

IBM XL Fortran Advanced Edition V10.1 for Linux



# Getting Started with XL Fortran



IBM XL Fortran Advanced Edition V10.1 for Linux



# Getting Started with XL Fortran

**Note!**

Before using this information and the product it supports, be sure to read the general information under “Notices” on page 31.

**First Edition (September 2005)**

This edition applies to IBM® XL Fortran Advanced Edition V10.1 for Linux™ (Program 5724-M17 ) 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.

IBM welcomes your comments. You can send your comments electronically to the network ID listed below. Be sure to include your entire network address if you wish a reply.

- Internet: [compinfo@ca.ibm.com](mailto:compinfo@ca.ibm.com)

When you send information to IBM, you grant IBM a nonexclusive right to use or distribute the information in any way it believes appropriate without incurring any obligation to you.

© Copyright International Business Machines Corporation 1990, 2005. All rights reserved.

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

---

# Contents

<b>About this document . . . . .</b>	<b>v</b>
Who should read this document. . . . .	v
How to use this document. . . . .	v
How this document is organized . . . . .	v
Conventions and terminology used in this document	vi
Typographical conventions . . . . .	vi
How to read syntax diagrams . . . . .	vi
Examples. . . . .	viii
Related information . . . . .	viii
IBM XL Fortran publications . . . . .	viii
Additional documentation. . . . .	x
Technical support. . . . .	x
How to send your comments. . . . .	x

## **Chapter 1. Overview of XL Fortran features . . . . . 1**

Commonality with other XL compilers. . . . .	1
Documentation, online help, and technical support . . . . .	1
Hardware and operating system support . . . . .	1
Highly configurable compiler. . . . .	2
Language standards compliance . . . . .	3
Source-code migration and conformance checking	3
Program optimization . . . . .	3
64-bit object capability . . . . .	4
Shared memory parallelization . . . . .	4
OpenMP directives . . . . .	5
Diagnostic listings . . . . .	5
Symbolic debugger support . . . . .	5

## **Chapter 2. What's new for V10.1 . . . . . 7**

Performance and optimization . . . . .	7
Architecture and processor-specific code tuning. . . . .	7
High performance libraries . . . . .	7
Other performance-related compiler options and directives . . . . .	8
Intrinsic procedures new for this release . . . . .	9
Support for language enhancements and APIs. . . . .	10
XL Fortran language enhancements . . . . .	10
OpenMP API V2.5 support for C, C++, and Fortran . . . . .	11
Ease of use . . . . .	11
New installation and configuration utilities. . . . .	11
Newly-supported filename extensions . . . . .	11
Support for IBM Tivoli License Manager. . . . .	11
New compiler options. . . . .	12
New command line options . . . . .	12

## **Chapter 3. Setting up and customizing XL Fortran . . . . . 15**

Environment variables and XL Fortran . . . . .	15
Setting the compiler working environment . . . . .	15
Setting the default runtime options . . . . .	15
Customizing the configuration file. . . . .	16
Determining what level of XL Fortran is installed. . . . .	16

## **Chapter 4. Editing, compiling, and linking programs with XL Fortran . . . . . 17**

The compiler phases . . . . .	17
Editing Fortran source files . . . . .	17
Compiling with XL Fortran . . . . .	18
Compiling Fortran 95 or Fortran 90 programs . . . . .	18
Compiling parallelized XL Fortran applications . . . . .	19
XL Fortran input files . . . . .	19
XL Fortran output files . . . . .	20
Specifying compiler options . . . . .	21
Linking XL Fortran programs . . . . .	22
Compiling and linking in separate steps. . . . .	22
Dynamic and static linking . . . . .	22

## **Chapter 5. Running XL Fortran programs . . . . . 25**

Canceling execution . . . . .	25
Setting runtime options . . . . .	25
Running compiled applications on other systems. . . . .	25

## **Chapter 6. XL Fortran compiler diagnostic aids . . . . . 27**

Compilation return codes. . . . .	27
XL Fortran compiler listings. . . . .	27
Debugging compiled applications . . . . .	28

## **Chapter 7. XL Fortran runtime environment information . . . . . 29**

External names in the runtime environment . . . . .	29
External names in XL Fortran libraries . . . . .	30

## **Notices . . . . . 31**

Programming interface information . . . . .	32
Trademarks and service marks . . . . .	33

## **Index . . . . . 35**



---

## About this document

*Getting Started with XL Fortran* provides a general overview of the XL Fortran compiler, its more significant features, and how those features can help you improve your software development productivity.

For the benefit of current XL Fortran users upgrading to this release, *Getting Started with XL Fortran* also includes a summary of features that are new or improved for V10.1.

*Getting Started with XL Fortran* is intended only to help familiarize you with the compiler. For detailed information on using the XL Fortran compiler, you will want to refer to other books in the XL Fortran Advanced Edition V10.1 for Linux library, described in “IBM XL Fortran publications” on page viii.

---

## Who should read this document

*Getting Started with XL Fortran* is intended for anyone who plans to work with IBM® XL Fortran Advanced Edition V10.1 for Linux, who is familiar with the Linux® operating system, and who has some previous Fortran programming experience.

---

## How to use this document

If you are new to XL Fortran, you should view Chapter 1, “Overview of XL Fortran features,” on page 1 to familiarize yourself with the key features of XL Fortran and how to begin using it to develop your applications.

If you are already an experienced XL Fortran user and are now upgrading to the latest release of XL Fortran, you may want to go directly to Chapter 2, “What’s new for V10.1,” on page 7 to review that latest changes and feature enhancements to the compiler.

The remaining sections of this guide provide a brief overview of basic program development tasks with XL Fortran.

---

## How this document is organized

This guide includes these topics:

- Chapter 1, “Overview of XL Fortran features,” on page 1 outlines the the key features of the XL Fortran compiler
- Chapter 2, “What’s new for V10.1,” on page 7 describes new and updated features offered by the latest version of XL Fortran.
- Chapter 3, “Setting up and customizing XL Fortran,” on page 15 provides brief overview information on the steps involved in setting up and customizing XL Fortran, together with pointers on where you can find more detailed information.
- Chapter 4, “Editing, compiling, and linking programs with XL Fortran,” on page 17 discusses the basic steps involved in creating and compiling your applications with XL Fortran.

- Chapter 5, “Running XL Fortran programs,” on page 25 describes how to run your compiled applications, including setting of run time options.
- Chapter 6, “XL Fortran compiler diagnostic aids,” on page 27 offers guidance on how to use XL Fortran compiler diagnostic aids to identify and correct compilation problems with your applications.

---

## Conventions and terminology used in this document

### Typographical conventions

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

Table 1. *Typographical conventions*

Typeface	Indicates	Example
<b>bold</b>	Commands, executable names, and compiler options.	By default, if you use the <b>-qsmp</b> compiler option in conjunction with one of these invocation commands, the option <b>-qdirective=IBM*:SMP\$:SOMP:IBMP:IBMT</b> will be on.
<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.	The maximum length of the <i>trigger_constant</i> in fixed source form is 4 for directives that are continued on one or more lines.
UPPERCASE	Fortran programming keywords, statements, directives, and intrinsic procedures.	The ASSERT directive applies only to the DO loop immediately following the directive, and not to any nested DO loops.
monospace	Programming keywords and library functions, compiler built-in functions, file and directory names, examples of program code, command strings, or user-defined names.	If you call <code>omp_destroy_lock</code> with an uninitialized lock variable, the result of the call is undefined.

### How to read syntax diagrams

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

If a variable or user-specified name ends in *\_list*, you can provide a list of these terms separated by commas.


You must enter punctuation marks, parentheses, arithmetic operators, and other special characters as part of the syntax.

- Read syntax diagrams from left to right and from top to bottom, following the path of the line:
  - The **▶** symbol indicates the beginning of a statement.
  - The **→** symbol indicates that the statement syntax continues on the next line.
  - The **▶** symbol indicates that a statement continues from the previous line.
  - The **→▶** symbol indicates the end of a statement.

- Program units, procedures, constructs, interface blocks and derived-type definitions consist of several individual statements. For such items, a box encloses the syntax representation, and individual syntax diagrams show the required order for the equivalent Fortran statements.
- IBM and Fortran 95 extensions are marked by a number in the syntax diagram with an explanatory note immediately following the diagram.
- Required items are shown on the horizontal line (the main path):




- Optional items are shown below the main path:




**Note:** Optional items (not in syntax diagrams) are enclosed by square brackets ([ and ]). For example, [UNIT=]u


- 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 repeat an item, and the separator character if it is other than a blank:



A repeat arrow above a stack indicates that you can make more than one choice from the items in the stack.



The following is an example of a syntax diagram with an interpretation:



Interpret the diagram as follows:

- Enter the keyword EXAMPLE.
- EXAMPLE is an IBM extension.
- 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.
- 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. (The *\_list* syntax is equivalent to the previous syntax for *e*.)

## Examples

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

---

## Related information

### IBM XL Fortran publications

XL Fortran provides product documentation in the following formats:

- Readme files

Readme files contain late-breaking information, including changes and corrections to the product documentation. Readme files are located by default in the `/opt/ibmcmp/xf/10.1/` directory and in the root directory of the installation CD.

- Installable man pages

Man pages are provided for the compiler invocations and all command-line utilities provided with the product. Instructions for installing and accessing the man pages are provided in the *IBM XL Fortran Advanced Edition V10.1 for Linux Installation Guide*.

- Information center

The information center of searchable HTML files can be launched on a network and accessed remotely or locally. Instructions for installing and accessing the information center are provided in the *IBM XL Fortran Advanced Edition V10.1 for Linux Installation Guide*. The information center is also viewable on the Web at:

[publib.boulder.ibm.com/infocenter/lxpcmp/index.jsp](http://publib.boulder.ibm.com/infocenter/lxpcmp/index.jsp)

- PDF documents

PDF documents are located by default in the `/opt/ibmcmp/xf/10.1/doc/language/PDF/` directory, and are also available on the Web at:

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

In addition to this document, the following files comprise the full set of XL Fortran product manuals:

*Table 2. XL Fortran PDF files*

Document title	PDF file name	Description
<i>IBM XL Fortran Advanced Edition V10.1 for Linux Installation Guide</i> , GC09-8020-00	install.pdf	Contains information for installing XL Fortran and configuring your environment for basic compilation and program execution.
<i>IBM XL Fortran Advanced Edition V10.1 for Linux Compiler Reference</i> , GC09-8019-00	cr.pdf	Contains information about the various compiler options and environment variables.
<i>IBM XL Fortran Advanced Edition V10.1 for Linux Language Reference</i> , GC09-8018-00	lr.pdf	Contains information about the Fortran programming language as supported by IBM, including language extensions for portability and conformance to non-proprietary standards, compiler directives and intrinsic procedures.
<i>IBM XL Fortran Advanced Edition V10.1 for Linux Optimization and Programming Guide</i> , GC09-8022-00	opg.pdf	Contains information on advanced programming topics, such as application porting, interlanguage calls, floating-point operations, input/output, application optimization and parallelization, and the XL Fortran high-performance libraries.

These PDF files are viewable and printable from Adobe Reader. If you do not have the Adobe Reader installed, you can download it from:

[www.adobe.com](http://www.adobe.com)

## Additional documentation

More documentation related to XL Fortran, including redbooks, whitepapers, tutorials, and other articles, is available on the Web at:

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

---

## Technical support

Additional technical support is available from the XL Fortran Support page. This page provides a portal with search capabilities to a large selection of technical support FAQs and other support documents. You can find the XL Fortran Support page on the Web at:

<http://www.ibm.com/software/awdtools/fortran/xlfortran/support>

If you cannot find what you need, you can e-mail:

[compinfo@ca.ibm.com](mailto:compinfo@ca.ibm.com)

For the latest information about XL Fortran, visit the product information site at:

<http://www.ibm.com/software/awdtools/fortran/xlfortran>

---

## How to send your comments

Your feedback is important in helping to provide accurate and high-quality information. If you have any comments about this document or any other XL Fortran documentation, send your comments by e-mail to:

[compinfo@ca.ibm.com](mailto:compinfo@ca.ibm.com)

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

---

## Chapter 1. Overview of XL Fortran features

XL Fortran Advanced Edition V10.1 for Linux can be used for large, complex, computationally intensive programs, including interlanguage calls with C and C++ programs. This section discusses the features of the XL Fortran compiler at a high level. It is intended for people who are evaluating XL Fortran and for new users who want to find out more about the product.

---

### Commonality with other XL compilers

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

The XL compilers are part of a larger family of IBM C, C++, and Fortran compilers that are derived from a common code base that shares compiler function and optimization technologies among a variety of platforms and programming languages, such as AIX, Linux distributions, OS/390, OS/400, z/OS, and z/VM operating systems. The common code base, along with compliance to international programming language standards, helps ensure consistent compiler performance and ease of program portability across multiple operating systems and hardware platforms.

The XL compilers are available for use on AIX and selected Linux distributions.

---

### Documentation, online help, and technical support

This guide provides an overview of XL Fortran and its features. You can also find more extensive product documentation in the following formats:

- Readme files.
- Installable man pages.
- An HTML-based information system.
- Portable Document Format (PDF) documents.
- Online technical support over the Web.

For more information about product documentation and technical support provided with XL Fortran, see:

- “IBM XL Fortran publications” on page viii
- “Additional documentation” on page x
- “Technical support” on page x

---

### Hardware and operating system support

XL Fortran Advanced Edition V10.1 for Linux supports several Linux distributions. See the README file and System prerequisites in the *XL Fortran Advanced Edition V10.1 for Linux Installation Guide* for a complete list of supported distributions and requirements.

The compiler, its libraries, and its generated object programs will run on all POWER3™, POWER4™, POWER5™, POWER5+™, PowerPC®, and PowerPC 970 systems with the required software and disk space.

To take maximum advantage of different hardware configurations, the compiler provides a number of options for performance tuning based on the configuration of the machine used for executing an application.

---

## Highly configurable compiler

XL Fortran offers you a wealth of features to let you tailor the compiler to your own unique compilation requirements.

### Compiler invocation commands

XL Fortran provides several different commands that you can use to invoke the compiler, for example, `xl`, `xlF95`, and `xlF90`. Each invocation command is unique in that it instructs the compiler to tailor compilation output to meet a specific language level specification. Compiler invocation commands are provided to support all standardized Fortran language levels, and many popular language extensions as well.

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

For more information about XL Fortran compiler invocation commands, see "Compiling with XL Fortran" on page 18 in this book or *Compiling XL Fortran programs in the XL Fortran Advanced Edition V10.1 for Linux Compiler Reference*.

### Compiler options

You can control the actions of the compiler through a large set of provided compiler options. Different categories of options help you to debug your applications, optimize and tune application performance, select language levels and extensions for compatibility with programs from other platforms, and do many other common tasks that would otherwise require changing the source code.

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

For more information about XL Fortran compiler options, see *Compiler-options reference in the XL Fortran Advanced Edition V10.1 for Linux Compiler Reference*.

### Custom compiler configuration files

The installation process creates a default compiler configuration file at `/etc/opt/ibmcmp/xlf/10.1/xlf.cfg`. This configuration file contains several stanzas that define compiler option default settings.

Your compilation needs may frequently call for specifying compiler option settings other than the defaults settings provided by XL Fortran. If so, XL Fortran provides the `xl_configure` utility that you can use to create additional configuration files. You can then modify these files with any text editor to contain your own frequently-used compiler option settings.

See *Customizing the configuration file in the XL Fortran Advanced Edition V10.1 for Linux Compiler Reference* for more information on creating and using custom configuration files.

---

## Language standards compliance

The compiler supports the following programming language specifications for Fortran:

- ISO/IEC 1539-1:1991(E) and ANSI X3.198-1992 (referred to as Fortran 90 or F90)
- ISO/IEC 1539-1:1997 (referred to as Fortran 95 or F95)
- Extensions to the Fortran 95 standard:
  - Industry extensions that are found in Fortran products from various compiler vendors
  - Extensions specified in SAA Fortran
- Partial support of the Fortran 2003 standard

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

- OpenMP extensions to support parallelized programming.
- Language extensions to support VMX vector programming.

See Language standards in the *XL Fortran Advanced Edition V10.1 for Linux Language Reference* for more information about Fortran language specifications and extensions.

## Source-code migration and conformance checking

XL Fortran helps protect your investment in your existing Fortran source code by providing compiler invocation commands that instruct the compiler to inspect your application for conformance to a specific language level and warn you if it finds constructs and keywords that do not conform to the specified language level. You can also use the `-qlanglvl` compiler option to specify a given language level, and the compiler will issue warnings if language elements in your program source do not conform to that language level. Additionally, you can name your source files with common filename extensions such as `.f77`, `.f90`, or `.f95`, then use the generic compiler invocations such as `xlfc` or `xlfc_r` to automatically select the appropriate language-level appropriate to the filename extension.

See `-qlanglvl` in the *XL Fortran Advanced Edition V10.1 for Linux Compiler Reference* for more information.

---

## Program optimization

XL Fortran provides several compiler options that can help you control the optimization of your programs. With these options, you can:

- 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 performance at run time. Fortran provides a portfolio of optimizing transformations tailored to various supported hardware. These transformations can:

- Reduce the number of instructions executed for critical operations.
- Restructure generated object code to make optimal use of the PowerPC architecture.
- Improve the usage of the memory subsystem.

- Exploit the ability of the architecture to handle large amounts of shared memory parallelization.

Significant performance improvements are possible with relatively little development effort because the compiler is capable of sophisticated program analysis and transformation. Moreover, XL Fortran enables programming models, such as OpenMP, which allow you to write high-performance code.

If possible, you should test and debug your code without optimization before attempting to optimize it.

For more information about optimization techniques, see *Optimizing XL compiler applications in the XL Fortran Advanced Edition V10.1 for Linux Optimization and Programming Guide*.

For a summary of optimization-related compiler options, see *Options for performance optimization in the XL Fortran Advanced Edition V10.1 for Linux Compiler Reference*.

---

## 64-bit object capability

The XL Fortran compiler's 64-bit object capability addresses increasing demand for larger storage requirements and greater processing power. The Linux operating system provides an environment that allows you to develop and execute programs that exploit 64-bit processors through the use of 64-bit address spaces.

To support larger executables that can be fit within a 64-bit address space, a separate, 64-bit object form is used to meet the requirements of 64-bit executables. The linker binds 64-bit objects to create 64-bit executables. Note that objects that are bound together must all be of the same object format. The following scenarios are not permitted and will fail to load, or 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 Fortran 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 XL Fortran in a 64-Bit Environment in the XL Fortran Advanced Edition V10.1 for Linux Compiler Reference*.

---

## Shared memory parallelization

XL Fortran Advanced Edition V10.1 for Linux supports application development for multiprocessor system architectures. You can use any of the following methods to develop your parallelized applications with XL Fortran:

- Directive-based shared memory parallelization (OpenMP)

- 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()`

For more information, see Parallel programming with XL Fortran in the *XL Fortran Advanced Edition V10.1 for Linux Optimization and Programming Guide*.

## OpenMP directives

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

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

OpenMP directives address three important issues of parallelizing an application:

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

XL Fortran supports the OpenMP API Version 2.5 specification. For more information, see [www.openmp.org](http://www.openmp.org).

---

## Diagnostic listings

The compiler output listing has optional sections that you can include or omit. For more information about the applicable compiler options and the listing itself, refer to “XL Fortran compiler listings” on page 27.

---

## Symbolic debugger support

You can use `gdb` or any other symbolic debugger when debugging your programs.



---

## Chapter 2. What's new for V10.1

The new features and enhancements in XL Fortran Advanced Edition V10.1 for Linux fall into four categories:

- “Performance and optimization”
- “Support for language enhancements and APIs” on page 10
- “Ease of use” on page 11
- “New compiler options” on page 12

---

### Performance and optimization

Many new features and enhancements fall into the category of optimization and performance tuning.

#### Architecture and processor-specific code tuning

The **-qarch** compiler option controls the particular instructions that are generated for the specified machine architecture. The **-qtune** compiler option adjusts the instructions, scheduling, and other optimizations to enhance performance on the specified hardware. These options work together to generate application code that gives the best performance for the specified architecture.

XL Fortran V10.1 augments the list of suboptions available to the **-qarch** compiler option to support processors that support the VMX instruction set and the newly-available POWER5+ processors. The following new **-qarch** options are available:

- **-qarch=ppc64v**
- **-qarch=pwr5x**

#### High performance libraries

XL Fortran includes highly-tuned mathematical functions that can greatly improve the performance of mathematically-intensive applications. These functions are provided through the following high-performance libraries:

##### Mathematical Acceleration Subsystem (MASS)

MASS libraries provide high-performance scalar and vector functions to perform common mathematical computations. The MASS libraries included with XL Fortran Advanced Edition V10.1 for Linux introduce new scalar and vector functions, and new support for the POWER5 processor architecture.

For more information about using the MASS libraries, see *Using the Mathematical Acceleration Subsystem in the XL Fortran Advanced Edition V10.1 for Linux Optimization and Programming Guide*.

##### Basic Linear Algebra Subprograms (BLAS)

XL Fortran Advanced Edition V10.1 for Linux introduces the BLAS set of high-performance algebraic functions. 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 Fortran Advanced Edition V10.1 for Linux Optimization and Programming Guide*.

## Other performance-related compiler options and directives

The entries in the following table describes new or changed compiler options and directives not already mentioned in the sections above.

Information presented here is just a brief overview. For more information about these compiler options, refer to Options for performance optimization in the *XL Fortran Advanced Edition V10.1 for Linux Compiler Reference*.

Table 3. Other Performance-Related Compiler Options and Directives

Option/directive	Description
<b>-qfloat</b>	<p><b>-qfloat</b> adds the following new suboptions:</p> <p><b>-qfloat=relax</b> This suboption relaxes strict-IEEE conformance in exchange for greater speed, typically by removing trivial floating-point arithmetic operations such as adds and subtracts involving a zero on the right.</p> <p><b>-qfloat=norelax</b> This is the default. Strict IEEE conformance is maintained.</p>
<b>-qipa</b>	<p><b>-qipa</b> adds the following new suboptions:</p> <p><b>-qipa=clonearch=arch{,arch}</b> Specifies one or more processor architectures for which multiple versions of the same instruction set are produced.</p> <p>XL Fortran lets you specify multiple specific processor architectures for which instruction sets will be generated. At run time, the application will detect the specific architecture of the operating environment and select the instruction set specialized for that architecture.</p> <p><b>-qipa=cloneproc=name{name}</b> Specifies the names of one or more functions to clone for the processor architectures specified by the <b>clonearch</b> suboption.</p>
<b>-O</b>	<p>Specifying the <b>-O3</b> compiler option now instructs the compiler to also assume the <b>-qhot=level=0</b> compiler option setting.</p> <p>Specifying the <b>-O4</b> or <b>-O5</b> compiler option now instructs the compiler to also assume the <b>-qhot=level=1</b> compiler option setting.</p>

Table 3. Other Performance-Related Compiler Options and Directives (continued)

Option/directive	Description
<b>-qsmallstack</b>	<p>The <b>-qsmallstack</b> compiler option adds the following new suboptions to control dynamic length variable allocation transformations:</p> <p><b>-qsmallstack=dynlenonheap</b>                      When this suboption is specified, certain automatically-sized objects are allocated from the heap. This suboption affects automatic objects that have nonconstant character lengths or a nonconstant array bound (DYNAMIC LENGTH ON HEAP). Specifying this suboption turns on both dynlenonheap and general smallstack transformations.</p> <p><b>-qsmallstack=nodynlenonheap</b>                      This is the default. If this suboption is not specified, those objects are allocated on the stack. This suboption affects automatic objects that have nonconstant character lengths or a nonconstant array bound (DYNAMIC LENGTH ON HEAP).</p>
<b>-qstacktemp</b>	<p>The <b>-qstacktemp</b> compiler option is new, and gives you the ability to control where certain compiler temporaries are stored. Available suboptions are:</p> <p><b>-qstacktemp=0</b>                      This is the default. Certain compiler temporaries are allocated to the heap instead of the stack at compiler's discretion, depending on the size of the compiler temporaries and the target operating system environment.</p> <p><b>-qstacktemp= -1</b>                      Certain compiler temporaries are always allocated on the stack, providing best performance but also using the most amount of stack space.</p> <p><b>-qstacktemp=num_bytes</b>                      Certain compiler temporaries less than <i>num_bytes</i> in size are allocated to the stack. Compiler temporaries greater than or equal to <i>num_bytes</i> are allocated to the heap.</p> <p>Programs that use large arrays may to use this option if they are running out of stack space at run time. SMP or OpenMP applications that are constrained by stack space may also find this option useful to move some compiler temporaries onto the heap from the stack.</p>

## Intrinsic procedures new for this release

The following table lists intrinsic procedures that are new for this release. For more information on intrinsic procedures provided by XL Fortran, see Intrinsic procedures in the *XL Fortran Advanced Edition V10.1 for Linux Language Reference*.

Table 4. Intrinsic procedures for XL Fortran

Function	Description
FRIM( <i>val</i> );	Takes an input <i>val</i> of REAL *8 format, rounds <i>val</i> down to the next lower integral value, and returns the result in REAL *8 format. Valid only for POWER5+ processors.
FRIMS( <i>val</i> );	Takes an input <i>val</i> in REAL *4 format, rounds <i>val</i> down to the next lower integral value, and returns the result in REAL *4 format. Valid only for POWER5+ processors.

Table 4. Intrinsic procedures for XL Fortran (continued)

Function	Description
FRIN( <i>val</i> );	Takes an input <i>val</i> in REAL *8 format, rounds <i>val</i> to the nearest integral value, and returns the result in REAL *8 format. Valid only for POWER5+ processors.
FRINS( <i>val</i> );	Takes an input <i>val</i> in REAL *4 format, rounds <i>val</i> to the nearest integral value, and returns the result in REAL *4 format. Valid only for POWER5+ processors.
FRIP( <i>val</i> );	Takes an input <i>val</i> in REAL *8 format, rounds <i>val</i> up to the next higher integral value, and returns the result in REAL *8 format. Valid only for POWER5+ processors.
FRIPS( <i>val</i> );	Takes an input <i>val</i> in REAL *4 format, rounds <i>val</i> up to the next higher integral value, and returns the result in REAL *4 format. Valid only for POWER5+ processors.
FRIZ( <i>val</i> );	Takes an input <i>val</i> in REAL *8 format, rounds <i>val</i> to the next integral value closest to zero, and returns the result in REAL *8 format. Valid only for POWER5+ processors.
FRIZS( <i>val</i> );	Takes an input <i>val</i> in REAL *4 format, rounds <i>val</i> to the next integral value closest to zero, and returns the result in REAL *4 format. Valid only for POWER5+ processors.

In addition, this release of XL Fortran features several new intrinsic procedures and a new VECTOR data type to support VMX vector programming. For more information about these new intrinsic procedures, see VMX intrinsic procedures in the *XL Fortran Advanced Edition V10.1 for Linux Language Reference*.

---

## Support for language enhancements and APIs

API and language enhancements can offer you additional ease of use and flexibility when developing your applications, as well as making it easier for you to develop code that more fully exploits the capabilities of your hardware platform.

### XL Fortran language enhancements

XL Fortran V10.1 implements new features compliant to the Fortran 2003 standard. These features, supported when `-qlanglvl=2003std` or `-qlanglvl=2003pure` is in effect, include:

#### ENUM statement

You can specify an ENUM statement to define and group a set of named integer constants called enumerators. The storage size of the enumerators can be set by the new `-qenum` compiler option.

#### READ statement - new BLANK and PAD specifiers

The BLANK specifier controls the default interpretation of blanks when you are using a format specification. The setting of the PAD specifier determines if input records are padded with blanks.

#### WRITE statement - new DELIM specifier

The DELIM specifier specifies what delimiter, if any, is used to delimit character constants written with list-directed or namelist formatting.

#### Relaxed rules for specification expression

XL Fortran now allows recursive functions to be specification functions.

### Procedure pointers

This release adds support for procedure pointers. A procedure pointer is a PROCEDURE entity that has the EXTERNAL and POINTER attributes. It may be a pointer associated with an external procedure, a module procedure, an intrinsic procedure, or a dummy procedure that is not a procedure pointer.

See *IBM XL Fortran Advanced Edition V10.1 for Linux Language Reference* for more information.

## OpenMP API V2.5 support for C, C++, and Fortran

XL Fortran now supports the OpenMP API V2.5 standard. This latest level of the OpenMP specification combines the previous C/C++ and Fortran OpenMP specifications into one single specification for both C/C++ and Fortran, and resolves previous inconsistencies between them.

The OpenMP Application Program Interface (API) is a portable, scalable programming model that provides a standard interface for developing user-directed shared-memory parallelization in C, C++, and Fortran applications. The specification is defined by the OpenMP organization, a group of computer hardware and software vendors, including IBM.

You can find more information about OpenMP specifications at:

[www.openmp.org](http://www.openmp.org)

---

## Ease of use

XL Fortran includes the following new features to help you more easily use the compiler for your application development.

### New installation and configuration utilities

This release of XL Fortran introduces the `xlfc_install` and `new_install` utilities to help you easily install and configure the compiler for initial use on your system.

### Newly-supported filename extensions

XL Fortran Advanced Edition V10.1 for Linux adds support for the `.f77`, `.f90`, and `.f95` filename extensions.

You can use generic XL Fortran compiler invocations, such as `xlfc` and `xlfc_r`, to compile program source files with these filename extensions. The compiler will recognize the filename extensions and apply the appropriate language standard defaults as if you were compiling using the `f77`, `xlfc90`, or `xlfc95` compiler invocations and their associated `_r` counterparts.

### Support for IBM Tivoli License Manager

IBM Tivoli License Manager (ITLM) is a Web-based solution that can help you manage software usage metering and license allocation services on supported systems. In general, ITLM recognizes and monitors the products that are installed and in use on your system.

IBM XL Fortran Advanced Edition V10.1 for Linux is ITLM-enabled for inventory and usage signature support, which means that ITLM is able to detect both product installation of XL Fortran and its usage.

**Note:** ITLM is not a part of the XL Fortran compiler offering, and must be purchased and installed separately.

Once installed and activated, ITLM scans your system for product inventory signatures that indicate whether a given product is installed on your system. ITLM also identifies that product's version, release, and modification levels. Signature files for XL Fortran are installed to the following directory:

**Default installations**

/opt/ibmcmp/xlf/10.1

**Non-default installations**

*compiler*/xlf/10.1 where *compiler* is the target directory for installation specified by the **--prefix** installation option.

For more information about IBM Tivoli License Manager Web, see:

[www.ibm.com/software/tivoli/products/license-mgr](http://www.ibm.com/software/tivoli/products/license-mgr)

---

## New compiler options

Compiler options can be specified on the command line or through directives embedded in your application source files.

## New command line options

The following table summarizes command line options new to XL Fortran. You can find detailed syntax and usage information for all compiler options in Compiler options reference in the *XL Fortran Advanced Edition V10.1 for Linux Compiler Reference*.

Option	Description and remarks
<b>-qenum</b>	The <b>-qenum</b> compiler option specifies the amount of storage used by enumerators defined with the ENUM statement.
<b>-qlanglvl</b>	This release adds the <b>2003std</b> and <b>2003pure</b> suboptions to the <b>-qlanglvl</b> compiler option.  <b>2003std</b> Accepts the language that the ISO Fortran 95 standard specifies, as well as all Fortran 2003 features supported by XL Fortran, and reports anything else as an error.  <b>2003pure</b> The same as 2003std except that it also reports errors for any obsolescent Fortran 2003 features used.
<b>-qlibansi</b>	This option is now recognized by the entire compiler, and not just by the IPA optimizer. It instructs the compiler to assume that all functions with the name of an ANSI C defined library function are in fact the library functions.
<b>-qlinedebug</b>	This new compiler option enables minimal generation (line number and source file name) of information for use by a debugger. This compiler option can be specified on the command line or in your program source code as a @PROCESS statement.

Option	Description and remarks
<b>-qlist</b>	The <b>-qlist</b> compiler option adds new <b>offset</b> and <b>nooffset</b> suboptions. Specifying <b>-qlist=offset</b> instructs the compiler to show object listing offsets from the start of a function rather than from the start of code generation.
<b>-qport</b>	This release adds <b>logicals</b> and <b>noclogicals</b> as new suboptions to the <b>-qport</b> compiler option. Specifying <b>-qport=logicals</b> together with <b>-qintlog</b> instructs the compiler to treat all non-zero integers used in logical expressions as TRUE. This option is useful when porting applications from other Fortran compilers that expect this behavior.



---

## Chapter 3. Setting up and customizing XL Fortran

This section provides brief overview information about setting up and customizing XL Fortran, together with pointers to other documentation that describes specific set-up and customization topics in greater detail.

---

### Environment variables and XL Fortran

XL Fortran uses a number of environment variables to control various aspects of compiler operation. Environment variables fall into two basic categories:

- Environment variables defining the basic working environment for the compiler.
- Environment variables defining run time compiler option defaults.

#### Setting the compiler working environment

These environment variables define the basic working environment for the compiler, including specifying your choice of national language or defining the location of libraries or temporary files. For complete information, refer to Correct settings for environment variables in the *XL Fortran Advanced Edition V10.1 for Linux Compiler Reference*.

##### LANG

Specifies the default national language *locale* used to display diagnostic messages and compiler listings. This environment variable also affects runtime behavior.

##### MANPATH

Specifies the search path for system, compiler, and third-party man pages.

##### NLSPATH

Specifies one or more directory locations where message catalogs can be found. This environment variable also affects runtime behavior.

##### PDFDIR

Specifies the directory location where profile-directed feedback information is stored when you compile with the **-qpdf** option.

##### TMPDIR

Specifies the directory location where the compiler will store temporary files created during program compilation. This environment variable also affects runtime behavior.

#### Setting the default runtime options

These environment variables define runtime compiler option defaults to be used by the compiler, unless explicitly overridden by compiler option settings specified on the command line or in directives located in your program source.

##### XL\_NOCLONEARCH

Instructs the program to execute only generic code, where generic code is compiled object code that is not versioned for a specific processor architecture. You can set the XL\_NOCLONEARCH environment variable to help you debug your application.

##### XLFRTEOPTS

The XLFRTEOPTS environment variable allows you to specify options that affect I/O, EOF error-handling, and the specification of random-number

generators. See Setting run time options in the XL Fortran Advanced Edition V10.1 for Linux Compiler Reference for more information.

**XLFSCRATCH\_unit**

Used to give a specific name to a scratch file.

**XLFUNIT\_unit**

Used to give a specific name to an implicitly connected file or a file opened with no **FILE=** specifier.

**XLSMPOPTS**

The **XLSMPOPTS** environment variable allows you to specify run time options that affect SMP execution. See Setting OMP and SMP runtime options in the XL Fortran Advanced Edition V10.1 for Linux Optimization and Programming Guide for more information.

**OMP\_DYNAMIC, OMP\_NESTED, OMP\_NUM\_THREADS, OMP\_SCHEDULE**

These environment variables, are part of the OpenMP standard. They let you specify how the application will execute sections of parallel code. See Setting OMP and SMP runtime options in the XL Fortran Advanced Edition V10.1 for Linux Optimization and Programming Guide for more information.

---

## Customizing the configuration file

The configuration file is a plain text file that specifies default settings for compiler options and invocations. XL Fortran provides a default configuration at file `/etc/opt/ibmcomp/xf/10.1/xf.cfg` during compiler installation.

If you are running on a single-user system, or if you already have a compilation environment with compilation scripts or makefiles, you may want to leave the default configuration file as it is.

As an alternative, you can create additional custom configuration files to meet special compilation requirements demanded by specific applications or groups of applications.

See Customizing the configuration file in the *XL Fortran Advanced Edition V10.1 for Linux Compiler Reference* for more information on creating and using custom configuration files.

---

## Determining what level of XL Fortran is installed

You may not be sure which level of XL Fortran is installed on a particular machine. You will need to know this information if contacting software support.

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

```
xf -qversion
```

---

## Chapter 4. Editing, compiling, and linking programs with XL Fortran

Basic Fortran program development consists of repeating cycles of editing, compiling and linking (by default a single step combined with compiling), and running.

### Prerequisite Information:

1. Before you can use the compiler, you must first ensure that all Linux settings (for example, certain environment variables and storage limits) are correctly configured. For more information see “Environment variables and XL Fortran” on page 15.
2. To learn more about writing Fortran programs, refer to the *XL Fortran Advanced Edition V10.1 for Linux Language Reference*.

---

### The compiler phases

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

1. A preprocessor
2. The compiler, which consists of the following phases:
  - a. Front-end parsing and semantic analysis
  - b. Loop transformations
  - c. Interprocedural analysis
  - d. Optimization
  - e. Register allocation
  - f. Final assembly
3. The assembler (for `.s` files and for `.S` files after they are preprocessed)
4. The linker `ld`

To see the compiler step through these phases, specify the `-qphsinfo` and `-v` compiler options when you compile your application.

---


### Editing Fortran source files

To create Fortran source programs, you can use any of the available text editors, such as `vi` or `emacs`. Source programs must use a recognized filename suffix unless the configuration file defines additional non-standard filename suffixes. See “XL Fortran input files” on page 19 for a list of filename suffixes recognized by XL Fortran.

For a Fortran source program to be a valid program, it must conform to the language definitions specified in the *XL Fortran Advanced Edition V10.1 for Linux Language Reference*.

## Compiling with XL Fortran

To compile a source program, use one of the compiler invocation commands with the syntax shown below:



The compiler invocation command performs all necessary steps to compile Fortran source files, assemble any `.s` and `.S` files, and link the object files and libraries into an executable program.

When working with source files whose filename extensions indicates a specific level of Fortran, such as `.f95`, `.f90`, or `f77`, compiling with the `xl` or `threadsafe` counterpart invocations will cause the compiler to automatically select the appropriate language-level defaults. The other base compiler invocation commands exist primarily to provide explicit compilation support for different levels and extensions of the Fortran language.

In addition to the base compiler invocation commands, XL Fortran also provides specialized variants of many base compiler invocations. A variation on a base compiler invocation is named by attaching a suffix to the name of that invocation command. Suffix meanings for invocation variants are:

`_r`      Threadsafe invocation variant that supports POSIX Pthread APIs for multithreaded applications, including applications compiled with `-qsmp` or with source code containing OpenMP program parallelization directives.

Table 5. XL Fortran compiler invocation commands

Base Invocation	Available Invocation Variants	Description
<code>xl</code> <code>f77</code> <code>fort77</code>	<code>xl_r</code>	Invokes the compiler so that source files are compiled as FORTRAN 77 source code.
<code>xl90</code> <code>f90</code>	<code>xl90_r</code>	Invokes the compiler so that source files are compiled as Fortran 90 source code.
<code>xl95</code> <code>f95</code>	<code>xl95_r</code>	Invokes the compiler so that source files are compiled as Fortran 95 source code.

## Compiling Fortran 95 or Fortran 90 programs

Use the following invocations (or their variants) to conform more closely to their corresponding Fortran language standards:

`f95`, `xl95`      Fortran 95

`f90`, `xl90`      Fortran 90

These are the preferred compiler invocation commands that you should use when creating and compiling new applications.

They all accept Fortran 90 free source form by default. To use fixed source form with these invocations, you must specify the `-qfixed` command line option.

I/O formats are slightly different between these commands and the other commands. I/O formats for the `xlF95` invocation are also different from those of `xlF90`. We recommend that you switch to the Fortran 95 formats for data files whenever possible.

By default, those invocation commands do not conform completely to their corresponding Fortran language standards. If you need full compliance, compile with the following additional compiler option settings:

```
-qnodirective -qnoescape -qextname -qfloat=nomaf:nofold  
-qnoswapomp -qlanglvl=90std -qlanglvl=95std
```

Also, specify the following runtime options before running the program, with a command similar to the following:

```
export XLF RTEOPTS="err_recovery=no:langlvl=90std"
```

The default settings are intended to provide the best combination of performance and usability, so you should change them only when absolutely required. Some of the options mentioned above are only required for compliance in very specific situations. For example, you would need to specify `-qextname` when an external symbol, such as a common block or subprogram, is named `main`.

## Compiling parallelized XL Fortran applications

XL Fortran provides threadsafe compilation invocations that you can use when compiling parallelized applications for use in multiprocessor environments.

- `xlF95_r`
- `xlF90_r`

These invocations are similar to their corresponding base compiler invocations, except that they link and bind compiled objects to threadsafe components and libraries.

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

### POSIX Pthreads API support

XL Fortran supports thread programming with the IEEE 1003.1-2001 (POSIX) standard Pthreads API.

---

## XL Fortran input files

The input files to the compiler are:

### Source files (.f .F ,f77 .F77 .f95 .f90 .F90 .F95 suffixes)

The compiler considers files with these suffixes as being Fortran source files for compilation.

The compiler compiles source files in the order you specify on the command line. If it cannot find a specified source file, the compiler produces an error message and proceeds to the next file, if one exists.

Files with a suffix of `.F`, `F77`, `F95`, or `F90` are passed through the C preprocessor (`cpp`) before being compiled.

Include files also contain source and often have suffixes different from those ordinarily used for Fortran source files.

#### **Object files (.o suffix)**

After the compiler compiles the source files, it uses the **ld** command to link the resulting .o files, any .o files that you specify as input files, and some of the .o and .a files in the product and system library directories. The compiler can then produce a single .o object file or a single executable output file from these object files.

#### **Assembler source files (.s and .S suffixes)**

The compiler sends assembler source files to the assembler (**as**). The assembler sends object files to the linker at link time.

**Note:** Assembler source files with a .S filename suffix are first preprocessed by the compiler, then sent to the assembler.

#### **Shared object or library files (.so suffix)**

These are object files that can be loaded and shared by multiple processes at run time. When a shared object is specified during linking, information about the object is recorded in the output file, but no code from the shared object is actually included in the output file.

#### **Configuration files (.cfg suffix)**

The contents of the configuration file determine many aspects of the compilation process, most commonly the default options for the compiler. You can use it to centralize different sets of default compiler options or to keep multiple levels of the XL Fortran compiler present on a system.

The default configuration file is `/etc/opt/ibmcomp/xf/10.1/xf.cfg`.

#### **Module symbol files: *modulename.mod***

A module symbol file is an output file from compiling a module and is an input file for subsequent compilations of files that **USE** that module. One .mod file is produced for each module, so compiling a single source file may produce multiple .mod files.

#### **Profile data files**

The **-qpdf1** option produces runtime profile information for use in subsequent compilations. This information is stored in one or more hidden files with names that match the pattern `.*pdf*`.

---

## **XL Fortran output files**

The output files that Fortran produces are:

#### **Executable files: *a.out***

By default, XL Fortran produces an executable file that is named **a.out** in the current directory.

#### **Object files: *filename.o***

If you specify the **-c** compiler option, instead of producing an executable file, the compiler produces an object file for each specified program source input file, and the assembler produces an object file for each specified assembler input file. By default, the output object files are saved to the current directory using the same file name prefixes as their corresponding source input files.

#### **Assembler source files: *filename.s***

If you specify the **-S** compiler option, instead of producing an executable

file, the XL Fortran compiler produces an equivalent assembler source file for each specified input source file. By default, the output assembler source files are saved to the current directory using the same file name prefixes as their corresponding source input files.

**Compiler listing files:** *filename.lst*

By default, no listing is produced unless you specify one or more listing-related compiler options. The listing file is placed in the current directory, with the same file name prefix as the source file.

**Module symbol files:** *modulename.mod*

Each module has an associated symbol file that holds information needed by program units, subprograms, and interface bodies that USE that module. By default, these symbol files must exist in the current directory. Compiling modules will also produce .o files that are needed when linking if you use the module.

**Preprocessed source files:** *Ffilename.f*

If you specify the **-d** option when compiling a file with a **.F** suffix, the intermediate file created by the C preprocessor (cpp) is saved rather than deleted.

**Profile data files (\*.pdf\*)**

These are the profile-directed feedback files that the **-qpdf1** compiler option produces. They are used in subsequent compilations to tune optimizations according to actual execution results.

---

## Specifying compiler options

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

You can specify compiler options:

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

When multiple compiler options have been specified, it is possible for option conflicts and incompatibilities to occur. To resolve these conflicts in a consistent fashion, the compiler usually applies the following general priority sequence:

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

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

**Note:** The **-I** compiler option is a special case. The compiler searches any directories specified with **-I** in the **xlfcfg** file before it searches the directories specified with **-I** on the command line. The option is cumulative rather than preemptive.

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

You can also pass compiler options to the linker, assembler, and preprocessor. See XL Fortran Compiler-option reference in the *XL Fortran Advanced Edition V10.1 for Linux Compiler Reference* for more information about compiler options and how to specify them.

---

## Linking XL Fortran programs

By default, you do not need to do anything special to link an XL Fortran program. The compiler invocation commands automatically call the linker to produce an executable output file. For example, running the following command:

```
xlf95 file1.f file2.o file3.f
```

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

After linking, follow the instructions in Chapter 5, “Running XL Fortran programs,” on page 25 to execute the program.

## Compiling and linking in separate steps

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

```
xlf95 -c file1.f           # Produce one object file (file1.o)
xlf95 -c file2.f file3.f   # Or multiple object files (file2.o, file3.o)
xlf95 file1.o file2.o file3.o # Link object files with appropriate libraries
```

It is often best to execute the linker through the compiler invocation command, because it passes some extra `ld` options and library names to the linker automatically.

## Dynamic and static linking

XL Fortran allows your programs to take advantage of the operating system facilities for both dynamic and static linking:

- Dynamic linking means that the code for some external routines is located and loaded when the program is first run. When you compile a program that uses shared libraries, the shared libraries are dynamically linked to your program by default.

Dynamically linked programs take up less disk space and less virtual memory if more than one program uses the routines in the shared libraries. During linking, they do not require any special precautions to avoid naming conflicts with library routines. They may perform better than statically linked programs if several programs use the same shared routines at the same time. They also allow you to upgrade the routines in the shared libraries without relinking.

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

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

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

See Linking XL Fortran programs in the *XL Fortran Advanced Edition V10.1 for Linux Compiler Reference* for more information about linking your programs.



---

## Chapter 5. Running XL Fortran programs

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

You can run a program by entering the name of a program executable file together with any runtime arguments on the command line.

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

---

### Canceling execution

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

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

---

### Setting runtime options

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

For more information on environment variables and how they can affect your applications at runtime, see “Environment variables and XL Fortran” on page 15.

---

### Running compiled applications on other systems

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

You can obtain the latest XL Fortran Runtime Environment install images, together with licensing and usage information, from the Download section on the XL Fortran Support page at:

[www.ibm.com/software/awdtools/fortran/xlfortran/support](http://www.ibm.com/software/awdtools/fortran/xlfortran/support)



---

## Chapter 6. XL Fortran compiler diagnostic aids

XL Fortran issues diagnostic messages when it encounters problems compiling your application. You can use these messages to help identify and correct such problems.

This section provides a brief overview of the main diagnostics messages offered by XL Fortran. For more information about related compiler options that can help you resolve problems with your application, see *Options for error checking and debugging* and *Options that control listings and messages* in the *XL Fortran Advanced Edition V10.1 for Linux Compiler Reference*.

---

### Compilation return codes

At the end of compilation, the compiler sets the return code to zero under any of the following conditions:

- No messages are issued.
- The highest severity level of all errors diagnosed is less than the setting of the **-qhalt** compiler option.

Otherwise, the compiler sets the return code to one of the following values:

Return Code	Error Type
1	An error with a severity level higher than the setting of the <b>-qhalt</b> compiler option has occurred.
40	An option error or unrecoverable error has occurred.
41	A configuration file error has occurred.
250	An out-of-memory has occurred. The compiler invocation command cannot allocate any more memory for its use.
251	A signal-received error has occurred. That is, an unrecoverable error or interrupt signal has occurred.
252	A file-not-found error has occurred.
253	An input/output error has occurred - files cannot be read or written to.
254	A fork error has occurred. A new process cannot be created.
255	An error has been detected while the process was running.

**Note:** Return codes may also be displayed for runtime errors.

---

### XL Fortran compiler listings

Diagnostic information is produced in the output listing according to the settings of the **-qlist**, **-qsource**, **-qxref**, **-qattr**, **-qreport**, and **-qlistopt** compiler options. The **-S** option generates an assembler listing in a separate file.

If the compiler encounters a programming error when compiling an application, the compiler issues a diagnostic message to the standard error device and, if the appropriate compiler options have been selected, to a listing file.

To locate the cause of a problem with the help of a listing, you can refer to:

- The source section (to see any compilation errors in the context of the source program)
- The attribute and cross-reference section (to find data objects that are misnamed or used without being declared or to find mismatched parameters)
- The transformation and object sections (to see if the generated code is similar to what you expect)

A heading identifies each major section of the listing. A string of "greater than" (>) symbols precede the section heading so that you can easily locate its beginning:

```
>>>> section name
```

You can select which sections appear in the listing by specifying the appropriate compiler options. For more information about these options see Options for error checking and debugging and Options that control listings and messages in the *XL Fortran Advanced Edition V10.1 for Linux Compiler Reference*.

---

## Debugging compiled applications

Specifying the **-g** or **-qlinedebug** compiler options at compile time instructs the XL Fortran compiler to include debugging information in compiled output. You can then use **gdb** or any other symbolic debugger to step through and inspect the behavior of your compiled application.

---

## Chapter 7. XL Fortran runtime environment information

Object code that the XL Fortran compiler produces often invokes compiler-supplied subprograms at run time to handle certain complex tasks. These subprograms are collected into several libraries.

The function of the XL Fortran Runtime Environment may be divided into these main categories:

- Support for Fortran I/O operations
- Mathematical calculation
- Operating-system services
- Support for SMP parallelization

The XL Fortran Runtime Environment also produces runtime diagnostic messages in the national language appropriate for your system.

Unless you bind statically, you cannot run object code produced by the XL Fortran compiler without the XL Fortran Runtime Environment. However, static binding is discouraged because it could cause problems when running applications on operating systems that require newer XL Fortran libraries than the one that was statically bound. The compiler defaults to dynamic binding, which is the most flexible solution for running your applications on multiple operating system levels.

The XL Fortran Runtime Environment is upward-compatible. Programs that are compiled and linked with a given level of the runtime environment and a given level of the operating system require the same or higher levels of both the runtime environment and the operating system to run.

---

### External names in the runtime environment

Runtime subprograms are collected into libraries. By default, the compiler invocation command also invokes the linker and gives it the names of the libraries that contain runtime subprograms called by Fortran object code.

The names of these runtime subprograms are external symbols. When object code that is produced by the XL Fortran compiler calls a runtime subprogram, the `.o` object code file contains an external symbol reference to the name of the subprogram. A library contains an external symbol definition for the subprogram. The linker resolves the runtime subprogram call with the subprogram definition.

You should avoid using names in your XL Fortran program that conflict with names of runtime subprograms. Conflict can arise under two conditions:

- The name of a subroutine, function, or common block that is defined in a Fortran program has the same name as a library subprogram.
- The Fortran program calls a subroutine or function with the same name as a library subprogram but does not supply a definition for the called subroutine or function.

---

## External names in XL Fortran libraries

To minimize naming conflicts between user-defined names and the names that are defined in the runtime libraries, the names of input/output routines in the runtime libraries are prefixed with an underscore(`_`), or `_xl`.

---

## Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing  
IBM Corporation  
North Castle Drive  
Armonk, NY 10504-1785  
U.S.A.

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

IBM World Trade Asia Corporation  
Licensing  
2-31 Roppongi 3-chome, Minato-ku  
Tokyo 106, Japan

**The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:**

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation  
Lab Director  
IBM Canada Limited  
8200 Warden Avenue  
Markham, Ontario, Canada  
L6G 1C7

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.

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 may contain sample application programs in source language, which illustrates programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. You may copy, modify, and distribute these sample programs in any form without payment to IBM for the purposes of developing, using, marketing, or distributing application programs conforming to IBM's application programming interfaces.

---

## Programming interface information

Programming interface information is intended to help you create application software using this program.

General-use programming interfaces allow the customer to write application software that obtain the services of this program's tools.

However, this information may also contain diagnosis, modification, and tuning information. Diagnosis, modification, and tuning information is provided to help you debug your application software.

**Note:** Do not use this diagnosis, modification, and tuning information as a programming interface because it is subject to change.

---

## Trademarks and service marks

The following terms, used in this publication, are trademarks or service marks of the International Business Machines Corporation in the United States or other countries or both:

AIX	IBM	OS/390
OS/400	POWER3	POWER4
POWER5	POWER5+	PowerPC
z/OS	z/VM	

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

OpenMP is a trademark of the OpenMP Architecture Review Board.

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

Windows is a trademark of Microsoft Corporation in the United States, other countries, or both.

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



---

# Index

## Special characters

- /etc/opt/ibmcmp/xf/10.1/xf.cfg
  - configuration file 16
- .a files 19
- .cfg files 19
- .f and .F files 19
- .lst files 20
- .mod files 19, 20
- .o files 19, 20
- .s files 19, 20
- .S files 19
- .so files 19

## Numerics

- 64-bit environment 4

## A

- a.out file 20
- ANSI
  - checking conformance to the Fortran 95 standard 3
  - checking conformance to the Fortran standard 3
- archive files 19
- assembler
  - source (.s) files 19, 20
  - source (.S) files 19

## C

- code optimization 3
- command-line options
  - See compiler options
- compiler listings 27
- compiler options
  - See the individual options listed under Special Characters at the start of the index
- compiling
  - description of how to compile a program 18
  - SMP programs 19
- configuration file 16, 19
- conformance checking 3
- customizing configuration file (including default compiler options) 16

## D

- debugger support 5, 28
- debugging 27, 28
- defaults
  - customizing compiler defaults 16
- documentation, online formats 1
- dynamic linking 22

## E

- editing source files 17
- environment variables
  - compile time 15
  - runtime 15
- example programs
  - See sample programs
- executable files 20
- executing a program 25
- executing the compiler 18
- external names
  - in the runtime environment 29

## F

- f77 command
  - description 18
  - level of Fortran standard compliance 18
- files
  - editing source 17
  - input 19
  - output 20
- FIPS FORTRAN standard, checking conformance to 3
- fort77 command
  - description 18
- Fortran 90
  - compiling programs written for 18

## H

- help system 1
- HTML documentation 1

## I

- I/O
  - See input/output
- input files 19
- invoking a program 25
- invoking the compiler 18
- ISO
  - checking conformance to the Fortran 2003 standard 3
  - checking conformance to the Fortran 90 standard 3
  - checking conformance to the Fortran 95 standard 3
  - checking conformance to the Fortran standard 3

## L

- language support 3
- level of XL Fortran, determining 16
- libraries 19
  - shared 30
- linking 22

- linking (*continued*)

- dynamic 22
  - static 22

- listing files 20

## M

- makefiles
  - configuration file as alternative for default options 16
- mod files 19, 20
- multiprocessor systems 4

## O

- object files 19, 20
- online compiler help 1
- online documentation 1
- OpenMP 5
- optimization 3
- output files 20

## P

- parallelization 4
- parameters
  - See arguments
- PDF documentation 1
- POSIX Pthreads
  - API support 19
- problem determination 27
- profiling data files 20

## R

- running a program 25
- running the compiler 18
- runtime
  - libraries 19
- runtime environment 25
  - external names in 29
- runtime options 25

## S

- SAA FORTRAN definition, checking conformance to 3
- setrteopts service and utility
  - procedure 15
- shared libraries 30
- shared memory parallelization 4
- shared object files 19
- SMP
  - programs, compiling 19
- SMP programs 4
- source files 19
- source-code conformance checking 3
- source-level debugging support 5
- static linking 22

symbolic debugger support 5

## T

text editors 17

## X

xlF command

description 18

level of Fortran standard

compliance 18

xlF\_r command

description 18

for compiling SMP programs 19

level of Fortran standard

compliance 18

xlF90 command

description 18

level of Fortran standard

compliance 18

xlF90\_r command

description 18

for compiling SMP programs 19

level of Fortran standard

compliance 18

xlF95 command

description 18

xlF95\_r command

description 18

for compiling SMP programs 19

level of Fortran standard

compliance 18

XLFRTEOPTS environment variable 15





Program Number: 5724-M17

SC09-8021-00

