AIX, or any other symbolic debugger that supports the AIX XCOFF executable format to help you debug your programs.

Chapter 2. What's new for IBM XL C for AIX, V12.1

This section describes features and enhancements added to the compiler in IBM XL C for AIX, V12.1.

C1X features

This release introduces support for selected features of C1X.

Note: C1X is a new version of the C programming language standard. IBM continues to develop and implement the features of the new standard. The implementation of the language level is based on IBM's interpretation of the standard. Until IBM's implementation of all the features of the C1X standard is complete, including the support of a new C standard library, the implementation may change from release to release. IBM makes no attempt to maintain compatibility, in source, binary, or listings and other compiler interfaces, with earlier releases of IBM's implementation of the new features of the C1X standard and therefore they should not be relied on as a stable programming interface.

The following features are introduced in IBM XL C for AIX, V12.1:

- Anonymous structures
- Complex type initialization
- New language level - **extc1x**
- The **_Noreturn** function specifier
- Static assertions

Anonymous structures

This feature enables the declaration of anonymous structures under the **extc1x** language level. For more information, see "Anonymous structures" in the *XL C Language Reference*.

Complex type initialization

Macros **CMPLX**, **CMPLXF**, and **CMPLXL** are defined inside the standard header file **complex.h** to enable the initialization of complex types under the **extc1x** language level. For more information, see "Initialization of complex types (C1X)" in the *XL C Language Reference*.

New language level - **extc1x**

A new suboption has been added to the **-qlanglvl** option in this release. You can use **-qlanglvl=extc1x** to enable C1X features that are currently supported by XL C.

The **_Noreturn** function specifier

The **_Noreturn** function specifier declares that a function does not return to its caller. You can define your own functions that do not return using this function specifier. The compiler can produce better code by ignoring what would happen if the function returns. For more information, see "The **_Noreturn** function specifier" in the *XL C Language Reference*.

Static assertions

The addition of static assertions to the C language has 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, improving usability.

You can declare static assertions to check important program invariants at compile time.

For more information, see "`_Static_assert` declaration (C1X)" in the *XL C Language Reference*.

OpenMP 3.1

IBM XL C for AIX, V12.1 supports the OpenMP Application Program Interface Version 3.1 specification. The XL C implementation is based on IBM's interpretation of the OpenMP Application Program Interface Version 3.1.

OpenMP 3.1 includes the following updates to OpenMP 3.0:

- Adds `final` and `mergeable` clauses to the task construct to support optimization.
- Adds the `taskyield` construct to allow users to specify where in the program can perform task switching.
- Adds the `omp_in_final` runtime library function to support specialization of final task regions.
- Extends the `atomic` construct to include `read`, `write`, and `capture` forms; adds the `update` clause to apply the existing form of the `atomic` construct.
- Adds two reduction operators: `min` and `max`.
- Allows `const`-qualified types to be specified on the `firstprivate` clause.
- Adds the `OMP_PROC_BIND` environment variable to control whether OpenMP threads are allowed to move between processors.
- Extends the `OMP_NUM_THREADS` environment variable to specify the number of threads to use for nested parallel regions.

Related information

- "OpenMP environment variables" in the *XL C Compiler Reference*
- "Pragma directives for parallel processing" in the *XL C Compiler Reference*
- www.openmp.org

Performance and optimization

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

Reports about compiler optimizations

There are a number of enhancements to the listing reports to give you more information about how the compiler optimized your code. You can use this information to get further benefits from the optimization capabilities of the compiler. For more details about these enhanced reports, see "New diagnostic reports" on page 11.

For additional information about performance tuning and program optimization, see "Optimizing your applications" in the *XL 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 HTML format

It is now possible to get information in XML or HTML 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 **-qlistfmt** option and its associated suboptions can be used to generate the XML or HTML report. By default, this option now generates all the available content if you do not specify the type of content.

To view the HTML version of an XML report that has been already generated, you can now use the **genhtml** tool. For more information about how to use this tool, see **genhtml** command in the *XL C Compiler Reference*.

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

Enhancements to profiling reports

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

Relevance of profiling data

This section shows the relevance of the profiling data to the source code during the **-qpdf1** phase. The relevance is indicated by a number in the range of 0 - 100. The larger the number is, the more relevant the profiling data is to the source code, and the more performance gain can be achieved by using the profiling data.

Missing profiling data

This section might include a warning message about missing profiling data. The warning message is issued for each function for which the compiler does not find profiling data.

Outdated profiling data

This section might include a warning message about outdated profiling data. The compiler issues this warning message for each function that is modified after the **-qpdf1** phase. The warning message is also issued when the optimization level changes from the **-qpdf1** phase to the **-qpdf2** phase.

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

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

Enhancements to showpdf reports

In addition to block-counter and call-counter profiling information currently provided, you can also use the **showpdf** utility to view cache-miss profiling and value profiling information. Value profiling and cache-miss profiling information can be displayed only in XML format. However, all the other types of profiling information can be displayed in either text or XML format. In this release, the profile-directed feedback (PDF) information is saved in two files. One is a PDF map file that is generated during the **-qpdf1** phase, and the other is a PDF file that is generated during the execution of the resulting application. You can run the **showpdf** utility to display the PDF information contained in these two files. For more information, see "Viewing profiling information with showpdf" in the *XL C Optimization and Programming Guide*.

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.

The 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 Compiler Reference*.

Table 4. Listings-related compiler options and directives

Option/directive	Description
-qlistfmt	<p>The -qlistfmt option has been enhanced to generate HTML reports as well as XML reports, containing information about optimizations performed by the compiler and missed optimization opportunities.</p> <p>The default behavior of this option has changed. Now, if you do not specify a particular type of content, the option generates all the available content, rather than generating none.</p>

New built-in functions

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

GCC atomic memory access built-in functions

New XL C built-in functions for atomic memory access, whose behavior corresponds to that provided by GNU Compiler Collection (GCC), are added in this release. In a program with multiple threads, you can use these functions to atomically and safely modify data in one thread without interference from another thread.

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

New or changed compiler options and pragma directives

This section describes new or changed compiler options and pragma directives.

You can specify compiler options on the command line. You can also modify compiler behavior through pragma directives embedded in your application source files. See the *XL C Compiler Reference* for detailed descriptions and usage information for these and other compiler options.

New or changed compiler options

-g The **-g** option is extended to have new levels to improve the debugging of optimized programs.

-qhaltmsg

The **-qhaltmsg** option, previously supported only by the C++ compiler, is now supported by XL C. It stops compilation before producing any object files, executable files, or assembler source files if a specified error message is generated. The negative form **-qnohaltmsg** has also been added.

-qinclude

The negative form **-qnoinclude** is added to ignore the previously specified **-qinclude** option.

-qinfo **-qinfo=all** now enables all diagnostic messages for all groups except **als** and **ppt**

-qlanglvl

The following suboptions are added or updated:

C1X **-qlanglvl=extc1x**

This suboption enables all the currently supported C1X features and other implementation-specific language extensions.

IBM **-qlanglvl=textafterendif**

This suboption suppresses the warning message that is emitted when you are porting code from a compiler that allows extra text after **#endif** or **#else** to IBM XL C compiler.

See “C1X features” on page 9 for more information about the C1X features.

-qlistfmt

The **-qlistfmt** option is enhanced to generate HTML reports as well as XML reports, containing information about optimizations performed by the compiler and missed optimization opportunities.

The default behavior of **-qlistfmt** has changed. In this release, if you do not specify a particular type of content, the option generates all the available content, rather than generating none.

-qoptfile

The new option **-qoptfile** specifies a file containing a list of additional command line options to be used for the compilation.

-qpic **-qpic=large** now enables large TOC access and prevents TOC overflow conditions when the Table of Contents is larger than 64 Kb.

-qshowpdf

The default value is changed from **-qnoshowpdf** to **-qshowpdf**.

New or changed pragma directives

#pragma ibm max_iterations

The **max_iterations** pragma is added. It specifies the approximate maximum number of loop iterations for the chosen loop.

#pragma ibm min_iterations

The **min_iterations** pragma is added. It specifies the approximate minimum number of loop iterations for the chosen loop.

#pragma simd_level

The **simd_level** pragma is added. It controls the compiler code generation of vector instructions for individual loops.

Chapter 3. Migrating from earlier versions

By migrating to the latest version of the compiler, you are able to take advantage of the new features of the compiler, including features to help boost the performance of the applications being compiled.

A later version of the compiler includes new and enhanced features. These features can provide the following benefits:

- Improved optimization of compiled applications with additional performance benefits
- Support for new language standards to facilitate code portability between multiple operating systems and hardware platforms
- Exploitation of the latest functionality in new hardware and operating systems

Migrating to the latest version of the compiler gives you the ability to exploit the new features with benefits of increased performance optimization, support of new language specifications, and exploitation of new hardware and software environments.

For detailed information about the new features in the current release, see Chapter 2, “What's new for IBM XL C for AIX, V12.1,” on page 9. The white paper *Upgrading XL C/C++ Compilers* at <https://www.ibm.com/support/docview.wss?uid=swg27022099> also provides additional information for compiler migration.

In general, a new version of the compiler is compatible with its earlier versions. However, there could be exceptions. For example, different diagnostic messages could be generated. Always back up your source code and other important data before migrating the compiler.

The following sections list the enhancements added to the compiler in earlier versions, which can assist you in migrating to a later version of the compiler.

Enhancements added in Version 11.1

This section describes features and enhancements added to the compiler in Version 11.1.

Operating system support

This section provides information about the operating systems supported in this release.

XL C V11.1 supports the following operating systems:

- AIX V5.3 TL 5300-07 or later
- AIX 6.1
- IBM i V6.1 PASE V6.1 with PTF SI30636 or later

Note: For operating systems supported by earlier versions of the compilers, see AIX OS levels supported by XL Compilers.

Support for POWER7 processors

XL C for AIX, 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 Language Reference* and Vector built-in functions in the *XL 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 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 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 24.

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 Compiler Reference* and **-qtune** in the *XL C Compiler Reference*.

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

In previous releases, if you modify the source files and compile them with the **-qpdf2** option, the compilation stops with an error. As of XL C for AIX, V11.1, the compiler issues a list of warnings but the compilation does not stop. This allows you to continue using the profiling data after modifying the source files.

Some new suboptions are added to the **-qpdf** option. You can use these new suboptions to get more control over performance improvements and extend **-qpdf** to support multiple-pass profiling, cache-miss profiling, and extended value profiling.

The new **-qpdf** suboptions are:

level Supports multiple-pass profiling, single-pass profiling, cache-miss profiling, value profiling, block-counter profiling, and call-counter profiling. You can compile your program with **-qpdf1=level=0|1|2** to specify the type of profiling information to be generated by the resulting application.

exename

Specifies the name of the generated PDF file according to the output file name specified by the **-o** option.

defname

Reverts the PDF file to its default file name.

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

Reports about compiler optimizations

There are a number of enhancements to the listing reports to give you more information about how the compiler optimizes your code. You can use this information to get further benefits from the optimization capabilities of the compiler. For more details about these enhanced reports, see “New diagnostic reports” on page 19.

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 options, directives, and other performance-related compiler options, see “Optimization and tuning options” in the *XL C Compiler Reference*.

Table 5. Performance-related compiler options and directives

-qfuncsect	An enhancement added to -qfuncsect is to improve linker garbage collection of functions with XL C programs. -qfuncsect places instructions for each function in a separate object file control section or CSECT which might reduce the size of your program. Placing each function in its own CSECT enables the linker to perform garbage collection on a per function basis rather than per object file. For details, see the -qfuncsect section in the <i>XL C Compiler Reference</i> .
-qhot	Two suboptions -qhot=fastmath and -qhot=nofastmath are added to -qhot , to tune your applications to use the fast scalar versions of the math routines or to use the default versions. -qhot=level=2 is also added for loop transformation analysis of nested loops. For details, see the -qhot section in the <i>XL C Compiler Reference</i> .
-qinline=level=number	A new option is added to -qinline 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 0 - 10 that indicates the level of inlining you want to use. For details, see -qinline in the <i>XL C Compiler Reference</i> .
-qipa	A new enhancement added to -qipa is -r -qipa=relink . You can generate relinkable objects while preserving IPA information by specifying -r -qipa=relink . This creates a nonexecutable package that contains all object files. By using this suboption, you can postpone linking until the last stage. -qipa=clonearch is no longer supported. Consider using -qtune=balanced . For detailed information, see -qipa section in the <i>XL C Compiler Reference</i> .
-qpdf	-qpdf provides suboptions to give you more control flexibility in controlling different PDF optimizations. For more information, see the -qpdf1 , -qpdf2 section in the <i>XL C Compiler Reference</i> .

Table 5. Performance-related compiler options and directives (continued)

-qprefetch	A new enhancement is added to -qprefetch for inserting prefetch instructions automatically where there are opportunities to improve code performance: -qprefetch=assistthread . For details, see -qprefetch in the <i>XL C Compiler Reference</i> .
-------------------	--

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

Enhancements to profiling reports

New sections have been added to your listing file to help you analyze your programs. When using the **-qreport** option with the **-qpdf2** option, you can get the following sections added to the 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 profile-directed feedback, see "Using profile-directed feedback" in the *XL C Optimization and Programming Guide*.

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

Table 6. Listings-related compiler options and directives

Option/directive	Description
<code>-qlistfmt</code>	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.
<code>-qreport</code>	The listing now contains a PDF report section when used with <code>-qpdf2</code> . Another new section in the listing files is a DATA REORGANIZATION section when used with <code>-qipa=level=2</code> or <code>-O5</code> .
<code>-qskipsrc</code>	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 Compiler Reference*.

New or changed compiler options and directives

This section describes new and changed compiler options and directives in XL C, V11.1.

You can specify compiler options on the command line. You can also modify compiler behavior through pragma directives embedded in your application source

files. See the *XL C Compiler Reference* for detailed descriptions and usage information for these and other compiler options.

Table 7. New or changed compiler options and directives

Option or 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. It is used to provide information about the characteristics of the files that can help to fine-tune optimizations.
-qconcurrentupdate	If you are building kernel extensions, you must use -qconcurrentupdate to enable hot patching. For details, see -qconcurrentupdate in the <i>XL C Compiler Reference</i> .
-qfunctrace	Traces the entry and exit points of functions in a compilation unit or only for a specific list of functions.
-qhot	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. 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.
-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 .
-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, profile-directed feedback, and data reorganization.
-qpdf1,-qpdf2	New suboptions 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.
-qsimd	Controls whether the compiler can automatically take advantage of vector instructions for processors that support them.

Table 7. New or changed compiler options and directives (continued)

Option or directive	Description
-qskipsrc	When a listing file is generated using the -qsource option, you can use -qskipsrc to control whether the source statements skipped by the compiler are shown in the source section of the listing file. Alternatively, you can use the -qskipsrc=hide option to hide the source statements skipped by the compiler.
-qstackprotect	Protects your applications against malicious code or programming errors that overwrite or corrupt the stack.
-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 8. Deprecated directives and options

Option or directive	Description
#pragma ibm critical	This directive is deprecated and might be removed in a future release. You can use the OpenMP equivalent.
#pragma ibm parallel_loop	This directive is deprecated and might be removed in a future release. You can use the OpenMP equivalent.
#pragma ibm schedule	This directive is deprecated and might be removed in a future release. You can use the OpenMP equivalent.
-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 .
-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 XL C, V11.1.

For more information about built-in functions provided by XL C, see Compiler built-in functions in the *XL 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 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`
- `__dcbt`
- `__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 Compiler Reference*.

POWER7 hardware built-in functions

New XL 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`

Compatibility of redistributable library `libxlopt.a`

Starting from this release, compatibility of the redistributable library, `libxlopt.a`, will be maintained. The `libxlopt.a` library will be compatible with the XL C for AIX, V11.1 compiler and its later releases.

Previously, the version of the redistributable library had to be the same as the version of the compiler with which the application was compiled.

You can download and use the latest redistributable library for multiple applications compiled with XL C for AIX, V11.1 or later.

For more information about the redistributable libraries, see Redistributable libraries in the *XL C Compiler Reference*.

Enhancements added in Version 10.1

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

Operating system support

This section provides information about the supported operating systems.

XL C V10.1 supports the following operating systems:

- AIX V5.3
- AIX V6.1
- AIX V7.1

This version of the compiler does not support AIX V5.2.

Note: For operating systems supported by earlier versions of the compilers, see AIX OS levels supported by XL Compilers.

XL C language-related updates

This section describes the language-related changes introduced in version 10.1.

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

Thread local storage

The thread local storage support has been enhanced to include `__attribute__((tls-model("string")))` where *string* is one of `local-exec`, `initial-exec`, `local-dynamic`, or `global-dynamic`.

OpenMP 3.0

IBM XL C for AIX, V10.1, supports the OpenMP API Version 3.0 specification. The XL 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 expected 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.

For more information, see:

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

Performance and optimization

XL C, V10.1 includes features and enhancements to assist with performance tuning and optimization of your applications.

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 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 Compiler Reference*.

Table 9. 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.
-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 26.

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

New or changed compiler options and pragma directives

This section describes new and changed compiler options and directives in XL C, V10.1.

You can specify compiler options on the command line. You can also modify compiler behavior through pragma directives embedded in your application source files. See the *XL C Compiler Reference* for detailed descriptions and usage information for these and other compiler options.

Table 10. New or changed compiler options and directives

Option or 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 17 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 26.
#pragma init and #pragma fini	Programmers can use #pragma init and #pragma fini to specify a list of functions to run before or after main() or when shared libraries are loaded or unloaded. These functions can be used to do initialization and cleanup. Note: A C++ invocation, such as xlC or the redistributable tools linkxlC or makeC++SharedLib must be used at link time.
-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.
-qunique	-qunique now applies to C as well.

Chapter 4. Setting up and customizing XL C

For complete prerequisite and installation information for XL C, refer to "Before installing" in the *XL 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 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 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 Compiler Reference*.

Chapter 5. Developing applications with XL 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 is properly installed and configured. For more information, see the *XL C Installation Guide*.
2. To learn about writing C programs, refer to the *XL 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 source files

To create C source programs, you can use any text editor available to your system, such as `vi` or `emacs`.

Source programs must be saved using a recognized file name suffix. See “XL C input and output files” on page 33 for a list of suffixes recognized by XL C.

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

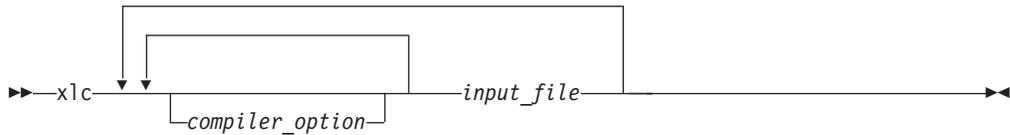
Compiling with XL C

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

Invoking the compiler

The compiler invocation commands perform all necessary steps to compile 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:



For most applications, you should compile with `xlc` or a thread safe counterpart.

Additional invocation commands are available to meet specialized compilation needs, primarily to provide explicit compilation support for different levels and extensions of the C language. See "Invoking the compiler" in the *XL 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.

Compiling parallelized XL C applications

XL 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 thread-safe compiler invocation is:

- `xlc_r`, `xlc_r7`, `xlc128_r`, `xlc128_r7`

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

Note: Using any of these commands alone does not imply parallelization. For the compiler to recognize SMP or 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 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 Compiler Reference* for more information about compiler options and their usage.

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

Reusing GNU C compiler options with `gxc`

XL C includes various features to help you transition from GNU C compilers to XL C, including the `gxc` command.

This utility accepts GNU C compiler options and translates them into comparable XL C options. It uses the XL C options to create an `xlc` invocation command, which the utility then uses to invoke the compiler. The `gxc` utility is provided to help you reuse makefiles created for applications previously developed with GNU C. However, to fully exploit the capabilities of XL C, you should use the `xlc` and `cc` invocation commands together with their associated options.

The actions of `gxc` are controlled by the configuration file `gxc.cfg`. The GNU C options that have an XL C counterpart are shown in this file. Not every GNU option has a corresponding XL C option. The `gxc` utility returns a warning for any GNU C option it cannot translate.

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

XL C input and output files

These file types are recognized by XL C.

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

Table 11. Input file types

Filename extension	Description
.a	Archive or library files
.c	C source files
.i	Preprocessed source files
.o	Object files
.s	Assembler files
.S	Unpreprocessed assembler files
.so	Shared object files

Table 12. Output file types

Filename extension	Description
a.out	Default name for executable file created by the compiler
.d	Target file suitable for inclusion in a makefile
.i	Preprocessed source files
.lst	Listing files
.o	Object files
.s	Assembler files
.so	Shared object files
.u	Make dependency files

Linking your compiled applications with XL C

By default, you do not need to do anything special to link an XL 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 (file2.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 Compiler Reference*

Relinking an existing executable file

The linker accepts executable files as input, so you can link an existing executable file with updated object files.

You cannot, however, relink executable files that were previously linked using the **-qipa** option.

If you have a program consisting of several source files and only make localized changes to some of the source files, you do not necessarily have to compile each file again. Instead, you can include the executable file as the last input file when compiling the changed files:

```
xlc -omansion front_door.c entry_hall.c parlor.c sitting_room.c \  
    master_bath.c kitchen.c dining_room.c pantry.c utility_room.c  
  
vi kitchen.c # Fix problem in OVEN function  
  
xlc -o newmansion kitchen.c mansion
```

Limiting the number of files to compile and link the second time reduces the compile time, disk activity, and memory use.

Note: You should avoid this type of linking unless you are experienced with linking. If done incorrectly, it can result in interface errors and other problems. If you do encounter problems, compiling with the **-qextchk** compiler option can help you diagnose problems with linking.

Dynamic and static linking

XL 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 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

After a program is compiled, you can run the generated executable file on the command line.

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

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`.

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

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 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 Installation Guide*.

Running compiled applications on other systems

In general, applications linked on a system using an earlier version of AIX will run with more recent versions of AIX. However, applications linked on a system using a newer version of AIX will not necessarily run with earlier versions of AIX.

XL C compiler diagnostic aids

XL 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 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.

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

You can then use **dbx**, the IBM Debugger for AIX, or any other symbolic debugger that supports the AIX XCOFF executable format 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 Optimization and Programming Guide*.

Determining what level of XL C is installed

When contacting software support for assistance, you will need to know what level of XL C is installed on your 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 websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites 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, 2012. All rights reserved.

Trademarks and service marks

IBM, the IBM logo, and [ibm.com](http://www.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, other countries, or both.

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 33
- .c files 33
- .i files 33
- .lst files 33
- .mod files 33
- .o files 33
- .s files 33
- .S files 33

Numerics

- 64-bit environment 6

A

- a.out file 33
- archive files 33
- assembler
 - source (.s) files 33
 - source (.S) files 33

B

- basic example, described viii
- built-in functions 12, 24

C

- code optimization 6
- compilation
 - sequence of activities 31
- compiler
 - controlling behavior of 32
 - invoking 32
 - running 32
- compiler directives
 - new or changed 13, 21
- compiler options
 - conflicts and incompatibilities 33
 - new or changed 13, 21
 - specification methods 32
- compiling
 - SMP programs 32
- customization
 - for compatibility with GNU 3

D

- dbx debugger 8, 37
- debugger support 37
 - output listings 36
 - symbolic 8
- debugging 37
- debugging compiled applications 36
- debugging information, generating 36
- dynamic linking 35

E

- editing source files 31
- executable files 33
- executing a program 36
- executing the linker 34

F

- files
 - editing source 31
 - input 33
 - output 33

G

- GNU
 - compatibility with 3

I

- input files 33
- invocation commands 32
- invoking a program 36
- invoking the compiler 32

L

- language standards 3
- language support 3
- level of XL C, determining 37
- libraries 33
- linking
 - dynamic 35
 - static 35
- linking process 34
- listings 33

M

- migrate 15
- migration 15
 - source code 32
- mod files 33
- multiprocessor systems 7, 26

O

- object files 33
 - creating 34
 - linking 34
- OMP directives 26
- OpenMP 7
- optimization
 - programs 6
- output files 33

P

- parallelization 7, 26
- performance
 - optimizing transformations 6
- problem determination 36
- programs
 - running 36

R

- running the compiler 32
- runtime
 - libraries 33
- runtime environment 36
- runtime options 36

S

- shared memory parallelization 7, 26
- shared object files 33
- SMP
 - programs, compiling 32
- SMP programs 7
- source files 33
- source-level debugging support 8
- static linking 35
- symbolic debugger support 8

T

- tools 4
 - cleanpdf utility 5
 - CreateExportList 5
 - custom installation 6
 - debugger 4
 - gxc utility 5
 - IBM Debugger 4
 - install 6
 - mergepdf utility 5
 - resetpdf utility 5
 - showpdf utility 5
 - urt 4
 - vacndi 6

U

- utilities 4
 - cleanpdf 5
 - CreateExportList 5
 - custom installation 6
 - gxc 5
 - IBM Debugger 4
 - install 6
 - mergepdf 5
 - resetpdf 5
 - showpdf 5
 - urt 4
 - vacndi 6

V

vac.cfg file 32



Product Number: 5765-J01; 5725-C71

Printed in USA

SC14-7323-00

