

IBM XL C for AIX, V13.1.3



Getting Started with XL C

Version 13.1.3

IBM XL C for AIX, V13.1.3



Getting Started with XL C

Version 13.1.3

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, V13.1.3 (Program 5765-J06; 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, 2015.

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

Contents

About this document	v	Support for POWER8 processors	14
Conventions	v	C11 features	15
Related information	viii	OpenMP 4.0	16
IBM XL C information	ix	Built-in functions	16
Standards and specifications	x	Compiler options and pragma directives	20
Other IBM information	x	Performance and optimization	22
Other information	x	Enhancements added in Version 12.1	22
Technical support	xi	C11 features	22
How to send your comments	xi	OpenMP 3.1	23
		Performance and optimization	24
		Diagnostic reports	24
		Built-in functions	26
		Compiler options and pragma directives	26
Chapter 1. Introducing XL C	1	Chapter 5. Setting up and customizing XL C.	29
Commonality with other IBM compilers	1	Using custom compiler configuration files	29
Operating system and hardware support	1	Configuring compiler utilization tracking and reporting	29
A highly configurable compiler	1		
Language standard compliance	3	Chapter 6. Developing applications with XL C	31
Compatibility with GNU	3	The compiler phases	31
Source-code migration and conformance checking	3	Editing C source files	31
Libraries	3	Compiling with XL C	31
Tools, utilities, and commands	4	Invoking the compiler	32
Program optimization	6	Compiling parallelized XL C applications	32
64-bit object capability	6	Specifying compiler options	32
Shared memory parallelization	7	XL C input and output files	33
Diagnostic reports	8	Linking your compiled applications with XL C	34
Symbolic debugger support	8	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 which level of XL C is being used	37
Chapter 2. What's new for IBM XL C for AIX, V13.1.3	9	Notices	39
		Trademarks	41
Chapter 3. Migration of your applications	11	Index	43
Migrating applications that use transactional memory built-in functions	11		
Mixing object files compiled with different compilers	11		
Chapter 4. Enhancements added in earlier releases	13		
Enhancements added in Version 13.1.2	13		
Built-in functions	13		
Commands	13		
Compiler options	14		
Enhancements added in Version 13.1	14		

About this document

This document contains overview and basic usage information for the IBM® XL C for AIX®, V13.1.3 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, 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 about 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 behavior of the compiler. You can, however, substitute other forms of the compiler invocation command if your particular environment requires it, and compiler option usage remains the same unless otherwise specified.

While this document covers information such as 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*.
- Compiler options: see the *XL C Compiler Reference* for detailed information about the syntax and usage of compiler options.
- The C programming language: see the *XL C Language Reference* for information about the syntax, semantics, and IBM implementation of the C programming language.
- Programming topics: see the *XL C Optimization and Programming Guide* for detailed information about developing applications with XL C, with a focus on program portability and optimization.

Conventions

Typographical conventions

The following table shows the typographical conventions used in the IBM XL C for AIX, V13.1.3 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 -03.

Qualifying elements (icons)







In descriptions of language elements where a feature is exclusive to the C11 standard, or where a feature is an IBM extension of the C standard, this information uses icons to delineate segments of text as follows:

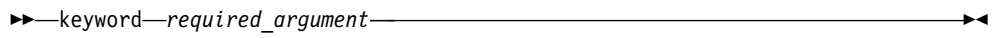
Table 2. Qualifying elements

Qualifier/Icon	Meaning
IBM extension begins   IBM extension ends	The text describes a feature that is an IBM extension to the standard language specifications.
C11 begins   C11 ends	The text describes a feature that is introduced into standard C as part of C11.

Syntax diagrams

Throughout this information, diagrams illustrate XL C syntax. This section helps you to interpret and use those diagrams.

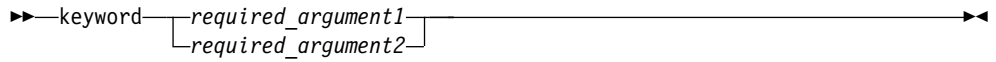
- Read the syntax diagrams from left to right, from top to bottom, following the path of the line.
 - The  symbol indicates the beginning of a command, directive, or statement.
 - The  symbol indicates that the command, directive, or statement syntax is continued on the next line.
 - The  symbol indicates 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):



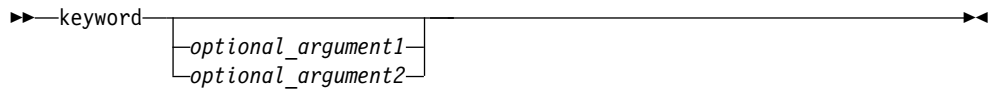
- Optional items are shown below the main path:



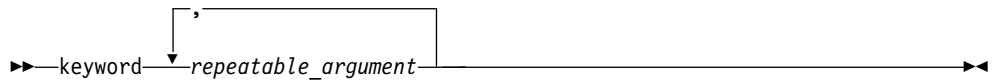
- If you can choose from two or more items, they are shown vertically, in a stack. If you *must* choose one of the items, one item of the stack is shown on the main path.



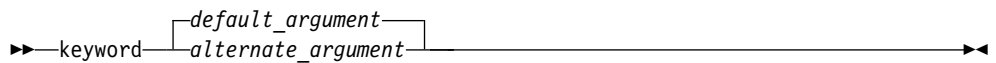
If choosing one of the items is optional, the entire stack is shown below the main path.



- 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:



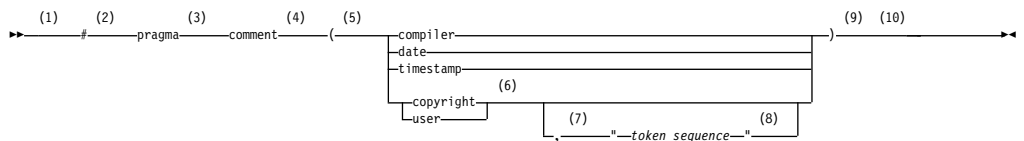
- The item that is the default is shown above the main path.



- 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")
```

Example of a syntax statement

EXAMPLE *char_constant* {*a|b*}[*c|d*]*e*[,*e*]... *name_list*{*name_list*}...

The following list explains the syntax statement:

- Enter the keyword EXAMPLE.
- Enter a value for *char_constant*.
- Enter a value for *a* or *b*, but not for both.
- Optionally, enter a value for *c* or *d*.
- Enter at least one value for *e*. If you enter more than one value, you must put a comma between each.
- Optionally, enter the value of at least one *name* for *name_list*. If you enter more than one value, you must put a comma between each *name*.

Note: The same example is used in both the syntax-statement and syntax-diagram representations.

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:

- Quick Start Guide

The Quick Start Guide (`quickstart.pdf`) is intended to get you started with IBM XL C for AIX, V13.1.3. It is located by default in the XL C directory and in the `\quickstart` directory of the installation DVD.

- 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 DVD.

- 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, V13.1.3 Installation Guide*.

- Online product documentation

The fully searchable HTML-based documentation is viewable in IBM Knowledge Center at http://www.ibm.com/support/knowledgecenter/SSGH2K_13.1.3/com.ibm.compilers.aix.doc/welcome.html.

- PDF documents

PDF documents are available on the web at <http://www.ibm.com/support/docview.wss?uid=swg27036590>.

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, V13.1.3 Installation Guide</i> , SC27-4238-02	<code>install.pdf</code>	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, V13.1.3</i> , SC27-4237-02	<code>getstart.pdf</code>	Contains an introduction to the XL C product, with information about setting up and configuring your environment, compiling and linking programs, and troubleshooting compilation errors.
<i>IBM XL C for AIX, V13.1.3 Compiler Reference</i> , SC27-4239-02	<code>compiler.pdf</code>	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, V13.1.3 Language Reference</i> , SC27-4240-02	<code>langref.pdf</code>	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, V13.1.3 Optimization and Programming Guide</i> , SC27-4241-02	<code>proguide.pdf</code>	Contains information about 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 Adobe Reader. If you do not have 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, and other articles, is available on the web at <http://www.ibm.com/support/docview.wss?uid=swg27036590>.

For more information about C/C++, see the C/C++ café at <https://www.ibm.com/developerworks/community/groups/service/html/communityview?communityUuid=5894415f-be62-4bc0-81c5-3956e82276f3>.

Standards and specifications

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

- *Information Technology - Programming languages - C, ISO/IEC 9899:1990*, also known as C89.
- *Information Technology - Programming languages - C, ISO/IEC 9899:1999*, also known as C99.
- *Information Technology - Programming languages - C, ISO/IEC 9899:2011*, also known as C11. (Partial support)
- *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 (full support)*, and *OpenMP Application Program Interface Version 4.0 (partial support)*, 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*

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/support/entry/portal/product/rational/xl_c_for_aix. 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 an email to compinfo@ca.ibm.com.

For the latest information about XL C, visit the product information site at <http://www.ibm.com/software/products/en/xlcaix>.

How to send your comments

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

Be sure to include the name of the manual, the part number of the manual, 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, V13.1.3 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 contains information about 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, V13.1.3 is part of a larger family of IBM C, C++, and Fortran compilers. XL C, together with XL C++ and XL Fortran, comprises 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[®]/Q, 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 code portability across multiple operating systems and hardware platforms.

Operating system and hardware support

This section describes the operating systems and hardware that IBM XL C for AIX, V13.1.3 supports.

IBM XL C for AIX, V13.1.3 supports the following operating systems:

- AIX V6.1 TL 2 Service Pack 5 or later
- AIX V7.1
- AIX V7.2
- IBM i V7.1 PASE V7.1
- IBM i V7.2 PASE V7.2

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 run on POWER5, POWER5+, POWER6[®], POWER7[®], POWER7+[™], and POWER8[®] 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 according to the hardware type that runs 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 commands to invoke the compiler, for example, `xc`, `c99`, and `c89`. Compiler invocation commands are provided to support most standardized C language levels and many popular language extensions.

The compiler also provides corresponding "`_r`" versions of most invocation commands, for example, `xc_r`. The "`_r`" invocations instruct the compiler to link and bind object files to threadsafe components and libraries, and produce threadsafe 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. You can benefit from using different options for the following tasks:

- Debugging your applications
- Optimizing and tuning application performance
- Selecting language levels and extensions for compatibility with nonstandard features and behaviors that are supported by other C compilers
- Performing many other common tasks that would otherwise require changing the source code

You can 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.

If you frequently specify compiler option settings other than the default settings of XL C, you can use makefiles to define your settings. Alternatively, you can create custom configuration files to define your own frequently used 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 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 topic describes the C programming language specifications that IBM XL C for AIX, V13.1.3 supports.

- Partial support for ISO/IEC 9899:2011 (referred to as C11)
- ISO/IEC 9899:1999 (referred to as C99)
- ISO/IEC 9899:1990 (referred to as C89)

In addition to the standard language levels, XL C supports the following language extensions:

- Partial support for OpenMP Application Program Interface V4.0
- OpenMP Application Program Interface V3.1
- Language extensions to support vector programming
- A subset of GNU C language extensions

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 that are developed with the `gcc` compilers.

This support is available when the `gxc` 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.

The 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 meet your unique compilation requirements. For more information, see "Reusing GNU C compiler options with `gxc`" in the *XL C Compiler Reference*.

Source-code migration and conformance checking

XL C provides 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 language level. If the language or language extension elements in your program source do not conform to the specified language level, the compiler issues diagnostic messages.

Related information

 `-qlanglvl`

Libraries

XL C includes a runtime environment that contains a number of libraries.

Mathematical Acceleration Subsystem library

The Mathematical Acceleration Subsystem (MASS) library consists of scalar and vector mathematical built-in 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 and offer improved performance over the default `libm` math library routines. These libraries are threadsafe and are called automatically when you request specific levels of optimization for your application. You can also make explicit calls to MASS library functions, whether optimization options are in effect or not.

For more information, see "Using the Mathematical Acceleration Subsystem" in the *XL C Optimization and Programming Guide*.

Basic Linear Algebra Subprograms

The Basic Linear Algebra Subprograms (BLAS) set of high-performance algebraic functions are shipped in the `libxlopt` library. You can use these functions to:

- 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 libraries 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 about support for libraries, search on the XL C Compiler support page at http://www.ibm.com/support/entry/portal/product/rational/xl_c_for_aix.

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 exportable 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 you find 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 the PDF files or the specified PDF files 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.

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*.

xlcmdi The **xlcmdi** script installs XL C to a nondefault directory location. For more information, see Updating an advanced installation using **xlcmdi** 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[®] processors
- Improving the usage of the memory subsystem
- Exploiting the ability of the architecture to handle large amounts of shared memory parallelization

Related information



Optimizing your applications



Optimization and tuning



Compiler built-in functions

64-bit object capability

The 64-bit object capability of the XL C compiler 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 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 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 block of code. The existence of the directives in the source removes the need for the compiler to perform any dependence analysis on the parallel code. OpenMP directives require the presence of Pthread libraries to provide the necessary infrastructure for parallelization.

OpenMP directives address the following important issues of parallelizing an application:

1. Clauses and directives are available for scoping variables. Generally, variables should not be shared; that is, each thread 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 threads.

3. Directives are available to control synchronization between threads.

IBM XL C for AIX, V13.1 supports OpenMP API Version 4.0 specification. For details, see "OpenMP 4.0" on page 16.

Related information



Optimizing your applications



The OpenMP API specification for parallel programming

Diagnostic reports

The compiler listings, XML reports, and 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*.

You can also obtain diagnostic information from the compiler in XML or HTML format. The XML and HTML reports provide information about optimizations that the compiler performed or could not perform. You can use this information 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 different levels of 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.

Related information



-g

Chapter 2. What's new for IBM XL C for AIX, V13.1.3

This section describes features and enhancements added to IBM XL C for AIX, V13.1.3.

Software License Metric (SLM) Tags logging support

IBM XL C for AIX, V13.1.3 fully supports IBM Software License Metric (SLM) Tags logging so that you can use IBM License Metric Tool (ILMT) to track compiler license usage.

For more information, see "Tracking compiler usage with Software License Metric Tags logging" in the *XL C Compiler Reference*.

Chapter 3. Migration of your applications

This section lists important considerations when you migrate your applications that were compiled with other versions of XL C.

Migrating applications that use transactional memory built-in functions

Starting from IBM XL C for AIX, V13.1.2, to use transactional memory built-in functions, you must include a header file in the source code. In addition, if you used numeric return values of the transaction begin and end built-in functions, you must replace numeric return values with macro return values that are provided by IBM XL C for AIX, V13.1.3.

New header file needed for transactional memory built-in functions

You must include the `htmxlintrin.h` file in the source code if you use any of the transactional memory built-in functions.

Changed return values of the transaction begin and end built-in functions

The return values of the transaction begin and end built-in functions are no longer numeric. You must update your program using the following return values:

`__TM_begin`

This function returns `_HTM_TBEGIN_STARTED` if successful; otherwise, it returns a different value.

`__TM_end`

This function returns `_HTM_TBEGIN_STARTED` if the thread is in the transactional state before the instruction starts; otherwise, it returns a different value.

`__TM_simple_begin`

This function returns `_HTM_TBEGIN_STARTED` if successful; otherwise, it returns a different value.

Related information

Transactional memory built-in functions



Transactional memory built-in functions

Mixing object files compiled with different compilers

Most object files that were compiled with different compilers can be linked together. However, under some circumstances, object files are not compatible and must be recompiled.

Note the following restrictions:

- There is no binary compatibility among AIX, Linux for big endian distributions, and Linux for little endian distributions.
- Do not mix object files that were compiled with the big endian compiler and object files that were compiled with the little endian compiler.

- Do not mix object and library files that were compiled with different versions of a compiler if the **-qipa** option was used during the compilation. The **-qipa** option instructs the compiler to perform an IPA link for these object and library files. An IPA link might not be able to handle mismatched versions.

Related information:

 -qipa

Chapter 4. Enhancements added in earlier releases

This section describes enhancements added in earlier releases. These enhancements also apply to the current release.

Enhancements added in Version 13.1.2

This section describes features and enhancements added in IBM XL C for AIX, V13.1.2. These features and enhancements apply to later releases as well.

Built-in functions

This section describes built-in functions that are new or changed for IBM XL C for AIX, V13.1.2.

New built-in functions

vec_mergee

Merges the values of even-numbered elements of two vectors.

vec_mergeo

Merges the values of odd-numbered elements of two vectors.

vec_revb

Returns a vector that contains the bytes of the corresponding element of the argument in the reverse byte order.

vec_reve

Returns a vector that contains the elements of the argument in the reverse element order.

vec_xl(a, b)

Loads a 16-byte vector from the memory address specified by the displacement *a* and the pointer *b*.

vec_xl_be(a, b)

Loads a 16-byte vector from the memory address specified by the displacement *a* and the pointer *b*.

vec_xst(a, b,c)

Stores the elements of the 16-byte vector *a* to the effective address obtained by adding the displacement provided in *b* with the address provided by *c*. The effective address is not truncated to a multiple of 16 bytes.

vec_xst_be(a, b,c)

Stores the elements of the 16-byte vector *a* in big endian element order to the effective address obtained by adding the displacement provided in *b* with the address provided by *c*. The effective address is not truncated to a multiple of 16 bytes.

Related information

 Compiler built-in functions

Commands

This section describes new, changed, or removed compiler commands.

resetpdf

This command has been removed. It is recommended that you use the **cleanpdf** command instead. The behavior of the **resetpdf** command is the same as that of the **cleanpdf** command. For more information, see `-qpdf1`, `-qpdf2` in the *XL C Compiler Reference*.

Compiler options

This section describes new or changed compiler options.

-qfloat

The following suboptions are added:

subnormals

This suboption asserts to the compiler that the code uses subnormal floating point values, also known as denormalized floating point values.

nosubnormals

This suboption asserts to the compiler that the code does not use subnormal floating point values, also known as denormalized floating point values.

Whether or not you specify this suboption, the behavior of your program will not change, but the compiler uses this information to gain possible performance improvements. The suboptions take effect only on POWER8 processors. To use **-qfloat=subnormals** or **-qfloat=nosubnormals**, you must also specify the **-qarch=pwr8** and **-qtune=pwr8** options.

-qstrict=guards

The actions that are performed by XL C if you specify the **-qstrict=guards** option have been increased. When the **-qstrict=guards** option is in effect, the compiler behavior is as follows:

- The compiler does not move operations past guards.
- When the compiler encounters if statements that contain pointer wraparound checks that can be resolved at compile time, it does not remove the checks or the enclosed operations.

Enhancements added in Version 13.1

This section describes features and enhancements added to the compiler in Version 13.1. These features and enhancements apply to later versions as well.

Support for POWER8 processors

XL C for AIX, V13.1 supports POWER8 processors.

The new features and enhancements introduced in support of the POWER8 processors, fall under the following categories:

- MASS libraries for POWER8 processors
- Compiler options for POWER8 processors
- Built-in functions for POWER8 processors

Mathematical Acceleration Subsystem (MASS) libraries for POWER8 processors

Vector libraries

The vector MASS library **libmassvp8.a** contains vector functions that have been tuned for the POWER8 architecture. The functions can be used in either 32-bit mode or 64-bit mode.

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_simdp8.a** contains an accelerated set of frequently used math built-in functions that provide improved performance over the corresponding standard system library functions.

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

Compiler options for POWER8 processors

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.

The new **-qarch=pwr8** suboption produces object code containing instructions that will run on the POWER8 hardware platforms. With the new **-qtune=pwr8** suboption, optimizations are tuned for the POWER8 hardware platforms.

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

Built-in functions for POWER8 processors

New hardware built-in functions are added to support the following POWER8 processor features:

- POWER8 functions for vector processing
- POWER8 binary-coded decimal functions
- POWER8 cryptography functions
- POWER8 quad-word arithmetic functions
- POWER8 load-and-reserve/store conditional instructions
- POWER8 cache and data prefetch control functions
- POWER8 transactional memory functions
- POWER8 prefetch functions

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

C11 features

In addition to the existing C11 features, new C11 features are supported in this release of XL C.



Note: IBM supports selected features of C11, known as C1X before its ratification. IBM will continue to develop and implement the features of this standard. The implementation of the language level is based on IBM's interpretation of the standard. Until IBM's implementation of all the C11 features is complete, including the support of a new C11 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 C11 features.

The following features are introduced in AIX, V13.1:

- typedef redeclaration
- Generic selection

typedef redeclaration

Using typedef redeclaration, you can redefine a name that is a previous typedef name in the same scope to refer to the same type.  The XL C compiler supports all types, including a variably modified type.  For more information, see "typedef definitions" in the *XL C Language Reference*.

Generic selection

Generic selection provides a mechanism to choose an expression according to a given type name at compile time. A common usage is to define type generic macros. For more information, see Generic selection (C11).

OpenMP 4.0

XL C for AIX, V13.1 partially supports the OpenMP Application Program Interface Version 4.0 specification. The XL C implementation is based on IBM's interpretation of the OpenMP Application Program Interface Version 4.0.

This version of XL C supports the following OpenMP 4.0 features:

- update and capture clauses enhancements
- OMP_DISPLAY_ENV environment variable

update and capture clauses enhancements

The update and capture clauses of the atomic construct are extended to support more expression forms.

OMP_DISPLAY_ENV environment variable

You can use the OMP_DISPLAY_ENV environment variable to display the values of the internal control variables (ICVs) associated with the environment variables and the build-specific information about the runtime library.

Related information

- "OpenMP environment variables" in the *XL C Compiler Reference*
- "Pragma directives for parallel processing" in the *XL C Compiler Reference*
- The OpenMP API specification for parallel programming

Built-in functions

The following major categories of built-in functions are new to this release.

Note: POWER8 built-in functions are valid only when **-qarch=pwr8** is set or implied.

POWER8 built-in functions for vector processing

The following vector built-in functions are added:

- The vector gather-bits-by-bytes doubleword function
 - `vec_gbb`
- The vector count leading zeros function
 - `vec_cntlz`
- The vector population count function
 - `vec_popcnt`
- Extended vector logical operations functions
 - `vec_eqv`
 - `vec_nand`
 - `vec_orc`
- 128-bit integer add subtract functions
 - `vec_add_u128`
 - `vec_sub_u128`
 - `vec_adde_u128`
 - `vec_sube_u128`
 - `vec_addc_u128`
 - `vec_subc_u128`
 - `vec_addec_u128`
 - `vec_subec_u128`
 - `vec_bperm`

The following built-in functions are extended to support doubleword types:

- Vector pack functions
 - `vec_pack`
 - `vec_packs`
 - `vec_packsu`
- Vector unpack functions
 - `vec_unpackh`
 - `vec_unpackl`
- Vector add and subtract functions
 - `vec_add`
 - `vec_sub`
- Vector max and min functions
 - `vec_max`
 - `vec_min`
- Vector shift and rotate functions
 - `vec_rl`
 - `vec_sl`
 - `vec_sr`
 - `vec_sra`
- Vector compare functions
 - `vec_cmpeq`
 - `vec_cmpgt`

- vec_cmpge
- vec_cmplt
- vec_cmple

Binary-coded decimal built-in functions

The following built-in functions are added to support binary-coded decimal (BCD) arithmetic and comparison. The first three types of built-in functions are POWER8 built-in functions. BCD load and store functions are valid when **-qarch** is set to target POWER8 or POWER7 processors:

- BCD add and subtract functions
 - __bcdadd
 - __bcdsub
- BCD test add and subtract for overflow functions
 - __bcdadd_ofl
 - __bcdsub_ofl
 - __bcd_invalid
- BCD comparison functions
 - __bcdcmpeq
 - __bcdcmpgt
 - __bcdcmpge
 - __bcdcmplt
 - __bcdcmple
- BCD load and store functions
 - __vec_ldrmb
 - __vec_strmb

POWER8 cryptography built-in functions

The following built-in functions are provided to perform cryptographic operations:

- Advanced Encryption Standard (AES) functions
 - __vcipher
 - __vcipherlast
 - __vncipher
 - __vncipherlast
 - __vsbox
- Secure Hash Algorithm (SHA) functions
 - __vshasigmad
 - __vshasigmaw
- Miscellaneous functions
 - __vpmsumb
 - __vpmsumh
 - __vpmsumw
 - __vpmsumd
 - __vpermxor

POWER8 non-vector built-in functions

The following built-in functions are added to improve the efficiency of cache:

- `__dcbtna`
- `__icbt`

Load and store built-in functions are extended with the following functions to support more types:

- `__lqarx`
- `__lharx`
- `__lbarx`
- `__stqcx`
- `__sthcx`
- `__stbcx`

POWER8 transactional memory built-in functions

Transactional memory is a model for parallel programming. In this model, you can designate a block of instructions or statements to be treated atomically.

You can use the following built-in functions to mark the beginning or end of transactions, and to diagnose the reasons for failure:

- Transaction begin and end functions
 - `__TM_begin`
 - `__TM_end`
 - `__TM_simple_begin`
- Transaction abort functions
 - `__TM_abort`
 - `__TM_named_abort`
- Transaction inquiry functions
 - `__TM_failure_address`
 - `__TM_failure_code`
 - `__TM_is_conflict`
 - `__TM_is_failure_persistent`
 - `__TM_is_footprint_exceeded`
 - `__TM_is_illegal`
 - `__TM_is_named_user_abort`
 - `__TM_is_nested_too_deep`
 - `__TM_is_user_abort`
 - `__TM_nesting_depth`

POWER8 prefetch built-in functions

The following built-in functions display the problem state control of the Data Stream Control Register (DSCR) in an intuitive, portable, and optimization-friendly way:

- Transient attribute enable functions
 - `__hardware_transient_enable`
 - `__load_transient_enable`

- __software_transient_enable
- __store_transient_enable
- Unit count enable and set functions
 - __hardware_unit_count_enable
 - __software_unit_count_enable
 - __set_prefetch_unit_count
- Prefetch depth functions
 - __default_prefetch_depth
 - __depth_attainment_urgency
- Load stream enable and disable functions
 - __load_stream_disable
 - __stride_n_stream_enable
- DSCR functions
 - __prefetch_get_dscr_register
 - __prefetch_set_dscr_register

Related information:

- Compiler built-in functions in the *XL C Compiler Reference*
- -qarch

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. For detailed descriptions and usage information for XL C compiler options, see the *XL C Compiler Reference*.

-qarch The option default is updated to **pwr4**. Suboptions denoting old hardware families are silently upgraded to newer architectures.

The following suboptions are added or updated:

-qarch=pwr7

This suboption produces object code containing instructions that run on the POWER7, POWER7+, or POWER8 hardware platforms.

-qarch=pwr8

This suboption produces object code containing instructions that run on the POWER8 hardware platforms.

-qcheck

The following suboptions are added or updated:

-qcheck=stacklobber

This suboption detects a certain type of stack corruption in your programs.

-qcheck=unset

This suboption checks for automatic variables that are used before they are set at run time.

-qdbfmt=dwarf4

This suboption generates debugging information in DWARF 4 format.

-qhelp This option displays the man page of the compiler.

-qinfo

The compiler does not issue informational messages for the following files:

- Files in the standard search paths for compiler and system header files.
- Files that are ultimately included by the files in the standard search paths for compiler and system header files.

The following suboptions are added or updated:

-qinfo=mt

This suboption notifies you about potential places where synchronization is needed.

-qinfo=unset

This suboption detects automatic variables that are used before they are set, and flags them with informational messages at compile time.

-qpdf1=unique

This suboption creates a unique PDF file for each process during run time.

-qprefetch=dscr

This suboption helps to improve the runtime performance of your applications. You can specify a value for `dscr` depending on your system architecture.

-qsimd=auto

This suboption controls the autosimdization, which was performed by the deprecated **-qhot=simd** option.

-qtune

The option default is updated.

The following suboptions are added or updated:

-qtune=pwr7

This suboption specifies that optimizations are tuned for the POWER7 or POWER7+ hardware platforms.

-qtune=pwr8

This suboption specifies that optimizations are tuned for the POWER8 hardware platforms.

SMT suboptions

The new **-qtune** simultaneous multithreading (SMT) suboptions allow you to specify a target SMT to direct optimization for best performance in that mode.

-qunroll=*n*

This suboption hints to the compiler to unroll loops by a factor of *n*. If the loop has fewer than *n* iterations, it is fully unrolled.

-qvisibility

This option specifies visibility attributes for entities. Entity visibility attributes describe whether and how entities defined in one module can be referenced or used in other modules. Visibility attributes affect entities with external linkage only, and cannot increase the visibility of other entities.

New or changed pragma directives

#pragma GCC visibility push, #pragma GCC visibility pop

This pair of pragma directives is the pragma equivalent of the **-qvisibility** option. The pragma directives are used to specify visibility attributes for external linkage symbols.

Performance and optimization

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



Visibility attributes of entities

Entity visibility attributes describe whether and how an entity that is defined in one module can be referenced or used in other modules. By using the visibility attributes for entities, you can get the following benefits:

- Decrease the size of shared libraries
- Reduce the chance of symbol collision
- Allow more optimization for the compile and link phases
- Improve the efficiency of dynamic linking

For more information, see "Using visibility attributes (IBM extension)" in the *XL C Optimization and Programming Guide*.



For more information about performance tuning and program optimization, see "Optimizing your applications" and "Coding your application to improve performance" in the *XL C Optimization and Programming Guide*.

Enhancements added in Version 12.1

This section describes features and enhancements added to the compiler in Version 12.1. These features and enhancements apply to later versions as well.

C11 features

XL C V12.1 introduces support for selected features of C11, which is a new C programming language standard.

Note: IBM supports selected features of C11, known as C1X before its ratification. IBM will continue to develop and implement the features of this standard. The implementation of the language level is based on IBM's interpretation of the standard. Until IBM's implementation of all the C11 features is complete, including the support of a new C11 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 C11 features.

The following C11 features are introduced in XL C, 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 (C11)" 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 C11 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 (C11)" in the *XL C Language Reference*.

OpenMP 3.1

XL C 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 the `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 in XL C V12.1 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 "Diagnostic reports."

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

Diagnostic reports

The new diagnostic reports added in XL C V12.1 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 the `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>

Built-in functions

This section describes the major categories of built-in functions that are new for V12.1.

GCC atomic memory access built-in functions (IBM extension)

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*.

Compiler options and pragma directives

This section describes new or changed compiler options and pragma directives in V12.1.

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.

New or changed compiler options

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

-qhaltonmsg

The **-qhaltonmsg** 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 **-qnohaltonmsg** 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:

C11 -qlanglvl=extc1x

This suboption enables all the currently supported C11 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.

For more information about the C11 features, see “C11 features” on page 22.

-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 5. Setting up and customizing XL C

This section describes how to set up and customize the compiler according to your own requirements.

For complete prerequisite and installation information for XL C, see "Before installing XL C" 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 creating your own configuration file.

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 that you specify in your custom configuration files with compiler settings that are specified in the default configuration file. Compiler updates that might later affect settings in the default configuration file do not affect the settings in your custom configuration files.

Related information



Using custom compiler configuration files

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 6. Developing applications with XL C

C application development consists of repeating cycles of editing, compiling, linking, and running. By default, compiling and linking are combined into a single step.

Notes:

- Before you use the compiler, ensure that XL C is properly installed and configured. For more information, see the *XL C Installation Guide*.
- 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 are executed more than once during a compilation. As each compilation component runs, the results are sent to the next step in the sequence.

1. Preprocessing of source files
2. Compilation, which might 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. Assembling 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 on 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*.

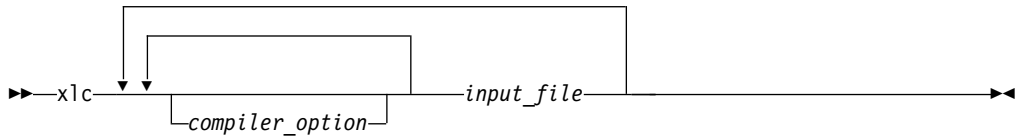
Compiling with XL C

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 or preprocessed files (.i or .ii), assemble any .s and .S files, and link the object files and libraries into an executable program.

To compile a C source program, use the following basic invocation syntax:



For most applications, compile with `xlc` or a threadsafe counterpart.

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

Compiling parallelized XL C applications

XL C provides threadsafe compiler invocation commands to compile 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 threadsafe components and libraries. The generic XL C threadsafe compiler invocations are as follows:

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

XL C provides additional threadsafe invocations to meet specific compilation requirements. For more information, see "Invoking the compiler" in the *XL C Compiler Reference*.

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 the `-qsmp` compiler option. In turn, you should specify the `-qsmp` option only when threadsafe invocations are used. When you specify `-qsmp`, the driver links the libraries that are specified on the `smp libraries` line in the active stanza of the configuration file.

For more information about 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 in one or any combination of the following ways:

- On the command line
- In your source code using directive statements
- In a makefile
- In the stanzas found in a compiler configuration file

You can also pass options to the linker, assembler, and preprocessor.

Priority sequence of compiler options


Option conflicts and incompatibilities might occur when multiple compiler options are specified. To resolve these conflicts in a consistent manner, the compiler applies the following general priority sequence to most options:

1. Directive statements in your source file override command line settings.
2. Compiler option settings on the command line override configuration file settings.
3. Configuration file settings override default settings.

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

Note: Some compiler options, such as the **-I** option, do not follow the priority sequence described above. The compiler searches any directories specified with **-I** in the `xl.c` file before it searches the directories specified with **-I** on the command line. The **-I** option is cumulative rather than preemptive.

Related information

 [Compiler options reference](#)

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.

The `gxc` 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 can use the `xlc` and `cc` invocation commands together with their associated options.

The actions of `gxc` are controlled by the `gxc.c` configuration file. 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 about 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

The topic describes the file types that 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 5. 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 6. 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, you can use `xlc` to compile `file1.c` and `file3.c` to produce object files `file1.o` and `file3.o`; after that, all object files, including `file2.o`, are submitted to the linker to produce one executable.


```
xlc file1.c file2.o file3.c
```

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
```

Related information

 [Linking](#)

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

You can use XL C 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 might perform better than statically linked programs if several programs use the same shared routines at the same time. By using dynamic linking, you can upgrade the routines in the shared libraries without relinking. This form of linking is the default and no additional options are needed.

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 might 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.

Note: Dynamically and statically linked programs might 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 and linked, 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 with runtime arguments on the command line.

Canceling execution

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

To cancel a running program, press **Ctrl+C** 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. Some environment variables do not control actual runtime behavior, but they can have an impact on how your applications run.

For more information about 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 can run with more recent versions of AIX. However, applications linked on a system using a newer version of AIX might 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.

At compile time, you can use the `-g` or `-qlinedebug` option to instruct the XL C compiler to include debugging information in compiled output. For `-g`, you can also use different levels to balance between debug capability and compiler optimization. For more information about the 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 you debug your applications. If you need to debug an optimized application, you can consider using the **-gN** form of the **-g** option along with any optimization options. This form of the **-g** option provides different levels of tradeoff between full optimization and full debugging support, depending on the value of N. For more information about debugging your optimized code, see "Debugging optimized code" in the *XL C Optimization and Programming Guide*.

Determining which level of XL C is being used

To display the version and release level of XL C that you are using, invoke the compiler with the **-qversion** compiler option.

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

```
xlc -qversion=verbose
```

Notices

Programming interfaces: Intended programming interfaces allow the customer to write programs to obtain the services of IBM XL C for AIX.

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, MD-NC119
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:

Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan, Ltd.
19-21, Nihonbashi-Hakozakicho, Chuo-ku
Tokyo 103-8510, 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 want 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:

Intellectual Property Dept. for Rational Software
IBM Corporation
5 Technology Park Drive
Westford, MA 01886
U.S.A.

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. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

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, 2015.

PRIVACY POLICY CONSIDERATIONS:

IBM Software products, including software as a service solutions, ("Software Offerings") may use cookies or other technologies to collect product usage information, to help improve the end user experience, or to tailor interactions with the end user, or for other purposes. In many cases no personally identifiable information is collected by the Software Offerings. Some of our Software Offerings can help enable you to collect personally identifiable information. If this Software Offering uses cookies to collect personally identifiable information, specific information about this offering's use of cookies is set forth below.

This Software Offering does not use cookies or other technologies to collect personally identifiable information.

If the configurations deployed for this Software Offering provide you as customer the ability to collect personally identifiable information from end users via cookies and other technologies, you should seek your own legal advice about any laws applicable to such data collection, including any requirements for notice and consent.

For more information about the use of various technologies, including cookies, for these purposes, see IBM's Privacy Policy at <http://www.ibm.com/privacy> and IBM's Online Privacy Statement at <http://www.ibm.com/privacy/details> in the section entitled "Cookies, Web Beacons and Other Technologies," and the "IBM Software Products and Software-as-a-Service Privacy Statement" at <http://www.ibm.com/software/info/product-privacy>.

Trademarks

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 and the Adobe 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.

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

Index

Special characters

- .a files 34
- .c files 34
- .i files 34
- .lst files 34
- .mod files 34
- .o files 34
- .s files 34
- .S files 34

Numerics

- 64-bit environment 6

A

- a.out file 34
- assembler
 - source (.s) files 34
 - source (.S) files 34

B

- basic example, described viii
- built-in functions 13, 16, 26

C

- code optimization 6
- commands 4
- compatibility with GNU 3
- compilation
 - sequence of activities 31
- compiler
 - invoking 32
 - running 32
- compiler directives
 - new or changed 14, 20, 26
- compiler options
 - conflicts and incompatibilities 33
 - new or changed 14, 20, 26
 - specification methods 33
- compiling
 - SMP programs 32
- customization
 - for compatibility with GNU 3

D

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

E

- editing source files 31
- executable files 34
- executing a program 35
- executing the linker 34

F

- files
 - editing source 31
 - input 34
 - output 34

I

- input files 34
- invocation commands 32
- invoking a program 35
- invoking the compiler 32

L

- language standards 3
- language support 3
- libraries 34
- linking
 - dynamic 35
 - static 35
- linking process 34
- listings 34

M

- migration
 - source code 33
- multiprocessor systems 7

O

- object files 34
 - creating 34
 - linking 34
- OpenMP 7
- optimization 22
 - programs 6
- output files 34

P

- parallelization 7
- performance 22
 - optimizing transformations 6
- problem determination 36
- programs
 - running 35

R

- running the compiler 32
- runtime environment 36
- runtime libraries 34
- runtime options 36

S

- shared memory parallelization 7
- shared object files 34
- SMP
 - programs, compiling 32
- SMP programs 7
- source files 34
- source-level debugging support 8
- static linking 35
- symbolic debugger support 8

T

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

U

- utilities
 - cleanpdf 5
 - CreateExportList 5
 - custom installation 6
 - gxc 5
 - IBM Debugger 4
 - install 6
 - mergepdf 5
 - showpdf 5
 - xlendi 6

X

- xl.cfg file 33



Product Number: 5765-J06; 5725-C71

Printed in USA

SC27-4237-02

