

IBM XL Fortran for Linux, V15.1.1



# Compiler Reference for Little Endian Distributions

*Version 15.1.1*



IBM XL Fortran for Linux, V15.1.1



# Compiler Reference for Little Endian Distributions

*Version 15.1.1*

**Note**

Before using this information and the product it supports, read the information in "Notices" on page 327.

**First edition**

This edition applies to IBM XL Fortran for Linux, V15.1.1 (Program 5765-J10; 5725-C75) 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 1990, 2014.

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

---

# Contents

<b>About this information . . . . .</b>	<b>vii</b>
Who should read this information. . . . .	vii
How to use this information . . . . .	vii
How this information is organized . . . . .	vii
Conventions. . . . .	viii
Related information . . . . .	xii
IBM XL Fortran information. . . . .	xii
Standards and specifications . . . . .	xiii
Other IBM information . . . . .	xiv
Technical support . . . . .	xiv
How to send your comments . . . . .	xiv

<b>Chapter 1. Introduction . . . . .</b>	<b>1</b>
--	----------

<b>Chapter 2. Overview of XL Fortran features . . . . .</b>	<b>3</b>
Hardware and operating-system support . . . . .	3
Language support . . . . .	3
Source-code conformance checking . . . . .	4
Highly configurable compiler. . . . .	4
Diagnostic listings . . . . .	5
Symbolic debugger support . . . . .	6
Program optimization . . . . .	6

<b>Chapter 3. Setting up and customizing XL Fortran. . . . .</b>	<b>7</b>
Where to find installation instructions . . . . .	7
Correct settings for environment variables . . . . .	7
Environment variable basics . . . . .	7
Setting library search paths . . . . .	8
Profile-directed feedback environment variables . . . . .	8
TMPDIR: Specifying a directory for temporary files . . . . .	9
XLFSCRATCH_unit: Specifying names for scratch files . . . . .	9
XLFUNIT_unit: Specifying names for implicitly connected files. . . . .	9
Using custom compiler configuration files . . . . .	9
Creating custom configuration files . . . . .	10
Using IBM XL Fortran for Linux, V15.1.1 with the Advance Toolchain . . . . .	13
Editing the default configuration file . . . . .	13
Configuration file attributes . . . . .	14
Determining which level of XL Fortran is installed . . . . .	16
Running two levels of XL Fortran . . . . .	16

<b>Chapter 4. Editing, compiling, linking, and running XL Fortran programs . . . . .</b>	<b>17</b>
Editing XL Fortran source files . . . . .	17
Compiling XL Fortran programs . . . . .	17
Compiling XL Fortran Version 2 programs . . . . .	19
Compiling Fortran 90 or Fortran 95 programs . . . . .	19
Compiling Fortran 2003 programs . . . . .	20
Compiling Fortran 2008 programs . . . . .	20

Compiling and linking a library . . . . .	21
Compiling XL Fortran SMP programs . . . . .	22
Compilation order for Fortran programs. . . . .	22
Canceling a compilation . . . . .	23
XL Fortran input files . . . . .	23
XL Fortran output files . . . . .	24
Scope and precedence of option settings. . . . .	26
Specifying options on the command line. . . . .	26
Specifying options in the source file . . . . .	27
Passing command-line options to the "ld" or "as" command . . . . .	28
Displaying information inside binary files (strings) . . . . .	28
Compiling for specific architectures . . . . .	29
Passing Fortran files through the C preprocessor . . . . .	29
cpp directives for XL Fortran programs . . . . .	30
Passing options to the C preprocessor . . . . .	31
Avoiding preprocessing problems . . . . .	31
Linking XL Fortran programs . . . . .	31
Compiling and linking in separate steps. . . . .	32
Passing options to the ld command . . . . .	32
Dynamic and static linking . . . . .	32
Avoiding naming conflicts during linking . . . . .	33
Running XL Fortran programs . . . . .	33
Canceling execution . . . . .	33
Compiling and executing on different systems. . . . .	33
Runtime libraries for POSIX pthreads support. . . . .	34
Setting runtime options . . . . .	34
Other environment variables that affect runtime behavior . . . . .	45
XL Fortran runtime exceptions . . . . .	45

<b>Chapter 5. Summary of compiler options by functional category . . . . .</b>	<b>47</b>
Output control . . . . .	47
Input control . . . . .	49
Language element control . . . . .	50
Floating-point and integer control . . . . .	52
Object code control . . . . .	53
Error checking and debugging . . . . .	54
Listings, messages, and compiler information . . . . .	56
Optimization and tuning . . . . .	57
Linking. . . . .	61
Portability and migration. . . . .	61
Compiler customization . . . . .	62

<b>Chapter 6. Detailed descriptions of the XL Fortran compiler options . . . . .</b>	<b>65</b>
-#. . . . .	66
-1. . . . .	67
-B . . . . .	67
-C . . . . .	68
-c. . . . .	69
-D . . . . .	69
-d . . . . .	70

-e . . . . .	70	-qinitauto . . . . .	154
-F . . . . .	71	-qinlgue . . . . .	157
-g . . . . .	73	-qinline . . . . .	157
-I . . . . .	76	-qintlog . . . . .	160
-k . . . . .	77	-qintsize . . . . .	161
-L . . . . .	77	-qipa . . . . .	163
-l . . . . .	78	-qkeepparm . . . . .	168
-MF . . . . .	79	-qlanglvl . . . . .	169
-MMD . . . . .	80	-qlibansi . . . . .	171
-MT . . . . .	81	-qlibmpi . . . . .	172
-NS . . . . .	82	-qlinedebug . . . . .	173
-O . . . . .	83	-qlist . . . . .	174
-o . . . . .	85	-qlistfmt . . . . .	175
-p . . . . .	86	-qlistopt . . . . .	178
-qalias . . . . .	87	-qlog4 . . . . .	179
-qalign . . . . .	90	-qmakedep . . . . .	179
-qaltivec . . . . .	92	-qmaxerr . . . . .	182
-qarch . . . . .	93	-qmaxmem . . . . .	183
-qassert . . . . .	95	-qmbcs . . . . .	185
-qattr . . . . .	96	-qmixed . . . . .	186
-qautodbl . . . . .	97	-qmksprobj . . . . .	186
-qbindcextname . . . . .	99	-qmoddir . . . . .	187
-qcache . . . . .	100	-qnoprint . . . . .	188
-qcclines . . . . .	103	-qnullterm . . . . .	189
-qcheck . . . . .	103	-qobject . . . . .	190
-qci . . . . .	106	-qonetrip . . . . .	191
-qcompact . . . . .	108	-qoptfile . . . . .	191
-qcr . . . . .	108	-qoptimize . . . . .	193
-qctyplss . . . . .	109	-qpath . . . . .	194
-qdbg . . . . .	110	-qpdf1, -qpdf2 . . . . .	195
-qddim . . . . .	112	-qphsinfo . . . . .	203
-qdescriptor . . . . .	113	-qpics . . . . .	205
-qdirective . . . . .	114	-qport . . . . .	205
-qdirectstorage . . . . .	116	-qposition . . . . .	208
-qdlines . . . . .	116	-qppsuborigarg . . . . .	209
-qdpc . . . . .	117	-qprefetch . . . . .	210
-qenum . . . . .	118	-qqcount . . . . .	212
-qescape . . . . .	119	-qrealsize . . . . .	213
-qessl . . . . .	120	-qrecur . . . . .	215
-qextern . . . . .	121	-qreport . . . . .	216
-qextername . . . . .	122	-qsaa . . . . .	218
-qfdpr . . . . .	124	-qsave . . . . .	218
-qfixed . . . . .	125	-qsaveopt . . . . .	220
-qflag . . . . .	126	-qsclk . . . . .	222
-qfloat . . . . .	127	-qshowpdf . . . . .	223
-qfpp . . . . .	131	-qsigtrap . . . . .	224
-qflttrap . . . . .	132	-qsimd . . . . .	225
-qfree . . . . .	134	-qsmallstack . . . . .	226
-qfullpath . . . . .	135	-qsmp . . . . .	227
-qfunctrace . . . . .	136	-qsource . . . . .	232
-qfunctrace_xlf_catch . . . . .	138	-qspillsize . . . . .	233
-qfunctrace_xlf_enter . . . . .	139	-qstackprotect . . . . .	233
-qfunctrace_xlf_exit . . . . .	140	-qstacktemp . . . . .	234
-qhalt . . . . .	140	-qstaticlink . . . . .	235
-qhaltonmsg . . . . .	142	-qstrict . . . . .	238
-qhelp . . . . .	143	-qstrictieemod . . . . .	242
-qhot . . . . .	143	-qstrict_induction . . . . .	242
-qieee . . . . .	146	-qsuffix . . . . .	244
-qinfo . . . . .	147	-qsuppress . . . . .	245
-qinit . . . . .	152	-qswapomp . . . . .	246
-qinitalloc . . . . .	152	-qtbtable . . . . .	248

-qthreaded . . . . .	249
-qtimestamps . . . . .	249
-qtune . . . . .	250
-qufmt . . . . .	252
-qundef . . . . .	253
-qunroll . . . . .	253
-qunwind . . . . .	255
-qversion . . . . .	255
-qwarn64 . . . . .	257
-qxflag=dvz . . . . .	258
-qxflag=oldtab . . . . .	259
-qxlf77 . . . . .	259
-qxlf90 . . . . .	261
-qxlf2003 . . . . .	263
-qxlf2008 . . . . .	267
-qxlines . . . . .	268
-qxref . . . . .	269
-qzerosize . . . . .	270
-r . . . . .	271
-S . . . . .	272
-t . . . . .	272
-U . . . . .	274
-u . . . . .	274
-v . . . . .	275
-V . . . . .	276
-W . . . . .	276
-w . . . . .	278
-y . . . . .	279

**Chapter 7. Using XL Fortran in a 64-bit environment . . . . . 281**

**Chapter 8. Problem determination and debugging . . . . . 283**

Understanding XL Fortran error messages . . . . .	283
Error severity . . . . .	283
Compiler return codes . . . . .	284
Runtime return codes . . . . .	284
Format of XL Fortran diagnostic messages. . . . .	284
Limiting the number of compile-time messages . . . . .	285
Fixing installation or system environment problems . . . . .	285

Fixing compile-time problems . . . . .	287
Fixing link-time problems . . . . .	288
Fixing runtime problems . . . . .	289
Debugging a Fortran program . . . . .	290

**Chapter 9. Understanding XL Fortran compiler listings . . . . . 291**

Header section . . . . .	291
Options section . . . . .	291
Source section . . . . .	292
Error messages . . . . .	292
PDF report section . . . . .	292
Transformation report section . . . . .	293
Data reorganization report section . . . . .	295
Attribute and cross reference section . . . . .	295
Object section . . . . .	296
File table section . . . . .	297
Compilation unit epilogue Section . . . . .	297
Compilation epilogue Section . . . . .	297

**Chapter 10. XL Fortran technical information . . . . . 299**

External names in XL Fortran libraries . . . . .	299
The XL Fortran runtime environment . . . . .	299
External names in the runtime environment . . . . .	299
Technical details of the -qfloat=hsflt option . . . . .	300
Implementation details for -qautodbl promotion and padding . . . . .	300
Terminology . . . . .	300
Examples of storage relationships for -qautodbl suboptions . . . . .	301

**XL Fortran internal limits . . . . . 305**

**Glossary . . . . . 307**

**Notices . . . . . 327**

Trademarks and service marks . . . . .	329
--	-----

**Index . . . . . 331**



---

## About this information

This document describes the IBM® XL Fortran for Linux, V15.1.1 compiler and explains how to set up the compilation environment and how to compile, link, and run programs written in the Fortran language. This guide also contains cross-references to relevant topics of other reference guides in the XL Fortran documentation suite.

---

## Who should read this information

This information is for anyone who wants to work with the IBM XL Fortran for Linux, V15.1.1 compiler, is familiar with the Linux operating system, and who has some previous Fortran programming experience. Users new to XL Fortran can also find information on the capabilities and features unique to XL Fortran. This information can help you understand what the features of the compiler are, especially the options, and how to use them for effective software development.

---

## How to use this information

While this information covers topics about configuring the compiler, and compiling, linking and running XL Fortran programs, it does not include information on the following topics, which are covered elsewhere:

- Installation, system requirements, last-minute updates: see the *XL Fortran Installation Guide* and product README.
- Overview of XL Fortran features: see the *Getting Started with XL Fortran*.
- Syntax, semantics, and implementation of the XL Fortran programming language: see the *XL Fortran Language Reference*.
- Optimizing, porting, OpenMP/ SMP programming: see the *XL Fortran Optimization and Programming Guide*.
- Operating system commands related to the use of the compiler: consult your Linux-specific distribution's man page help and documentation.

---

## How this information is organized

This information starts with an overview of the compiler and then outlines the tasks you need to do before invoking the compiler. It then continues with reference information about the compiler options and debugging problems.

This reference includes the following topics:

- Chapter 1, “Introduction,” on page 1 through Chapter 4, “Editing, compiling, linking, and running XL Fortran programs,” on page 17 discuss setting up the compilation environment and the environment variables that you need for different compilation modes, customizing the configuration file, the types of input and output files, compiler listings and messages and information specific to invoking the preprocessor and linkage editor.
- Chapter 5, “Summary of compiler options by functional category,” on page 47 organizes the compiler options by their functional category. You can search for options by their name, or alternatively use the links in the functional category tables and look up options according to their functionality. Chapter 6, “Detailed

descriptions of the XL Fortran compiler options,” on page 65 includes individual descriptions of the compiler options sorted alphabetically. Descriptions provide examples and list related topics.

- Chapter 7, “Using XL Fortran in a 64-bit environment,” on page 281 discusses application development for the 64-bit environment.
- Chapter 8, “Problem determination and debugging,” on page 283 addresses debugging and understanding compiler listings.
- Chapter 10, “XL Fortran technical information,” on page 299 and “XL Fortran internal limits” on page 305 provide information that advanced programmers may need to diagnose unusual problems and run the compiler in a specialized environment.

## Conventions

### Typographical conventions

The following table shows the typographical conventions used in the IBM XL Fortran for Linux, V15.1.1 information.

Table 1. Typographical conventions

Typeface	Indicates	Example
<b>lowercase bold</b>	Invocation commands, executable names, and compiler options.	The compiler provides basic invocation commands, <code>xlf</code> , along with several other compiler invocation commands to support various Fortran language levels and compilation environments.  The default file name for the executable program is <b>a.out</b> .
<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.	<code>nomaf   <u>maf</u></code>
monospace	Examples of program code, reference to program code, file names, path names, command strings, or user-defined names.	To compile and optimize <code>myprogram.f</code> , enter: <code>xlf myprogram.f -03</code> .
<b>UPPERCASE bold</b>	Fortran programming keywords, statements, directives, and intrinsic procedures. Uppercase letters may also be used to indicate the minimum number of characters required to invoke a compiler option/suboption.	The <b>ASSERT</b> directive applies only to the <b>DO</b> loop immediately following the directive, and not to any nested <b>DO</b> loops.

### Qualifying elements (icons and bracket separators)

In descriptions of language elements, this information uses icons and marked bracket separators to delineate the Fortran language standard text as follows:

Table 2. Qualifying elements

Icon	Bracket separator text	Meaning
	N/A	The text describes an IBM XL Fortran implementation of the Fortran 2008 standard.
	Fortran 2003 begins / ends	The text describes an IBM XL Fortran implementation of the Fortran 2003 standard, and it applies to all later standards.
	IBM extension begins / ends	The text describes a feature that is an IBM XL Fortran extension to the standard language specifications.
	N/A	The text describes a feature in a Technical Specification that is not part of the current Fortran standard.

**Note:** If the information is marked with a Fortran language standard icon or bracket separators, it applies to this specific Fortran language standard and all later ones. If it is not marked, it applies to all Fortran language standards.

## Syntax diagrams

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

- Read the syntax diagrams from left to right, from top to bottom, following the path of the line.

The  symbol indicates the beginning of a command, directive, or statement.

The  symbol indicates that the command, directive, or statement syntax is continued on the next line.

The  symbol indicates that a command, directive, or statement is continued from the previous line.

The  symbol indicates the end of a command, directive, or statement.

Fragments, which are diagrams of syntactical units other than complete commands, directives, or statements, start with the  symbol and end with the  symbol.


IBM XL Fortran extensions are marked by a number in the syntax diagram with an explanatory note immediately following the diagram.

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.

- 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 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:



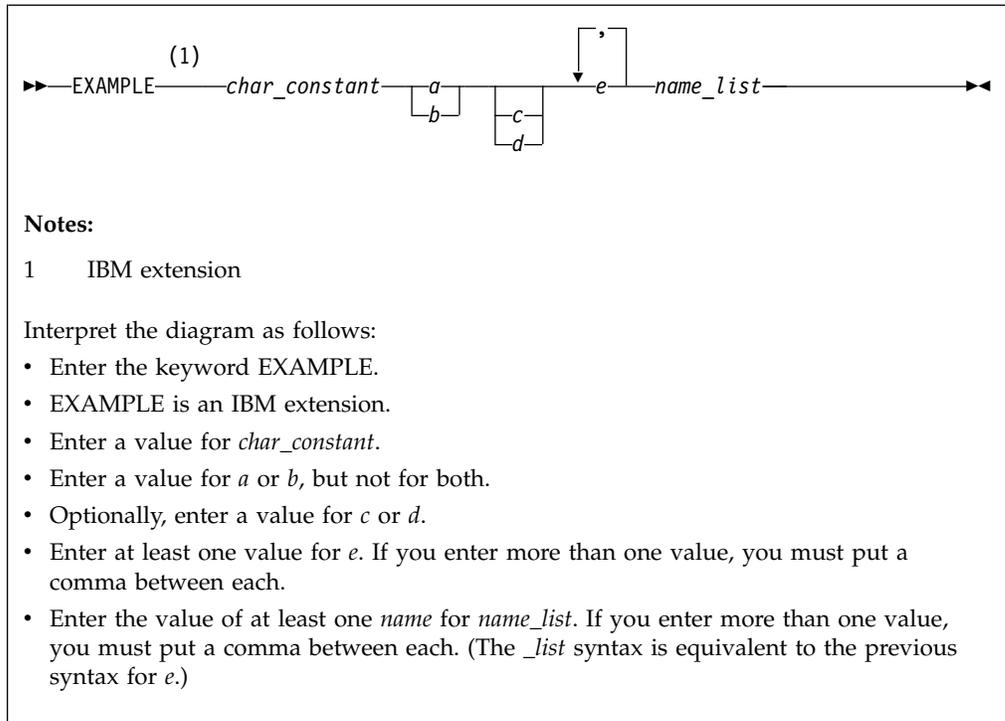
- The item that is the default is shown above the main path.



- 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 a variable or user-specified name ends in *\_list*, you can provide a list of these terms separated by commas.
- 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 is an example of a syntax diagram with an interpretation:



## How to read syntax statements

Syntax statements are read from left to right:

- Individual required arguments are shown with no special notation.
- When you must make a choice between a set of alternatives, they are enclosed by { and } symbols.
- Optional arguments are enclosed by [ and ] symbols.
- When you can select from a group of choices, they are separated by | characters.
- Arguments that you can repeat are followed by ellipses (...).

### 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.

## Notes on the terminology used

Some of the terminology in this information is shortened as follows:

- The term *free source form format* often appears as *free source form*.
- The term *fixed source form format* often appears as *fixed source form*.
- The term *XL Fortran* often appears as *XL F*.

---

## Related information

The following sections provide related information for XL Fortran:

### IBM XL Fortran information

XL Fortran provides product information in the following formats:

- README files  
README files contain late-breaking information, including changes and corrections to the product information. README files are located by default in the XL Fortran directory, and in the root directory and subdirectories 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 Fortran for Linux, V15.1.1 Installation Guide*.
- Online product documentation  
The fully searchable HTML-based documentation is viewable in IBM Knowledge Center at [http://www.ibm.com/support/knowledgecenter/SSAT4T\\_15.1.1/com.ibm.compilers.linux.doc/welcome.html](http://www.ibm.com/support/knowledgecenter/SSAT4T_15.1.1/com.ibm.compilers.linux.doc/welcome.html).
- PDF documents  
PDF documents are available on the web at <http://www.ibm.com/support/docview.wss?uid=swg27036672>.

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

Table 3. XL Fortran PDF files

Document title	PDF file name	Description
<i>IBM XL Fortran for Linux, V15.1.1 Installation Guide, GC27-6580-00</i>	install.pdf	Contains information for installing XL Fortran and configuring your environment for basic compilation and program execution.

Table 3. XL Fortran PDF files (continued)

Document title	PDF file name	Description
<i>Getting Started with IBM XL Fortran for Linux, V15.1.1, SC27-6620-00</i>	getstart.pdf	Contains an introduction to the XL Fortran product, with information on setting up and configuring your environment, compiling and linking programs, and troubleshooting compilation errors.
<i>IBM XL Fortran for Linux, V15.1.1 Compiler Reference, SC27-6610-00</i>	compiler.pdf	Contains information about the various compiler options and environment variables.
<i>IBM XL Fortran for Linux, V15.1.1 Language Reference, SC27-6590-00</i>	langref.pdf	Contains information about the Fortran programming language as supported by IBM, including language extensions for portability and conformance to nonproprietary standards, compiler directives and intrinsic procedures.
<i>IBM XL Fortran for Linux, V15.1.1 Optimization and Programming Guide, SC27-6600-00</i>	proguide.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.

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 Fortran including IBM Redbooks® publications, white papers, tutorials, documentation errata, and other articles, is available on the web at:

<http://www.ibm.com/support/docview.wss?uid=swg27036672>

For more information about Fortran, see the Fortran café at <https://www.ibm.com/developerworks/mydeveloperworks/groups/service/html/communityview?communityUuid=b10932b4-0edd-4e61-89f2-6e478ccb9aa>.

## Standards and specifications

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

- *American National Standard Programming Language FORTRAN, ANSI X3.9-1978.*
- *American National Standard Programming Language Fortran 90, ANSI X3.198-1992.*
- *ANSI/IEEE Standard for Binary Floating-Point Arithmetic, ANSI/IEEE Std 754-1985.*
- *Federal (USA) Information Processing Standards Publication Fortran, FIPS PUB 69-1.*
- *Information technology - Programming languages - Fortran, ISO/IEC 1539-1:1991. (This information uses its informal name, Fortran 90.)*
- *Information technology - Programming languages - Fortran - Part 1: Base language, ISO/IEC 1539-1:1997. (This information uses its informal name, Fortran 95.)*
- *Information technology - Programming languages - Fortran - Part 1: Base language, ISO/IEC 1539-1:2004. (This information uses its informal name, Fortran 2003.)*

- *Information technology - Programming languages - Fortran - Part 1: Base language, ISO/IEC 1539-1:2010.* (This information uses its informal name, Fortran 2008. We currently provide partial support to this standard.)
- *Military Standard Fortran DOD Supplement to ANSI X3.9-1978, MIL-STD-1753* (United States of America, Department of Defense standard). Note that XL Fortran supports only those extensions documented in this standard that have also been subsequently incorporated into the Fortran 90 standard.
- *OpenMP Application Program Interface Version 4.0 (Partial support)*, available at <http://www.openmp.org>

## Other IBM information

- *ESSL product documentation* available at [http://www.ibm.com/support/knowledgecenter/SSFHY8/essl\\_welcome.html](http://www.ibm.com/support/knowledgecenter/SSFHY8/essl_welcome.html)

---

## Technical support

Additional technical support is available from the XL Fortran Support page at [http://www.ibm.com/support/entry/portal/overview/software/rational/xl\\_fortran\\_for\\_linux](http://www.ibm.com/support/entry/portal/overview/software/rational/xl_fortran_for_linux). This page provides a portal with search capabilities to a large selection of Technotes and other support information.

If you cannot find what you need, you can send email to [compinfo@cn.ibm.com](mailto:compinfo@cn.ibm.com).

For the latest information about XL Fortran, visit the product information site at <http://www.ibm.com/software/products/us/en/xlfortran-linux>.

---

## How to send your comments

Your feedback is important in helping to provide accurate and high-quality information. If you have any comments about this information or any other XL Fortran information, send your comments by email to [compinfo@cn.ibm.com](mailto:compinfo@cn.ibm.com).

Be sure to include the name of the manual, the part number of the manual, 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. Introduction

IBM XL Fortran for Linux, V15.1.1 is an optimizing, standards-based, command-line compiler for the Linux operating system, running on Power<sup>®</sup> hardware with the Power Architecture. The XL Fortran compiler enables application developers to create and maintain optimized 64-bit applications for the Linux operating system. The compiler also offers a diversified portfolio of optimization techniques that allow an application developer to exploit the multi-layered architecture of the Power processor.

The implementation of the Fortran programming language is intended to promote portability among different environments by enforcing conformance to language standards. A program that conforms strictly to its language specification has maximum portability among different environments. In theory, a program that compiles correctly with one standards-conforming compiler will compile and execute correctly under all other conforming compilers, insofar as hardware differences permit. A program that correctly exploits the extensions to the programming language in which it is written can improve the efficiency of its object code.

XL Fortran can be used for large, complex, and computationally intensive programs. It also supports interlanguage calls with C. For applications that require SIMD (single-instruction, multiple data) parallel processing, performance improvements can be achieved through optimization techniques, which may be less labor-intensive than vector programming. Many of the optimizations developed by IBM are controlled by compiler options and directives.



---

## Chapter 2. Overview of XL Fortran features

This section discusses the features of the XL Fortran compiler, language, and development environment 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.

---

### Hardware and operating-system support

IBM XL Fortran for Linux, V15.1.1 is supported on several Linux distributions. See the *XL Fortran for Linux Installation Guide* and README file for a list of supported distributions and requirements.

The compiler, its generated object programs, and runtime library can run on POWER8<sup>®</sup> systems with the required software, disk space, and virtual storage.

The POWER8 processor is a type of PowerPC<sup>®</sup> processor. In this document, any statement or reference to the PowerPC processor also applies to the POWER8 processor.

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.

---

### Language support

The XL Fortran language consists of the following:

- Partial ISO Technical specification 29113 support (referred to as TS 29113), defined in the document *Information technology -- Further interoperability of Fortran with C, ISO/IEC TS 29113:2012*.
- Partial ISO Fortran 2008 language standard (referred to as Fortran 2008 or F2008), defined in the document *Information technology - Programming languages - Part 1: Base language, ISO/IEC 1539-1:2010*.
- The full ISO Fortran 2003 language standard (referred to as Fortran 2003 or F2003), defined in the document *Information technology - Programming languages - Part 1: Base language, ISO/IEC 1539-1:2004*.
- The full ISO Fortran 95 language standard (referred to as Fortran 95 or F95), defined in the document *Information technology - Programming languages - Fortran - Part 1: Base language, ISO/IEC 1539-1:1997*.
- The full American National Standard Fortran 90 language (referred to as Fortran 90 or F90), defined in the documents *American National Standard Programming Language Fortran 90, ANSI X3.198-1992* and *Information technology - Programming languages - Fortran, ISO/IEC 1539-1:1991 (E)*. This language has a superset of the features found in the FORTRAN 77 standard. It adds many more features that are intended to shift more of the tasks of error checking, array processing, memory allocation, and so on from the programmer to the compiler.
- Extensions to the Fortran standard:
  - Common Fortran language extensions defined by other compiler vendors, in addition to those defined by IBM
  - Industry extensions that are found in Fortran products from various compiler vendors

- Extensions specified in SAA Fortran
- Extensions of the Vector Multimedia Extension (VMX) and the Vector Scalar Extension (VSX) intrinsic functions

In the *XL Fortran Language Reference*, extensions to the Fortran 2003 language and Fortran 2008 language are marked as described in the *Conventions: Qualifying elements* section.

---

## Source-code conformance checking

To help you find possible problems when you port your application to or from different FORTRAN 77, Fortran 90, Fortran 95, Fortran 2003, or Fortran 2008 standards, the XL Fortran compiler provides options to warn you about features that no longer conform to certain Fortran definitions.

If you specify the appropriate compiler options, the XL Fortran compiler checks source statements for conformance to the following Fortran language definitions:

- Partial Fortran 2008 Standard (**-qlanglvl=2008std** option)
- Full Fortran 2003 Standard (**-qlanglvl=2003std** option)
- Full Fortran 95 Standard (**-qlanglvl=95std** option)
- Full American National Standard Fortran 90 Standard (**-qlanglvl=90std** option)
- Full American National Standard FORTRAN 77 Standard (**-qlanglvl=77std** option)
- Fortran 2008, less any obsolescent features (**-qlanglvl=2008pure** option)
- Fortran 2003, less any obsolescent features (**-qlanglvl=2003pure** option)
- Fortran 95, less any obsolescent features (**-qlanglvl=95pure** option)
- Fortran 90, less any obsolescent features (**-qlanglvl=90pure** option)
- Partial support for Technical Specification 29113 for further interoperability with C (**-qlanglvl=ts** option)
- Technical Specification 29113 supplements to Fortran 2008, less any obsolescent features (**-qlanglvl=tspure** option)
- IBM SAA FORTRAN (**-qsaa** option)

You can also use the **langlvl** option for conformance checking.

**Note:** Fortran 2008 conformance checking is based on the current subset implementation of this standard.

---

## Highly configurable compiler

You can invoke the compiler by using the following commands:

- xlf
- xlf\_r
- f77
- fort77
- xlf90
- xlf90\_r
- f90
- xlf95
- xlf95\_r
- f95
- xlf2003
- xlf2003\_r
- f2003

- `xlF2008`
- `xlF2008_r`
- `f2008`

The following commands maintain maximum compatibility with the behavior and I/O formats of XL Fortran Version 2:

- For `.f`, `.F`, `.f77` and `.F77` files: `xlF` and `xlF_r`
- For any source files: `f77` and `fort77`

The `xlF90`, `xlF90_r`, and `f90` commands provide more Fortran 90 conformance and some implementation choices for efficiency and usability. The `f95`, `xlF95` and `xlF95_r` commands provide more Fortran 95 conformance and some implementation choices for efficiency and usability. The `xlF2003`, `xlF2003_r`, and `f2003` commands provide more Fortran 2003 conformance and some implementation choices for efficiency and usability. The `xlF2008`, `xlF2008_r`, and `f2008` commands provide more Fortran 2008 conformance and some implementation choices for efficiency and usability. The `f77` or `fort77` command provides maximum compatibility with the XPG4 behavior.

The main difference between the set of `xlF_r`, `xlF90_r`, `xlF95_r`, `xlF2003_r`, and `xlF2008_r` commands and the set of `xlF`, `xlF90`, `f90`, `xlF95`, `f95`, `xlF2003`, `f2003`, `xlF2008`, `f2008`, `f77`, and `fort77` commands is that the first set links and binds the object files to the threadsafe components (libraries, and so on). You can have this behavior with the second set of commands by using the `-F` compiler option to specify the configuration file stanza to use. For example:

```
xlF -F/opt/ibm/xlF/15.1.1/etc/xlF.cfg:xlF_r
```

You can control the actions of the compiler through a set of options. The different categories of options help you to debug, to optimize and tune program performance, to select extensions for compatibility with programs from other platforms, and to do other common tasks that would otherwise require changing the source code.

To simplify the task of managing many different sets of compiler options, you can edit the default configuration file or use a customized configuration file instead of creating many separate aliases or shell scripts.

### Related information

- “Using custom compiler configuration files” on page 9
- “Compiling XL Fortran programs” on page 17
- Chapter 5, “Summary of compiler options by functional category,” on page 47 and Chapter 6, “Detailed descriptions of the XL Fortran compiler options,” on page 65

---

## Diagnostic listings

The compiler output listing has optional sections that you can include or omit. For information about the applicable compiler options and the listing itself, refer to “Listings, messages, and compiler information” on page 56 and Chapter 9, “Understanding XL Fortran compiler listings,” on page 291.

The `-S` option gives you a true assembler source file.

---

## Symbolic debugger support

You can instruct XL Fortran to include debugging information in your compiled objects by using different levels of the **-g** or **-qdbg** compiler option.

For details, see **-g** or **-qdbg**.

The debugging information can be examined by **gdb** or any other symbolic debugger to help you debug your programs.

---

## Program optimization

The XL Fortran compiler helps you control the optimization of your programs:

- You can select different levels of compiler optimizations.
- You can turn on separate optimizations for loops, floating point, and other categories.
- You can optimize a program for a particular class of machines or for a very specific machine configuration, depending on where the program will run.

The *XL Fortran Optimization and Programming Guide* provides a road map and optimization strategies for these features.

---

## Chapter 3. Setting up and customizing XL Fortran

This section explains how to customize XL Fortran settings for yourself or all users. The full installation procedure is beyond the scope of this section, which refers you to the documents that cover the procedure in detail.

This section can also help you to diagnose problems that relate to installing or configuring the compiler.

Some of the instructions require you to be a superuser, and so they are only applicable if you are a system administrator.

---

### Where to find installation instructions

To install the compiler, refer to these documents (preferably in this order):

1. Read the file called `/opt/ibm/xlf/15.1.1/README`, and follow any directions it gives. It contains information that you must know and possibly distribute to other people who use XL Fortran.
2. Read the *XL Fortran Installation Guide* to see if there are any important notices you need to be aware of or any updates you might need to apply to your system before doing the installation.

#### Using default installation tools

- You should be familiar with the RPM Package Manager (RPM) for installing this product on SLES. For information on using RPM, visit the RPM Web page at URL <http://www.rpm.org/>, or type `rpm --help` at the command line.  
If you are already experienced with software installation, you can use the `rpm` command to install all the images from the distribution medium.
- You should be familiar with the Debian Package Manager (dpkg) for installing this product on Ubuntu. For information on using dpkg, you can view the dpkg man page at URL <http://manpages.ubuntu.com/manpages/trusty/en/man1/dpkg.1.html>, or type `dpkg --help` at the command line.

---

### Correct settings for environment variables

You can set and export a number of environment variables for use with the operating system. The following sections deal with the environment variables that have special significance to the XL Fortran compiler, application programs, or both.

#### Environment variable basics

You can set the environment variables from shell command lines or from within shell scripts. (For more information about setting environment variables, see the man page help for the shell you are using.) If you are not sure which shell is in use, a quick way to find out is to issue `echo $SHELL` to show the name of the current shell.

To display the contents of an environment variable, enter the command `echo $var_name`.

**Note:** For the remainder of this document, most examples of shell commands use **Bash** notation instead of repeating the syntax for all shells.

## Setting library search paths

If your executable program is linked with shared libraries, you must set the runtime library search paths. You can use one of the following three ways to set runtime library search paths:

- When linking the shared library into the executable, use the **-R** (or **-rpath**) compiler/link option.
- Before linking the shared library into the executable, set the **LD\_RUN\_PATH** environment variable.
- Set the **LD\_LIBRARY\_PATH** environment variable.

For example:

```
# Compile and link
xlf95 -L/usr/lib/mydir1 -R/usr/lib/mydir1 -L/usr/lib/mydir2 -R/usr/lib/mydir2
      -lmylib1 -lmylib2 test.f

# -L directories are searched at link time for both static and shared libraries.
# -R directories are searched at run time for shared libraries.
```

For more information about the linker option **-R** (or **-rpath**), and environment variables **LD\_RUN\_PATH** and **LD\_LIBRARY\_PATH**, see the man pages for the **ld** command.

## Profile-directed feedback environment variables

The following list includes profile-directed feedback (PDF) environment variables that you can use with the **-qpdf** compiler option:

- **PDF\_BIND\_PROCESSOR:**

If you want to bind your application to the specified processor for cache-miss profiling, set the **PDF\_BIND\_PROCESSOR** environment variable. Processor 0 is set by default.

- **PDFDIR:**

When you compile a Fortran program with the **-qpdf** compiler option, you can specify the directory where profiling information is stored by setting the **PDFDIR** environment variable to the name of the directory. The compiler creates files to hold the profile information. XL Fortran updates the files when you run an application that is compiled with the **-qpdf1** option.

Problems can occur if the profiling information is stored in a wrong place or is updated by more than one application. To avoid these problems, you must follow these guidelines:

- Always set the **PDFDIR** environment variable when using the **-qpdf** option. Make sure that the directory specified by the **PDFDIR** environment variable exists; otherwise, the compiler issues a warning message.
- Store the profiling information for each application in a different directory, or use the **-qpdf1=pdfname**, **-qpdf1=exename** option to explicitly name the temporary profiling files according to the template provided.
- Leave the value of the **PDFDIR** environment variable unchanged until you have completed the PDF process (compiling, running, and compiling again) for the application.

- **PDF\_PM\_EVENT:**

When you run an application compiled with **-qpdf1=level=2** and want to gather different levels of cache-miss profiling information, set the **PDF\_PM\_EVENT** environment variable to **L1MISS**, **L2MISS**, or **L3MISS** (if applicable) accordingly.

- **PDF\_WL\_ID:**

The PDF\_WL\_ID environment variable is used to distinguish the sets of PDF counters that are generated by multiple training runs of the user program. Each run receives distinct input.

By default, PDF counters for training runs after the first training run are added to the first and the only set of PDF counters. This behavior can be changed by setting the PDF\_WL\_ID environment variable before each PDF training run. You can set PDF\_WL\_ID to an integer value in the range 1 - 65535. The PDF run time then uses this number to tag the set of PDF counters that are generated by this training run. After all the training runs complete, the PDF profile file contains multiple sets of PDF counters, each set with an ID number.

## **TMPDIR: Specifying a directory for temporary files**

The XL Fortran compiler creates a number of temporary files for use during compilation. An XL Fortran application program creates a temporary file at run time for a file opened with `STATUS='SCRATCH'`. By default, these files are placed in the directory `/tmp`.

If you want to change the directory where these files are placed, perhaps because `/tmp` is not large enough to hold all the temporary files, set and export the `TMPDIR` environment variable before running the compiler or the application program.

If you explicitly name a scratch file by using the `XLFSCRATCH_unit` method described below, the `TMPDIR` environment variable has no effect on that file.

## **XLFSCRATCH\_unit: Specifying names for scratch files**

To give a specific name to a scratch file, you can set the runtime option `scratch_vars=yes`; then set one or more environment variables with names of the form `XLFSCRATCH_unit` to file names to use when those units are opened as scratch files. See *Naming scratch files* in the *XL Fortran Optimization and Programming Guide* for examples.

## **XLFUNIT\_unit: Specifying names for implicitly connected files**

To give a specific name to an implicitly connected file or a file opened with no `FILE=` specifier, you can set the runtime option `unit_vars=yes`; then set one or more environment variables with names of the form `XLFUNIT_unit` to file names. See *Naming files that are connected with no explicit name* in the *XL Fortran Optimization and Programming Guide* for examples.

---

## **Using custom compiler configuration files**

The XL Fortran compiler generates a default configuration file `/opt/ibm/xlf/15.1.1/etc/xlf.cfg.$OSRelease.gcc$gccVersion`. For example, `/opt/ibm/xlf/15.1.1/etc/xlf.cfg.sles.12.gcc.4.8.2` or `/opt/ibm/xlf/15.1.1/etc/xlf.cfg.ubuntu.14.04.gcc.4.8.2` at installation time. (See the *XL Fortran Installation Guide* for more information on the various tools you can use to generate the configuration file during installation.) The configuration file specifies information that the compiler uses when you invoke it.

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

If you want users to be able to choose among several sets of compiler options, you might want to use custom configuration files for specific needs. For example, you might want to enable `-qlist` by default for compilations using the `xlf` compiler invocation command. This is to avoid forcing your users to specify this option on the command line for every compilation, because `-qnolist` is automatically in effect every time the compiler is called with the `xlf` command.

You have several options for customizing configuration files:

- You can directly edit the default configuration file. In this case, the customized options will apply for all users for all compilations. The disadvantage of this option is that you will need to reapply your customizations to the new default configuration file that is provided every time you install a compiler update.
- You can use the default configuration file as the basis of customized copies that you specify at compile time with the `-F` option. In this case, the custom file overrides the default file on a per-compilation basis.

**Note:** This option requires you to reapply your customization after you apply service to the compiler.

- You can create custom, or user-defined, configuration files that are specified at compile time with the `XL_FUSR_CONFIG` environment variable. In this case, the custom user-defined files complement, rather than override, the default configuration file, and they can be specified on a per-compilation or global basis. The advantage of this option is that you do not need to modify your existing, custom configuration files when a new system configuration file is installed during an update installation. Procedures for creating custom, user-defined configuration files are provided below.

**Related information:**

- “-F” on page 71

## Creating custom configuration files

If you use the `XL_FUSR_CONFIG` environment variable to instruct the compiler to use a custom user-defined configuration file, the compiler examines and processes the settings in that user-defined configuration file before looking at the settings in the default system configuration file.

To create a custom user-defined configuration file, you add stanzas which specify multiple levels of the `use` attribute. The user-defined configuration file can reference definitions specified elsewhere in the same file, as well as those specified in the system configuration file. For a given compilation, when the compiler looks for a given stanza, it searches from the beginning of the user-defined configuration file and follows any other stanza named in the `use` attribute, including those specified in the system configuration file.

If the stanza named in the `use` attribute has a name different from the stanza currently being processed, the search for the `use` stanza starts from the beginning of the user-defined configuration file. This is the case for stanzas A, C, and D which you see in the following example. However, if the stanza in the `use` attribute has the same name as the stanza currently being processed, as is the case of the two B stanzas in the example, the search for the `use` stanza starts from the location of the current stanza.

The following example shows how you can use multiple levels for the **use** attribute. This example uses the **options** attribute to help show how the **use** attribute works, but any other attributes, such as **libraries** can also be used.

```
A: use =DEFLT
   options=<set of options A>
B: use =B
   options=<set of options B1>
B: use =D
   options=<set of options B2>
C: use =A
   options=<set of options C>
D: use =A
   options=<set of options D>
DEFLT:
   options=<set of options Z>
```

Figure 1. Sample configuration file

In this example:

- stanza A uses option sets *A* and *Z*
- stanza B uses option sets *B1*, *B2*, *D*, *A*, and *Z*
- stanza C uses option sets *C*, *A*, and *Z*
- stanza D uses option sets *D*, *A*, and *Z*

Attributes are processed in the same order as the stanzas. The order in which the options are specified is important for option resolution. Ordinarily, if an option is specified more than once, the last specified instance of that option wins.

By default, values defined in a stanza in a configuration file are added to the list of values specified in previously processed stanzas. For example, assume that the `XLF_USR_CONFIG` environment variable is set to point to the user-defined configuration file at `~/userconfig1`. With the user-defined and default configuration files shown in the following example, the compiler references the `xlf` stanza in the user-defined configuration file and uses the option sets specified in the configuration files in the following order: *A1*, *A*, *D*, and *C*.

```
xlf: use=xlf
     options= <A1>

DEFLT: use=DEFLT
       options=<D>
```

Figure 2. Custom user-defined configuration file `~/userconfig1`

```
xlf: use=DEFLT
     options=<A>

DEFLT:
     options=<C>
```

Figure 3. Default configuration file `xlf.cfg`

## Overriding the default order of attribute values

You can override the default order of attribute values by changing the assignment operator(=) for any attribute in the configuration file.

Table 4. Assignment operators and attribute ordering

Assignment Operator	Description
<code>--</code>	Prepend the following values before any values determined by the default search order.
<code>:=</code>	Replace any values determined by the default search order with the following values.
<code>+=</code>	Append the following values after any values determined by the default search order.

For example, assume that the `XLF_USR_CONFIG` environment variable is set to point to the custom user-defined configuration file at `~/userconfig2`.

#### Custom user-defined configuration file

`~/userconfig2`

```
xlf_prepend: use=xlf
              options--<B1>
xlf_replace: use=xlf
              options:=<B2>
xlf_append:  use=xlf
              options+=<B3>
```

```
DEFLT: use=DEFLT
        options=<D>
```

#### Default configuration file xlf.cfg

```
xlf: use=DEFLT
      options=<B>
DEFLT:
      options=<C>
```

The stanzas in the preceding configuration files use the following option sets, in the following orders:

1. stanza `xlf` uses `B`, `D`, and `C`
2. stanza `xlf_prepend` uses `B1`, `B`, `D`, and `C`
3. stanza `xlf_replace` uses `B2`
4. stanza `xlf_append` uses `B`, `D`, `C`, and `B3`

You can also use assignment operators to specify an attribute more than once. For example:

```
xlf:
  use=xlf
  options--Isome_include_path
  options+=some options
```

Figure 4. Using additional assignment operations

## Examples of stanzas in custom configuration files

```
DEFLT: use=DEFLT
        options = -g
```

This example specifies that the `-g` option is to be used in all compilations.

```
xlf: use=xlf
      options+=-qlist
xlf_r: use=xlf_r
        options+=-qlist
```

This example specifies that `-qlist` is to be used for any compilation called by the `xlf` and `xlf_r` commands. This `-qlist` specification overrides the default setting of `-qlist` specified in the system configuration file.

```
DEFLT: use=DEFLT
        libraries=-L/home/user/lib,-lmylib
```

This example specifies that all compilations should link with `/home/user/lib/libmylib.a`.

## Using IBM XL Fortran for Linux, V15.1.1 with the Advance Toolchain

IBM XL Fortran for Linux, V15.1.1 supports IBM Advance Toolchain 8.0, which is a set of open source development tools and runtime libraries. With IBM Advance Toolchain 8.0, you can take advantage of the latest POWER® hardware features on Linux, especially the tuned libraries. For more information about the Advance Toolchain 8.0, see IBM Advance Toolchain for PowerLinux™ Documentation.

To use IBM XL Fortran for Linux, V15.1.1 with the Advance Toolchain, take the following steps:

1. Install the **at8.0** packages into the default installation location. For instructions, see IBM Advance Toolchain for PowerLinux Documentation.
2. Run the **xlf\_configure** utility to create the **xlf.at.cfg** configuration file. In the **xlf.at.cfg** configuration file, all other entities except the XL Fortran compiler are directed to those of the Advance Toolchain. The entities include the linker, headers, and runtime libraries.

**Note:** To run the **xlf\_configure** utility, you must either become the root user or use the `sudo` command.

- If you installed the compiler in the default location, issue the following command:

```
xlf_configure -at
```

- If you installed the compiler in a nondefault installation (NDI) location, issue the following command:

```
xlf_configure -at -ibmcmp $ndi_path
```

where *\$ndi\_path* is the directory in which you installed the compiler.

3. Invoke the XL compiler with the Advance Toolchain support.
  - If you installed the compiler in the default location, issue the following command to invoke the Fortran compiler:

```
/opt/ibm/xlf/15.1.1/bin/xlf_at
```
  - If you installed the compiler in an NDI location, issue the following command:

```
$ndi_path/xlf/15.1.1/bin/xlf_at
```

**Note:** If you use the XL compiler with the Advance Toolchain support to build your application, your application can run only under the Advance Toolchain environment because the application depends on the runtime library of the Advance Toolchain. If you copy the application to run on other machines, ensure that the Advance Toolchain, or at least the runtime library of the Advance Toolchain is available on those machines.

---

## Editing the default configuration file

The configuration file specifies information that the compiler uses when you invoke it. XL Fortran provides the default configuration file `/opt/ibm/xlf/15.1.1/etc/xlf.cfg` at installation time.

If you want many users to be able to choose among several sets of compiler options, you may want to add new named stanzas to the configuration file and to create new commands that are links to existing commands. For example, you could specify something similar to the following to create a link to the **xlf95** command:

```
ln -s /opt/ibm/xlf/15.1.1/bin/xlf95 /home/username/bin/xlf95
```

When you run the compiler under another name, it uses whatever options, libraries, and so on, that are listed in the corresponding stanza.

**Note:**

1. The configuration file contains other named stanzas to which you may want to link.
2. If you make any changes to the configuration file and then move or copy your makefiles to another system, you will also need to copy the changed configuration file.
3. You cannot use tabs as separator characters in the configuration file. If you modify the configuration file, make sure that you use spaces for any indentation.

## Configuration file attributes

The configuration file contains the following attributes:

**use** The named and local stanzas provide the values for attributes. For single-valued attributes, values in the **use** attribute apply if there is no value in the local, or default, stanza. For comma-separated lists, the values from the **use** attribute are added to the values from the local stanza. You can only use a single level of the **use** attribute. Do not specify a **use** attribute that names a stanza with another **use** attribute.

**crt\_64** The path name of the object file which contains the startup code. This object file is passed as the first parameter to the linkage editor.

**mcrt\_64** Same as for **crt\_64**, but the object file contains profiling code for the **-p** option.

**gcrt\_64** Same as **crt\_64**, but the object file contains profiling code for the **-pg** option.

**gcc\_libs\_64** The linker options to specify the path to the GCC libraries and to link the GCC library.

**gcc\_path\_64** Specifies the path to the 64-bit tool chain.

**cpp** The absolute path name of the C preprocessor, which is automatically called for files ending with a specific suffix (usually **.F**).

**xlf** The absolute path name of the main compiler executable file. The compiler commands are driver programs that execute this file.

**hot** Absolute path name of the program that does array language transformations.

**ipa** Absolute path name of the program that performs interprocedural optimizations, loop optimizations, and program parallelization.

**code** The absolute path name of the optimizing code generator.

**as\_64** The absolute path name of the assembler.

**ld\_64** The absolute path name of the linker.

**bolt** Absolute path name of the binder.

**options**

A string of options that are separated by commas. The compiler processes these options as if you entered them on the command line before any other option. This attribute lets you shorten the command line by including commonly used options in one central place.

**cppoptions**

A string of options that are separated by commas, to be processed by **cpp** (the C preprocessor) as if you entered them on the command line before any other option. This attribute is needed because some **cpp** options are usually required to produce output that can be compiled by XL Fortran. The default is **-C**, which preserves any C-style comments in the output. Also, refer to the “-qfpp” on page 131 and “-qppsuborigarg” on page 209 options for other useful **cpp** options.

**Note:** You can specify **-C!** preprocessor option on the command line (**-WF**, **-C!**) to override the default setting.

**xlfopt** Lists names of options that are assumed to be compiler options, for cases where, for example, a compiler option and a linker option use the same letter. The list is a concatenated set of single-letter flags. Any flag that takes an argument is followed by a colon, and the whole list is enclosed by double quotation marks.

**asopt** Lists names of options that are assumed to be assembler options for cases where, for example, a compiler option and an assembler option use the same letter. The list is a concatenated set of single-letter flags. Any flag that takes an argument is followed by a colon, and the whole list is enclosed by double quotation marks. You may find it more convenient to set up this attribute than to pass options to the assembler through the **-W** compiler option.

**ldopt** Lists names of options that are assumed to be linker options for cases where, for example, a compiler option and a linker option use the same letter. The list is a concatenated set of single-letter flags. Any flag that takes an argument is followed by a colon, and the whole list is enclosed by double quotation marks.

You might find it more convenient to set up this attribute than to pass options to the linker through the **-W** compiler option. However, most unrecognized options are passed to the linker anyway.

**cppsuffix**

The suffix that indicates a file must be preprocessed by the C preprocessor (**cpp**) before being compiled by XL Fortran. The default is **F**.

**fsuffix** The allowed suffix for Fortran source files. The default is **f**. The compiler requires that all source files in a single compilation have the same suffix. Therefore, to compile files with other suffixes, such as **f95**, you must change this attribute in the configuration file or use the **-qsuffix** compiler option. For more information on **-qsuffix**, see “-qsuffix” on page 244.

**osuffix**

The suffix used to recognize object files that are specified as input files. The default is **o**.

**ssuffix**

The suffix used to recognize assembler files that are specified as input files. The default is **s**.

**defaultmsg**

Absolute path name of the default message files.

**include\_64**

Indicates the search path that is used for compilation include files, module symbol files, and submodule symbol files.

**Notes:**

- To specify multiple search paths for compilation include files, separate each path location with a comma as follows:  
include = -I/path1, -I/path2, ...
- You can use the “-F” on page 71 option to select a different configuration file, a specific stanza in the configuration file, or both.

**Related information**

- “XL Fortran input files” on page 23
- “XL Fortran output files” on page 24

---

## Determining which level of XL Fortran is installed

Sometimes, you may not be sure which level of XL Fortran is installed on a particular machine. You would need to know this information before contacting software support.

To check whether the latest level of the product has been installed through the system installation procedure, issue the command:

**On SLES**

```
rpm -qa | grep xlf.15.1.1 | xargs rpm -qi
```

**On Ubuntu**

```
dpkg -l xlf.15.1.1
```

The result includes the version, release, modification, and fix level of the compiler image installed on the system.

You can also use the **-qversion** compiler option to display the version, release, and level of the compiler and its components.

---

## Running two levels of XL Fortran

It is possible for two different levels of the XL Fortran compiler to coexist on one system. This allows you to invoke one level by default and to invoke the other one whenever you explicitly choose to.

To do this, consult the *XL Fortran Installation Guide* for details.

---

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

Most Fortran program development consists of a repeating cycle of editing, compiling and linking (which is by default a single step), and running. For problems encountered during the development cycle, refer to the following sections for help.

### Prerequisite information:

1. Before you can use the compiler, all the required Linux settings (for example, certain environment variables and storage limits) must be correct for your user ID; for details, see “Correct settings for environment variables” on page 7.
2. To learn more about writing and optimizing XL Fortran programs, refer to the *XL Fortran Language Reference* and *XL Fortran Optimization and Programming Guide*.

---

### Editing XL 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 have a suffix of **.f** unless the **fsuffix** attribute in the configuration file specifies a different suffix, or the **-qsuffix** compiler option is used. You can also use a suffix of **.F** if the program contains C preprocessor (**cpp**) directives that must be processed before compilation begins. Source files with the **.f77**, **.f90**, **.f95**, **.f03**, or **.f08** suffix are also valid.

For the Fortran source program to be a valid program, it must conform to the language definition that is specified in the *XL Fortran Language Reference*.

---

### Compiling XL Fortran programs

To compile a source program, use one of the **xlf90**, **xlf90\_r**, **f90**, **xlf95**, **xlf95\_r**, **f95**, **xlf2003**, **xlf2003\_r**, **f2003**, **xlf2008**, **xlf2008\_r**, **f2008**, **xlf**, **xlf\_r**, **f77**, or **fort77** commands, which have the form:



These commands all accept essentially the same Fortran language. The main difference is that they use different default options (which you can see by reading the configuration file `/opt/ibm/xlf/15.1.1/etc/xlf.cfg`).

The invocation command performs the necessary steps to compile the Fortran source files, assemble any `.s` files, and link the object files and libraries into an executable program. In particular, the `xlf_r`, `xlf90_r`, `xlf95_r`, `xlf2003_r`, and `xlf2008_r` commands use the components for multi-threading (libraries, and so on) to link and bind object files.

The following table summarizes the invocation commands that you can use:

*Table 5. XL Fortran Invocation commands*

Driver Invocation	Path or Location	Chief Functionality	Linked Libraries
<code>xlf</code>	<code>/opt/ibm/xlf/15.1.1/bin</code>	Selected Fortran language level	<code>libxlf90.so</code>
<code>xlf_r</code>	<code>/opt/ibm/xlf/15.1.1/bin</code>	The threadsafe version of the selected language level	<code>libxlf90_r.so</code>
<code>f77</code> , <code>fort77</code>	<code>/opt/ibm/xlf/15.1.1/bin</code>	FORTRAN 77	<code>libxlf90.so</code>
<code>xlf90</code> , <code>f90</code>	<code>/opt/ibm/xlf/15.1.1/bin</code>	Fortran 90	<code>libxlf90.so</code>
<code>xlf90_r</code>	<code>/opt/ibm/xlf/15.1.1/bin</code>	Threadsafe Fortran 90	<code>libxlf90_r.so</code>
<code>xlf95</code> , <code>f95</code>	<code>/opt/ibm/xlf/15.1.1/bin</code>	Fortran 95	<code>libxlf90.so</code>
<code>xlf95_r</code>	<code>/opt/ibm/xlf/15.1.1/bin</code>	Threadsafe Fortran 95	<code>libxlf90_r.so</code>
<code>xlf2003</code>	<code>/opt/ibm/xlf/15.1.1/bin</code>	Fortran 2003	<code>libxlf90.so</code>
<code>xlf2003_r</code>	<code>/opt/ibm/xlf/15.1.1/bin</code>	Threadsafe Fortran 2003	<code>libxlf90.so</code>
<code>f2003</code>	<code>/opt/ibm/xlf/15.1.1/bin</code>	Fortran 2003	<code>libxlf90.so</code>
<code>xlf2008</code>	<code>/opt/ibm/xlf/15.1.1/bin</code>	Fortran 2008	<code>libxlf90.so</code>
<code>xlf2008_r</code>	<code>/opt/ibm/xlf/15.1.1/bin</code>	Threadsafe Fortran 2008	<code>libxlf90.so</code>
<code>f2008</code>	<code>/opt/ibm/xlf/15.1.1/bin</code>	Fortran 2008	<code>libxlf90.so</code>

The `xlf` and `xlf_r` invocation commands select the appropriate language level according to the suffix of the source file names. The other invocation commands behave consistently no matter what the Fortran source file name suffix is. For example, when you compile the following source files:

```
xlf program1.f program2.f90 program3.f95 program4.f03 program5.f08
```

the behavior is as follows:

- The `program1.f` file is compiled as if the invocation command was `f77`.
- The `program2.f90` file is compiled as if the invocation command was `xlf90`.

- The program3.f95 file is compiled as if the invocation command was **xlF95**.
- The program4.f03 file is compiled as if the invocation command was **xlF2003**.
- The program5.f08 file is compiled as if the invocation command was **xlF2008**.

**libxlf90.so** is provided for both threaded and non-threaded applications. XL Fortran determines at run time whether your application is threaded.

XL Fortran provides the library **libxlf90\_t.so**, in addition to **libxlf90\_r.so**. **libxlf90\_t.so** exports the same entry points as **libxlf90\_r.so** does. The library **libxlf90\_r.so** is a superset of **libxlf90\_t.so**. The file **xlF.cfg** is set up to link to **libxlf90\_r.so** automatically when you use the **xlF90\_r**, **xlF95\_r**, and **xlF\_r** commands. **libxlf90\_t.so** is a partial thread-support runtime library. Unlike **libxlf90\_r.so**, **libxlf90\_t.so** does not provide thread synchronization and routines in **libxlf90\_t.so** are not thread-reentrant. Therefore, only one Fortran thread at a time can perform I/O operations or invoke Fortran intrinsics. You can use **libxlf90\_t.so** instead of **libxlf90\_r.so** in multithread applications where there is only one Fortran thread, to avoid the thread synchronization overhead in **libxlf90\_r.so**.

When you bind a multithreaded executable with multiple Fortran threads, **libxlf90\_r.so** should be used. Note that using the **xlF\_r**, **xlF90\_r**, **xlF95\_r**, **xlF2003\_r**, or **xlF2008\_r** invocation command ensures the proper linking.

The invocation commands have the following implications for directive triggers:

- For **f77**, **fort77**, **f90**, **f95**, **f2003**, **xlF**, **xlF90**, **xlF95**, **xlF2003**, and **xlF2008**, the directive trigger is **IBM\*** by default.
- For all other commands, the directive triggers are **IBM\*** and **IBMT** by default.

If you specify the **-qsmp** option, the following occurs:

- The compiler turns on automatic parallelization.
- In addition to **IBM\*** and **IBMT**, the compiler also recognizes the **IBMP**, **SMP\$**, and **\$OMP** directive triggers.
- If you specify the **-qsmp=omp** option, the compiler only recognizes the **\$OMP** directive trigger.

## Compiling XL Fortran Version 2 programs

When used to compile **.f**, **.F**, **.f77**, or **.F77** files, **xlF** maintains, wherever possible, compatibility with existing programs by using the same I/O formats as earlier versions of XL Fortran and some implementation behavior compatible with FORTRAN 77.

If the configuration file has not been customized, **f77** is identical to **xlF** when **xlF** is used to compile the **.f**, **.F**, **.f77**, or **.F77** files.

You may find that you need to continue using these commands for compatibility with existing makefiles and build environments. However, be aware that programs that you compile with these commands may not conform to the newer standard in subtle ways.

## Compiling Fortran 90 or Fortran 95 programs

The **f90**, **xlF90**, and **xlF90\_r** commands make your programs conform more closely to the Fortran 90 standard than do the other invocation commands. The **f95**, **xlF95**, and **xlF95\_r** commands make your programs conform more closely to the Fortran 95 standard than do the other invocation commands. **f90**, **xlF90**, **xlF90\_r**, **f95**, **xlF95**, and **xlF95\_r** are the preferred commands for compiling any new programs. They all

accept Fortran 90 free source form by default; to use them for fixed source form, you must use the **-qfixed** option. I/O formats are slightly different between these commands and the other commands. I/O formats also differ between the set of **f90**, **xlf90** and **xlf90\_r** commands and the set of **f95**, **xlf95** and **xlf95\_r** commands. We recommend that you switch to the Fortran 95 formats for data files whenever possible.

By default, the **f90**, **xlf90**, and **xlf90\_r** commands do not conform completely to the Fortran 90 standard. Also, by default, the **f95**, **xlf95**, and **xlf95\_r** commands do not conform completely to the Fortran 95 standard. If you need full Fortran 90 or Fortran 95 compliance, compile with any of the following additional compiler options (and suboptions):

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

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

```
export XLF RTEOPTS="err_recovery=no:langlvl=90std"  
export XLF RTEOPTS="err_recovery=no:langlvl=95std"
```

The default settings are intended to provide the best combination of performance and usability. Therefore, it is usually a good idea to change them only when required. Some of the options above are only required for compliance in very specific situations.

## Compiling Fortran 2003 programs

The **f2003**, **xlf2003**, and **xlf2003\_r** commands make your programs conform more closely to the Fortran 2003 Standard than do the other invocation commands. The Fortran 2003 commands accept free source form by default. I/O formats for the Fortran 2003 commands are the same as for the **f95**, **xlf95**, and **xlf95\_r** commands. The Fortran 2003 commands format infinity and NaN floating-point values differently from previous commands. The Fortran 2003 commands enable polymorphism by default.

By default, the **f2003**, **xlf2003**, and **xlf2003\_r** commands do not conform completely to the Fortran 2003 standard. If you need full compliance, compile with the following additional compiler suboptions:

```
-qlanglvl=2003std -qnodirective -qnoescape -qfloat=nomaf:rndsngl:nofold  
-qnoswapomp -qstrictieemod
```

Also specify the following runtime options:

```
XLF RTEOPTS="err_recovery=no:langlvl=2003std:iostat_end=2003std:  
internal_nldelim=2003std"
```

## Compiling Fortran 2008 programs

The **f2008**, **xlf2008**, and **xlf2008\_r** commands make your programs conform more closely to the Fortran 2008 Standard than do the other invocation commands. The Fortran 2008 commands accept free source form by default. I/O formats for the Fortran 2008 commands are the same as for the **f95**, **xlf95**, **xlf95\_r**, **xlf95\_r7**, **f2003**, **xlf2003**, and **xlf2003\_r** commands. The Fortran 2008 commands format infinity and NaN floating-point values in the same way as the Fortran 2003 commands. The Fortran 2008 commands enable polymorphism by default.

By default, the `f2008`, `xlF2008`, and `xlF2008_r` commands do not conform completely to the Fortran 2008 standard. If you need full compliance, compile with the following additional compiler suboptions:

```
-qlanglvl=2008std -qnodirective -qnoescape -qfloat=nomaf:rndsngl:nofold  
-qnoswapomp -qstrictieemod
```

Also specify the following runtime options:

```
XLFRTEOPTS="err_recovery=no:langlvl=2008std:iostat_end=2003std:  
internal_nldelim=2003std"
```

## Related information

- “Compiling Fortran 2003 programs” on page 20

## Compiling and linking a library

### Compiling a static library

To compile a static library:

1. Compile each source file into an object file, with no linking. For example:

```
xlF -c bar.f example.f
```

2. Use the `ar` command to add the generated object files to an archive library file. For example:

```
ar -rv libfoo.a bar.o example.o
```

### Compiling a shared library

To compile a shared library, you must use the `-qpic` option.

Use the following steps to compile a shared library:

1. Compile your source files into an object file, with no linking. For example:

```
xlF -qpic -c foo.f
```

2. Use the `-qmkshrobj` compiler option to create a shared object from the generated object files. For example:

```
xlF -qmkshrobj -o libfoo.so foo.o
```

### Related information in the *XL Fortran Compiler Reference*



`-qpic`



`-qmkshrobj`

### Linking a library to an application

You can use the same command string to link a static or shared library to your main program. For example:

```
xlF -o myprogram main.f -Ldirectory [-Rdirectory] -ltest
```

At compile time, you instruct the linker to search for `libtest.so` in the first directory specified via the `-L` option. If `libtest.so` is not found, the linker searches for `libtest.a`. If neither file is found, the search continues with the next directory specified via the `-L` option.

At run time, the runtime linker searches for `libtest.so` in the first directory specified via the `-R` option. If `libtest.so` is not found, the search continues with the next directory specified via the `-R` option. The path specified by the `-R` option can be overridden at run time via the `LD_LIBRARY_PATH` environment variable.

For additional linkage options, including options that modify the default behavior, see the operating system **ld** documentation .

**Related information in the *XL Fortran Compiler Reference***

 -l

 -L

### Linking a shared library to another shared library

Just as you link modules into an application, you can create dependencies between shared libraries by linking them together. For example:

```
xlf -qmksprobj -o mylib.so myfile.o -Ldirectory -Rdirectory -lfoo
```

**Related information in the *XL Fortran Compiler Reference***

 -qmksprobj

 -L

## Compiling XL Fortran SMP programs

You can use the `xlf_r`, `xlf90_r`, `xlf95_r`, `xlf2003_r`, or `xlf2008_r` command to compile XL Fortran SMP programs. The `xlf_r` command is similar to the `xlf` command; the `xlf90_r` command is similar to the `xlf90` command; the `xlf95_r` command is similar to the `xlf95` command; the `xlf2003_r` command is similar to the `xlf2003` command; the `xlf2008_r` command is similar to the `xlf2008` command. The main difference is that the components for multi-threading are used to link and bind the object files if you specify the `xlf_r`, `xlf90_r`, `xlf95_r`, `xlf2003_r`, or `xlf2008_r` command.

Note that using any of these commands alone does not imply parallelization. For the compiler to recognize the SMP directives and activate parallelization, you must also specify `-qsmp`. In turn, you can only specify the `-qsmp` option in conjunction with one of these invocation commands. When you specify `-qsmp`, the driver links in the libraries specified on the `smplibraries` 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.

To compile and then link your program with the standard interface libraries, use the `xlf_r`, `xlf90_r`, `xlf95_r`, `xlf2003_r`, or `xlf2008_r` command. For example, you could specify:

```
xlf95_r test.f
```

## Compilation order for Fortran programs

If you have a program unit, subprogram, or interface body that uses a module, you must first compile the module. If the module and the code that uses the module are in separate files, you must first compile the file that contains the module. If they are in the same file, the module must come before the code that uses it in the file. If you change any entity in a module, you must recompile any files that use that module.

► **F2008** If only the implementation of a separate module procedure is changed, but the interface remains the same, you do not need to recompile the file that contains the module in which the corresponding module procedure interface body is declared. ◀ **F2008**

## Canceling a compilation

To stop the compiler before it finishes compiling, press **Ctrl+C** in interactive mode, or use the **kill** command.

## XL Fortran input files

The input files to the compiler are:

### Source Files (.f or .F suffix)

All .f, .f77, .f90, .f95, .f03, .f08, and .F, .F77, .F90, .F95, .F03, and .F08 files are 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 are passed through the C preprocessor (**cpp**) before being compiled.

Include files also contain source and often have different suffixes from .f.

**Related information:** See “Passing Fortran files through the C preprocessor” on page 29.

The **fsuffix** and **cppsuffix** attributes in “Editing the default configuration file” on page 13 and “-qsuffix” on page 244 let you select a different suffix.

### Object Files (.o suffix)

All .o files are object files. After the compiler compiles the source files, it uses the **ld** command to link-edit 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. It then produces a single executable output file.

**Related information:** See “Linking” on page 61 and “Linking XL Fortran programs” on page 31.

The **osuffix** attribute, which is described in “Editing the default configuration file” on page 13 and “-qsuffix” on page 244, lets you select a different suffix.

### Assembler Source Files (.s suffix)

The compiler sends any specified .s files to the assembler (**as**). The assembler output consists of object files that are sent to the linker at link time.

**Related information:** The **ssuffix** attribute, which is described in “Editing the default configuration file” on page 13 and “-qsuffix” on page 244, lets you select a different suffix.

### 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 /opt/ibm/xlf/15.1.1/etc/xlf.cfg.

**Related information:** See “Using custom compiler configuration files” on page 9 and “-F” on page 71 for information about selecting the configuration file.

**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.

**Related information:** See “-I” on page 76 and “-qmoddir” on page 187.

▶ **F2008** **Submodule Symbol Files:** *ancestormodulename\_submodulename.smod*

A submodule symbol file is an output file from compiling a submodule and is an input file for subsequent compilations of the descendant submodules. One **.smod** file is produced for each submodule, so compiling a single source file might produce multiple **.smod** files.

The submodule symbol file is not needed for compiling the ancestor module or any compilation units that access the ancestor module via use association.

**Related information:** See “-qmoddir” on page 187.

◀ **F2008**

**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\*” or “\*.pdf\_map\*”.

**Related information:** See “-qpdf1, -qpdf2” on page 195.

## XL Fortran output files

The output files that XL Fortran produces are:

**Executable Files:** *a.out*

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

**Related information:** See “-o” on page 85 for information on selecting a different name and “-c” on page 69 for information on generating only an object file.

**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 source file, and the assembler produces an object file for each specified assembler source file. By default, the object files have the same file name prefixes as the source files and appear in the current directory.

**Related information:** See “-c” on page 69 and “Linking XL Fortran programs” on page 31. For information on renaming the object file, see “-o” on page 85.

**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 source file. By default, the assembler source files have the same file name prefixes as the source files and appear in the current directory.

**Related information:** See “-S” on page 272 and “Linking XL Fortran programs” on page 31. For information on renaming the assembler source file, see “-o” on page 85.

**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 and a suffix of **.lst**.

**Related information:** See “Listings, messages, and compiler information” on page 56.

**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.

**Related information:** For information on putting **.mod** files in a different directory, see “-qmoddir” on page 187.

**F2008** **Submodule Symbol Files:** *ancestormodulename\_submodulename.smod*

Each submodule has an associated symbol file that holds information needed by the descendant submodules. By default, these symbol files must exist in the current directory.

The submodule symbol file is not needed for compiling the ancestor module or any compilation units that access the ancestor module via use association.

**Related information:** For information on putting **.smod** files in a different directory, see “-qmoddir” on page 187.

**F2008**

**cpp-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.

**Related information:** See “Passing Fortran files through the C preprocessor” on page 29 and “-d” on page 70.

**Profile Data Files (.pdf\*, .pdf\_map\*)**

These are the files that the **-qpdf1** option produces. They are used in subsequent compilations to tune optimizations that are based on actual execution results.

**Related information:** See “-qpdf1, -qpdf2” on page 195.

**Dependency Files:** *filename.d*

A dependency file contains source file dependency information. Dependency files are used by the **make** command to determine the order in which files should be compiled and the minimum set of files that must be recompiled when a file is changed. You can specify the “-MMD” on page 80 or “-qmakedep” on page 179 option to generate a dependency file.

**Related information:** For information on setting the name of the dependency file, see “-MF” on page 79. For information on specifying the target name of the object file in the dependency file, see “-MT” on page 81.

## Scope and precedence of option settings

You can specify compiler options in any of three locations. Their scope and precedence are defined by the location you use. (XL Fortran also has comment directives, such as **SOURCEFORM**, that can specify option settings. There is no general rule about the scope and precedence of such directives.)

Location	Scope	Precedence
In a stanza of the configuration file.	All compilation units in all files compiled with that stanza in effect.	Lowest
On the command line.	All compilation units in all files compiled with that command.	Medium
In an <b>@PROCESS</b> directive. (XL Fortran also has comment directives, such as <b>SOURCEFORM</b> , that can specify option settings. There is no general rule about the scope and precedence of such directives.)	The following compilation unit.	Highest

If you specify an option more than once with different settings, the last setting generally takes effect. Any exceptions are noted in the individual descriptions in the Chapter 6, “Detailed descriptions of the XL Fortran compiler options,” on page 65 and are indexed under “conflicting options”.

## Specifying options on the command line

XL Fortran supports the traditional UNIX method of specifying command-line options, with one or more letters (known as flags) following a minus sign:

```
xlf95 -c file.f
```

You can often concatenate multiple flags or specify them individually:

```
xlf95 -cv file.f      # These forms
xlf95 -c -v file.f   # are equivalent
```

(There are some exceptions, such as **-pg**, which is a single option and not the same as **-p -g**.)


Some of the flags require additional argument strings. Again, XL Fortran is flexible in interpreting them; you can concatenate multiple flags as long as the flag with an argument appears at the end. The following example shows how you can specify flags:

```
# All of these commands are equivalent.
xlf95 -g -v -o montecarlo -p montecarlo.f
xlf95 montecarlo.f -g -v -o montecarlo -p
xlf95 -g -v montecarlo.f -o montecarlo -p
xlf95 -g -v -omontecarlo -p montecarlo.f
# Because -o takes a blank-delimited argument,
```

```
# the -p cannot be concatenated.
xlf95 -gvomontecarlo -p montecarlo.f
# Unless we switch the order.
xlf95 -gvpomontecarlo montecarlo.f
```

If you are familiar with other compilers, particularly those in the XL family of compilers, you may already be familiar with many of these flags.

You can also specify many command-line options in a form that is intended to be easy to remember and make compilation scripts and makefiles relatively self-explanatory:




This format is more restrictive about the placement of blanks; you must separate individual **-q** options by blanks, and there must be no blank between a **-q** option and a following argument string. Unlike the names of flag options, **-q** option names are not case-sensitive except that the **q** must be lowercase. Use an equal sign to separate a **-q** option from any arguments it requires, and use colons to separate suboptions within the argument string.

For example:

```
xlf95 -qddim -qXREF=full -qfloat=nomaf:rsqrt -03 -qcache=type=c:level=1 file.f
```

## Specifying options in the source file

By putting the **@PROCESS** compiler directive in the source file, you can specify compiler options to affect an individual compilation unit. The **@PROCESS** compiler directive can override options specified in the configuration file, in the default settings, or on the command line.



*option* is the name of a compiler option without the **-q**.

*suboption*

is a suboption of a compiler option.

In fixed source form, **@PROCESS** can start in column 1 or after column 6. In free source form, the **@PROCESS** compiler directive can start in any column.

You cannot place a statement label or inline comment on the same line as an **@PROCESS** compiler directive.

By default, option settings you designate with the **@PROCESS** compiler directive are effective only for the compilation unit in which the statement appears. If the file has more than one compilation unit, the option setting is reset to its original

state before the next unit is compiled. Trigger constants specified by the **DIRECTIVE** option are in effect until the end of the file (or until **NODIRECTIVE** is processed).

The **@PROCESS** compiler directive must usually appear before the first statement of a compilation unit. The only exceptions are when specifying **SOURCE** and **NOSOURCE**; you can put them in **@PROCESS** directives anywhere in the compilation unit.

## Passing command-line options to the "ld" or "as" command

Because the compiler automatically executes other commands, such as **ld** and **as**, as needed during compilation, you usually do not need to concern yourself with the options of those commands. If you want to choose options for these individual commands, you can do one of the following:

- Include linker options on the compiler command line. When the compiler does not recognize a command-line option other than a **-q** option, it passes the option on to the linker:

```
xlf95 --print-map file.f # --print-map is passed to ld
```

- Use the **-W** compiler option to construct an argument list for the command:

```
xlf95 -Wl,--print-map file.f # --print-map is passed to ld
```

In this example, the **ld** option **--print-map** is passed to the linker (which is denoted by the **l** in the **-Wl** option) when the linker is executed.

This form is more general than the previous one because it works for the **as** command and any other commands called during compilation, by using different letters after the **-W** option.

- Edit the configuration file `/opt/ibm/xlf/15.1.1/etc/xlf.cfg`, or construct your own configuration file. You can customize particular stanzas to allow specific command-line options to be passed through to the assembler or linker.

For example, if you include these lines in the **xlf95** stanza of `/opt/ibm/xlf/15.1.1/etc/xlf.cfg`:

```
asopt = "W"  
ldopt = "M"
```

and issue this command:

```
xlf95 -Wa,-Z -Wl,-s -w produces_warnings.s uses_many_symbols.f
```

the file **produces\_warnings.s** is assembled with the options **-W** and **-Z** (issue warnings and produce an object file even if there are compilation errors), and the linker is invoked with the options **-s** and **-M** (strip final executable file and produce a load map).

**Related information:** See “-W” on page 276 and “Using custom compiler configuration files” on page 9.

## Displaying information inside binary files (strings)

The **strings** command reads information encoded into some binary files, as follows:

- Information about the compiler version is encoded in the compiler binary executables and libraries.
- Information about the parent module, bit mode, the compiler that created the **.mod** file, the date and time the **.mod** file was created, and the source file is encoded in each **.mod** file.

For example, to see the information embedded in `/opt/ibm/xlf/15.1.1/exe/xlfentry`, issue the following command:

```
strings /opt/ibm/xlf/15.1.1/exe/xlfentry | grep "@(#)"
```

## Compiling for specific architectures

You can use `-qarch` and `-qtune` to instruct the compiler to generate and tune code for a particular architecture. This allows the compiler to take advantage of machine-specific instructions that can improve performance. The `-qarch` option determines the architectures on which the resulting programs can run. The `-qtune` and `-qcache` options refine the degree of platform-specific optimization performed.

By default, the `-qarch` setting produces code using only instructions common to all supported architectures, with resultant settings of `-qtune` and `-qcache` that are relatively general. To tune performance for a particular processor set or architecture, you may need to specify different settings for one or more of these options. The natural progression to try is to use `-qarch`, and then add `-qtune`, and then add `-qcache`. Because the defaults for `-qarch` also affect the defaults for `-qtune` and `-qcache`, the `-qarch` option is often all that is needed.

If the compiling machine is also the target architecture, `-qarch=auto` will automatically detect the setting for the compiling machine. For more information on this compiler option setting, see “`-qarch`” on page 93 and also `-O4` and `-O5` under the `-O` option.

If the compiling machine is also the target architecture, `-qtune=auto` will automatically detect the setting for the compiling machine. For more information on this compiler option setting, see “`-qtune`” on page 250 and also `-O4` and `-O5` under the `-O` option.

If your programs are intended for execution mostly on particular architectures, you may want to add one or more of these options to the configuration file so that they become the default for all compilations.

## Passing Fortran files through the C preprocessor

A common programming practice is to pass files through the C preprocessor (`cpp`). `cpp` can include or omit lines from the output file based on user-specified conditions (“conditional compilation”). It can also perform string substitution (“macro expansion”).

XL Fortran can use `cpp` to preprocess a file before compiling it.

To call `cpp` for a particular file, use a file suffix of `.F`, `.F77`, `.F90`, `.F95`, `.F03`, or `.F08`. Each `.F*` file `filename.F*` is preprocessed into an intermediate file. You can save the intermediate file by specifying the `-d` compiler option; otherwise, the file is deleted. If you specify the `-d` option, the intermediate file name is `Ffilename.F*`. Otherwise, the intermediate file name is `/tmpdir/F8xxxxxx`, where `x` is an alphanumeric character and `tmpdir` is the contents of the `TMPDIR` environment variable or, if you have not specified a value for `TMPDIR`, `/tmp`. If you only want to preprocess and do not want to produce object or executable files, specify the `-qnoobject` option also.

When XL Fortran uses `cpp` for a file, the preprocessor will emit `#line` directives unless you also specify the `-d` option. The `#line` directive associates code that is created by `cpp` or any other Fortran source code generator with input code that you create. The preprocessor may cause lines of code to be inserted or deleted.

Therefore, the **#line** directives that it emits can be useful in error reporting and debugging, because they identify the source statements found in the preprocessed code by listing the line numbers that were used in the original source.

The **\_OPENMP** C preprocessor macro can be used to conditionally include code. This macro is defined when the C preprocessor is invoked and when you specify the **-qsmp=omp** compiler option. An example of using this macro follows:

```
program par_mat_mul
  implicit none
  integer(kind=8)                :: i,j,nthreads
  integer(kind=8),parameter      :: N=60
  integer(kind=8),dimension(N,N) :: Ai,Bi,Ci
  integer(kind=8)                :: Sumi
#ifdef _OPENMP
  integer omp_get_num_threads
#endif

  common/data/ Ai,Bi,Ci
!$OMP threadprivate (/data/)

!$omp parallel
  forall(i=1:N,j=1:N) Ai(i,j) = (i-N/2)**2+(j+N/2)
  forall(i=1:N,j=1:N) Bi(i,j) = 3-((i/2)+(j-N/2)**2)
!$omp master
#ifdef _OPENMP
  nthreads=omp_get_num_threads()
#else
  nthreads=8
#endif
!$omp end master
!$omp end parallel

!$OMP parallel default(private),copyin(Ai,Bi),shared(nthreads)
!$omp do
  do i=1,nthreads
    call imat_mul(Sumi)
  enddo
!$omp end do
!$omp end parallel

end
```

See *Conditional compilation* in the *Language Elements* section of the *XL Fortran Language Reference* for more information on conditional compilation.

To customize **cpp** preprocessing, the configuration file accepts the attributes **cpp**, **cppsuffix**, and **cppoptions**.

The letter **F** denotes the C preprocessor with the **-t** and **-W** options.

#### **Related information:**

- “-d” on page 70
- “-t” on page 272
- “-W” on page 276
- “-qfpp” on page 131
- “-qppsuborigarg” on page 209
- “Using custom compiler configuration files” on page 9

## **cpp directives for XL Fortran programs**

Macro expansion can have unexpected consequences that are difficult to debug, such as modifying a **FORMAT** statement or making a line longer than 72

characters in fixed source form. Therefore, we recommend using **cpp** primarily for conditional compilation of Fortran programs. The **cpp** directives that are most often used for conditional compilation are **#if**, **#ifdef**, **#ifndef**, **#elif**, **#else**, and **#endif**.

## Passing options to the C preprocessor

Because the compiler does not recognize **cpp** options other than **-I** directly on the command line, you must pass them through the **-W** option. For example, if a program contains **#ifdef** directives that test the existence of a symbol named **LNXXV1**, you can define that symbol to **cpp** by compiling with a command like:

```
xlf95 conditional.F -WF,-DLNXXV1
```

## Avoiding preprocessing problems

Because Fortran and C differ in their treatment of some sequences of characters, be careful when using **/\*** or **\*/**. These might be interpreted as C comment delimiters, possibly causing problems even if they occur inside Fortran comments. Also be careful when using three-character sequences that begin with **??** (which might be interpreted as C trigraphs).

Consider the following example:

```
program testcase
character a
character*4 word
a = '?'
word(1:2) = '??'
print *, word(1:2)
end program testcase
```

If the preprocessor matches your character combination with the corresponding trigraph sequence, your output may not be what you expected.

If your code does *not* require the use of the XL Fortran compiler option **-qnoescape**, a possible solution is to replace the character string with an escape sequence **word(1:2) = '\?\?'**. However, if you are using the **-qnoescape** compiler option, this solution will not work. In this case, you require a **cpp** that will ignore the trigraph sequence. XL Fortran uses the **cpp** that is shipped as part of the compiler. It is **ISO C** compliant and therefore recognizes trigraph sequences.

---

## 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 compiles and produces object files **file1.o** and **file3.o**.

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

Then all object files are submitted to the linker to produce one executable file.

After linking, follow the instructions in “Running XL Fortran programs” on page 33 to execute the program.

To link a library, follow the instructions in “Compiling and linking a library” on page 21.

**Note:** If you want to use a nondefault linker, you can use either of the following options:

- Use **-t** and **-B** to specify the nondefault linker, for example,  
`-t1 -Blinker_path`
- Customize the configuration file of the compiler to use the nondefault linker. For more information about how to customize the configuration file, see Using custom compiler configuration files and Creating custom configuration files.

## 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 (file1.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.

## Passing options to the ld command

For the detailed information about passing options to the **ld** command, see Passing command-line options to the "ld" or "as" command.

## 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, there are less chances for naming conflicts with library routines or external data objects because only exported symbols are visible outside a shared library. 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. There are more chances for naming conflicts with library routines or external data objects because all global symbols are visible outside a static library. 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.

To link statically, add the **-qstaticlink** option to the linker command. For example:

```
xlf95 -qstaticlink test.f
```

## Avoiding naming conflicts during linking

If you define an external subroutine, external function, or common block with the same name as a runtime or system library routine, your definition of that name may be used in its place, or it may cause a link-edit error.

Try the following general solution to help avoid these kinds of naming conflicts:

- Compile all files with the **-qextname** option. It adds an underscore to the end of the name of each global entity, making it distinct from any names in the system libraries.

**Note:** When you use this option, you do not need to use the final underscore in the names of Service and Utility Subprograms, such as **dtime\_** and **flush\_**.

- Link your programs dynamically, which is the default.

If you do not use the **-qextname** option, you must take the following extra precautions to avoid conflicts with the names of the external symbols in the XL Fortran and system libraries:

- Do not name a subroutine or function **main**, because XL Fortran defines an entry point **main** to start your program.
- Do not use *any* global names that begin with an underscore. In particular, the XL Fortran libraries reserve all names that begin with **\_xl**.
- Do not use names that are the same as names in the XL Fortran library or one of the system libraries. To determine which names are not safe to use in your program, run the **nm** command on any libraries that are linked into the program and search the output for names you suspect might also be in your program.

Be careful not to use the names of subroutines or functions without defining the actual routines in your program. If the name conflicts with a name from one of the libraries, the program could use the wrong version of the routine and not produce any compile-time or link-time errors.

---

## Running XL Fortran programs

The default file name for the executable program is **a.out**. You can select a different name with the **-o** compiler option. You should avoid giving your programs the same names as system or shell commands (such as **test** or **cp**), as you could accidentally execute the wrong command. If a name conflict does occur, you can execute the program by specifying a path name, such as **./test**.

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

### Canceling execution

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

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

### Compiling and executing on different systems

If you want to move an XL Fortran executable file to a different system (running the same or a compatible operating system) for execution, you can link statically and copy the program, and optionally the runtime message catalogs. Alternatively, you can link dynamically and copy the program as well as the XL Fortran libraries

if needed and optionally the runtime message catalogs. For non-SMP programs, **libxlf90.so**, **libxlfmath.so**, and **libxlfomp\_ser.so** are usually the only XL Fortran libraries needed. For SMP programs, you will usually need at least the **libxlf90.so**, **libxlfmath.so**, and **libxlfomp.so** libraries. **libxlfomp\*.so** is only needed if the program is compiled with the **-qautodbl** option.

For a dynamically linked program to work correctly, the XL Fortran libraries and the operating system on the execution system must be at either the same level or a more recent level than on the compilation system.

For a statically linked program to work properly, the operating system level may need to be the same on the execution system as it is on the compilation system.

**Related information:** See “Dynamic and static linking” on page 32.

## Runtime libraries for POSIX pthreads support

There are two runtime libraries that are connected with POSIX thread support. The **libxlf90\_r.so** library is a threadsafe version of the Fortran runtime library. The **libxlfomp.so** library is the SMP runtime library.

Depending on the invocation command, and in some cases, the compiler option, the appropriate set of libraries for thread support is bound in. For example:


Command	Libraries Used	Include Directory
xlf90_r xlf95_r	/opt/ibm/xlf/15.1.1/lib/libxlf90_r.so	/opt/ibm/xlf/15.1.1/include
xlf_r	/opt/ibm/xlf/15.1.1/lib/libxlfomp.so	

## Setting runtime options

Internal switches in an XL Fortran program control runtime behavior, similar to the way compiler options control compile-time behavior. You can set the runtime options through either environment variables or a procedure call within the program. You can specify XL Fortran runtime option settings by using the following environment variables: **XLFRTEOPTS** and **XLSMPOPTS**.

### The XLFRTEOPTS environment variable

The XLFRTEOPTS environment variable allows you to specify options that affect the runtime behavior of items such as I/O, EOF error-handling, the specification of random-number generators, and more. You can declare XLFRTEOPTS by using the following **bash** command format:



You can specify option names and settings in uppercase or lowercase. You can add blanks before and after the colons and equal signs to improve readability. However, if the XLFRTEOPTS option string contains imbedded blanks, you must enclose the entire option string in double quotation marks (").

The environment variable is checked when the program first encounters one of the following conditions:

- An I/O statement is executed.

- The **RANDOM\_SEED** procedure is executed.
- An **ALLOCATE** statement needs to issue a runtime error message.
- A **DEALLOCATE** statement needs to issue a runtime error message.
- The multithreaded implementation of the **MATMUL** procedure is executed.

Changing the **XLFRTEOPTS** environment variable during the execution of a program has no effect on the program.

The **SETRTEOPTS** procedure (which is defined in the *XL Fortran Language Reference*) accepts a single-string argument that contains the same name-value pairs as the **XLFRTEOPTS** environment variable. It overrides the environment variable and can be used to change settings during the execution of a program. The new settings remain in effect for the rest of the program unless changed by another call to **SETRTEOPTS**. Only the settings that you specified in the procedure call are changed.

You can specify the following runtime options with the **XLFRTEOPTS** environment variable or the **SETRTEOPTS** procedure:

**aggressive\_array\_io={yes | no}**

Controls whether or not the XL Fortran run time will take advantage of descriptor information when deciding to apply slower or faster algorithms to do array I/O operations. Descriptor information that specifies an array or array section as contiguous can be used to apply the faster algorithms which would otherwise be unsafe if the array or array section was not contiguous. The default is to perform aggressive array I/O operations.

Code executing under the current XL Fortran runtime but compiled with older XL Fortran compilers can cause the aggressive array I/O operations to be unsafe if the older compilers did not set the XL Fortran descriptor information correctly. This can be a problem with code built with old XL Fortran compilers no longer in service or built with XL Fortran compilers not at the latest service levels. Older code should be recompiled, if possible, with the current compiler instead of relying on the use of this option.

**buffering={enable | disable\_preconn | disable\_all}**

Determines whether the XL Fortran runtime library performs buffering for I/O operations.

The library reads data from, or writes data to the file system in chunks for **READ** or **WRITE** statements, instead of piece by piece. The major benefit of buffering is performance improvement.

If you have applications in which Fortran routines work with routines in other languages or in which a Fortran process works with other processes on the same data file, the data written by Fortran routines may not be seen immediately by other parties (and vice versa), because of the buffering. Also, a Fortran **READ** statement may read more data than it needs into the I/O buffer and cause the input operation performed by a routine in other languages or another process that is supposed to read the next data item to fail. In these cases, you can use the **buffering** runtime option to disable the buffering in the XL Fortran runtime library. As a result, a **READ** statement will read in exactly the data it needs from a file and the data written by a **WRITE** statement will be flushed out to the file system at the completion of the statement.

Note: I/O buffering is always enabled for files on sequential access devices (such as pipes, terminals, sockets). The setting of the **buffering** option has no effect on these types of files.

If you disable I/O buffering for a logical unit, you do not need to flush the contents of the I/O buffer for that logical unit with the **FLUSH** statement or the Fortran service routine **flush\_**.

The suboptions for **buffering** are as follows:

**enable**

The Fortran runtime library maintains an I/O buffer for each connected logical unit. The current read-write file pointers that the runtime library maintains might not be synchronized with the read-write pointers of the corresponding files in the file system.

**disable\_preconn**

The Fortran runtime library does not maintain an I/O buffer for each preconnected logical unit (0, 5, and 6). However, it does maintain I/O buffers for all other connected logical units. The current read-write file pointers that the runtime library maintains for the preconnected units are the same as the read-write pointers of the corresponding files in the file system.

**disable\_all**

The Fortran runtime library does not maintain I/O buffers for any logical units. You should not specify the **buffering=disable\_all** option with Fortran programs that perform asynchronous I/O.

In the following example, Fortran and C routines read a data file through redirected standard input. First, the main Fortran program reads one integer. Then, the C routine reads one integer. Finally, the main Fortran program reads another integer.

Fortran main program:

```
integer(4) p1,p2,p3
print *, 'Reading p1 in Fortran...'
read(5,*) p1
call c_func(p2)
print *, 'Reading p3 in Fortran...'
read(5,*) p3
print *, 'p1 p2 p3 Read: ', p1,p2,p3
end
```

C subroutine (c\_func.c):

```
#include <stdio.h>
void
c_func(int *p2)
{
    int n1 = -1;

    printf("Reading p2 in C...\n");
    setbuf(stdin, NULL); /* Specifies no buffering for stdin */
    fscanf(stdin,"%d", &n1);
    *p2=n1;
    fflush(stdout);
}
```

Input data file (infile):

```
11111
22222
33333
44444
```

The main program runs by using infile as redirected standard input, as follows:

```
$ main < infile
```

If you turn on **buffering=disable\_preconn**, the results are as follows:

```
Reading p1 in Fortran...
Reading p2 in C...
Reading p3 in Fortran...
p1 p2 p3 Read: 11111 22222 33333
```

If you turn on **buffering=enable**, the results are unpredictable.

#### **buffer\_size=size**

Specifies the size of I/O buffers in bytes instead of using the block size of devices. *size* must be either -1 or an integer value that is greater than or equal to 4096. The default, -1, uses the block size of the device where the file resides.

Using this option can reduce the amount of memory used for I/O buffers when an application runs out of memory because the block size of devices is very large and the application opens many files at the same time.

Note the following when using this runtime option:

- Preconnected units remain unaffected by this option. Their buffer size is the same as the block size of the device where they reside except when the block size is larger than 64KB, in which case the buffer size is set to 64KB.
- This runtime option does not apply to files on a tape device or logical volume.
- Specifying the buffer size with the **SETRTEOPTS** procedure overrides any value previously set by the **XLFRTEOPTS** environment variable or **SETRTEOPTS** procedure. The resetting of this option does not affect units that have already been opened.

#### **cnvrr={yes | no}**

If you set this runtime option to **no**, the program does not obey the **IOSTAT=** and **ERR=** specifiers for I/O statements that encounter conversion errors. Instead, it performs default recovery actions (regardless of the setting of **err\_recovery**) and may issue warning messages (depending on the setting of **xrf\_messages**).

**Related information:** For more information about conversion errors, see *Data transfer statements* in the *XL Fortran Language Reference*. For more information about **IOSTAT** values, see *Conditions and IOSTAT values* in the *XL Fortran Language Reference*.

#### **cpu\_time\_type={usertime | systime | alltime | total\_usertime | total\_systime | total\_alltime}**

Determines the measure of time returned by a call to **CPU\_TIME(TIME)**.

The suboptions for **cpu\_time\_type** are as follows:

##### **usertime**

Returns the user time of a process.

##### **systime**

Returns the system time of a process.

##### **alltime**

Returns the sum of the user and system time of a process.

##### **total\_usertime**

Returns the total user time of a process. The total user time is the sum of the user time of a process and the total user times of its child processes, if any.

**total\_systemtime**

Returns the total system time of a process. The total system time is the sum of the system time of the current process and the total system times of its child processes, if any.

**total\_alltime**

Returns the total user and system time of a process. The total user and system time is the sum of the user and system time of the current process and the total user and system times of their child processes, if any.

**default\_recl={64 | 32}**

Allows you to determine the default record size for sequential files opened without a **RECL=** specifier. The suboptions are as follows:

**64** Uses a 64-bit value as the default record size.

**32** Uses a 32-bit value as the default record size.

Use **default\_recl** when porting 32-bit programs to 64-bit mode where a 64-bit record length will not fit into the specified integer variable. Consider the following:

```
INTEGER(4) I
OPEN (11)
INQUIRE (11, RECL=i)
```

A runtime error occurs in the above code sample when **default\_recl=64**, since the default record length of  $2^{63}-1$  does not fit into the 4-byte integer **I**. Specifying **default\_recl=32** ensures a default record size of  $2^{31}-1$ , which fits into **I**.

For more information on the **RECL=** specifier, see the **OPEN** *statement* in the *XL Fortran Language Reference*.

**errloc={yes | no}**

Controls whether the file name and line number are displayed with an error message if a runtime error condition occurs during an I/O, **ALLOCATE** or **DEALLOCATE** statement. By default, the line number and file name appear prepended to the runtime error messages. If **errloc=no** is specified, runtime error messages are displayed without the source location information.

The **errloc** runtime option can be specified with the **SETRTEOPTS** procedure, as well.

**erroeof={yes | no}**

Determines whether the label specified by the **ERR=** specifier is to be branched to if no **END=** specifier is present when an end-of-file condition is encountered.

**err\_recovery={yes | no}**

If you set this runtime option to **no**, the program stops if there is a recoverable error while executing an I/O statement with no **IOSTAT=** or **ERR=** specifiers. By default, the program takes some recovery action and continues when one of these statements encounters a recoverable error. Setting **cnverr** to **yes** and **err\_recovery** to **no** can cause conversion errors to halt the program.

**errthrdnum={yes | no}**

When **errthrdnum=yes** is in effect, XL Fortran appends to all error messages the thread number of the running thread that is specified by the **omp\_get\_thread\_num** routine. For single-threaded programs, the thread number is 0.

If you specify **errloc=yes**, the thread number is displayed in front of the file name and line number. If the **IOMSG=** specifier is present in an I/O statement, the thread number is prefixed to the error message and the other part of the message uses the same format as displayed on Standard error.

**errtrace={yes | no}**

Controls whether a traceback is displayed with an error message if a runtime error condition occurs during an I/O, **ALLOCATE** or **DEALLOCATE** statement. Specifying **errtrace=no** means runtime error messages are displayed without tracebacks.

To show more detailed information in tracebacks, compile with the **-q1inedebug** or **-g** option.

Tracebacks are not displayed if either of the following conditions is true:

- You use the **IOSTAT=**, **ERR=**, **END=**, or **EOR=** specifier in an I/O statement.
- You use the **STAT=** specifier in an **ALLOCATE** or **DEALLOCATE** statement.

For example, in the following sample code, the **ERR=** specifier is used to branch to some error handling code. Because the I/O statement causing the error specifies **ERR=**, no traceback is generated. The output is `Open error.`, although **errtrace=yes** is specified.

```
program open_error
open(unit=11, file='doesnotexist', status='old', err=200) ! no traceback
close(11)
200 print *, 'Open error.'
end
```

**iostat\_end={extended | 2003std}**

Sets the **IOSTAT** values based on the XL Fortran definition or the Fortran 2003 Standard when end-of-file and end-of-record conditions occur. The suboptions are as follows:

**extended**

Sets the **IOSTAT** variables based on XL Fortran's definition of values and conditions.

**2003std**

Sets the **IOSTAT** variables based on Fortran 2003's definition of values and conditions.

For example, setting the **iostat\_end=2003std** runtime option results in a different **IOSTAT** value from extensions being returned for the end-of-file condition

```
export XLFRT_OPTS=iostat_end=2003std
character(10) ifl
integer(4) aa(3), ios
ifl = "12344321 "
read(ifl, '(3i4)', iostat=ios) aa ! end-of-file condition occurs and
! ios is set to -1 instead of -2.
```

For more information on setting and using **IOSTAT** values, see the **READ**, **WRITE**, and *Conditions and IOSTAT values* sections in the *XL Fortran Language Reference*.

**intrinths={num\_threads}**

Specifies the number of threads for parallel execution of the **MATMUL** and **RANDOM\_NUMBER** intrinsic procedures. The default value for **num\_threads** when using the **MATMUL** intrinsic equals the number of processors online. The default value for **num\_threads** when using the **RANDOM\_NUMBER** intrinsic is equal to the number of processors online\*2.

Changing the number of threads available to the **MATMUL** and **RANDOM\_NUMBER** intrinsic procedures can influence performance.

**langlvl={ | 90std | 95std | 2003std | 2008std | extended}**

Determines the level of support for Fortran standards and extensions to the standards. The values of the suboptions are as follows:

**90std** Instructs the compiler to flag any extensions to the Fortran 90 standard I/O statements and formats as errors.

**95std** Instructs the compiler to flag any extensions to the Fortran 95 standard I/O statements and formats as errors.

#### **2003std**

Instructs the compiler to flag any extensions to the Fortran 2003 standard I/O statements and formats as errors.

For example, setting the **langlvl=2003std** runtime option results in a runtime error message.

```
integer(4) aa(100)
call setrteopts("langlvl=2003std")
...           ! Write to a unit without explicitly
...           ! connecting the unit to a file.
write(10, *) aa ! The implicit connection to a file does not
...           ! conform with Fortran 2003 behavior.
```

#### **2008std**

Instructs the compiler to accept all standard I/O statements and formats that the Fortran 2003 standard specifies, as well as any Fortran 2008 formats that XL Fortran supports. Anything else is flagged as an error.

#### **extended**

Instructs that the compiler to accept the Fortran 95 language standard, Fortran 2003 features, the Fortran 2008 features supported by XL Fortran, and extensions, effectively turning off language-level checking.

To obtain support for items that are part of the Fortran 95 standard and are available in XL Fortran (such as namelist comments), you must specify one of the following suboptions:

- **95std**
- **2003std**
- **2008std**
- **extended**

The following example contains a Fortran 95 extension (the *file* specifier is missing from the **OPEN** statement):

```
program test1

call setrteopts("langlvl=95std")
open(unit=1,access="sequential",form="formatted")

10 format(I3)

write(1,fmt=10) 123

end
```

Specifying **langlvl=95std** results in a runtime error message.

The following example contains a Fortran 95 feature (namelist comments) that was not part of Fortran 90:

```

program test2

INTEGER I
LOGICAL G
NAMELIST /TODAY/G, I

call setrteopts("langlvl=95std:namelist=new")

open(unit=2,file="today.new",form="formatted", &
      & access="sequential", status="old")

read(2,nml=today)
close(2)

end

today.new:

&TODAY ! This is a comment
I = 123, G=.true. /

```

If you specify **langlvl=95std**, no runtime error message is issued. However, if you specify **langlvl=90std**, a runtime error message is issued.

The **err\_recovery** setting determines whether any resulting errors are treated as recoverable or severe.

#### **multconn={yes | no}**

Enables you to access the same file through more than one logical unit simultaneously. With this option, you can read more than one location within a file simultaneously without making a copy of the file.

You can only use multiple connections within the same program for files on random-access devices, such as disk drives. In particular, you cannot use multiple connections within the same program for:

- Files have been connected for write-only (**ACTION='WRITE'**)
- Asynchronous I/O
- Files on sequential-access devices (such as pipes, terminals, sockets)

To avoid the possibility of damaging the file, keep the following points in mind:

- The second and subsequent **OPEN** statements for the same file can only be for reading.
- If you initially opened the file for both input and output purposes (**ACTION='READWRITE'**), the unit connected to the file by the first **OPEN** becomes read-only (**ACCESS='READ'**) when the second unit is connected. You must close all of the units that are connected to the file and reopen the first unit to restore write access to it.
- Two files are considered to be the same file if they share the same device and i-node numbers. Thus, linked files are considered to be the same file.

#### **multconnio={tty | nulldev | combined | no }**

Enables you to connect a device to more than one logical unit. You can then write to, or read from, more than one logical unit that is attached to the same device. The suboptions are as follows:

##### **combined**

Enables you to connect a combination of null and TTY devices to more than one logical unit.

##### **nulldev**

Enables you to connect the null device to more than one logical unit.

**tty**

Enables you to connect a TTY device to more than one logical unit.

**Note:** Using this option can produce unpredictable results.

In your program, you can now specify multiple **OPEN** statements that contain different values for the **UNIT** parameters but the same value for the **FILE** parameters. For example, if you have a symbolic link called **mytty** that is linked to TTY device **/dev/tty**, you can run the following program when you specify the **multconnio=tty** option:

```
PROGRAM iotest
OPEN(UNIT=3, FILE='mytty', ACTION="WRITE")
OPEN(UNIT=7, FILE='mytty', ACTION="WRITE")
END PROGRAM iotest
```

Fortran preconnects units 0, 5, and 6 to the same TTY device. Normally, you cannot use the **OPEN** statement to explicitly connect additional units to the TTY device that is connected to units 0, 5, and 6. However, this is possible if you specify the **multconnio=tty** option. For example, if units 0, 5, and 6 are preconnected to TTY device **/dev/tty**, you can run the following program if you specify the **multconnio=tty** option:

```
PROGRAM iotest
! /dev/pts/2 is your current tty, as reported by the 'tty' command.
! (This changes every time you login.)
CALL SETRTEOPTS ('multconnio=tty')
OPEN (UNIT=3, FILE='/dev/pts/2')
WRITE (3, *) 'hello' ! Display 'hello' on your screen
END PROGRAM
```

**namelist={new | old}**

Determines whether the program uses the XL Fortran new or old **NAMELIST** format for input and output. The Fortran 90 and Fortran 95 standards require the new format.

**Note:** You may need the **old** setting to read existing data files that contain **NAMELIST** output. However, use the standard-compiler new format in writing any new data files.

With **namelist=old**, the nonstandard **NAMELIST** format is not considered an error by the **langlvl=90std**, **langlvl=95std**, or **langlvl=2003std** setting.

**Related information:** For more information about **NAMELIST** I/O, see *Namelist formatting* in the *XL Fortran Language Reference*.

**naninfoutput={2003std | old | default}**

Controls whether the display of IEEE exceptional values conform to the Fortran 2003 standard or revert to the old XL Fortran behavior. This runtime option allows object files created with different compilation commands to output all IEEE exceptional values based on the old behavior, or the Fortran 2003 standard. The suboptions are:

**default**

Exceptional values output depends on how the program is compiled.

**old**

Exceptional values output conforms to the old XL Fortran behavior.

**2003std**

Exceptional values output conforms to the Fortran 2003 standard.

**nlinewidth=record\_width**

By default, a **NAMELIST** write statement produces a single output record long

enough to contain all of the written **NAMELIST** items. To restrict **NAMELIST** output records to a given width, use the **nlwidth** runtime option.

**Note:** The **RECL=** specifier for sequential files has largely made this option obsolete, because programs attempt to fit **NAMELIST** output within the specified record length. You can still use **nlwidth** in conjunction with **RECL=** as long as the **nlwidth** width does not exceed the stated record length for the file.

**random={generator1 | generator2}**

Specifies the generator to be used by **RANDOM\_NUMBER** if **RANDOM\_SEED** has not yet been called with the **GENERATOR** argument. The value **generator1** (the default) corresponds to **GENERATOR=1**, and **generator2** corresponds to **GENERATOR=2**. If you call **RANDOM\_SEED** with the **GENERATOR** argument, it overrides the random option from that point onward in the program. Changing the random option by calling **SETRTEOPTS** after calling **RANDOM\_SEED** with the **GENERATOR** option has no effect.

**scratch\_vars={yes | no}**

To give a specific name to a scratch file, set the **scratch\_vars** runtime option to **yes**, and set the environment variable **XLFSCRATCH\_unit** to the name of the file you want to be associated with the specified unit number. See *Naming scratch files* in the *XL Fortran Optimization and Programming Guide* for examples.

**ufmt\_bigendian={units\_list}**

Specifies unit numbers of unformatted data files on which big-endian I/O is to be performed. The big-endian format data in the specified unformatted files is converted, on-the-fly, during the I/O operation to and from the little-endian format used on machines where XL Fortran applications are running.

This runtime option does not work with internal files; internal files are always **FORMATTED**. Units specified must be connected by an explicit or implicit **OPEN** for the **UNFORMATTED** form of I/O.

The syntax for this option is as follows:

```
ufmt_bigendian=units_list
```

where:

```
units_list = units | units_list, units
```

```
units = unit | unit- | -unit | unit1-unit2 | newunit | *
```

The unit number must be an integer, whose value is in the range 1 through 2 147 483 647.

**unit** Specifies the number of the logical unit.

**unit-** Specifies the range of units, starting from unit number *unit* to the highest possible unit number

**-unit** Specifies the range of units, starting from unit number 1 to unit number *unit*.

**unit1-unit2**

Specifies the range of units, starting from unit number *unit1* to unit number *unit2*.

**newunit**

Specifies the range of all **NEWUNIT** values. For details about **NEWUNIT** values, see **NEWUNIT** value.

\* Specifies all units. It is equivalent to 1-,newunit.

**Note:**

1. The byte order of data of type **CHARACTER** is not swapped.
2. The compiler assumes that the internal representation of values of type **REAL\*4** or **REAL\*8** is IEEE floating-point format compliant. I/O may not work properly with an internal representation that is different.
3. The internal representation of values of type **REAL\*16** is inconsistent among different vendors. The compiler treats the internal representation of values of type **REAL\*16** to be the same as XL Fortran's. I/O may not work properly with an internal representation that is different.
4. Conversion of derived type data is not supported. The alignment of derived types is inconsistent among different vendors.
5. Discrepancies in implementations from different vendors may cause problems in exchanging the big-endian unformatted data files between XL Fortran applications running on Linux and Fortran applications running on big-endian systems. XL Fortran provides a number of options that help users port their programs to XL Fortran. If there are problems exchanging big-endian data files, check these options to see if they can help with the problem.
6. XL Fortran provides several methods of specifying the byte order of unformatted data transfer. The **ufmt\_bigendian** option has the highest precedence, followed by the **CONVERT=** specifier and the **-qufmt** option. For more information, see **CONVERT=** specifier in the **OPEN** statement and “-qufmt” on page 252.

**unit\_vars={yes | no}**

To give a specific name to an implicitly connected file or to a file opened with no **FILE=** specifier, you can set the runtime option **unit\_vars=yes** and set one or more environment variables with names of the form **XLFUNIT\_unit** to file names. See *Naming files that are connected with no explicit name* in the *XL Fortran Optimization and Programming Guide* for examples.

**uwidth={32 | 64}**

To specify the width of record length fields in unformatted sequential files, specify the value in bits. When the record length of an unformatted sequential file is greater than  $(2^{*31} - 1)$  bytes minus 8 bytes (for the record terminators surrounding the data), you need to set the runtime option **uwidth=64** to extend the record length fields to 64 bits. This allows the record length to be up to  $(2^{*63} - 1)$  minus 16 bytes (for the record terminators surrounding the data).

**xrf\_messages={yes | no}**

To prevent programs from displaying runtime messages for error conditions during I/O operations, **RANDOM\_SEED** calls, and **ALLOCATE** or **DEALLOCATE** statements, set the **xrf\_messages** runtime option to **no**. Otherwise, runtime messages for conversion errors and other problems are sent to the standard error stream.

The following examples set the **cnvrr** runtime option to **yes** and the **xrf\_messages** option to **no**.

```
# Basic format
XLFRTEOPTS=cnvrr=yes:xrf_messages=no
export XLFRTEOPTS
```

```
# With imbedded blanks
XLFRTLOPTS="xrf_messages = NO : cnvrr = YES"
export XLFRTLOPTS
```

As a call to SETRTEOPTS, this example could be:

```
CALL setrteopts('xrf_messages=NO:cnvrr=yes')
! Name is in lowercase in case -U (mixed) option is used.
```

## Setting OMP and SMP run time options

The **XLSMPOPTS** environment variable allows you to specify options that affect SMP execution. The OpenMP environment variables, **OMP\_DYNAMIC**, **OMP\_NESTED**, **OMP\_NUM\_THREADS**, and **OMP\_SCHEDULE**, allow you to control the execution of parallel code. For details on using these, see *XLSMPOPTS* and *OpenMP environment variables* sections in the *XL Fortran Optimization and Programming Guide*.

## BLAS/ESSL environment variable

By default, the `libxlopt` library is linked with any application you compile with XL Fortran. However, if you are using a third-party Basic Linear Algebra Subprograms (BLAS) library or want to ship a binary file that includes ESSL routines, you must specify these using the `XL_BLAS_LIB` environment variable. For example, if your own BLAS library is called `libblas.a`, set the environment variable as follows:

```
export XL_BLAS_LIB=/usr/lib/libblas.a
```

When the compiler generates calls to BLAS routines, the ones defined in the `libblas` library will be used at run time instead of those defined in `libxlopt`.

## XLF\_USR\_CONFIG

Use the **XLF\_USR\_CONFIG** environment variable to specify the location of a custom configuration file to be used by the compiler. The file name must be given with its absolute path. The compiler will first process the definitions in this file before processing those in the default system configuration file, or those in a customized file specified by the `-F` option; for more information, see “Using custom compiler configuration files” on page 9.

---

## Other environment variables that affect runtime behavior

The `LD_LIBRARY_PATH`, `LD_RUN_PATH`, and `TMPDIR` environment variables have an effect at run time, as explained in “Correct settings for environment variables” on page 7. They are not XL Fortran runtime options and cannot be set in either `XLFRTLOPTS` or `XLSMPOPTS`.

---

## XL Fortran runtime exceptions

The following operations cause runtime exceptions in the form of **SIGTRAP** signals, which typically result in a “Trace/breakpoint trap” message:

- Character substring expression or array subscript out of bounds after you specified the `-C` option at compile time.
- Lengths of character pointer and target do not match after you specified the `-C` option at compile time.
- The flow of control in the program reaches a location for which a semantic error with severity of **S** was issued when the program was compiled.
- Floating-point operations that generate NaN values and loads of the NaN values after you specify the `-qfloat=nanq` option at compile time.

- Fixed-point division by zero.
- Calls to the TRAP hardware-specific intrinsic procedure.

The following operations cause runtime exceptions in the form of **SIGFPE** signals:

- Floating-point exceptions provided you specify the appropriate **-qflttrap** suboptions at compile time.

If you install one of the predefined XL Fortran exception handlers before the exception occurs, a diagnostic message and a traceback showing the offset within each routine called that led to the exception are written to standard error after the exception occurs. The file buffers are also flushed before the program ends. If you compile the program with the **-g** option, the traceback shows source line numbers in addition to the address offsets.

You can use a symbolic debugger to determine the error. **gdb** provides a specific error message that describes the cause of the exception.

**Related information:**

- “-C” on page 68
- “-qflttrap” on page 132
- “-qsigtrap” on page 224

Also see the following topics in the *XL Fortran Optimization and Programming Guide*:

- *Detecting and trapping floating-point exceptions* for more details about these exceptions
- *Controlling the floating-point status and control register* for a list of exception handlers.

---

## Chapter 5. Summary of compiler options by functional category

The XL Fortran options available on the Linux platform are grouped into the following categories:

- “Output control”
- “Input control” on page 49
- “Language element control” on page 50
- “Floating-point and integer control” on page 52
- “Object code control” on page 53
- “Error checking and debugging” on page 54
- “Listings, messages, and compiler information” on page 56
- “Optimization and tuning” on page 57
- “Linking” on page 61
- “Portability and migration” on page 61
- “Compiler customization” on page 62

If the option supports an equivalent @PROCESS directive, this is indicated. To get detailed information on any option listed, see the full description page for that option.

You can enter compiler options that start with **-q**, suboptions, and @PROCESS directives in either uppercase or lowercase. However, note that if you specify the **-qmixed** option, procedure names that you specify for the **-qextern** option are case-sensitive.

In general, this document uses the convention of lowercase for **-q** compiler options and suboptions and uppercase for @PROCESS directives.

Understanding the significance of the options you use and knowing the alternatives available can save you considerable time and effort in making your programs work correctly and efficiently.

For detailed information about each compiler option, see Chapter 6, “Detailed descriptions of the XL Fortran compiler options,” on page 65.

---

### Output control

The options in this category control the type of file output the compiler produces, as well as the locations of the output. These are the basic options that determine the compiler components that will be invoked, the preprocessing, compilation, and linking steps that will (or will not) be taken, and the kind of output to be generated.

Table 6. Compiler output options

Option name	@PROCESS directive	Description
“-c” on page 69	None.	Instructs the compiler to compile or assemble the source files only but do not link. With this option, the output is a .o file for each source file.

Table 6. Compiler output options (continued)

Option name	@PROCESS directive	Description
"-d" on page 70	None.	Causes preprocessed source files that are produced by <b>cpp</b> to be kept rather than deleted.
"-MMD" on page 80	None.	Produces a dependency output file containing targets suitable for inclusion in a description file for the <b>make</b> command. <b>-MMD</b> is the short form of <b>-qmakedep</b> on page 179.
"-MF" on page 79	None.	Specifies the name or location for the dependency output files that are generated by the <b>-qmakedep</b> or <b>-MMD</b> option.
"-MT" on page 81	None.	Specifies the target name of the object file in the <b>make</b> rule in the dependency output file that is generated by the <b>-qmakedep</b> or <b>-MMD</b> option.
"-o" on page 85	None.	Specifies a name for the output object, assembler, or executable file.
"-qmakedep" on page 179	None.	Produces a dependency output file containing targets suitable for inclusion in a description file for the <b>make</b> command. <b>-qmakedep</b> is the long form of <b>-MMD</b> on page 80.
"-qmkshrobj" on page 186	None.	Creates a shared object from generated object files.
"-qmoddir" on page 187	None.	Specifies the location for any module (.mod) or <b>F2008</b> submodule <b>F2008</b> (.smod) files that the compiler writes.
"-qtimestamps" on page 249	None.	Controls whether or not implicit time stamps are inserted into an object file.
"-S" on page 272	None.	Generates an assembler language file for each source file.

## Input control

The options in this category specify the type and location of your source files.

Table 7. Compiler input options

Option name	@PROCESS directive	Description
"-D" on page 69, "-qdlines" on page 116	DLINES	Specifies whether the compiler compiles fixed source form lines with a D in column 1 or treats them as comments.
"-I" on page 76	None.	Adds a directory to the search path for include files, .mod files, and .smod files.
"-qcclines" on page 103	CCLINES	Determines whether the compiler recognizes conditional compilation lines in fixed source form and F90 free source form. This option is not supported with IBM free source form.
"-qci" on page 106	CI	Specifies the identification numbers (from 1 to 255) of the INCLUDE lines to process.
"-qcr" on page 108	None.	Controls how the compiler interprets the CR (carriage return) character.
"-qdirective" on page 114	DIRECTIVE	Specifies sequences of characters, known as trigger constants, that identify comment lines as compiler comment directives.
"-qfixed" on page 125	FIXED	Indicates that the input source program is in fixed source form and optionally specifies the maximum line length.
"-qfpp" on page 131	None.	Controls Fortran-specific preprocessing in the C preprocessor.  This is a C preprocessor option, and must therefore be specified using the <b>-WF</b> option.
"-qfree" on page 134	FREE	Indicates that the source code is in free source form.
"-qmixed" on page 186, "-U" on page 274	MIXED	Makes the compiler sensitive to the case of letters in names.

Table 7. Compiler input options (continued)

Option name	@PROCESS directive	Description
"-qppsuborigarg" on page 209	None.	Instructs the C preprocessor to substitute original macro arguments before further macro expansion.  This is a C preprocessor option, and must therefore be specified using the <b>-WF</b> option.
"-qsuffix" on page 244	None.	Specifies the source-file suffix on the command line.
"-qxlines" on page 268	XLINES	Specifies whether fixed source form lines with an X in column 1 are compiled or treated as comments.

## Language element control

The options in this category allow you to specify the characteristics of the source code. You can also use these options to enforce or relax language restrictions and enable or disable language extensions.

Table 8. Language element control options

Option name	@PROCESS directive	Description
"-qaltivec" on page 92	ALTIVEC	Specifies the order of vector elements in vector registers.
"-qinit" on page 152	INIT(F90PTR)	Makes the initial association status of pointers disassociated.
"-qlanglvl" on page 169	LANGLVL	Determines which language standard (or superset, or subset of a standard) to consult for nonconformance. It identifies nonconforming source code and also options that allow such nonconformances.
"-qmbs" on page 185	MBCS	Indicates to the compiler whether character literal constants, Hollerith constants, H edit descriptors, and character string edit descriptors can contain Multibyte Character Set (MBCS) or Unicode characters.
"-qnullterm" on page 189	NULLTERM	Appends a null character to each character constant expression that is passed as a dummy argument, making it more convenient to pass strings to C functions.

Table 8. Language element control options (continued)

Option name	@PROCESS directive	Description
"-1" on page 67, "-qonetrip" on page 191	ONETRIP	Executes each <b>DO</b> loop in the compiled program at least once if its <b>DO</b> statement is executed, even if the iteration count is 0. This option provides compatibility with FORTRAN 66.
"-qposition" on page 208	POSITION	Positions the file pointer at the end of the file when data is written after an <b>OPEN</b> statement with no <b>POSITION=</b> specifier and the corresponding <b>STATUS=</b> value ( <b>OLD</b> or <b>UNKNOWN</b> ) is specified.
"-qqcount" on page 212	QCOUNT	Accepts the <b>Q</b> character-count edit descriptor ( <b>Q</b> ) as well as the extended-precision <b>Q</b> edit descriptor ( <b>Qw.d</b> ).
"-qsaa" on page 218	SAA	Checks for conformance to the SAA FORTRAN language definition. It identifies nonconforming source code and also options that allow such nonconformances.
"-qsave" on page 218	SAVE	Specifies the default storage class for local variables.
"-qsclk" on page 222	None.	Specifies the resolution that the <b>SYSTEM_CLOCK</b> intrinsic procedure uses in a program.
"-u" on page 274, "-qundef" on page 253	UNDEF	Specifies that no implicit typing of variable names is permitted.  <b>-qundef</b> is the long form of the "-u" on page 274 option.
"-qxf77" on page 259	XLF77	Provides compatibility with FORTRAN 77 aspects of language semantics and I/O data format that have changed.
"-qxf90" on page 261	XLF90	Provides compatibility with the Fortran 90 standard for certain aspects of the Fortran language.

Table 8. Language element control options (continued)

Option name	@PROCESS directive	Description
"-qxf2003" on page 263	XLF2003	Provides the ability to use language features specific to the Fortran 2003 standard when compiling with compiler invocations that conform to earlier Fortran standards, as well as the ability to disable these features when compiling with compiler invocations that conform to the Fortran 2003 standard.
"-qxf2008" on page 267	XLF2008	Provides the ability to use language features specific to the Fortran 2008 standard when compiling with compiler invocations that conform to earlier Fortran standards, as well as the ability to disable these features when compiling with compiler invocations that conform to the Fortran 2008 standard.
"-qzerosize" on page 270	ZEROSIZE	Determines whether checking for zero-sized character strings and arrays takes place in programs that might process such objects.

## Floating-point and integer control

Specifying the details of how your applications perform calculations can allow you to take better advantage of your system's floating-point performance and precision, including how to direct rounding. However, keep in mind that strictly adhering to IEEE floating-point specifications can impact the performance of your application. Using the options in the following table, you can control trade-offs between floating-point performance and adherence to IEEE standards. Some of these options also allow you to control certain aspects of integer calculations.

Table 9. Floating-point and integer control options

Option name	@PROCESS directive	Description
"-qautodbl" on page 97	AUTODBL	Provides an automatic means of converting single-precision floating-point calculations to double-precision and of converting double-precision calculations to extended-precision.
"-qdpc" on page 117	DPC	Increases the precision of real constants for maximum accuracy, when assigning real constants to <b>DOUBLE PRECISION</b> variables.
"-qenum" on page 118	None.	Specifies the range of the enumerator constant and enables storage size to be determined.

Table 9. Floating-point and integer control options (continued)

Option name	@PROCESS directive	Description
"-qfloat" on page 127	FLOAT	Selects different strategies for speeding up or improving the accuracy of floating-point calculations.
"-qieee" on page 146, "-y" on page 279	IEEE	Specifies the rounding mode that the compiler will use when it evaluates constant floating-point expressions at compile time.
"-qintlog" on page 160	INTLOG	Specifies that you can mix integer and logical data entities in expressions and statements.
"-qintsize" on page 161	INTSIZE	Sets the size of default <b>INTEGER</b> and <b>LOGICAL</b> data entities that have no length or kind specified.
"-qrealsize" on page 213	REALSIZE	Sets the default size of <b>REAL</b> , <b>DOUBLE PRECISION</b> , <b>COMPLEX</b> , and <b>DOUBLE COMPLEX</b> values.
"-qstrictieemod" on page 242	STRICTIEEMOD	Specifies whether the compiler will adhere to the Fortran 2003 IEEE arithmetic rules for the <b>ieee_arithmetic</b> and <b>ieee_exceptions</b> intrinsic modules.

## Object code control

These options affect the characteristics of the object code, preprocessed code, or other output generated by the compiler.

Table 10. Object code control options

Option name	@PROCESS directive	Description
"-qinlglue" on page 157	INLGLUE	When used with <b>-O2</b> or higher optimization, inlines glue code that optimizes external function calls in your application.
"-qplic" on page 205	None.	Generates position-independent code required for use in shared libraries.

Table 10. Object code control options (continued)

Option name	@PROCESS directive	Description
"-qsaveopt" on page 220	None.	Saves the command-line options used for compiling a source file, the user's configuration file name and the options specified in the configuration files, the version and level of each compiler component invoked during compilation, and other information to the corresponding object file.
"-qstackprotect" on page 233	None.	Provides protection against malicious input data or programming errors that overwrite or corrupt the stack.
"-qtbtable" on page 248	None.	Controls the amount of debugging traceback information that is included in the object files.
"-qthreaded" on page 249	None.	Indicates to the compiler whether it must generate threadsafe code.

## Error checking and debugging

The options in the following table allow you to detect and correct problems in your source code. In some cases, these options can alter your object code, increase your compile time, or introduce runtime checking that can slow down the execution of your application. The option descriptions indicate how extra checking can impact performance.

To control the amount and type of information you receive regarding the behavior and performance of your application, consult "Listings, messages, and compiler information" on page 56.

For information on debugging optimized code, see the *XL Fortran Optimization and Programming Guide*.

Table 11. Error checking and debugging options

Option name	@PROCESS directive	Description
"-#" on page 66	None.	Previews the compilation steps specified on the command line, without actually invoking any compiler components.
"-C" on page 68, "-qcheck" on page 103	CHECK	Generates code that performs certain types of runtime checking.
"-g" on page 73, "-qdbg" on page 110	DBG	Generates debugging information for use by a symbolic debugger, and makes the program state available to the debugging session at selected source locations.

Table 11. Error checking and debugging options (continued)

Option name	@PROCESS directive	Description
"-qflttrap" on page 132	FLTTRAP	Determines what types of floating-point exception conditions to detect at run time.
"-qfullpath" on page 135	None.	When used with the <b>-g</b> or <b>-qlinedebug</b> option, this option records the full, or absolute, path names of source and include files in object files compiled with debugging information, so that debugging tools can correctly locate the source files.
"-qfunctrace" on page 136	None.	Traces entry and exit points of procedures in your program. If your program contains C++ compilation units, this option also traces C++ catch blocks.
"-qfunctrace_xlf_catch" on page 138	None.	Specifies the name of the catch tracing subroutine.
"-qfunctrace_xlf_enter" on page 139	None.	Specifies the name of the entry tracing subroutine.
"-qfunctrace_xlf_exit" on page 140	None.	Specifies the name of the exit tracing subroutine.
"-qhalt" on page 140	HALT	Stops compilation before producing any object, executable, or assembler source files if the maximum severity of compile-time messages equals or exceeds the severity you specify.
"-qhaltonmsg" on page 142	HALTONMSG	Stops compilation before producing any object files, executable files, or assembler source files if a specified error message is generated.
"-qinfo" on page 147	None.	Produces or suppresses groups of informational messages.
"-qinitialloc" on page 152	INITALLOC	Initializes allocatable and pointer variables that are allocated but not initialized to a specific value, for debugging purposes.
"-qinitauto" on page 154	None.	Initializes uninitialized automatic variables to a specific value, for debugging purposes.
"-qkeepparm" on page 168	None.	When used with <b>-O2</b> or higher optimization, specifies whether procedure parameters are stored on the stack.
"-qlinedebug" on page 173	None.	Generates only line number and source file name information for a debugger.
"-qobject" on page 190	OBJECT	Specifies whether to produce an object file or to stop immediately after checking the syntax of the source files.
"-qsigtrap" on page 224	None.	Sets up the specified trap handler to catch <b>SIGTRAP</b> and <b>SIGFPE</b> exceptions when compiling a file that contains a main program.

Table 11. Error checking and debugging options (continued)

Option name	@PROCESS directive	Description
"-qwarn64" on page 257	None.	Displays informational messages identifying statements that may cause problems with 32-bit to 64-bit migration.
"-qxflag=dvz" on page 258	None.	Causes the compiler to generate code to detect floating-point divide-by-zero operations.

## Listings, messages, and compiler information

The options in the following table allow you control over the listing file, as well as how and when to display compiler messages. You can use these options in conjunction with those described in "Error checking and debugging" on page 54 to provide a more robust overview of your application when checking for errors and unexpected behavior.

Table 12. Listings and messages options

Option name	@PROCESS directive	Description
"-qattr" on page 96	ATTR	Produces a compiler listing that includes the attribute component of the attribute and cross-reference section of the listing.
"-qflag" on page 126	FLAG	Limits the diagnostic messages to those of a specified severity level or higher.
"-qhelp" on page 143	None.	Displays the man page of the compiler.
"-qlist" on page 174	LIST	Produces a compiler listing file that includes object and constant area sections.
"-qlistfmt" on page 175	None.	Creates an XML or HTML report to assist with finding optimization opportunities.
"-qlistopt" on page 178	None.	Produces a compiler listing file that includes all options in effect at the time of compiler invocation.
"-qmaxerr" on page 182	MAXERR	Stops compilation when the number of error messages of a specified severity level or higher reaches a specified number.
"-qphsinfo" on page 203	PHSINFO	Reports the time taken in each compilation phase to standard output.

Table 12. Listings and messages options (continued)

Option name	@PROCESS directive	Description
"-qnoprint" on page 188	None.	Prevents the compiler from creating the listing file, regardless of the settings of other listing options.
"-qreport" on page 216	None.	Produces listing files that show how sections of code have been optimized.
"-qsource" on page 232	SOURCE	Produces a compiler listing file that includes the source section of the listing and provides additional source information when printing error messages.
"-qsuppress" on page 245	None.	Prevents specific informational or warning messages from being displayed or added to the listing file, if one is generated.
"-qversion" on page 255	None.	Displays the version and release of the compiler being invoked.
"-V" on page 276	None.	The same as <code>-v</code> except that you can cut and paste directly from the display to create a command.
"-v" on page 275	None.	Reports the progress of compilation, by naming the programs being invoked and the options being specified to each program.
"-w" on page 278	None.	Suppresses informational, language-level and warning messages (equivalent to <code>-qflag=e:e</code> ).
"-qxref" on page 269	XREF	Produces a compiler listing that includes the cross-reference component of the attribute and cross-reference section of the listing.

## Optimization and tuning

You can control the optimization and tuning process, which can improve the performance of your application at run time, using the options in the following table. Remember that not all options benefit all applications. Trade-offs sometimes occur between an increase in compile time, a reduction in debugging capability, and the improvements that optimization can provide. In addition to the option

descriptions in this section, consult the *XL Fortran Optimization and Programming Guide* for details on the optimization and tuning process as well as writing optimization friendly source code.

Some of the options in “Floating-point and integer control” on page 52 can also improve performance, but you must use them with care to ensure your application retains the floating-point semantics it requires.

*Table 13. Optimization and tuning options*

Option name	@PROCESS directive	Description
“-qalias” on page 87	ALIAS( <i>argument_list</i> )	Indicates whether a program contains certain categories of aliasing or does not conform to Fortran standard aliasing rules. The compiler limits the scope of some optimizations when there is a possibility that different names are aliases for the same storage location.
“-qarch” on page 93	None.	Specifies the processor architecture, or family of architectures, where the code may run. This allows the compiler to take maximum advantage of the machine instructions specific to an architecture, or common to a family of architectures.
“-qassert” on page 95	ASSERT	Provides information about the characteristics of your code that can help the compiler fine-tune optimizations.
“-qcache” on page 100	None.	Specifies the cache configuration for a specific execution machine.
“-qcompact” on page 108	COMPACT	Avoids optimizations that increase code size.
“-qdirectstorage” on page 116	None.	Informs the compiler that a given compilation unit may reference write-through-enabled or cache-inhibited storage.
“-qessl” on page 120	None.	Allows the compiler to substitute the Engineering and Scientific Subroutine Library (ESSL) routines in place of Fortran 90 intrinsic procedures.
“-qfdpr” on page 124	None.	Provides object files with information that the IBM Feedback Directed Program Restructuring (FDPR®) performance-tuning utility needs to optimize the resulting executable file.

Table 13. Optimization and tuning options (continued)

Option name	@PROCESS directive	Description
“-qhot” on page 143	HOT( <i>suboptions</i> )	Performs high-order loop analysis and transformations (HOT) during optimization.
-qinline	None.	Attempts to inline procedures instead of generating calls to those procedures, for improved performance.
“-qipa” on page 163	None.	Enables or customizes a class of optimizations known as interprocedural analysis (IPA).
“-qlibansi” on page 171	None.	Assumes that all functions with the name of an ANSI C library function are, in fact, the library functions and not a user function with different semantics.
-qlibmpi	None.	Asserts that all functions with Message Passing Interface (MPI) names are in fact MPI functions and not a user function with different semantics.
“-qmaxmem” on page 183	MAXMEM	Limits the amount of memory that the compiler allocates while performing specific, memory-intensive optimizations to the specified number of kilobytes.
“-O” on page 83	OPTIMIZE	Specifies whether to optimize code during compilation and, if so, at which level.
“-p” on page 86	None.	Prepares the object files produced by the compiler for profiling.
“-qpdf1, -qpdf2” on page 195	None.	Tunes optimizations through <i>profile-directed feedback</i> (PDF), where results from sample program execution are used to improve optimization near conditional branches and in frequently executed code sections.
“-qprefetch” on page 210	None.	Inserts prefetch instructions automatically where there are opportunities to improve code performance.

Table 13. Optimization and tuning options (continued)

Option name	@PROCESS directive	Description
"-qshowpdf " on page 223	None.	When used with <b>-qpdf1</b> and a minimum optimization level of <b>-02</b> at compile and link steps, creates a PDF map file that contains additional profiling information for all procedures in your application.
"-qsimd" on page 225	None.	Controls whether the compiler can automatically take advantage of vector instructions for processors that support them.
"-qsmallstack" on page 226	None.	Minimizes stack usage where possible.
"-qsmp" on page 227	None.	Enables parallelization of program code.
"-qstacktemp" on page 234	None.	Determines where to allocate certain XL Fortran compiler temporaries at run time.
"-qstrict" on page 238	STRICT	Ensures that optimizations done by default at the <b>-03</b> and higher optimization levels, and, optionally at <b>-02</b> , do not alter certain program semantics mostly related to strict IEEE floating-point conformance.
"-qstrict_induction" on page 242	None.	Prevents the compiler from performing induction (loop counter) variable optimizations. These optimizations may be unsafe (may alter the semantics of your program) when there are integer overflow operations involving the induction variables.
"-qtune" on page 250	None.	Tunes instruction selection, scheduling, and other architecture-dependent performance enhancements to run best on a specific hardware architecture. Allows specification of a target SMT mode to direct optimizations for best performance in that mode.
"-qunroll" on page 253	None.	Specifies whether unrolling <b>DO</b> loops is allowed in a program. Unrolling is allowed on outer and inner <b>DO</b> loops.

Table 13. Optimization and tuning options (continued)

Option name	@PROCESS directive	Description
"-qunwind" on page 255	None.	Specifies that the compiler will preserve the default behavior for saves and restores to volatile registers during a procedure call.

## Linking

Though linking occurs automatically, the options in the following table allow you to direct input and output to the linker, controlling how the linker processes your object files.

You can actually include **ld** options on the compiler command line, because the compiler passes unrecognized options on to the linker.

Table 14. Linking options

Option name	@PROCESS directive	Description
"-e" on page 70	None.	When used together with the <b>-qmksbobj</b> , specifies an entry point for a shared object.
"-L" on page 77	None.	At link time, searches the directory path for library files specified by the <b>-l</b> option.
"-l" on page 78	None.	Searches for the specified library file. The linker searches for <i>libkey.so</i> , and then <i>libkey.a</i> if <i>libkey.so</i> is not found.
"-qstaticlink" on page 235	None.	Controls whether static or shared runtime libraries are linked into an application.

## Portability and migration

The options in this category can help you maintain application behavior compatibility on past, current, and future hardware, operating systems and compilers, or help move your applications to an XL compiler with minimal change.

Table 15. Portability and migration options

Option name	@PROCESS directive	Description
"-qalign" on page 90	ALIGN	Specifies the alignment of data objects in storage, which avoids performance problems with misaligned data.
"-qbindcextname" on page 99	BINDCEXTNAME	Controls whether the <b>-qextname</b> option affects <b>BIND(C)</b> entities.
"-qctyplss" on page 109	CTYPLSS	Specifies whether character constant expressions are allowed wherever typeless constants may be used.

Table 15. Portability and migration options (continued)

Option name	@PROCESS directive	Description
"-qddim" on page 112	DDIM	Specifies that the bounds of pointee arrays are re-evaluated each time the arrays are referenced and removes some restrictions on the bounds expressions for pointee arrays.
"-qdescriptor" on page 113	None.	Specifies the XL Fortran internal descriptor data structure format to use for non object-oriented entities in your compiled applications.
"-qescape" on page 119	ESCAPE	Specifies how the backslash is treated in character strings, Hollerith constants, H edit descriptors, and character string edit descriptors.
"-qextern" on page 121	None.	Allows user-written procedures to be called instead of XL Fortran intrinsics.
"-qextname" on page 122	EXTNAME	Adds an underscore to the names of all global entities.
"-qlog4" on page 179	LOG4	Specifies whether the result of a logical operation with logical operands is a <b>LOGICAL(4)</b> or is a <b>LOGICAL</b> with the maximum length of the operands.
"-qport" on page 205	PORT	Provides options to accommodate other Fortran language extensions when porting programs to XL Fortran.
"-qswapomp" on page 246	SWAPOMP	Specifies that the compiler should recognize and substitute OpenMP routines in XL Fortran programs.
"-qufmt" on page 252	UFMT	Sets the byte order for I/O operations on unformatted data files.
"-qxflag=oldtab" on page 259	XFLAG(OLDTAB)	Interprets a tab in columns 1 to 5 as a single character (for fixed source form programs).

## Compiler customization

The options in the following table allow you to specify alternative locations for compiler components, configuration files, standard include directories, and internal compiler operation. These options are useful for specialized installations, testing scenarios, and the specification of additional command-line options.

Table 16. Compiler customization options

Option name	@PROCESS directive	Description
"-B" on page 67	None.	Specifies substitute path names for XL Fortran components such as the assembler, C preprocessor, and linker.
"-F" on page 71	None.	Specifies an alternative configuration file, which stanza to use within the configuration file, or both.
"-NS" on page 82, "-qspillsize" on page 233	SPILLSIZE	Specifies the size (in bytes) of the register spill space; the internal program storage areas used by the optimizer for register spills to storage.
"-qoptfile" on page 191	None.	Specifies a response file that contains a list of additional command line options to be used for the compilation. Response files typically have the .rsp suffix.
"-qpath" on page 194	None.	Specifies substitute path names for XL Fortran components such as the assembler, C preprocessor, and linker.
"-t" on page 272	None.	Applies the prefix specified by the <b>-B</b> option to the designated components.
"-W" on page 276	None.	Passes the listed options to a component that is executed during compilation.



---

## Chapter 6. Detailed descriptions of the XL Fortran compiler options

This section contains descriptions of the individual options available in XL Fortran.

For each option, the following information is provided:

### Category

The functional category to which the option belongs is listed here.

### @PROCESS

For many compiler options, you can use an equivalent @PROCESS directive to apply the option's functionality within the source code, limiting the scope of the option's application to a single source file or compilation unit, or even selected sections of code.

### Purpose

This section provides a brief description of the effect of the option (and equivalent directives), and why you might want to use it.

### Syntax

This section provides the syntax for the command-line option and for the equivalent @PROCESS directive, if applicable. Syntax is shown first in command-line form, and then in @PROCESS form. For an explanation of the notations used to represent command-line syntax, see "Conventions" on page viii.

Uppercase letters are sometimes used to indicate the minimum number of characters for an option. For example, in **-qassert=CONTIGuous**, the uppercase letters **CONTIG** indicate the minimum number of characters you must use for this option. Therefore if you use **-qassert=contig** or **-qassert=contigu**, the compiler recognizes both as valid.

For @PROCESS syntax, the following notations are used:

- Defaults for each option are underlined and in boldface type.
- Individual required arguments are shown with no special notation.
- When you must make a choice between a set of alternatives, they are enclosed by { and } symbols.
- Optional arguments are enclosed by [ and ] symbols.
- When you can select from a group of choices, they are separated by | characters.
- Arguments that you can repeat are followed by ellipses (...).

### Defaults

In most cases, the default option setting is clearly indicated in the syntax diagram. However, for many options, there are multiple default settings, depending on other compiler options in effect. This section indicates the different defaults that may apply.

### Parameters

This section describes the suboptions that are available for the option.

**Usage** This section describes any rules or usage considerations you should be aware of. These can include restrictions on the option's applicability, precedence rules for multiple option specifications, and so on.

## Examples

Where appropriate, examples of the command-line syntax and use are provided in this section.

---

**-#**

## Category

Error checking and debugging

## @PROCESS

None.

## Purpose

Previews the compilation steps specified on the command line, without actually invoking any compiler components.

At the points where the compiler executes commands to perform different compilation steps, this option displays a simulation of the commands it would run and the system argument lists it would pass, but it does not actually perform these actions.

## Syntax

### Option:

►► -# ◄◄

## Defaults

Not applicable.

## Usage

Examining the output of this option can help you quickly and safely determine the following information for a particular compilation:

- What files are involved
- What options are in effect for each step

It avoids the overhead of compiling the source code and avoids overwriting any existing files, such as `.lst` files. (If you are familiar with the `make` command, it is similar to `make -n`.)

This option produces the same output as `-v` and `-V`, but does not perform the compilation.

Note that if you specify this option with `-qipa`, the compiler does not display linker information subsequent to the IPA link step. This is because the compiler does not actually call IPA.

## Related information

- “-v” on page 275
- “-V” on page 276

---

## -1

### Category

Language element control

### Purpose

Executes each **DO** loop in the compiled program at least once if its **DO** statement is executed, even if the iteration count is 0. This option provides compatibility with FORTRAN 66.

**-qonetrip** is the long form of **-1**.

### Syntax

#### Option:

►► -1 ◀◀

#### @PROCESS:

@PROCESS ONETRIP | NOONETRIP

### Defaults

The default is to follow the behavior of later Fortran standards, where **DO** loops are not performed if the iteration count is 0.

### Restrictions

It has no effect on **FORALL** statements, **FORALL** constructs, or array constructor implied-**DO** loops.

---

## -B

### Category

Compiler customization

### @PROCESS

None.

### Purpose

Specifies substitute path names for XL Fortran components such as the assembler, C preprocessor, and linker.

It can be used in combination with the **-t** option, which determines which of these components are affected by **-B**. However, it is preferred that you use the **-qpath** option to accomplish this instead.

## Syntax

► — -B —————►  
          └── prefix ─┘

## Defaults

The default paths for the compiler components are defined in the compiler configuration file.

## Parameters

*prefix*

The name of a directory where the components reside. It must end in a / (slash).

## Usage

To form the complete path name for each component, the driver program adds *prefix* to the standard program names. You can restrict the components that are affected by this option by also including one or more **-t***mnemonic* options.

You can also specify default path names for these commands in the configuration file.

This option allows you to keep multiple levels of some or all of the XL Fortran components or to try out an upgraded component before installing it permanently. When keeping multiple levels of XL Fortran available, you might want to put the appropriate **-B** and **-t** options into a configuration-file stanza and to use the **-F** option to select the stanza to use.

## Related information

- “-qpath” on page 194
- “-t” on page 272
- “-F” on page 71
- “Using custom compiler configuration files” on page 9
- “Running two levels of XL Fortran” on page 16

---

## -C

### Category

Error checking and debugging

### Purpose

Generates code that performs certain types of runtime checking.

**-qcheck** is the long form of **-C**.

### Syntax

**Option:**

▶▶ -C ◀◀

@PROCESS:

@PROCESS CHECK | NOCHECK

## Defaults

-qnocheck

---

## -C

### Category

Object code control

### @PROCESS

None.

### Purpose

Instructs the compiler to compile or assemble the source files only but do not link. With this option, the output is a .o file for each source file.

### Syntax

▶▶ -c ◀◀

### Defaults

Not applicable.

### Usage

Using the **-o** option in combination with **-c** selects a different name for the .o file. In this case, you can only compile one source file at a time.

### Related information

- “-o” on page 85.
- 

## -D

### Category

Input control

### Purpose

Specifies whether the compiler compiles fixed source form lines with a D in column 1 or treats them as comments.

**-qdlines** is the long form of **-D**.



## @PROCESS

None.

### Purpose

When used together with the **-qmkshrobj**, specifies an entry point for a shared object.

### Syntax

► — **-e** *entry\_name* —►

### Defaults

None.

### Parameters

*name*

The name of the entry point for the shared executable.

### Usage

Specify the **-e** option only with the **-qmkshrobj** option.

**Note:** When you link object files, do not use the **-e** option. The default entry point of the executable output is `__start`. Changing this label with the **-e** flag can produce errors.

### Related information

- “-qmkshrobj” on page 186

---

## -F

### Category

Compiler customization

## @PROCESS


None.

### Purpose

Specifies an alternative configuration file, which stanza to use within the configuration file, or both.

The configuration file specifies different kinds of defaults, such as options for particular compilation steps and the locations of various files that the compiler requires.

## Syntax



## Defaults

By default, the compiler uses the configuration file that is configured at installation time, and the stanza defined in that file for the invocation command currently being used (for example, `xlf2003`, `xlf90_r`, `xlf90`, and so on.).

## Parameters

### *config\_file*

The full path name of the alternate compiler configuration file to use.

### *stanza*

The name of the configuration file stanza to use for compilation. This directs the compiler to use the entries under that *stanza* regardless of the invocation command being used. For example, if you are compiling with `xlf2003`, but you specify the `xlf95` stanza, the compiler will use all the settings specified in the `xlf95` stanza.

## Usage

A simple way to customize the way the compiler works, as an alternative to writing complicated compilation scripts, is to add new stanzas to `/opt/ibm/xlf/15.1.1/etc/xlf.cfg`, giving each stanza a different name and a different set of default compiler options. Or, you can specify a user-defined configuration file with the `XL_FUSR_CONFIG` environment variable rather than editing the default configuration file. You may find a single, centralized file easier to maintain than many scattered compilation scripts and makefiles.

By running the compiler with an appropriate `-F` option, you can select the set of options that you want. You might have one set for full optimization, another set for full error checking, and so on. Note that the settings in any user-defined configuration file are processed before the ones specified by the `-F` option.

## Restrictions

Because the default configuration file is replaced each time a new compiler release is installed, make sure to save a copy of any new stanzas or compiler options that you add.

Alternatively, you can store customized settings in the user-defined configuration file specified by the `XL_FUSR_CONFIG` environment variable. This file will not be replaced during reinstallation.

## Examples

```
# Use stanza debug in default xlf.cfg.  
xlf95 -F:debug t.f  
  
# Use stanza xlf95 in /home/fred/xlf.cfg.
```

```
xlf95 -F/home/fred/xlf.cfg t.f
# Use stanza myxlf in /home/fred/xlf.cfg.
xlf95 -F/home/fred/xlf.cfg:myxlf t.f
```

## Related information

- “Creating custom configuration files” on page 10 explains the contents of a custom, user-defined configuration file and shows how to select different stanzas in the file without using the **-F** option.
- “Editing the default configuration file” on page 13 explains how to edit the contents of a configuration file for use with the **-F** option.
- “-B” on page 67
- “-t” on page 272
- “-W” on page 276

---

## -g

### Category

Error checking and debugging

### @PROCESS

DBG

### Purpose

Generates debugging information for use by a symbolic debugger, and makes the program state available to the debugging session at selected source locations.

You can use different **-g** levels to balance between debug capability and compiler optimization. Higher **-g** levels provide a more complete debug support, at the cost of runtime or possible compile-time performance, while lower **-g** levels provide higher runtime performance, at the cost of some capability in the debugging session.

When the **-O2** optimization level is in effect, the debug capability is completely supported.

**Note:** When an optimization level higher than **-O2** is in effect, the debug capability is limited.

**-g** is the short form of **-qdbg**.

## Syntax



### @PROCESS:

@PROCESS DBG | NODBG

## Defaults

`-g0`, equivalent to `-qnodbg`

## Parameters

`-g`

- When no optimization is enabled (`-qnoopt`), `-g` is equivalent to `-g9` or `-qdbg=level=9`.
- When the `-O2` optimization level is in effect, `-g` is equivalent to `-g2` or `-qdbg=level=2`.

`-g0` Generates no debugging information. No program state is preserved.

`-g1` Generates minimal read-only debugging information about line numbers and source file names. No program state is preserved. This option is equivalent to `-q1 inedebug`.

`-g2` Generates read-only debugging information about line numbers, source file names, and variables.

When the `-O2` optimization level is in effect, no program state is preserved.

`-g3, -g4`

Generates read-only debugging information about line numbers, source file names, and variables.

When the `-O2` optimization level is in effect:

- No program state is preserved.
- Procedure parameter values are available to the debugger at the beginning of each procedure.

`-g5, -g6, -g7`

Generates read-only debugging information about line numbers, source file names, and variables.

When the `-O2` optimization level is in effect:

- Program state is available to the debugger at **IF** constructs, loop constructs, procedure definitions, and procedure calls. For details, see "Usage" on page 75.
- Procedure parameter values are available to the debugger at the beginning of each procedure.

- g8** Generates read-only debugging information about line numbers, source file names, and variables.

When the **-O2** optimization level is in effect:

- Program state is available to the debugger at the beginning of every executable statement.
- Procedure parameter values are available to the debugger at the beginning of each procedure.

- g9** Generates debugging information about line numbers, source file names, and variables. You can modify the value of the variables in the debugger.

When the **-O2** optimization level is in effect:

- Program state is available to the debugger at the beginning of every executable statement.
- Procedure parameter values are available to the debugger at the beginning of each procedure.

## Usage

When no optimization is enabled, the debugging information is always available if you specify **-g2** or a higher level. When the **-O2** optimization level is in effect, the debugging information is available at selected source locations if you specify **-g5** or a higher level.

When you specify **-g8** or **-g9** with **-O2**, the debugging information is available at every source line with an executable statement.

When you specify **-g5**, **-g6**, or **-g7** with **-O2**, the debugging information is available for the following language constructs:

- **IF** constructs

The debugging information is available at the beginning of every **IF** statement, namely at the line where the **IF** keyword is specified. It is also available at the beginning of the next executable statement right after the **IF** construct.

- Loop constructs

The debugging information is available at the beginning of every **DO** statement, namely at the line where the **DO** keyword is specified. It is also available at the beginning of the next executable statement right after the **DO** construct.

- Procedure definitions

The debugging information is available at the first executable statement in the body of the procedure.

- Procedure calls

The debugging information is available at the beginning of every statement where a user-defined procedure is called. It is also available at the beginning of the next executable statement right after the statement that contains the procedure call.

## Examples

Use the following command to compile `myprogram.f` and generate an executable program called `testing` for debugging:

```
xlf myprogram.f -o testing -g
```

The following command uses a specific **-g** level with **-O2** to compile `myprogram.f` and generate debugging information:

```
xlf myprogram.f -O2 -g8
```

### Related information

- SNAPSHOT
- “Debugging a Fortran program” on page 290
- “Symbolic debugger support” on page 6
- “-qdbg” on page 110
- “-qlinedebug” on page 173
- “-qfullpath” on page 135
- “-O” on page 83
- “-qkeeparm” on page 168

---

**-I**

### Category

Input control

### @PROCESS

None.

### Purpose

Adds a directory to the search path for include files, `.mod` files, and `.smod` files.

### Syntax

►► `-I` *path\_name* ◀◀

### Defaults

Not applicable.

### Parameters

*path\_name*

A valid path name (for example, `/home/dir`, `/tmp`, or `./subdir`).

### Usage

If XL Fortran calls `cpp`, this option adds a directory to the search path for `#include` files. Before checking the default directories for include, `.mod`, and `.smod` files, the compiler checks each directory in the search path. For include files, this path is only used if the file name in an `INCLUDE` line is not provided with an absolute path. For `#include` files, refer to the `cpp` documentation for the details of the `-I` option.

## Rules

The compiler appends a / to *path\_name* and then concatenates that with the file name before making a search. If you specify more than one **-I** option on the command line, files are searched in the order of the *path\_name* names as they appear on the command line.

The following directories are searched, in this order, after any paths that are specified by **-I** options:

1. The current directory (from which the compiler is executed)
2. The directory where the source file is (if different from 1.)
3. /usr/include

Also, the compiler will search /opt/ibm/xlf/15.1.1/include where include and .mod files shipped with the compiler are located.

## Related information

- “-qmoddir” on page 187
- “-qfullpath” on page 135

---

## -k

### Category

Input control

### Purpose

Indicates that the source code is in free source form.

This option is the short form of **-qfree=f90**.

### Syntax

**Option:**

►► -k ◄◄

**@PROCESS:**

@PROCESS FREE(F90)

### Related information

- “-qfree” on page 134
- *Free source form* in the *XL Fortran Language Reference*.

---

## -L

### Category

Linking

**@PROCESS**

None.

## Purpose

At link time, searches the directory path for library files specified by the `-l` option.

## Syntax

### Option:

►► `-L—Directory` ◀◀

## Defaults

Not applicable.

## Usage

Adds *Directory* to the list of search directories that are used for finding libraries designated by the `-l` flag (lowercase letter **L**). If you use libraries other than the default ones specified in `/opt/ibm/xlf/15.1.1/lib`, you can specify one or more `-L` options that point to the locations of the other libraries.

## Rules

This option is passed directly to the `ld` command and is not processed by XL Fortran at all.

## Related information

- “Linking” on page 61
- “Linking XL Fortran programs” on page 31

---

## -l

### Category

Linking

### @PROCESS

None.

## Purpose

Searches for the specified library file. The linker searches for *libkey.so*, and then *libkey.a* if *libkey.so* is not found.

## Syntax

►► `-l—key` ◀◀

## Defaults

The compiler default is to search only for some of the compiler runtime libraries. The default configuration file specifies the default library names to search for with

the `-l` compiler option, and the default search path for libraries with the `-L` compiler option.

## Parameters

*key*

The name of the library minus the `lib` and `.a` or `.so` characters.

## Rules

This option is passed directly to the `ld` command and is not processed by XL Fortran at all.

## Examples

To compile `myprogram.f` and link it with library `libmylibrary.so` or `libmylibrary.a` that is found in the `/usr/mylibdir` directory, enter the following command. Preference is given to `libmylibrary.so` over `libmylibrary.a`.

```
xlc myprogram.f -lmylibrary -L/usr/mylibdir
```

## Related information

- “Linking” on page 61
- “Linking XL Fortran programs” on page 31

---

## -MF

### Category

Output control

### @PROCESS

None.

### Purpose

Specifies the name or location for the dependency output files that are generated by the `-qmakedep` or `-MMD` option.

For more information about the `-qmakedep` and `-MMD` options, see “`-qmakedep`” on page 179 and “`-MMD`” on page 80.

### Syntax

▶▶ `-MF` *file\_path* ▶▶

### Defaults

If `-MF` is not specified, the dependency output file is generated in the current working directory. The file has the same base name as the object file, but with a `.d` suffix.

## Parameters

### *file\_path*

The name or location of the generated dependency output file. *file\_path* can be the path name of a file or a directory. It can be relative or absolute.

If *file\_path* is the name of an existing directory, the generated dependency output file is placed in the specified directory. It will have the same base name as the generated object file, but with a `.d` suffix. Otherwise, *file\_path* is the name of the generated dependency output file. For more information, see the Examples section.

## Usage

If the file specified by `-MF` option already exists, it will be overwritten.

If you specify a single file name for the `-MF` option when you compile multiple source files, only a single dependency file will be generated. The dependency file contains the **make** rule for the last file specified on the command line.

## Examples

To compile `source.f` and create a dependency output file named `mysource.d`, enter:

```
xlf -c -qmakedep source.f -MF mysource.d
```

To compile `source1.f` and `source2.f` and create two dependency output files, named `source1.d` and `source2.d`, respectively, in the `/project/deps/` directory, enter:

```
xlf -c -qmakedep source1.f source2.f -MF /project/deps/
```

To compile `source.f` and create a dependency output file named `mysource.d` in the `deps/` directory, enter:

```
xlf -c -qmakedep source.f -MF deps/mysource.d
```

To compile `source.f` and create an object file named `object.o` and a dependency output file named `mysource.d`, enter:

```
xlf -c -qmakedep source.f -o object.o -MF mysource.d
```

## Related information

- “-qmakedep” on page 179
- “-MMD”
- “-MT” on page 81

---

## -MMD

### Category

Output control

### @PROCESS

None.

## Purpose

Produces a dependency output file containing targets suitable for inclusion in a description file for the **make** command.

The dependency output file is named with a `.d` suffix and specifies a separate rule for each of the main source file's dependencies.

**-MMD** is the short form of “`-qmakedep`” on page 179.

## Syntax

▶▶ `-MMD` ◀◀

## Defaults

Not applicable.

## Related information

- “`-qmakedep`” on page 179

---

## -MT

### Category

Output control

### @PROCESS

None.

## Purpose

Specifies the target name of the object file in the **make** rule in the dependency output file that is generated by the **-qmakedep** or **-MMD** option.

For more information about the **-qmakedep** and **-MMD** options, see “`-qmakedep`” on page 179 and “`-MMD`” on page 80.

## Syntax

▶▶ `-MT` *target* ◀◀

## Defaults

If **-MT** is not specified, the target name is the base name of the object file.

## Parameters

*target* The name that you specify for the object file in the generated dependency file.

## Usage

Only one target is accepted. If you specify multiple **-MT** options, the target from the last **-MT** option is used.

## Examples

A source file named `mysource.f` contains the following code:

```
#include "options.h"
MODULE m
  USE n

CONTAINS
  SUBROUTINE sub
    IMPLICIT NONE
    INCLUDE 'constants.h'
    CALL my_print(pi)
  END SUBROUTINE
END MODULE
```

To compile `mysource.f` and create a dependency output file named `mysource.d`, and also to include path information `/home/user/sample/` as part of the target name of the object file in the `mysource.d` file, enter:

```
xlf -c -qmakedep mysource.f -MT '/home/user/sample/mysource.o'
```

The generated `mysource.d` file is as follows:

```
/home/user/sample/mysource.o m.mod: n.mod
/home/user/sample/mysource.o m.mod: option.h
/home/user/sample/mysource.o m.mod: mysource.f
/home/user/sample/mysource.o m.mod: constants.h
```

## Related information

- “-qmakedep” on page 179
- “-MMD” on page 80
- “-MF” on page 79

---

## -NS

### Category

Compiler customization

### Purpose

Specifies the size (in bytes) of the register spill space; the internal program storage areas used by the optimizer for register spills to storage.

**-qspillsize** is the long form of **-NS**.

### Syntax

#### Option:

▶▶ **-NS** *bytes* ◀◀

#### @PROCESS:

@PROCESS SPILLSIZE(*bytes*)

## Defaults

By default, each subprogram stack has 512 bytes of spill space reserved.

If you need this option, a compile-time message informs you of the fact.

## Parameters

*bytes*

The number of bytes of stack space to reserve in each subprogram, in case there are too many variables to hold in registers and the program needs temporary storage for register contents.

## Related information

- “-qspillsize” on page 233

---

## -O

### Category

Optimization and tuning


### Purpose

Specifies whether to optimize code during compilation and, if so, at which level.

**-qOPTimize** is the long form of **-O**.

### Syntax

**Option:**



**@PROCESS:**

@PROCESS OPTimize[(*level*)] | **NOOPTimize**

### Defaults

**nooptimize** or **-O0** or **optimize=0**

### Parameters

**not specified**

Almost all optimizations are disabled. This is equivalent to specifying **-O0** or **-qnoopt**.

**-O** For each release of XL Fortran, **-O** enables the level of optimization that represents the best tradeoff between compilation speed and runtime performance. If you need a specific level of optimization, specify the appropriate numeric value. **-O** is equivalent to **-O2**.

- O0** Almost all optimizations are disabled. This option is equivalent to **-qnoopt**.
- O1** Reserved for future use. This form is ignored and has no effect on the outcome of the compilation.
- O2** Performs a set of optimizations that are intended to offer improved performance without an unreasonable increase in time or storage that is required for compilation.
- O3** Performs additional optimizations that are memory intensive, compile-time intensive, and may change the semantics of the program slightly, unless **-qstrict** is specified. We recommend these optimizations when the desire for runtime speed improvements outweighs the concern for limiting compile-time resources.

This level of optimization also affects the setting of the **-qfloat** option, turning on the **rsqrt** suboption by default, and sets **-qmaxmem=-1**.

Specifying **-O3** implies **-qhot=level=0**, unless you explicitly specify **-qhot**, **-qhot=level=1**, or **-qnohot**.

- O4** Aggressively optimizes the source program, trading off additional compile time for potential improvements in the generated code. You can specify the option at compile time or at link time. If you specify it at link time, it will have no effect unless you also specify it at compile time for at least the file that contains the main program.

**-O4** implies the following other options:

- **-qhot**
- **-qipa**
- **-O3** (and all the options and settings that it implies)
- **-qarch=auto**
- **-qtune=auto**
- **-qcache=auto**

Note that the **auto** setting of **-qarch**, **-qtune**, and **-qcache** implies that the execution environment will be the same as the compilation environment.

This option follows the "last option wins" conflict resolution rule, so any of the options that are modified by **-O4** can be subsequently changed.

- O5** Provides all of the functionality of the **-O4** option, but also provides the functionality of the **-qipa=level=2** option.

#### Note:

To obtain the same floating-point accuracy for optimized and non-optimized applications, you must specify the **-qfloat=nomaf** compiler option. In cases where differences in floating-point accuracy still occur after specifying **-qfloat=nomaf**, the **-qstrict** compiler option allows you to exert greater control over changes that optimization can cause in floating-point semantics.

## Usage

Generally, use the same optimization level for both the compile and link steps. This is important when using either the **-O4** or **-O5** optimization level to get the best runtime performance. For the **-O5** level, all loop transformations (as specified via the **-qhot** option) are done at the link step.

Increasing the level of optimization may or may not result in additional performance improvements, depending on whether the additional analysis detects any further optimization opportunities.

An optimization level of **-03** or higher can change the behavior of the program and potentially cause exceptions that would not otherwise occur. Use of the **-qstrict** option maintains the same program behavior as with **-02**, at the cost of optimization opportunity. Refer to the **-qstrict** option for the list of optimizations it disables.

When you use **-0** or higher optimization, **-qtbtable=small** is in effect. The traceback table generated has no function name or parameter information.

If the **-0** option is used in an **@PROCESS** statement, only an optimization level of 0, 2, or 3 is allowed. Note that unlike using **-03** in command line invocation, specifying **@PROCESS OPT(3)** does not imply **-qhot=level=0**. If optimization level **-03** or higher is specified on the command line, the **-qhot** and **-qipa** options that are set by the optimization level cannot be overridden by **@PROCESS OPT(0)** or **@PROCESS OPT(2)**.

Compilations with optimization may require more time and machine resources than other compilations.

The more the compiler optimizes a program, the more difficult it is to debug the program with a symbolic debugger.

### Related information

- “-qtbtable” on page 248 controls the amount of debugging traceback information that is included in the object files.
- “-qstrict” on page 238 shows how to turn off the effects of **-03** that might change the semantics of a program.
- “-qipa” on page 163, “-qpdf1, -qpdf2” on page 195, and “-qhot” on page 143 turn on additional optimizations that may improve performance for some programs.
- "Optimizing your applications" in the *XL Fortran Optimization and Programming Guide* discusses technical details of the optimization techniques the compiler uses and some strategies you can use to get maximum performance from your code.

---

## **-O**

### Category

Output control

### **@PROCESS**

None.

### Purpose

Specifies a name for the output object, assembler, or executable file.

## Syntax

►► `-o name` ◀◀

## Defaults

The default name for an executable file is **a.out**. The default name for an object or assembler source file is the same as the source file except that it has a **.o** or **.s** extension.

## Usage

To choose the name for an object file, use this option in combination with the **-c** option. For an assembler source file, use it in combination with the **-S** option.

## Rules

Except when you specify the **-c** or **-S** option, the **-o** option is passed directly to the **ld** command, instead of being processed by XL Fortran.

## Examples

```
xlf95 t.f                # Produces "a.out"
xlf95 -c t.f            # Produces "t.o"
xlf95 -o test_program t.f # Produces "test_program"
xlf95 -S -o t2.s t.f    # Produces "t2.s"
```

---

## -p

### Category

Optimization and tuning

### @PROCESS

None.

### Purpose

Prepares the object files produced by the compiler for profiling.

The compiler produces monitoring code that counts the number of times each routine is called. The compiler inserts a call to the monitor subroutine at the start of each subprogram.

### Syntax

►► `-p` ◀◀

### Defaults

Not applicable.



## Parameters

### aryovrlp | **noaryovrlp**

Indicates whether the compilation units contain any array assignments between storage-associated arrays. If not, specify **noaryovrlp** to improve performance.

### intptr | **nointptr**

Indicates whether the compilation units contain any integer **POINTER** statements. If so, specify **intptr**.

### pteovrlp | **nopteovrlp**

Indicates whether any pointee variables may be used to refer to any data objects that are not pointee variables, or whether two pointee variables may be used to refer to the same storage location. If not, specify **nopteovrlp**.

### std | **nostd**

Indicates whether the compilation units contain any nonstandard aliasing (which is explained below). If so, specify **nostd**.

## Usage

An alias exists when an item in storage can be referred to by more than one name. The Fortran 90, Fortran 95, Fortran 2003, and Fortran 2008 standards allow some types of aliasing and disallow some others. The sophisticated optimizations that the XL Fortran compiler performs increase the likelihood of undesirable results when nonstandard aliasing is present, as in the following situations:

- The same data object is passed as an actual argument two or more times in the same subprogram reference. The aliasing is not valid if either of the actual arguments becomes defined, undefined, or redefined.
- A subprogram reference associates a dummy argument with an object that is accessible inside the referenced subprogram. The aliasing is not valid if any part of the object associated with the dummy argument becomes defined, undefined, or redefined other than through a reference to the dummy argument.
- A dummy argument becomes defined, undefined, or redefined inside a called subprogram in some other way than through the dummy argument.
- A subscript to an array within a common block exceeds that array's bounds.

## Restrictions

Because this option inhibits some optimizations of some variables, using it can lower performance.

Programs that contain nonstandard or integer **POINTER** aliasing may produce incorrect results if you do not compile them with the correct **-qalias** settings. When you use the **xl** or **xl\_r** invocation command to compile the **.f**, **.F**, **.f77**, or **.F77** files, or when you use the **f77** or **fort77** invocation command, the compiler assumes that integer **POINTERS** might be present (**-qalias=aryovrlp:pteovrlp:std:intptr**). All the other invocation commands assume that a program contains only standard aliasing (**-qalias=aryovrlp:pteovrlp:std:ointptr**).

## Examples

If the following subroutine is compiled with `-qalias=nopteovr1p`, the compiler may be able to generate more efficient code. You can compile this subroutine with `-qalias=nopteovr1p`, because the integer pointers, `ptr1` and `ptr2`, point at dynamically allocated memory only.

```
subroutine sub(arg)
  real arg
  pointer(ptr1, pte1)
  pointer(ptr2, pte2)
  real pte1, pte2

  ptr1 = malloc(%val(4))
  ptr2 = malloc(%val(4))
  pte1 = arg*arg
  pte2 = int(sqrt(arg))
  arg = pte1 + pte2
  call free(%val(ptr1))
  call free(%val(ptr2))
end subroutine
```

If most array assignments in a compilation unit involve arrays that do not overlap but a few assignments do involve storage-associated arrays, you can code the overlapping assignments with an extra step so that the `NOARYOVRLP` suboption is still safe to use.

```
@PROCESS ALIAS(NOARYOVRLP)
! The assertion that no array assignments involve overlapping
! arrays allows the assignment to be done without creating a
! temporary array.
program test
  real(8) a(100)
  integer :: j=1, k=50, m=51, n=100

  a(1:50) = 0.0d0
  a(51:100) = 1.0d0

  ! Timing loop to achieve accurate timing results
  do i = 1, 1000000
    a(j:k) = a(m:n)    ! Here is the array assignment
  end do

  print *, a
end program

! We cannot assert that this unit is free
! of array-assignment aliasing because of the assignments below.
subroutine sub1
  integer a(10), b(10)
  equivalence (a, b(3))
  a = b          ! a and b overlap.
  a = a(10:1:-1) ! The elements of a are reversed.
end subroutine

! When the overlapping assignment is recoded to explicitly use a
! temporary array, the array-assignment aliasing is removed.
! Although ALIAS(NOARYOVRLP) does not speed up this assignment,
! subsequent assignments of non-overlapping arrays in this unit
! are optimized.
@PROCESS ALIAS(NOARYOVRLP)
subroutine sub2
  integer a(10), b(10), t(10)
  equivalence (a, b(3))
  t = b; a = t
  t = a(10:1:-1); a = t
end subroutine
```

When **SUB1** is called, an alias exists between **J** and **K**. **J** and **K** refer to the same item in storage. In Fortran, this aliasing is not permitted if **J** or **K** are updated, and, if it is left undetected, it can have unpredictable results.

```
CALL SUB1(I,I)
...
SUBROUTINE SUB1(J,K)
```

In the following example, the program might store 5 instead of 6 into **J** unless **-qalias=nostd** indicates that an alias might exist.

```
INTEGER BIG(1000)
INTEGER SMALL(10)
COMMON // BIG
EQUIVALENCE(BIG,SMALL)
...
BIG(500) = 5
SMALL (I) = 6 ! Where I has the value 500
J = BIG(500)
```

### Related information

- See *Optimizing your applications* in the *XL Fortran Optimization and Programming Guide* for information on aliasing strategies you should consider.

## -qalign


### Category

Portability and migration

### Purpose

Specifies the alignment of data objects in storage, which avoids performance problems with misaligned data.

### Format



@PROCESS:

```
@PROCESS ALIGN({[NO]4K|STRUCT{(suboption)}|BINDC{(suboption)}})
```

### Defaults

`-qalign= no4k:struct=natural:bindc=linuxppc.`

### Parameters

The `[no]4k`, `bindc`, and `struct` options can be specified and are not mutually exclusive. The `[no]4k` option is useful primarily in combination with logical volume I/O and disk striping.

**[no]4k**

Specifies whether to align large data objects on page (4 KB) boundaries, for improved performance with data-stripped I/O. Objects are affected depending on their representation within the object file. The affected objects are arrays and structures that are 4 KB or larger and are in static or bss storage and also CSECTs (typically **COMMON** blocks) that are 8 KB or larger. A large **COMMON** block, equivalence group containing arrays, or structure is aligned on a page boundary, so the alignment of the arrays depends on their position within the containing object. Inside a structure of non-sequence derived type, the compiler adds padding to align large arrays on page boundaries.

**bindc={suboption}**

Specifies that the alignment and padding for an XL Fortran derived type with the **BIND(C)** attribute is compatible with a C struct type that is compiled with the corresponding XL C alignment option. The compatible alignment options include:

XL Fortran Option	Corresponding XL C Option
<b>-qalign=bindc=bit_packed</b>	<b>-qalign=bit_packed</b>
<b>-qalign=bindc=linuxppc</b>	<b>-qalign=linuxppc</b>

**struct={suboption}**

The **struct** option specifies how objects or arrays of a derived type declared using a record structure are stored, and whether or not padding is used between components. All program units must be compiled with the same settings of the **-qalign=struct** option. The three suboptions available are:

**packed**

If the **packed** suboption of the **struct** option is specified, objects of a derived type are stored with no padding between components, other than any padding represented by **%FILL** components. The storage format is the same as would result for a sequence structure whose derived type was declared using a standard derived type declaration.

**natural**

If the **natural** suboption of the **struct** option is specified, objects of a derived type are stored with sufficient padding such that components will be stored on their natural alignment boundaries, unless storage association requires otherwise. The natural alignment boundaries for objects of a type that appears in the left-hand column of the following table is shown in terms of a multiple of some number of bytes in the corresponding entry in the right-hand column of the table.

Type	Natural Alignment (in multiples of bytes)
INTEGER(1), LOGICAL(1), BYTE, CHARACTER	1
INTEGER(2), LOGICAL(2)	2
INTEGER(4), LOGICAL(4), REAL(4)	4
INTEGER(8), LOGICAL(8), REAL(8), COMPLEX(4)	8
REAL(16), COMPLEX(8), COMPLEX(16)	16

Type	Natural Alignment (in multiples of bytes)
Derived	Maximum alignment of its components

If the **natural** suboption of the **struct** option is specified, arrays of derived type are stored so that each component of each element is stored on its natural alignment boundary, unless storage association requires otherwise.

### port

If the **port** suboption of the **struct** option is specified,

- Storage padding is the same as described above for the **natural** suboption, with the exception that the alignment of components of type complex is the same as the alignment of components of type real of the same kind.
- The padding for an object that is immediately followed by a union is inserted at the beginning of the first map component for each map in that union.

### Restrictions

The **port** suboption does not affect any arrays or structures with the **AUTOMATIC** attribute or arrays that are allocated dynamically. Because this option may change the layout of non-sequence derived types, when compiling programs that read or write such objects with unformatted files, use the same setting for this option for all source files.

You can tell if an array has the **AUTOMATIC** attribute and is thus unaffected by **-qalign=4k** if you look for the keywords **AUTOMATIC** or **CONTROLLED AUTOMATIC** in the listing of “-qattr” on page 96. This listing also shows the offsets of data objects.

## -qaltivec

### Category

Language element control

### @PROCESS

ALTIVEC(LE | BE)

### Purpose

Specifies the order of vector elements in vector registers.

### Syntax

```

▶▶ -q—altivec—[=le]
                    [=be]
▶▶

```

## Defaults

`-qaltivec=le`

## Parameters

- be** Specifies big endian element order. Vectors are laid out in vector registers from left to right, so that element 0 is the leftmost element in the register.
- le** Specifies little endian element order. Vectors are laid out in vector registers from right to left, so that element 0 is the rightmost element in the register.

## Usage

The `-qaltivec` option affects the following categories of procedures:

- Vector Multimedia Extension (VMX) load and store intrinsic procedures
- Vector Scalar Extension (VSX) load and store intrinsic procedures
- The nonload and nonstore intrinsic procedures referring to the vector element order

The following list shows all the procedures affected:

- Load procedures
  - VMX load procedures: **VEC\_LD**, **VEC\_LDE**, and **VEC\_LDL**
  - VSX load procedures: **VEC\_XLD2**, **VEC\_XLW4**, and **VEC\_XL**
- Store procedures
  - VMX store procedures: **VEC\_ST**, **VEC\_STE**, and **VEC\_STL**
  - VSX store procedures: **VEC\_XSTD2**, **VEC\_XSTW4**, and **VEC\_XST**
- Nonload and nonstore procedures: **VPERMXOR**, **VEC\_EXTRACT**, **VEC\_INSERT**, **VEC\_MERGEH**, **VEC\_MERGEL**, **VEC\_MULE**, **VEC\_MULO**, **VEC\_PACK**, **VEC\_PERM**, **VEC\_PROMOTE**, **VEC\_SPLAT**, **VEC\_SUM2S**, **VEC\_SUMS**, **VEC\_UNPACKH**, and **VEC\_UNPACKL**

## Examples

- To change the vector element sequence to big endian element order in registers, enter the following command:  

```
xlf95 myprogram.f -qaltivec=be
```

## Related information

- “-qarch”
- Vector intrinsic procedures (IBM extension)
- Vector data types
- Vector layout in register (big endian or little endian element order)
- “-qsimd” on page 225
- *AltiVec Technology Programming Interface Manual*, available at [http://www.freescale.com/files/32bit/doc/ref\\_manual/ALTIVECPIM.pdf](http://www.freescale.com/files/32bit/doc/ref_manual/ALTIVECPIM.pdf)

---

## -qarch

### Category

Optimization and tuning

## @PROCESS

None.

### Purpose

Specifies the processor architecture, or family of architectures, where the code may run. This allows the compiler to take maximum advantage of the machine instructions specific to an architecture, or common to a family of architectures.

### Syntax

►► -q-arch =pwr8  
=auto ►►

### Defaults

- `-qarch=pwr8`
- `-qarch=auto` when `-04` or `-05` is in effect

### Parameters

#### **auto**

Automatically detects the specific architecture of the compiling machine. It assumes that the execution environment will be the same as the compilation environment. This option is implied if the `-04` or `-05` option is set or implied.

#### **pwr8**

Produces object code containing instructions that run on the POWER8 hardware platforms.

### Usage

For any given or `-qarch` setting, the compiler defaults to a specific, matching or `-qtune` setting, which can provide additional performance improvements. For detailed information about using `-qarch` and `-qtune` together, see “`-qtune`” on page 250.

The POWER8 architecture supports graphics, square root, Vector Multimedia Extension (VMX) processing, Vector Scalar Extension (VSX) processing, hardware transactional memory, and cryptography.

### Examples

To specify that the executable program `testing` compiled from `myprogram.f` is to run on a computer with VSX instruction support, enter:

```
xlf -o testing myprogram.f -qarch=pwr8
```

### Related information

- `-qprefetch`
- `-qfloat`
- “`-qtune`” on page 250
- “Compiling for specific architectures” on page 29



When **-qassert=contig** is specified, the compiler can perform optimizations according to the memory layout of the objects occupying contiguous blocks of memory.

Using **-qassert=contig** does not have the same effect as specifying the **CONTIGUOUS** attribute for array pointers and assumed-shape arrays. **-qassert=contig** does not validate contiguity assertion. [▶ F2008](#) To ensure Fortran 2008 semantics, use the **CONTIGUOUS** attribute. [▶ F2008](#)

**Notes:**

- **-qassert=contig** is not supported through the **ASSERT** directive.
- Using this suboption might produce unexpected results without warning.

**REFalign | NOREFalign**

Specifies that all pointers inside the compilation unit only point to data that is naturally aligned according to the length of the pointer types. With this assertion, the compiler might generate more efficient code. This assertion is particularly useful when you target a SIMD architecture with **-qhot=level=0** or **-qhot=level=1** with **-qsimd=auto**.

## Usage

The **itercnt**, **minitercnt**, and **maxitercnt** values are not required to be accurate. The values only affect performance, never correctness. Specify the values following the rule **minitercnt <= itercnt <= maxitercnt**. Otherwise, messages are issued to indicate that the values are inconsistent and the inconsistent value is ignored.

## Related information

- *High-order transformation* in the *XL Fortran Optimization and Programming Guide* for background information and instructions on using these assertions
- The **ASSERT** directive in the *XL Fortran Language Reference*
- [▶ F2008](#) The **CONTIGUOUS** attribute [▶ F2008](#)

## -qattr

### Category

Listings, messages, and compiler information

### Purpose

Produces a compiler listing that includes the attribute component of the attribute and cross-reference section of the listing.

### Syntax



**@PROCESS:**

**@PROCESS ATTR[(FULL)] | NOATTR**

## Defaults

-qnoattr

## Parameters

### full

Reports all identifiers in the program, whether they are referenced or not. If you specify **-qattr** without this suboption, reports only those identifiers that are used.

## Usage

If you specify **-qattr** after **-qattr=full**, the full attribute listing is still produced.

You can use the attribute listing to help debug problems caused by incorrectly specified attributes or as a reminder of the attributes of each object while writing new code.

## Related information

- “Listings, messages, and compiler information” on page 56
- “Attribute and cross reference section” on page 295

---

## -qautodbl

### Category

Floating-point and integer control

### Purpose

Provides an automatic means of converting single-precision floating-point calculations to double-precision and of converting double-precision calculations to extended-precision.

### Syntax

►► -qautodbl=

none
dbl4
dbl8
dbl

►►

@PROCESS:

@PROCESS AUTODBL(*setting*)

### Defaults

**-qautodbl=none**

### Parameters

The **-qautodbl** suboptions offer different strategies to preserve storage relationships between objects that are promoted or padded and those that are not.

The settings you can use are as follows:

**none** Does not promote or pad any objects that share storage. This setting is the default.

**dbl4** Promotes floating-point objects that are single-precision (4 bytes in size) or that are composed of such objects (for example, **COMPLEX** or array objects):

- **REAL(4)** is promoted to **REAL(8)**.
- **COMPLEX(4)** is promoted to **COMPLEX(8)**.

This suboption requires the **libxlfpm4.a** library during linking.

**dbl8** Promotes floating-point objects that are double-precision (8 bytes in size) or that are composed of such objects:

- **REAL(8)** is promoted to **REAL(16)**.
- **COMPLEX(8)** is promoted to **COMPLEX(16)**.

This suboption requires the **libxlfpm8.a** library during linking.

**dbl** Combines the promotions that **dbl4** and **dbl8** perform.

This suboption requires the **libxlfpm4.a** and **libxlfpm8.a** libraries during linking.

## Usage

You might find this option helpful in porting code where storage relationships are significant and different from the XL Fortran defaults. For example, programs that are written for the IBM VS FORTRAN compiler may rely on that compiler's equivalent option.

If the appropriate **-qautodbl** option is specified during linking, the program is automatically linked with the necessary extra libraries. Otherwise, you must link them in manually.

When you have both **REAL(4)** and **REAL(8)** calculations in the same program and want to speed up the **REAL(4)** operations without slowing down the **REAL(8)** ones, use **dbl4**.

If you want maximum precision of all results, you can use **dbl**. **dbl4** and **dbl8** select a subset of real types that have their precision increased.

By using **dbl4**, you can increase the size of **REAL(4)** objects without turning **REAL(8)** objects into **REAL(16)**s. **REAL(16)** is less efficient in calculations than **REAL(8)** is.

The **-qautodbl** option handles calls to intrinsics with arguments that are promoted; when necessary, the correct higher-precision intrinsic function is substituted. For example, if single-precision items are being promoted, a call in your program to **SIN** automatically becomes a call to **DSIN**.

You must not specify the **-qautodbl** option if your program contains vector types.

## Restrictions

- Because character data is not promoted or padded, its relationship with storage-associated items that are promoted or padded may not be maintained.
- If the storage space for a pointee is acquired through the system routine **malloc**, the size specified to **malloc** should take into account the extra space needed to represent the pointee if it is promoted or padded.

- If an intrinsic function cannot be promoted because there is no higher-precision specific name, the original intrinsic function is used, and the compiler displays a warning message.
- You must compile every compilation unit in a program with the same **-qautodbl** setting.

## Related information

For background information on promotion, padding, and storage/value relationships and for some source examples, see “Implementation details for -qautodbl promotion and padding” on page 300.

“-qrealsize” on page 213 describes another option that works like **-qautodbl**, but it only affects items that are of default kind type and does not do any padding. If you specify both the **-qrealsize** and the **-qautodbl** options, only **-qautodbl** takes effect. Also, **-qautodbl** overrides the **-qdpc** option.

## -qbindcextname

### Category

Portability and migration

### Purpose

Controls whether the **-qextname** option affects **BIND(C)** entities.

### Syntax



@PROCESS:

@PROCESS **BINDCEXTNAME** | NOBINDCEXTNAME

### Defaults

-qbindcextname

### Usage

The **-qextname** option and the **BIND(C)** attribute are two ways of modifying the names of Fortran global entities to facilitate use in C.

If you explicitly specify a **BIND(C)** binding label in an interface block using the **NAME=** specifier, the compiler uses this binding label in calls to the procedure regardless of the **-qextname** and **-qbindcextname** options.

If your interface does not explicitly specify a **BIND(C)** binding label using the **NAME=** specifier, the compiler creates an implicit binding label. If you also specify the **-qextname** option, the compiler appends an underscore to the implicit binding label only when the **-qbindcextname** option is in effect.

If you specify the **-qextname** and **-qbindcextname** options for a compilation unit declaring a **BIND(C)** procedure, the compiler appends an underscore to the binding label, even when the binding label is explicitly specified.

**Notes:**

- You must ensure that the names of a **BIND(C)** entity are identical. Accordingly, if two compilation units access the same **BIND(C)** entity that does not have an explicitly-specified binding label, you must not compile one unit with the **-qbindcextname** option and the other with the **-qnobindcextname** option.
- The **-q[no]bindcextname** option has effect only if the **-qextname** option is also specified. If the **-qextname** option is specified with a list of named entities, the **-q[no]bindcextname** option only affects these named entities.

**Examples**

```
interface
  integer function foo() bind(c)
  end function
  integer function bar()
  end function
end interface

print *, foo()
print *, bar()
end
```

```
xlf90 x.f -qextname -qbindcextname      # calls "foo_", and "bar_"
xlf90 x.f -qextname -qnobindcextname    # calls "foo", and "bar"
xlf90 x.f -qextname=foo -qbindcextname  # calls "foo_", and "bar_"
xlf90 x.f -qextname=foo -qnobindcextname # calls "foo", and "bar"
xlf90 x.f                               # calls "foo", and "bar"
xlf90 x.f -qnobindcextname              # calls "foo", and "bar"
```

**Related information**

- “-qextname” on page 122
- "BIND (Fortran 2003)"
- "Binding labels"

---

## **-qcache**

**Category**

Optimization and tuning

**@PROCESS**


None.

**Purpose**

Specifies the cache configuration for a specific execution machine.

The compiler uses this information to tune program performance, especially for loop operations that can be structured (or *blocked*) to process only the amount of data that can fit into the data cache.

## Syntax



## Defaults

Not applicable.

## Parameters

**assoc=number**

Specifies the set associativity of the cache:

- 0** Direct-mapped cache
- 1** Fully associative cache
- n > 1** *n*-way set-associative cache

**auto** Automatically detects the specific cache configuration of the compiling machine. It assumes that the execution environment will be the same as the compilation environment.

**cost=cycles**

Specifies the performance penalty that results from a cache miss so that the compiler can decide whether to perform an optimization that might result in extra cache misses.

**level=level**

Specifies which level of cache is affected:

- 1** Basic cache
- 2** Level-2 cache or the table lookaside buffer (TLB) if the machine has no level-2 cache
- 3** TLB in a machine that does have a level-2 cache

Other levels are possible but are currently undefined. If a system has more than one level of cache, use a separate **-qcache** option to describe each level.

**line=bytes**

Specifies the line size of the cache.

**size=Kbytes**

Specifies the total size of this cache.

**type={C|c|D|d|I|i}**

Specifies the type of cache that the settings apply to, as follows:

- **C** or **c** for a combined data and instruction cache
- **D** or **d** for the data cache
- **I** or **i** for the instruction cache

## Usage

If you know exactly what type of system a program is intended to be executed on and that system has its instruction or data cache configured differently from the default case (as governed by the **-qtune** setting), you can specify the exact characteristics of the cache to allow the compiler to compute more precisely the benefits of particular cache-related optimizations.

For the **-qcache** option to have any effect, you must include the **level** and **type** suboptions and specify the **-qhot** option or an option that implies **-qhot**.

- If you know some but not all of the values, specify the ones you do know.
- If a system has more than one level of cache, use a separate **-qcache** option to describe each level. If you have limited time to spend experimenting with this option, it is more important to specify the characteristics of the data cache than of the instruction cache.
- If you are not sure of the exact cache sizes of the target systems, use relatively small estimated values. It is better to have some cache memory that is not used than to have cache misses or page faults from specifying a cache that is larger than the target system has.

If you specify the wrong values for the cache configuration or run the program on a machine with a different configuration, the program may not be as fast as possible but will still work correctly. Remember, if you are not sure of the exact values for cache sizes, use a conservative estimate.

## Examples

To tune performance for a system with a combined instruction and data level-1 cache where the cache is two-way associative, 8 KB in size, and has 64-byte cache lines:

```
xlf95 -O3 -qhot -qcache=type=c:level=1:size=8:line=64:assoc=2 file.f
```

To tune performance for a system with two levels of data cache, use two **-qcache** options:

```
xlf95 -O3 -qhot -qcache=type=D:level=1:size=256:line=256:assoc=4 \  
-qcache=type=D:level=2:size=512:line=256:assoc=2 file.f
```

To tune performance for a system with two types of cache, again use two **-qcache** options:

```
xlf95 -O3 -qhot -qcache=type=D:level=1:size=256:line=256:assoc=4 \  
-qcache=type=I:level=1:size=512:line=256:assoc=2 file.f
```

## Related information

- “-qarch” on page 93
- “-qhot” on page 143
- “-qtune” on page 250

---

## -qcclines

### Category

Input control

### Purpose

Determines whether the compiler recognizes conditional compilation lines in fixed source form and F90 free source form. This option is not supported with IBM free source form.

### Syntax

►► -q cclines  
nocclines ►►

@PROCESS:

@PROCESS CCLINES | NOCCLINES

### Defaults

The default is **-qcclines** if the **-qsmp=omp** option is turned on; otherwise, the default is **-qnocclines**.

### Related information

- *Conditional compilation in the XL Fortran Language Reference*

---

## -qcheck

### Category

Error checking and debugging

### Purpose

Generates code that performs certain types of runtime checking.

**-qcheck** is the long form of the **-C** option.

### Syntax

►► -q nocheck  
check all  
bounds  
nobounds  
stacklobber  
nostacklobber  
unset  
nounset ►►

@PROCESS:

@PROCESS CHECK[(*suboptions*)] | **NOCHECK**

## Defaults

-qnocheck

## Parameters

### **all**

Enables all suboptions.

### **bounds | nobounds**

Checks each reference to an array element, array section, or character substring to ensure the reference stays within the defined bounds of the entity.

### **stacklobber | nostacklobber**

Detects stack corruption of nonvolatile registers in the save area in user programs. This type of corruption happens only if any of the nonvolatile registers in the save area of the stack is modified.

If the **-qstackprotect** option and this suboption are both on, this suboption catches the stack corruption first.

### **unset | nounset**

Checks for automatic variables that are used before they are set. A trap will occur at run time if an automatic variable is not set before it is used.

The **-qinitauto** option initializes automatic variables. As a result, the **-qinitauto** option hides uninitialized variables from the **-qcheck=unset** option.

The **-qsave** option changes the storage class of automatic variables to `STATIC`. As a result, the **-qsave** option hides variables that are used before they are set from the **-qcheck=unset** option.

Specify the **-qsigtrap** option to get a traceback that shows the procedure calls leading to the use of the variable. If you also specify the **-g** option, the traceback will include the line numbers and procedure names.

Specifying **-qcheck** with no suboption is equivalent to specifying **-qcheck=all**.

## Usage

You can specify the **-qcheck** option more than once. The suboption settings are accumulated, but the later suboptions override the earlier ones.

You can use the **all** suboption along with the **no...** form of one or more of the other options as a filter. For example, using:

```
xlf myprogram.f -qcheck=all:nobounds
```

provides checking for everything except for references that go out of bounds. If you use **all** with the **no...** form of the suboptions, **all** should usually be the first suboption.

At compile time, if the compiler can determine that a reference goes out of bounds, the severity of the error reported is increased to **S** (severe) when **-qcheck=bounds** is enabled.

At run time, if a reference goes out of bounds or if a certain type of stack corruption is detected, the program generates a **SIGTRAP** signal. By default, this signal ends the program and produces a core dump. This is an expected behavior and does not indicate there is a defect in the compiler product.

Because runtime checking can slow execution, you should decide which is the more important factor for each program: the performance impact or the possibility of incorrect results if an error goes undetected. You might decide to use this option only while testing and debugging a program (if performance is more important) or also for compiling the production version (if safety is more important).

The **-qcheck** option prevents some optimizations. You may want to remove the **-qcheck** option after the debugging of your code is complete and then add any wanted optimization options for better performance.

The valid bounds for character substring expressions differ depending on the setting of the **-qzerosize** option.

## Examples

The following code example shows an out-of-bounds access of an allocatable array that is detected by **-qcheck=bounds**:

```
program outofbounds
  real, allocatable :: x(:)
  allocate(x(5:10))
  x(1) = 3.0 ! Out of bound access.
  print *, x
end
```

In the following example, a random assignment of an integer pointer corrupts the save area of the call stack. **-qcheck=stackcllobber** detects the problem and causes a trap:

```
subroutine update(i, off, value)
  implicit none
  integer, value :: i, off
  character, value :: value
  integer pointee
  pointer(p, pointee)
  p = i + off
  pointee = ichar(value)
end subroutine

subroutine sub1()
  implicit none
  interface
    subroutine update(i, off, value)
      integer, value :: i, off
      character, value :: value
    end subroutine
  end interface
  character(9) buffer
  buffer = ""
  ! This call to update will corrupt the call stack.
  ! Offset 48 is within the save area for this program.
  call update(loc(buffer), 48, 'a')
  print *, buffer
end subroutine

program main
  call sub1
end program
```

**Note:** The offset of the save area of the call stack depends on the program source and optimization level. When **-O2** or a lower optimization level is in effect, offset 48 falls within the save area.

In the following function `factorial`, `temp` is not initialized when `n<=1`, and `result` is accessed before it is set when `n>1`. To compile `factorial.f` to detect this issue and causes traps during run time, enter the following command:

```
xlf95 -qcheck=unset -O -g -qsigtrap factorial.f
```

`factorial.f` contains the following code:

```
module m
contains
  recursive function factorial(n) result(result)
    integer, value :: n
    integer result, temp

    if (n > 1) then
      temp = n * factorial(n - 1)
      print *, result ! line 9
    endif

    result = temp ! line 12
  end function
end module

use m
integer x
x = factorial(1)
end
```

The compiler issues the following informational messages and traps occur near line 12 and line 18 during run time:

```
1500-098 (I) "factorial.f", line 9: "result" is used before it is set.
1500-098 (I) "factorial.f", line 12: "temp" might be used before it is set.
1501-510 Compilation successful for file factorial.f.
$ ./a.out
```

```
Signal received: SIGTRAP - Trace trap
Fortran language trap: reference to unset memory location
```

```
Traceback:
Offset 0x00000014 in procedure __m_NMOD_factorial, near line 12 in file factorial.f
Offset 0x00000028 in procedure __main, near line 18 in file factorial.f
--- End of call chain ---
```

**Note:** If you set `-qcheck=unset` at `noopt`, the compiler does not issue informational messages at compile time.

### Related information

- “`-qstackprotect`” on page 233
- “`-qhot`” on page 143
- “`-qzerosize`” on page 270
- “`-qsigtrap`” on page 224 and *Installing an exception handler* in the *XL Fortran Optimization and Programming Guide* describe how to detect and recover from `SIGTRAP` signals without ending the program.

---

## -qci

### Category

Input control

## Purpose

Specifies the identification numbers (from 1 to 255) of the **INCLUDE** lines to process.

## Syntax

►► -qci= number

**@PROCESS:**

@PROCESS CI(*number*,...,*number*)

## Defaults

Not applicable.

## Usage

This option allows a kind of conditional compilation because you can put code that is only sometimes needed (such as debugging **WRITE** statements, additional error-checking code, or XLF-specific code) into separate files and decide for each compilation whether to process them.

If an **INCLUDE** line has a number at the end, the file is only included if you specify that number in a **-qci** option. The set of identification numbers that is recognized is the union of all identification numbers that are specified on all occurrences of the **-qci** option.

### Note:

1. Because the optional number in **INCLUDE** lines is not a widespread XL Fortran feature, using it may restrict the portability of a program.
2. This option works only with the XL Fortran **INCLUDE** directive and not with the **#include** C preprocessor directive.

## Examples

```
REAL X /1.0/  
INCLUDE 'print_all_variables.f' 1  
X = 2.5  
INCLUDE 'print_all_variables.f' 1  
INCLUDE 'test_value_of_x.f' 2  
END
```

In this example, compiling without the **-qci** option simply declares **X** and assigns it a value. Compiling with **-qci=1** includes two instances of an include file, and compiling with **-qci=1:2** includes both include files.

## Related information

- The **INCLUDE** directive in the *XL Fortran Language Reference*

---

## -qcompact

### Category

Optimization and tuning

### Purpose

Avoids optimizations that increase code size.

### Syntax

►► -q nocompact  
compact ►►

**@PROCESS:**

@PROCESS COMPACT | NOCOMPACT

### Defaults

-qnocompact

### Usage

By default, some techniques the optimizer uses to improve performance, such as loop unrolling and array vectorization, may also make the program larger. For systems with limited storage, you can use **-qcompact** to reduce the expansion that takes place. If your program has many loop and array language constructs, using the **-qcompact** option will affect your application's overall performance. You may want to restrict using this option to those parts of your program where optimization gains will remain unaffected.

### Rules

With **-qcompact** in effect, other optimization options still work; the reductions in code size come from limiting code replication that is done automatically during optimization.

This option takes effect only when it is specified at the **-O2** optimization level, or higher.

---

## -qcr

### Category

Input control

**@PROCESS**

None.

### Purpose

Controls how the compiler interprets the CR (carriage return) character.

This option allows you to compile code written using a Mac OS or DOS/Windows editor.

## Syntax

►► -q no cr

## Defaults

By default, the CR (Hex value X'0d') or LF (Hex value X'0a') character, or the CRLF (Hex value X'0d0a') combination indicates line termination in a source file.

## Usage

If you specify **-qno cr**, the compiler recognizes only the LF character as a line terminator. You must specify **-qno cr** if you use the CR character for a purpose other than line termination.

---

## -qctyp1ss

### Category

Portability and migration

### Purpose

Specifies whether character constant expressions are allowed wherever typeless constants may be used.

This language extension might be needed when you are porting programs from other platforms.

### Syntax

►► -q noctyp1ss ctyp1ss noarg arg

@PROCESS:

@PROCESS CTYPLSS[([NO]ARG)] | NOCTYPLSS

### Defaults

-qnoctyp1ss

### Parameters

arg | noarg

Suboptions retain the behavior of **-qctyp1ss**. Additionally, **arg** specifies that Hollerith constants used as actual arguments will be treated as integer actual arguments.

## Usage

With **-qctyp1ss**, character constant expressions are treated as if they were Hollerith constants and thus can be used in logical and arithmetic expressions.

- If you specify the **-qctyp1ss** option and use a character-constant expression with the **%VAL** argument-list keyword, a distinction is made between Hollerith constants and character constants. Character constants are placed in the rightmost byte of the register and padded on the left with zeros, while Hollerith constants are placed in the leftmost byte and padded on the right with blanks. All of the other **%VAL** rules apply.
- The option does not apply to character expressions that involve a constant array or subobject of a constant array at any point.

## Examples

**Example 1:** In the following example, the compiler option **-qctyp1ss** allows the use of a character constant expression.

```
@PROCESS CTYPLSS
  INTEGER I,J
  INTEGER, PARAMETER :: K(1) = (/97/)
  CHARACTER, PARAMETER :: C(1) = (/ 'A' /)

  I = 4HABCD      ! Hollerith constant
  J = 'ABCD'      ! I and J have the same bit representation

! These calls are to routines in other languages.
  CALL SUB(%VAL('A')) ! Equivalent to CALL SUB(97)
  CALL SUB(%VAL(1HA)) ! Equivalent to CALL SUB(1627389952)

! These statements are not allowed because of the constant-array
! restriction.
!   I = C // C
!   I = C(1)
!   I = CHAR(K(1))
  END
```

**Example 2:** In the following example, the variable *J* is passed by reference. The suboption **arg** specifies that the Hollerith constant is passed as if it were an integer actual argument.

```
@PROCESS CTYPLSS(ARG)
  INTEGER :: J

  J = 3HIBM
! These calls are to routines in other languages.
  CALL SUB(J)
  CALL SUB(3HIBM) ! The Hollerith constant is passed as if
                  ! it were an integer actual argument
```

## Related information

- *Hollerith constants in the XL Fortran Language Reference*
- *Passing arguments by reference or by value in the XL Fortran Optimization and Programming Guide*

---

## -qdbg


### Category

Error checking and debugging

## Purpose

`-qdbg` is the long form of `-g`.

## Syntax



**@PROCESS:**

@PROCESS DBG | NODBG

## Defaults

`-qnodbg`, or `-qdbg=level=0`

## Parameters

**`-qdbg=level=0`**

Equivalent to `-qnodbg` or `-g0`.

**`-qdbg=level=1`**

Equivalent to `-g1` or `-qlinedebug`.

**`-qdbg=level=2`**

Equivalent to `-g2`.

**`-qdbg=level=3`**

Equivalent to `-g3`.

**`-qdbg=level=4`**

Equivalent to `-g4`.

**`-qdbg=level=5`**

Equivalent to `-g5`.

**`-qdbg=level=6`**

Equivalent to `-g6`.

**`-qdbg=level=7`**

Equivalent to `-g7`.

**`-qdbg=level=8`**

Equivalent to `-g8`.

**`-qdbg=level=9`**

Equivalent to `-g9`.

## Related information

- “-g” on page 73
- “-qlinedebug” on page 173

---

## -qddim

### Category

Portability and migration

### Purpose

Specifies that the bounds of pointee arrays are re-evaluated each time the arrays are referenced and removes some restrictions on the bounds expressions for pointee arrays.

### Syntax

►► -q noddim  
ddim ◀◀

@PROCESS:

@PROCESS DDIM | NODDIM

### Defaults

-qnoddim

### Usage

By default, a pointee array can only have dimension declarators containing variable names if the array appears in a subprogram, and any variables in the dimension declarators must be dummy arguments, members of a common block, or use- or host-associated. The size of the dimension is evaluated on entry to the subprogram and remains constant during execution of the subprogram.

With the **-qddim** option:

- The bounds of a pointee array are re-evaluated each time the pointee is referenced. This process is called *dynamic dimensioning*. Because the variables in the declarators are evaluated each time the array is referenced, changing the values of the variables changes the size of the pointee array.
- The restriction on the variables that can appear in the array declarators is lifted, so ordinary local variables can be used in these expressions.
- Pointee arrays in the main program can also have variables in their array declarators.

### Examples

```
@PROCESS DDIM
INTEGER PTE, N, ARRAY(10)
POINTER (P, PTE(N))
DO I=1, 10
  ARRAY(I)=I
END DO
N = 5
P = LOC(ARRAY(2))
```

```
PRINT *, PTE ! Print elements 2 through 6.  
N = 7 ! Increase the size.  
PRINT *, PTE ! Print elements 2 through 8.  
END
```

---

## -qdescriptor

### Category

Portability and migration

### @PROCESS

None.

### Purpose

Specifies the XL Fortran internal descriptor data structure format to use for non object-oriented entities in your compiled applications.

### Syntax

►► -qdescriptor=v1/v2 ►►

### Defaults

- -qdescriptor=v1

### Parameters

- v1** Use an internal descriptor data structure format that is compact, but incapable of representing some of the new Fortran language features such as object-orientation.
- v2** Use an internal descriptor data structure format that is extensible. This setting allows your programs to take advantage of Fortran's object-oriented features, as well as parameterized derived types.

### Usage

Regardless of what **-qdescriptor** setting is in effect, applications containing object-oriented constructs or parameterized derived types will use the **v2** data structure format for those constructs.

The choice of **-qdescriptor** setting is an important consideration when building libraries or modules for distribution. Users of these libraries and modules will need to be aware of the **-qdescriptor** setting and compile the code that uses them in a compatible way. It is suggested that such libraries and modules be built with the **-qsaveopt** option so that the objects themselves will encode the compilation options in a user-readable form.

If you are building modules that contain user-visible derived types, consider building them with the **-qxlf2003=polymorphic** suboption. This allows users of the module to use or extend the derived types in a Fortran object-oriented context that uses polymorphism.

In the Fortran 2003 object-oriented programming model, the XL Fortran compiler supports using types and type extensions from types defined in modules not compiled with `-qxlf2003=polymorphic`, as long as the types are not used in a context that requires polymorphism. However, if the compiler detects the attempted use of a type or a type extension from a module not compiled with `-qxlf2003=polymorphic` in a context that requires polymorphism, an error message will be issued and compilation halted.

If a module built with the `-qdescriptor=v1` setting is used in a compilation where `-qdescriptor=v2` has been specified, the compiler will diagnose this mismatch and halt compilation after issuing an error message.

When using the `-qdescriptor=v2` option, the compiler is unable to diagnose unsafe usage where objects built with the `v2` setting are mixed with those built with the `v1` setting or with XL Fortran 10.1 or older compilers. Even if your program appears to function properly, this usage is unsupported. The descriptor formats are different sizes and, when used with certain constructs, data layouts will change resulting in undefined and unsupported behavior. For example, the sizes of allocatable and pointer entities within derived types will be different resulting a differing size for the derived type itself.

### Related information

- “-qsaveopt” on page 220
- “-qxlf2003” on page 263

## -qdirective

### Category


Input control

### Purpose

Specifies sequences of characters, known as trigger constants, that identify comment lines as compiler comment directives.

A compiler comment directive is a line that is not a Fortran statement but is recognized and acted on by the compiler.

### Format



**@PROCESS:**

@PROCESS DIRECTIVE[(directive\_list)] | NODIRECTIVE[(directive\_list)]

### Defaults

The compiler recognizes the default trigger constant **IBM\***.

Specifying `-qsmp` implies `-qdirective=smp\$: $omp:ibmp`, and, by default, the trigger constants **SMP\$**, **\$OMP**, and **IBMP** are also turned on. If you specify

**-qsmp=omp**, the compiler ignores all trigger constants that you have specified up to that point and recognizes only the **\$OMP** trigger constant. Specifying **-qthreaded** implies **-qdirective=ibmt**, and, by default, the trigger constant **IBMT** is also turned on.

## Parameters

The **-qdirective** option with no *directive\_list* turns off all previously specified directive identifiers; with a *directive\_list*, it turns off only the selected identifiers.

**-qdirective** with no *directive\_list* turns on the default trigger constant **IBM\*** if it has been turned off by a previous **-qdirective**.

## Usage

Note the following:

- Multiple **-qdirective** and **-qnodirective** options are additive; that is, you can turn directive identifiers on and off again multiple times.
- One or more *directive\_lists* can be applied to a particular file or compilation unit; any comment line beginning with one of the strings in the *directive\_list* is then considered to be a compiler comment directive.
- The trigger constants are not case-sensitive.
- The characters (, ), ', ", :, =, comma, and blank cannot be part of a trigger constant.
- To avoid wildcard expansion in trigger constants that you might use with these options, you can enclose them in single quotation marks on the command line. For example:

```
xlf95 -03 -qhot -qcache=type=D:level=1 -qdirective='dbg*' -qnodirective='IBM*' directives.f
```

- This option only affects Fortran directives that the XL Fortran compiler provides, not those that any preprocessors provide.
- As the use of incorrect trigger constants can generate warning messages, error messages, or both, you should check the particular directive statement for the suitable associated trigger constant.

## Examples

```
! This program is written in Fortran free source form.
PROGRAM DIRECTV
INTEGER A, B, C, D, E, F
A = 1 ! Begin in free source form.
B = 2
!OLDSTYLE SOURCEFORM(FIXED)
! Switch to fixed source form for this include file.
      INCLUDE 'set_c_and_d.inc'
!IBM* SOURCEFORM(FREE)
! Switch back to free source form.
E = 5
F = 6
END
```

For this example, compile with the option **-qdirective=oldstyle** to ensure that the compiler recognizes the **SOURCEFORM** directive before the **INCLUDE** line. After processing the include-file line, the program reverts back to free source form, after the **SOURCEFORM(FREE)** statement.

- The **SOURCEFORM** directive in the *XL Fortran Language Reference*
- The *Directives* section in the *XL Fortran Language Reference*

---

## **-qdirectstorage**

### **Category**

Optimization and tuning

### **@PROCESS**

None.

### **Context**

None.

### **Purpose**

Informs the compiler that a given compilation unit may reference write-through-enabled or cache-inhibited storage.

### **Format**

►► -q 

nodirectstorage
directstorage

 ◄◄

### **Defaults**

-qnodirectstorage

### **Usage**

Use this option with discretion. It is intended for programmers who know how the memory and cache blocks work, and how to tune their applications for optimal performance. For a program to execute correctly on all Power implementations of cache organization, the programmer should assume that separate instruction and data caches exist, and should program to the separate cache model.

**Note:** Using the **-qdirectstorage** option together with the **CACHE\_ZERO** directive may cause your program to fail, or to produce incorrect results.

### **Related information**

- **CACHE\_ZERO** in the *XL Fortran Language Reference*.

---

## **-qdlines**

### **Category**

Input control

### **Purpose**

**-qdlines** is the long form of **-D**.

## Format

►► -q nodlines  
dlines ►►

@PROCESS:

@PROCESS D LINES | NODLINES

## Defaults

-qnodlines

## -qdpc

### Category

Floating-point and integer control

### Purpose

Increases the precision of real constants for maximum accuracy, when assigning real constants to **DOUBLE PRECISION** variables.

This language extension might be needed when you are porting programs from other platforms.

### Format

►► -q nodpc  
dpc =e ►►

@PROCESS:

@PROCESS DPC[(E)] | NODPC

## Defaults

-qnodpc

### Usage

If you specify **-qdpc**, all basic real constants (for example, 1.1) are treated as double-precision constants; the compiler preserves some digits of precision that would otherwise be lost during the assignment to the **DOUBLE PRECISION** variable. If you specify **-qdpc=e**, all single-precision constants, including constants with an e exponent, are treated as double-precision constants.

This option does not affect constants with a kind type parameter specified.

**-qautodbl** and **-qrealsize** are more general-purpose options that can also do what **-qdpc** does. **-qdpc** has no effect if you specify either of these options.

## Examples

```
@process nodpc
  subroutine nodpc
    real x
    double precision y
    data x /1.000000000001/ ! The trailing digit is lost
    data y /1.000000000001/ ! The trailing digit is lost

    print *, x, y, x .eq. y ! So x is considered equal to y
  end

@process dpc
  subroutine dpc
    real x
    double precision y
    data x /1.000000000001/ ! The trailing digit is lost
    data y /1.000000000001/ ! The trailing digit is preserved

    print *, x, y, x .eq. y ! So x and y are considered different
  end

  program testdpc
    call nodpc
    call dpc
  end
```

When compiled, this program prints the following:

```
1.000000000    1.000000000000000000    T
1.000000000    1.0000000000100009      F
```

showing that with **-qdp** the extra precision is preserved.

- “-qautodbl” on page 97
- “-qrealsize” on page 213

---

## -qenum

### Category

Floating-point and integer control


### @PROCESS

None.

### Purpose

Specifies the range of the enumerator constant and enables storage size to be determined.

### Syntax



### Defaults

-qenum=4



- Porting code from another Fortran compiler that uses the backslash as an escape character.
- Including “unusual” characters, such as tabs or newlines, in character data. Without this option, the alternative is to encode the ASCII values (or EBCDIC values, on mainframe systems) directly in the program, making it harder to port.

If you are writing or porting code that depends on backslash characters being passed through unchanged, specify **-qnoescape** so that they do not get any special interpretation. You could also write `\\` to mean a single backslash character under the default setting.

## Examples

```
$ # Demonstrate how backslashes can affect the output
$ cat escape.f
      PRINT *,'a\bcd\fg'
      END
$ xlf95 escape.f
** _main === End of Compilation 1 ===
1501-510  Compilation successful for file escape.f.
$ a.out
cde
  g
$ xlf95 -qnoescape escape.f
** _main === End of Compilation 1 ===
1501-510  Compilation successful for file escape.f.
$ a.out
a\bcd\fg
```

In the first compilation, with the default setting of **-qescape**, `\b` is printed as a backspace, and `\f` is printed as a formfeed character.

With the **-qnoescape** option specified, the backslashes are printed like any other character.

## Related information

The list of escape sequences that XL Fortran recognizes is shown in *Escape sequences for character strings* in the *XL Fortran Optimization and Programming Guide*.

## -qessl

### Category

Optimization and tuning

### @PROCESS

None.

### Purpose

Allows the compiler to substitute the Engineering and Scientific Subroutine Library (ESSL) routines in place of Fortran 90 intrinsic procedures.

The ESSL is a collection of subroutines that provides a wide range of mathematical functions for various scientific and engineering applications. The subroutines are tuned for performance on specific architectures. Some of the Fortran 90 intrinsic procedures have similar counterparts in ESSL. Performance is improved when

these Fortran 90 intrinsic procedures are linked with ESSL. In this case, you can keep the interface of Fortran 90 intrinsic procedures, and get the added benefit of improved performance using ESSL.

## Syntax

►► `-q` noessl  
essl \_\_\_\_\_►►

## Defaults

`-qnoessl`

## Usage

Use the ESSL Serial Library when linking with `-lessl`. Use the ESSL SMP Library when linking with `-lessl smp`.

`-lessl` or `-lessl smp` must be used whenever code is being compiled with `-qessl`.

Also, since `libessl.so` and `libessl smp.so` have a dependency on `libxlf90_r.so`, compile with `xlf_r`, `xlf90_r`, or `xlf95_r`, which use `libxlf90_r.so` as the default to link. You can also specify `-lxlf90_r` on the link command line if you use the linker directly, or other commands to link.

The following MATMUL function calls may use ESSL routines when `-qessl` is enabled:

```
real a(10,10), b(10,10), c(10,10)
c=MATMUL(a,b)
```

## Related information

The ESSL libraries are not shipped with the XL Fortran compiler. For more information about these libraries, see the Engineering and Scientific Subroutine Library (ESSL) and Parallel ESSL web page.

---

## **-qextern**

### Category

Portability and migration

### @PROCESS

None.

### Purpose

Allows user-written procedures to be called instead of XL Fortran intrinsics.

### Syntax

►► `-q` `extern` `=` `names` \_\_\_\_\_►►

## Defaults

Not applicable.

## Parameters

*names*

A list of procedure names separated by colons.

## Usage

The procedure names are treated as if they appear in an **EXTERNAL** statement in each compilation unit being compiled. If any of your procedure names conflict with XL Fortran intrinsic procedures, use this option to call the procedures in the source code instead of the intrinsic ones.

Because of the many Fortran 90 and Fortran 95 intrinsic functions and subroutines, you might need to use this option even if you did not need it for FORTRAN 77 programs.

## Examples

```
subroutine matmul(res, aa, bb, ext)
  implicit none
  integer ext, i, j, k
  real aa(ext, ext), bb(ext, ext), res(ext, ext), temp
  do i = 1, ext
    do j = 1, ext
      temp = 0
      do k = 1, ext
        temp = temp + aa(i, k) * bb(k, j)
      end do
      res(i, j) = temp
    end do
  end do
end subroutine

implicit none
integer i, j, irand
integer, parameter :: ext = 100
real ma(ext, ext), mb(ext, ext), res(ext, ext)

do i = 1, ext
  do j = 1, ext
    ma(i, j) = float(irand())
    mb(i, j) = float(irand())
  end do
end do

call matmul(res, ma, mb, ext)
end
```

Compiling this program with no options fails because the call to **MATMUL** is actually calling the intrinsic subroutine, not the subroutine defined in the program. Compiling with **-qextern=matmul** allows the program to be compiled and run successfully.

---

## -qextname


### Category

Portability and migration

## Purpose

Adds an underscore to the names of all global entities.

## Syntax



**@PROCESS:**

**@PROCESS** EXTNAME[(*name1*, *name2*,...)] | **NOEXTNAME**

## Defaults

-qnoextname

## Parameters

*name*

Identifies a specific global entity (or entities). For a list of named entities, separate each name with a colon. For example: *name1: name2:....*

The name of a main program is not affected.

## Usage

The **-qextname** option helps to port mixed-language programs to XL Fortran without modifications.

Use of this option avoids naming problems that might otherwise be caused by:

- Fortran subroutines, functions, or common blocks that are named **main**, **MAIN**, or have the same name as a system subroutine
- Non-Fortran routines that are referenced from Fortran and contain an underscore at the end of the routine name

**Note:** XL Fortran Service and Utility Procedures, such as **flush\_** and **dtime\_**, have these underscores in their names already. By compiling with the **-qextname** option, you can code the names of these procedures without the trailing underscores.

- Non-Fortran routines that call Fortran procedures and use underscores at the end of the Fortran names
- Non-Fortran external or global data objects that contain an underscore at the end of the data name and are shared with a Fortran procedure

You must compile all the source files for a program, including the source files of any required module files, with the same **-qextname** setting.

If you use the **xlfortility** module to ensure that the Service and Utility subprograms are correctly declared, you must change the name to **xlfortility\_extname** when compiling with **-qextname**.

If there is more than one Service and Utility subprogram referenced in a compilation unit, using **-qextname** with no names specified and the **xlfortility\_extname** module may cause the procedure declaration check not to work accurately.

This option also affects the names that are specified in the **-qextern**, **-qinline**, and **-qsigtrap** options. You do not have to include underscores in their names on the command line.

## Examples

```
@PROCESS EXTNAME
  SUBROUTINE STORE_DATA
    CALL FLUSH(10) ! Using EXTNAME, we can drop the final underscore.
  END SUBROUTINE

@PROCESS(EXTNAME(sub1))
program main
  external :: sub1, sub2
  call sub1()      ! An underscore is added.
  call sub2()      ! No underscore is added.
end program
```

## Related information

- “-qextern” on page 121
- “-qinline” on page 157
- “-qsigtrap” on page 224
- “-qbindcextname” on page 99

---

## -qfdpr

### Category

Optimization and tuning

### @PROCESS

None.

### Purpose

Provides object files with information that the IBM Feedback Directed Program Restructuring (FDPR) performance-tuning utility needs to optimize the resulting executable file.

When **-qfdpr** is in effect, optimization data is stored in the object file.

### Syntax

```
►► -q nofdpr / fdpr ◀◀
```

### Defaults

-qnofdpr

## Usage

For best results, use **-qfdpr** for all object files in a program; FDPR will perform optimizations only on the files compiled with **-qfdpr**, and not library code, even if it is statically linked.

The optimizations that the FDPR utility performs are similar to those that the **-qpdf** option performs.

The FDPR performance-tuning utility has its own set of restrictions, and it is not guaranteed to speed up all programs or produce executables that produce exactly the same results as the original programs.

## Examples

To compile `myprogram.f` so it includes data required by the FDPR utility, enter:

```
xlf myprogram.f -qfdpr
```

## Related information

- “-qpdf1, -qpdf2” on page 195

---

## -qfixed

### Category

Input control

### Purpose

Indicates that the input source program is in fixed source form and optionally specifies the maximum line length.

### Syntax

```
►► -qfixed [==right_margin] ◄◄
```

@PROCESS:

```
@PROCESS FIXED[(right_margin)]
```

### Defaults

**-qfixed=72** is the default for the **xlf** and **xlf\_r** invocation commands when they are used to compile the `.f`, `.F`, `.f77`, or `.F77` files.

**-qfixed=72** is also the default for the **f77** and **fort77** invocation commands.

**Note:**

**-qfree=f90** is the default for the **f90**, **xlf90**, **xlf90\_r**, **f95**, **xlf95**, **xlf95\_r**, **f2003**, **xlf2003**, **xlf2003\_r**, **f2008**, **xlf2008**, and **xlf2008\_r** commands.

## Usage

The source form specified when executing the compiler applies to all of the input files, although you can switch the form for a compilation unit by using a **FREE** or **FIXED @PROCESS** directive, or switch the form for the rest of the file by using a **SOURCEFORM** comment directive (even inside a compilation unit).

For source code from some other systems, you may find you need to specify a right margin larger than the default. This option allows a maximum right margin of 132.

## Related information

- “-qfree” on page 134
- See *Fixed source form* in the *XL Fortran Language Reference*.

---

## -qflag


### Category

Listings, messages, and compiler information

### Purpose

Limits the diagnostic messages to those of a specified severity level or higher.

### Syntax



### Notes:

- 1 Minimum severity level of messages reported in listing
- 2 Minimum severity level of messages reported on terminal

### @PROCESS:

@PROCESS FLAG(*listing\_severity*,*terminal\_severity*)

### Defaults

`-qflag=i:i`, which shows all compiler messages.

### Parameters

The severity levels (from lowest to highest) are:

- i** Informational messages. They explain things that you should know, but they usually do not require any action on your part.
- l** Language-level messages, such as those produced under the `-qlanglvl` option. They indicate possible nonportable language constructs.

- w** Warning messages. They indicate error conditions that might require action on your part, but the program is still correct.
- e** Error messages. They indicate error conditions that require action on your part to make the program correct, but the resulting program can probably still be executed.
- s** Severe error messages. They indicate error conditions that require action on your part to make the program correct, and the resulting program will fail if it reaches the location of the error. You must change the **-qhalt** setting to make the compiler produce an object file when it encounters this kind of error.
- u** Unrecoverable error messages. They indicate error conditions that prevent the compiler from continuing. They require action on your part before you can compile your program.
- q** No messages. A severity level that can never be generated by any defined error condition. Specifying it prevents the compiler from displaying messages, even if it encounters unrecoverable errors.

## Usage

You must specify both *listing\_severity* and *terminal\_severity*.

Only messages with severity *listing\_severity* or higher are written to the listing file. Only messages with severity *terminal\_severity* or higher are written to the terminal.

The **-qflag** option overrides any **-qlanglvl** or **-qsaa** options specified.

The **-w** option is a short form for **-qflag=e:e**.

The **-qhaltmsg** option has precedence over the **-qflag** option. If you specify both **-qhaltmsg** and **-qflag**, messages that **-qflag** does not select are also printed and compilation stops.

**Note:** If **-qflag=u:u** or **-qflag=q:q** is specified, the message specified by **-qhaltmsg** is not shown.

## Related information

- “-qhalt” on page 140
- “-qlanglvl” on page 169
- “-qmaxerr” on page 182
- “-qhaltmsg” on page 142
- “-qsaa” on page 218
- “-qsuppress” on page 245
- “-w” on page 278
- “Understanding XL Fortran error messages” on page 283

---

## **-qfloat**


### Category

Floating-point and integer control

## Purpose

Selects different strategies for speeding up or improving the accuracy of floating-point calculations.

## Syntax



**@PROCESS:**

@PROCESS FLOAT(*suboptions*)

## Defaults

- **-qfloat=nofenv:fold:nogcclongdouble:nohscmplx:nohsflt:maf:nonans:rngchk:norrm:norsqrt:nostrictnmaf**
- **-qfloat=rsqrt:norngchk** when **-qnostrict**, **-qstrict=nooperationprecision:noexceptions**, or the **-O3** or higher optimization level is in effect.

## Parameters

**fenv** | **nofenv**

Specifies whether the code depends on the hardware environment and whether to suppress optimizations that could cause unexpected results due to this dependency.

Certain floating-point operations rely on the status of Floating-Point Status and Control Register (FPSCR), for example, to control the rounding mode or to detect underflow. In particular, many compiler built-in functions read values directly from the FPSCR.

When **nofenv** is in effect, the compiler assumes that the program does not depend on the hardware environment, and that aggressive compiler optimizations that change the sequence of floating-point operations are allowed. When **fenv** is in effect, such optimizations are suppressed.

You should use **fenv** for any code containing statements that read or set the hardware floating-point environment, to guard against optimizations that could cause unexpected behavior.

**fold** | **nofold**

Evaluates constant floating-point expressions at compile time, which may yield slightly different results from evaluating them at run time. The compiler always evaluates constant expressions in specification statements, even if you specify **nofold**.

**gcclongdouble** | **nogcclongdouble**

Specifies whether the compiler uses GCC-supplied or IBM-supplied library functions for 128-bit REAL(16) and COMPLEX(32) operations.

**gcclongdouble** ensures binary compatibility with GCC for mathematical calculations. If this compatibility is not important in your application, you should use **nogcclongdouble** for better performance.

**Note:** Passing results from modules compiled with **nogcclongdouble** to modules compiled with **gcclongdouble** may produce different results for numbers such as Inf, NaN and other rare cases. To avoid such incompatibilities, the compiler provides built-in functions to convert IBM long double types to GCC long double types.

**hscmplx** | **nohscmplx**

Speeds up operations involving complex division and complex absolute value. This suboption, which provides a subset of the optimizations of the **hsflt** suboption, is preferred for complex calculations.

**hsflt** | **nohsflt**

Speeds up calculations by preventing rounding for single-precision expressions and by replacing floating-point division by multiplication with the reciprocal of the divisor. **hsflt** implies **hscmplx**.

The **hsflt** suboption overrides the **nans** and **spnans** suboptions.

**Note:** Use **-qfloat=hsflt** on applications that perform complex division and floating-point conversions where floating-point calculations have known characteristics. In particular, all floating-point results must be within the defined range of representation of single precision. Use with discretion, as this option may produce unexpected results without warning. For complex computations, it is recommended that you use the **hscmplx** suboption (described above), which provides equivalent speed-up without the undesirable results of **hsflt**.

**maf** | **nomaf**

Makes floating-point calculations faster and more accurate by using floating-point multiply-add instructions where appropriate. The results may not be exactly equivalent to those from similar calculations performed at compile time or on other types of computers. Negative zero results may be produced. Rounding towards negative infinity or positive infinity will be reversed for these operations. This suboption may affect the precision of floating-point intermediate results. If **-qfloat=nomaf** is specified, no multiply-add instructions will be generated unless they are required for correctness.

**nans** | **nonans**

Allows you to use the **-qflttrap=invalid:enable** option to detect and deal with exception conditions that involve signaling NaN (not-a-number) values.

Use this suboption only if your program explicitly creates signaling NaN values, because these values never result from other floating-point operations.

### **rngchk | norngchk**

At optimization level **-O3** and above, and without **-qstrict**, controls whether range checking is performed for input arguments for software divide and inlined square root operations. Specifying **norngchk** instructs the compiler to skip range checking, allowing for increased performance where division and square root operations are performed repeatedly within a loop.

Note that with **norngchk** in effect the following restrictions apply:

- The dividend of a division operation must not be +/-INF.
- The divisor of a division operation must not be 0.0, +/- INF, or denormalized values.
- The quotient of dividend and divisor must not be +/-INF.
- The input for a square root operation must not be INF.

If any of these conditions are not met, incorrect results may be produced. For example, if the divisor for a division operation is 0.0 or a denormalized number (absolute value  $< 2^{-1022}$  for double precision, and absolute value  $< 2^{-126}$  for single precision), NaN, instead of INF, may result; when the divisor is +/- INF, NaN instead of 0.0 may result. If the input is +INF for a sqrt operation, NaN, rather than INF, may result.

**norngchk** is only allowed when **-qnostrict** is in effect. If **-qstrict**, **-qstrict=infinities**, **-qstrict=operationprecision**, or **-qstrict=exceptions** is in effect, **norngchk** is ignored.

### **rrm | norrm**

Prevents floating-point optimizations that require the rounding mode to be the default, round-to-nearest, at run time, by informing the compiler that the floating-point rounding mode may change or is not round-to-nearest at run time. You should use **rrm** if your program changes the runtime rounding mode by any means; otherwise, the program may compute incorrect results.

### **rsqrt | norsqrt**

Speeds up some calculations by replacing division by the result of a square root with multiplication by the reciprocal of the square root.

If you compile with the **-O3** or higher optimization level, **rsqrt** is enabled automatically. To disable it, also specify **-qstrict**, **-qstrict=nans**, **-qstrict=infinities**, **-qstrict=zerosigns**, or **-qstrict=exceptions**.

### **strictnmaf | nostrictnmaf**

Turns off floating-point transformations that are used to introduce negative MAF instructions, as these transformations do not preserve the sign of a zero value. By default, the compiler enables these types of transformations.

To ensure strict semantics, specify both **-qstrict** and **-qfloat=strictnmaf**.

**Note:** For details about the relationship between **-qfloat** suboptions and their **-qstrict** counterparts, see “-qstrict” on page 238.

## **Usage**

Using **-qfloat** suboptions other than the default settings might produce incorrect results in floating-point computations if the system does not meet all required conditions for a given suboption. Therefore, use this option only if the

floating-point calculations involving IEEE floating-point values are manipulated and can properly assess the possibility of introducing errors in the program.

If the **-qstrict** | **-qnostrict** and **float** suboptions conflict, the last setting specified is used.

## Examples

To compile `myprogram.f` so that the constant floating-point expressions are evaluated at compile time and multiply-add instructions are not generated, enter:

```
xlf myprogram.f -qfloat=fold:nomaf
```

## Related information

- “-qarch” on page 93
- “-qflttrap” on page 132
- “-qstrict” on page 238
- “Implementation details of XL Fortran floating-point processing” in the *XL Fortran Optimization and Programming Guide*.

---

## -qfpp

### Category

Input control


### @PROCESS

None.

### Purpose

Controls Fortran-specific preprocessing in the C preprocessor.

### Syntax



### Defaults

- -qnofpp

### Parameters

#### comment | nocomment

Instructs the C preprocessor (**cpp**) to recognize the `!` character as a comment delimiter in macro expansion. When this suboption is enabled, the `!` character and all characters following it on that line will be ignored by **cpp** when performing macro expansion.

## linecont | nolinecont

Instructs **cpp** to recognize the & character as a line continuation character. When this suboption is enabled, **cpp** treats the & character and the C-style \ line continuation character equivalently.

Specifying **-qfpp** without any suboptions is equivalent to **-qfpp=comment:linecont**.

## Usage

**-qfpp** is a C preprocessor option, and must therefore be specified using the **-WF** option.

## Related information

- “-W” on page 276
- “-qppsborigarg” on page 209
- “Passing Fortran files through the C preprocessor” on page 29

---

## -qflttrap

### Category


Error checking and debugging

### Purpose

Determines what types of floating-point exceptions to detect at run time.

The program receives a **SIGFPE** signal when the corresponding exception occurs.

### Syntax



@PROCESS:

FLTTRAP[(*suboptions*)] | NOFLTTRAP

### Defaults

**-qnoflttrap** or

Specifying **-qflttrap** option with no suboptions is equivalent to **-qflttrap=invalid:inexact:overflow:undflow:zerodivide**

### Parameters

**ENable**

Turn on checking for the specified exceptions in the main program so that

the exceptions generate **SIGFPE** signals. You must specify this suboption if you want to turn on exception trapping without modifying your source code.

#### **INEXact**

Detect and trap on floating-point inexact if exception-checking is enabled. Because inexact results are very common in floating-point calculations, you usually should not need to turn this type of exception on.

#### **INValid**

Detect and trap on floating-point invalid operations if exception-checking is enabled.

#### **NANQ**

Detect and trap on all quiet not-a-number values (NaNQs) and signaling not-a-number values (NaNSs). Trapping code is generated regardless of specifying the **enable** suboption. This suboption detects all NaN values handled by or generated by floating-point instructions, including those not created by invalid operations. This option can impact performance.

#### **OVerflow**

Detect and trap on floating-point overflow if exception-checking is enabled.

#### **UNDerflow**

Detect and trap on floating-point underflow if exception-checking is enabled.

#### **ZERODivide**

Detect and trap on floating-point division by zero if exception-checking is enabled.

### **Usage**

Exceptions will be detected by the hardware, but trapping is not enabled. Because this default does not include **enable**, it is probably only useful if you already use **fpsets** or similar subroutines in your source.

If you specify **-qfltrap** more than once, both with and without suboptions, the **-qfltrap** without suboptions is ignored.

The **-qfltrap** option is recognized during linking with IPA. Specifying the option at the link step overrides the compile-time setting.

**Note:** Due to the transformations performed and the exception handling support of some vector instructions, use of **-qsimd=auto** may change the location where an exception is caught or even cause the compiler to miss catching an exception.

For full instructions on how and when to use the **-qfltrap** option, especially if you are just starting to use it, see the *Detecting and trapping floating-point exceptions* topic in *XL Fortran Optimization and Programming Guide*.

### **Example**

```
REAL :: x, y, z
DATA x /5.0/, y /0.0/
z = x / y
PRINT *, z
END
```

When you compile this program with the following command, the program stops when the division is performed.

```
xlf -qfltrap=zerodivide:enable -qsigtrap divide_by_zero.f
```

The **zerodivide** suboption identifies the type of exception to guard against. The **enable** suboption causes a **SIGFPE** signal to be generated when the exception occurs. The **-qsigtrap** option produces informative output when the signal stops the program.

### Related information

- *Detecting and trapping floating-point exceptions*
- “-qfloat” on page 127
- “-qarch” on page 93
- “-qsigtrap” on page 224

---

## -qfree

### Category

Input control

### Purpose

Indicates that the source code is in free source form.

### Syntax

→ -q—free →

```
└─ f90 ─┘  
└─ ibm ─┘
```

@PROCESS:

```
@PROCESS FREE[({F90|IBM})]
```

### Defaults

**-qfree** by itself specifies Fortran 90 free source form.

**-qfixed=72** is the default for the **xlf** and **xlf\_r** invocation commands when they are used to compile the **.f**, **.F**, **.f77**, or **.F77** files. **-qfixed=72** is also the default for the **f77** and **fort77** invocation commands.

**-qfree=f90** is the default for the **f90**, **xlf90**, **xlf90\_r**, **f95**, **xlf95**, **xlf95\_r**, **f2003**, **xlf2003**, **xlf2003\_r**, **f2008**, **xlf2008**, and **xlf2008\_r** commands.

### Parameters

**ibm**

Specifies compatibility with the free source form defined for VS FORTRAN.

**f90**

Specifies compatibility with the free source form defined for Fortran 90.

Note that the free source form defined for Fortran 90 also applies to Fortran 95, Fortran 2003, and Fortran 2008.

## Usage

The source form specified when executing the compiler applies to all of the input files, although you can switch the form for a compilation unit by using a **FREE** or **FIXED @PROCESS** directive or for the rest of the file by using a **SOURCEFORM** comment directive (even inside a compilation unit).

Fortran 90 free source form is the format to use for maximum portability across compilers that support Fortran 90 and Fortran 95 features now and in the future.

IBM free source form is equivalent to the free format of the IBM VS FORTRAN compiler, and it is intended to help port programs from the z/OS® platform.

**-k** is equivalent to **-qfree=f90**.

## Related information

- “-qfixed” on page 125
- “-k” on page 77
- Free source form in the *XL Fortran Language Reference*

---

## -qfullpath

### Category

Error checking and debugging

### @PROCESS

None.

### Purpose

When used with the **-g** or **-q1inedebug** option, this option records the full, or absolute, path names of source and include files in object files compiled with debugging information, so that debugging tools can correctly locate the source files.

### Syntax

►► -q nofullpath  
fullpath ►►

### Defaults

By default, the compiler records the relative path names of the original source file in each **.o** file. It may also record relative path names for include files.

### Usage

If you need to move an executable file into a different directory before debugging it or have multiple versions of the source files and want to ensure that the debugger uses the original source files, use the **-qfullpath** option in combination with the **-g** or **-q1inedebug** option so that source-level debuggers can locate the correct source files.

Although **-qfullpath** works without the **-g** or **-qlinedebug** option, you cannot do source-level debugging unless you also specify the **-g** or **-qlinedebug** option.

## Examples

In this example, the executable file is moved after being created, but the debugger can still locate the original source files:

```
$ xlf95 -g -qfullpath file1.f file2.f file3.f -o debug_version
...
$ mv debug_version $HOME/test_bucket
$ cd $HOME/test_bucket
$ gdb debug_version
```

## Related information

- “-g” on page 73
- “-qlinedebug” on page 173

---

## -qfunctrace

### Category

Error checking and debugging


### @PROCESS

None.

### Purpose

Traces entry and exit points of procedures in your program. If your program contains C++ compilation units, this option also traces C++ catch blocks.

### Syntax



### Defaults

-qnofunctrace

### Parameters

- + Instructs the compiler to trace the specified program, procedures, or module procedures. All their internal procedures are traced by default.
- Instructs the compiler not to trace the specified program, procedures, module procedures, or any of their internal procedures.


#### **procedure\_name**

The name of a program, external procedure, or module procedure. The name is case sensitive when **-qmixed** is in effect. **BIND(C)** binding labels and mangled

module procedure names are allowed, but they must have the correct case. If **-qextname** is in effect, *procedure\_name* is the name of the procedure without the additional underscore.


**module\_name**

The name of a module. The name is case sensitive when **-qmixed** is in effect.





**submodule\_name**

The name of a submodule. The name is case sensitive when **-qmixed** is in effect.



## Usage

**-qfunctrace** enables tracing for all procedures in your program. **-qnofunctrace** disables tracing that was enabled by **-qfunctrace**.

The **-qfunctrace+** and **-qfunctrace-** suboptions enable tracing for a specific list of procedures and are not affected by **-qnofunctrace**. The list of procedures is cumulative. When a module procedure and its containing module  or submodule  are both specified, the procedure specification takes precedence.

This option inserts calls to the tracing procedures that you have defined. These procedures must be provided at the link step. For details about the interface of tracing procedures, as well as when they are called, see the Trace procedures in your code section in the *XL Fortran Optimization and Programming Guide*.

## Examples

The following table provides some examples of using the **-qfunctrace** option to achieve different purposes.

Purpose	Usage example
Tracing all procedures	<b>-qfunctrace</b>
Tracing procedures <i>x</i> , <i>y</i> , and <i>z</i>	<b>-qfunctrace+x:y:z</b>
Tracing all procedures except <i>x</i>	<b>-qfunctrace -qfunctrace-x</b> or <b>-qfunctrace-x -qfunctrace</b>
Tracing only procedures <i>x</i> and <i>y</i>	<b>-qfunctrace+x -qfunctrace+y</b> or <b>-qfunctrace+x -qnofunctrace</b> <b>-qfunctrace+y</b>
Tracing only procedure <i>y</i>	<b>-qfunctrace+y -qnofunctrace</b> or <b>-qfunctrace+y</b>
Tracing all the module procedures, except procedure <i>x</i> , in module <i>y</i>	<b>-qfunctrace-x -qfunctrace+y</b> or <b>-qfunctrace+y -qfunctrace-x</b>

## Related information

- “**-qfunctrace\_xlf\_catch**” on page 138
- “**-qfunctrace\_xlf\_enter**” on page 139
- “**-qfunctrace\_xlf\_exit**” on page 140

- For details about the directives that you can use to specify the name of the tracing procedures, see the **FUNCTRACE\_XLF\_CATCH**, **FUNCTRACE\_XLF\_ENTER**, **FUNCTRACE\_XLF\_EXIT** sections in the *XL Fortran Language Reference*.
- For details about the rules for using the **NOFUNCTRACE** directive, see **NOFUNCTRACE** in the *XL Fortran Language Reference*.
- For detailed information about how to implement procedure tracing routines in your code, as well as detailed examples and a list of rules for using them, see **Tracing procedures in your code** in the *XL Fortran Optimization and Programming Guide*.

---

## **-qfunctrace\_xlf\_catch**

### **Category**

Error checking and debugging

### **@PROCESS**

None.

### **Purpose**

Specifies the name of the catch tracing subroutine.

### **Syntax**

▶▶ `-qfunctrace_xlf_catch=catch_routine` ◀◀

### **Defaults**

Not applicable.

### **Parameters**

*catch\_routine*

Indicates the name of the catch tracing subroutine.

### **Usage**

You use the **-qfunctrace\_xlf\_catch** option to specify that the external or module procedure being compiled must be used as a catch tracing procedure.

#### **Note:**

- If you write a tracing subroutine, make sure that the program does not contain any user procedures called `__func_trace_catch`.
- You must not specify the name of an internal subroutine when you use the **-qfunctrace\_xlf\_catch** option.

### **Related information**

- The **FUNCTRACE\_XLF\_CATCH** directive in the *XL Fortran Language Reference*.
- “-qfunctrace” on page 136
- “-qfunctrace\_xlf\_enter” on page 139
- “-qfunctrace\_xlf\_exit” on page 140

- For detailed information about how to implement the tracing procedures in your code, see **Tracing procedures in your code** in the *XL Fortran Optimization and Programming Guide*.

---

## **-qfunctrace\_xlf\_enter**

### **Category**

Error checking and debugging

### **@PROCESS**

None.

### **Purpose**

Specifies the name of the entry tracing subroutine.

### **Syntax**

► `--qfunctrace_xlf_enter=enter_routine` ◀

### **Defaults**

Not applicable.

### **Parameters**

*enter\_routine*

Indicates the name of the entry tracing subroutine.

### **Usage**

You use the **-qfunctrace\_xlf\_enter** option to specify that the external or module procedure being compiled must be used as an entry tracing procedure.

#### **Note:**

- If you write a tracing subroutine, make sure that the program does not contain any user procedures called `__func_trace_enter`.
- You must not specify the name of an internal subroutine when you use the **-qfunctrace\_xlf\_enter** option.

### **Related information**

- The **FUNCTRACE\_XLF\_ENTER** directive in the *XL Fortran Language Reference*.
- “-qfunctrace” on page 136
- “-qfunctrace\_xlf\_catch” on page 138
- “-qfunctrace\_xlf\_exit” on page 140
- For detailed information about how to implement the tracing procedures in your code, see **Tracing procedures in your code** in the *XL Fortran Optimization and Programming Guide*.

---

## **-qfunctrace\_xlf\_exit**

### **Category**

Error checking and debugging

### **@PROCESS**

None.

### **Purpose**

Specifies the name of the exit tracing subroutine.

### **Syntax**

▶▶—qfunctrace\_xlf\_exit—=*exit\_routine*—▶▶

### **Defaults**

Not applicable.

### **Parameters**

*exit\_routine*

Indicates the name of the exit tracing subroutine.

### **Usage**

You use the **-qfunctrace\_xlf\_exit** option to specify that the external or module procedure being compiled must be used as an exit tracing procedure.

#### **Note:**

- If you write a tracing subroutine, make sure that the program does not contain any user procedures called `__func_trace_exit`.
- You must not specify the name of an internal subroutine when you use the **-qfunctrace\_xlf\_exit** option.

### **Related information**

- The **FUNCTTRACE\_XLF\_EXIT** directive in the *XL Fortran Language Reference*.
- “-qfunctrace” on page 136
- “-qfunctrace\_xlf\_catch” on page 138
- “-qfunctrace\_xlf\_exit”
- For detailed information about how to implement the tracing procedures in your code, see **Tracing procedures in your code** in the *XL Fortran Optimization and Programming Guide*.

---

## **-qhalt**

### **Category**

Error checking and debugging

## Purpose

Stops compilation before producing any object, executable, or assembler source files if the maximum severity of compile-time messages equals or exceeds the severity you specify.

## Syntax



## Notes:

- 1 Minimum severity level of messages that will prevent an object file from being created

## @PROCESS:

@PROCESS HALT(*severity*)

## Defaults

`-qhalt=s`, which prevents the compiler from generating an object file when compilation fails.

## Parameters

The severity levels (from lowest to highest) are:

- i** Informational messages. They explain things that you should know, but they usually do not require any action on your part.
- l** Language-level messages, such as those produced under the `-qlanglvl` option. They indicate possible nonportable language constructs.
- w** Warning messages. They indicate error conditions that might require action on your part, but the program is still correct.
- e** Error messages. They indicate error conditions that require action on your part to make the program correct, but the resulting program can probably still be executed.
- s** Severe error messages. They indicate error conditions that require action on your part to make the program correct, and the resulting program will fail if it reaches the location of the error. You must change the `-qhalt` setting to make the compiler produce an object file when it encounters this kind of error.
- u** Unrecoverable error messages. They indicate error conditions that prevent the compiler from continuing. They require action on your part before you can compile your program.

## Usage

The `-qhalt` option can override the `-qobject` option, and `-qnoobject` can override `-qhalt`.

## Related information

- “-qflag” on page 126
- “-qhaltonmsg”
- “-qmaxerr” on page 182
- “-qobject” on page 190

---

## -qhaltonmsg


### Category

Error checking and debugging

### Purpose

Stops compilation before producing any object files, executable files, or assembler source files if a specified error message is generated.

### Syntax



### @PROCESS:

@PROCESS HALTONMSG(*message\_identifier* [, *message\_identifier* [, ...]]) | NOHALTONMSG

### Defaults

-qnohaltonmsg

### Parameters

*message\_identifier*[:*message\_identifier* ...]

Stops compilation before producing any object files, executable files, or assembler source files if a specified error message (*nnnn-mmm*) or a list of messages (*nnnn-mmm[:nnnn-mmm ...]*) is generated. To specify a list of messages, separate each message number with a colon.

*nnnn-mmm* is the message number, where:

- *nnnn* must be a four-digit integer between 1500 and 1585; this is the range of XL Fortran message numbers.
- *mmm* must be any three-digit integer (with leading zeros if necessary).

### Usage

When the compiler stops as a result of the **-qhaltonmsg** option, the compiler return code is nonzero. The severity level of a message that is specified by **-qhaltonmsg** is changed to S if its original severity level is lower than S.

If **-qflag=u:u** or **-qflag=q:q** is specified, the message specified by **-qhaltonmsg** is not shown.

**-qhaltonmsg** has precedence over **-qsuppress** and **-qflag**.

### Related information

- “-qhalt” on page 140
  - “-qflag” on page 126
  - “-qsuppress” on page 245
- 

## -qhelp

### Category

Listings, messages, and compiler information

### @PROCESS

None.

### Purpose

Displays the man page of the compiler.

### Syntax

▶▶ — -q—help —————▶▶

### Usage

If you specify the **-qhelp** option, regardless of whether you provide input files, the compiler man page is displayed and the compilation stops.

### Related information

- “-qversion” on page 255
- 

## -qhot

### Category


Optimization and tuning

### Purpose

Performs high-order loop analysis and transformations (HOT) during optimization.

The **-qhot** compiler option is a powerful alternative to hand tuning that provides opportunities to optimize loops and array language. This compiler option will always attempt to optimize loops, regardless of the suboptions you specify.

### Syntax



#### @PROCESS:

@PROCESS HOT[=suboptions] | **NOHOT**

### Defaults

- **-qnohot**
- **-qhot=noarraypad:level=0:novector:fastmath** when **-O3** is in effect.
- **-qhot=noarraypad:level=1:vector:fastmath** when **-qsmp**, **-O4** or **-O5** is in effect.
- Specifying **-qhot** without suboptions is equivalent to **-qhot=noarraypad:level=1:vector:fastmath**.

### Parameters

#### arraypad | noarraypad

Permits the compiler to increase the dimensions of arrays where doing so might improve the efficiency of array-processing loops. (Because of the implementation of the cache architecture, array dimensions that are powers of two can lead to decreased cache utilization.) Specifying **-qhot=arraypad** when your source includes large arrays with dimensions that are powers of 2 can reduce cache misses and page faults that slow your array processing programs. This can be particularly effective when the first dimension is a power of 2. If you use this suboption with no *number*, the compiler will pad any arrays where it infers there may be a benefit and will pad by whatever amount it chooses. Not all arrays will necessarily be padded, and different arrays may be padded by different amounts. If you specify a *number*, the compiler will pad every array in the code.

**Note:** Using **arraypad** can be unsafe, as it does not perform any checking for reshaping or equivalences that may cause the code to break if padding takes place.

#### *number*

A positive integer value representing the number of elements by which each array will be padded in the source. The pad amount must be a positive integer value. It is recommended that pad values be multiples of the largest array element size, typically 4, 8, or 16.

#### level=0

Performs a subset of the high-order transformations and sets the default to **novector:noarraypad:fastmath**.

#### level=1

Performs the default set of high-order transformations.

### **level=2**

Performs the default set of high-order transformations and some more aggressive loop transformations. **-qhot=level=2** must be used with **-qsmp**. This option performs aggressive loop analysis and transformations to improve cache reuse and exploit loop parallelization opportunities.

### **vector | novector**

When specified with **-qnostrict** and an optimization level of **-03** or higher, **vector** causes the compiler to convert certain operations that are performed in a loop on successive elements of an array (for example, square root, reciprocal square root) into a call to a routine in the Mathematical Acceleration Subsystem (MASS) library in libxlopt. The **vector** suboption supports single and double-precision floating-point mathematics, and is useful for applications with significant mathematical processing demands.

**novector** disables the conversion of loop array operations into calls to MASS library routines.

Since vectorization can affect the precision of your program's results, if you are using **-03** or higher, you should specify **-qhot=novector** if the change in precision is unacceptable to you.

### **fastmath | nofastmath**

You can use this suboption to tune your application to either use fast scalar versions of math functions or use the default versions.

**-qhot=fastmath** enables the replacement of math routines with available math routines from the XLOPT library only if **-qstrict=nolibrary** is enabled.

**-qhot=nofastmath** disables the replacement of math routines by the XLOPT library. **-qhot=fastmath** is enabled by default if **-qhot** is specified regardless of the hot level.

## **Usage**

If you do not also specify an optimization level when specifying **-qhot** on the command line, the compiler assumes **-02**.

If you want to override the default **level** setting of **1** when using **-qsmp**, **-04** or **-05**, be sure to specify **-qhot=level=0** or **-qhot=level=2** *after* the other options.

If **-02**, **-qnohot**, or **-qnoopt** is used on the command line, specifying HOT options in an **@PROCESS** directive will have no effect on the compilation unit.

The **-C** option turns off some array optimizations.

You can use the **-qreport** option in conjunction with **-qhot** or any optimization option that implies **-qhot** to produce a pseudo-Fortran report showing how the loops were transformed. The loop transformations are included in the listing report if either the **-qreport** or **-qlistfmt** option is also specified. This LOOP TRANSFORMATION SECTION of the listing file also contains information about data prefetch insertion locations. In addition, when you use **-qprefetch=assistthread** to generate prefetching assist threads, a message Assist thread for data prefetching was generated also appears in the LOOP TRANSFORMATION SECTION of the listing file. Specifying **-qprefetch=assistthread** guides the compiler to generate aggressive data prefetching at optimization level **-03 -qhot** or higher. For more information, see “-qreport” on page 216.



Specifying the same rounding mode for compile-time and runtime operations avoids inconsistencies in floating-point results.

**Note:** Compile-time arithmetic is most extensive when you also specify the **-0** option.

If you change the rounding mode from the default (round-to-nearest) at run time, be sure to also specify **-qfloat=rrm** to turn off optimizations that only apply in the default rounding mode.

If your program contains operations involving real(16) values, the rounding mode must be set to **-qieee=near**, round-to-nearest.

### Related information

- *Selecting the rounding mode* in the *XL Fortran Optimization and Programming Guide*
- “-O” on page 83
- “-qfloat” on page 127

---

## -qinfo

### Category

Error checking and debugging

### @PROCESS

None.


### Purpose

Produces or suppresses groups of informational messages.

The messages are written to standard output and, optionally, to the listing file if one is generated.

### Syntax

#### Option syntax



#### @PROCESS:

@PROCESS INFO[(suboptions)] | **NOINFO**

## Defaults

-qnoinfo

## Parameters

### all

Enables diagnostic messages for all groups except **mt**.

### noall (option only)

Disables all diagnostic messages for all groups.

### mt | nomt

Reports potential synchronization issues in parallel code. This suboption detects the Global Thread Flag pattern where one thread uses a shared volatile flag variable to notify other threads that it has completed a computation and stored its result to memory. Other threads check the flag variable in a loop that does not change the flag variable. When the value of the flag variable changes, the waiting threads access the computation result from memory. The PowerPC storage model requires synchronization before the flag variable is set in the first thread, and after the flag variable is checked in the waiting threads. Synchronization can be done by the **LIGHT\_SYNC** or **ISYNC** directive.

The type of synchronization directives you need to use depends on your code. Usually, it is enough to use the **LIGHT\_SYNC** directive, as it preserves the storage access order to system memory. However, if the loops in the waiting threads are written in such a way that might cause instruction prefetching to start executing code that accesses the computation result before the flag variable is updated, you need to use the **ISYNC** directive to preserve order. Such patterns are typically as follows:

```
10 CALL sleep_(value)
   IF (.NOT. flag) GOTO 10
   ! The SYNC directive is needed here.
   x = shared_computation_result
```

Some patterns that do not require synchronization are similar to the patterns described above. The messages generated by this suboption are only suggestions about potential synchronization issues.

To use the **-qinfo=mt** suboption, you must enable the **-qthreaded** option and specify at least one of the following options:

- **-03**
- **-04**
- **-05**
- **-qipa**
- **-qhot**
- **-qsmp**

The default option is **-qinfo=nomt**.

### **HOSTASSOCIation** | **NOHOSTASSOCIation**

Issues an information message for an entity that is accessed by host association for the first time. However, if the entity is accessed by an **IMPORT** statement, no information message is issued.

You can use **-qinfo=all / noall** to control whether to enable the **-qinfo=HOSTASSOCIation** option. The default option is **-qinfo=NOHOSTASSOCIation**.

### **stp | nostp**

Issues warnings for procedures that are not protected against stack corruption. **-qinfo=stp** has no effects unless the **-qstackprotect** option is also enabled. Like other **-qinfo** options, **-qinfo=stp** is enabled or disabled through **-qinfo=all / noall**. **-qinfo=nostp** is the default option.

### **unset | nounset**

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

**-qinfo=unset** uses program information, such as control flow information, available during optimization. As a result, detection accuracy improves with the optimization level. For example, some uses of unset variables that are not flagged at **-O0** are flagged at **-O2**. Aggressive optimization levels, such as **-O4** and **-O5**, can cause the line numbers in the informational messages to become inaccurate. In very rare cases, these optimization levels reorder the program sufficiently to cause false positive messages from static analysis. You can get the best results with the **-O2** optimization level.

The **-qinitauto** option initializes automatic variables. As a result, the **-qinitauto** option hides variables that are used before they are set from the **-qinfo=unset** option.

The **-qsave** option changes the storage class of automatic variables to **STATIC**. As a result, the **-qsave** option hides variables that are used before they are set from the **-qinfo=unset** option.

## **Usage**

Specifying **-qinfo** with no suboptions is equivalent to **-qinfo=all**.

Specifying **-qnoinfo** is equivalent to **-qinfo=noall**.

## **Examples**

To compile a program to produce informational messages about stack protection, enter the following command:

```
xlf90 myprogram.f -qinfo=stp -qstackprotect
```

To compile `t.f` to produce informational messages about host-associated variables, enter the following command:

```
xlf2008 t.f -qinfo=HOSTASSOCIATION
```

Suppose that `t.f` contains the following code:

```
PROGRAM p
  IMPLICIT none

  INTEGER :: var

  var = 3
  CALL sub()

  CONTAINS
    SUBROUTINE sub()
      PRINT *, var ! The compiler issues an information message when entity 'var' is
                  ! accessed by host association for the first time.

      PRINT *, var ! No message is issued here.
    END

  INTEGER FUNCTION func()
```

```

func = var ! Entity 'var' is in a different scope. The compiler issues an
           ! information message when the entity is accessed by host association
           ! for the first time in each scope.

```

```

END
END

```

The compiler issues the following information messages:

```

"t.f", line 11.20: 1521-004 (I) Entity var is accessed by host association.
"t.f", line 16.18: 1521-004 (I) Entity var is accessed by host association.

```

To compile `sync.f` to produce informational messages about potential synchronization issues in parallel code, enter the following command:

```
xlf95_r -qinfo=mt -03 sync.f
```

Suppose that `sync.f` contains the following code:

```

MODULE m
  IMPLICIT NONE
  LOGICAL, VOLATILE :: done ! shared flag
  INTEGER, VOLATILE :: shared_result
CONTAINS
  SUBROUTINE setter(id)
    IMPLICIT NONE
    INTEGER, INTENT(IN) :: id

    CALL sleep_(5)
    shared_result = 7
    ! !ibm* light_sync
    done = .TRUE. ! line 13
  END SUBROUTINE

  SUBROUTINE waiter(id)
    IMPLICIT NONE
    INTEGER, INTENT(IN) :: id

    DO WHILE (.NOT. done)
      CALL sleep_(1)
    END DO ! line 22
    ! !ibm* light_sync
    PRINT *, shared_result
  END SUBROUTINE
END MODULE

PROGRAM MAIN
  USE m, ONLY: waiter, setter
  USE f_thread
  IMPLICIT NONE

  TYPE(f_thread_t) threads(2)
  TYPE(f_thread_attr_t) attr
  INTEGER(4) flag, result

  ! Initialization
  result = f_thread_attr_init(attr)
  IF (result /= 0) ERROR STOP 1
  flag = FLAG_DEFAULT

  ! Create threads
  result = f_thread_create(threads(1), attr, flag, waiter, 1)
  IF (result /= 0) ERROR STOP 2

  result = f_thread_create(threads(2), attr, flag, setter, 2)
  IF (result /= 0) ERROR STOP 3

```

```

    result = f_pthread_join(threads(1))
    result = f_pthread_join(threads(2))
END PROGRAM

```

The compiler issues the following informational message:

```

** m === End of Compilation 1 ===
** main === End of Compilation 2 ===

```

1586-669 (I) "sync.f", line 22: If this loop is used as a synchronization point, additional synchronization via a directive or built-in function might be needed.

1586-670 (I) "sync.f", line 13: If this statement is used as a synchronization point, additional synchronization via a directive or built-in function might be needed.

1501-510 Compilation successful for file sync.f.

The following function `factorial.f` does not initialize `temp` when `n<=1`. `factorial.f` also accesses `result` before it is set when `n>1`. With **-qinfo=unset** at **-qnoopt**, this issue is not detected. To compile `factorial.f` to produce informational messages about the uninitialized variables, enter the following command:

```
xlf95 -qinfo=unset -O factorial.f
```

`factorial.f` contains the following code:

```

module m
contains
  recursive function factorial(n) result(result)
    integer, value :: n
    integer result, temp

    if (n > 1) then
      temp = n * factorial(n - 1)
      print *, result ! line 9
    endif

    result = temp ! line 12
  end function
end module

use m
integer x
x = factorial(1)
end

```

The compiler issues the following informational messages:

```

1500-098 (I) "factorial.f", line 9: "result" is used before it is set.
1500-099 (I) "factorial.f", line 12: "temp" might be used before it is set.
1501-510  Compilation successful for file factorial.f.

```

## Related information

- “-qreport” on page 216
- `isync` in the *XL Fortran Language Reference*
- `light_sync` in the *XL Fortran Language Reference*
- For more information about synchronization and the PowerPC storage model, see the article at <http://www.ibm.com/developerworks/systems/articles/powerpc.html>.

---

## -qinit

### Category

Language element control

### Purpose

Makes the initial association status of pointers disassociated.

Note that this option applies to Fortran 90 and above.

### Syntax

►► -q—init—==f90ptr—————►►

@PROCESS:

@PROCESS INIT(F90PTR)

### Defaults

Not applicable.

### Usage

You can use this option to help locate and fix problems that are due to using a pointer before you define it.

### Related information

- See *Pointer association* in the *XL Fortran Language Reference*.

---

## -qinitialloc

### Category

Error checking and debugging

### Purpose

Initializes allocatable and pointer variables that are allocated but not initialized to a specific value, for debugging purposes.

The **-qinitialloc** option applies to the following uninitialized variables:

- Variables that have the **ALLOCATABLE** attribute and are allocated using the **ALLOCATE** statement
- Variables that have the **POINTER** attribute and are allocated using the **ALLOCATE** statement

### Syntax

►► -q—

noinitialloc
initialloc

—

==hex_value
-------------

—————►►

**@PROCESS:**

@PROCESS INITALLOC[(*hex\_value*)] | **NOINITALLOC**

## Defaults

### **-qnoinitalloc**

By default, the compiler does not initialize allocated storage to any particular value.

## Parameters

### **hex\_value**

A one- to eight-digit hexadecimal number.

- If you do not specify a *hex\_value*, the compiler initializes the value of each byte of allocated storage to zero.
- To initialize each byte of storage to a specific value, specify one or two digits for the *hex\_value*. If you specify only one digit, the compiler pads the *hex\_value* on the left with a zero.
- To initialize each word of storage to a specific value, specify three to eight digits for the *hex\_value*. If you specify more than two but fewer than eight digits, the compiler pads the *hex\_value* on the left with zeros.
- In the case of word initialization, if allocatable variables are not a multiple of 4 bytes in length, the *hex\_value* is truncated on the left to fit. For example, if an allocatable variable is only 1 byte and you specify five digits for the *hex\_value*, the compiler truncates the three digits on the left and assigns the other two digits on the right to the variable.
- You can specify alphabetic digits as either uppercase or lowercase.
- For derived type variables with default initialization, initialization with *hex\_value* occurs before the default initialization. See Example 2.

## Usage

The **-qinitalloc** option provides the following benefits:

- Setting *hex\_value* to zero (the default value) ensures that all allocatable variables are cleared before being used.
- You can use this option to initialize variables of real or complex type to a signaling or quiet NaN, which helps locate uninitialized variables in your program.

The usage of this option is similar to that of the **-qinitauto** option. For more information, see “-qinitauto” on page 154.

**Restriction:** Objects that are equivalenced, structure components, and array elements are not initialized individually. Instead, the entire storage sequence is initialized collectively.

## Examples

### **Example 1:**

The following example shows how the **-qinitalloc** option works.

```

SUBROUTINE Sub()
  REAL(4), ALLOCATABLE :: a, b
  CHARACTER, ALLOCATABLE :: c
  REAL(8), ALLOCATABLE :: d

  ALLOCATE(a)           ! a is allocated but not initialized.
  ALLOCATE(b, SOURCE = 3.0) ! b is allocated and initialized to 3.0.
  ALLOCATE(c)
  ALLOCATE(d)
END SUBROUTINE

```

If you compile your program with **-qinitalloc=0cf**, for example, the compiler performs the following initialization:

- Pads 0cf with five zeros and initializes a to 000000CF
- Keeps the original initialization for b
- Truncates the first digit of 0cf and initializes c to CF
- Pads 0cf with five zeros, repeats the value, and initializes d to 000000CF000000CF

### Example 2:

The following example shows how the **-qinitalloc** option works when a derived type contains a component with default initialization.

```

TYPE dt
  INTEGER :: i = 1    ! i has default initialization
  INTEGER :: j
END TYPE
TYPE(dt), ALLOCATABLE :: dt1
ALLOCATE(dt1)

```

If you compile your program with **-qinitalloc**, the compiler keeps the default initialization for i, and initializes j to zero.

### Related information

- The **ALLOCATABLE** attribute in the *XL Fortran Language Reference*
- The **ALLOCATE** statement in the *XL Fortran Language Reference*
- The **POINTER** statement in the *XL Fortran Language Reference*

---

## -qinitauto

### Category

Error checking and debugging

### @PROCESS

None.

### Purpose

Initializes uninitialized automatic variables to a specific value, for debugging purposes.

## Syntax



## Defaults

-qnoinitauto

By default, the compiler does not initialize automatic storage to any particular value. However, it is possible that a region of storage contains all zeros.

## Parameters

### hex\_value

A 1 to 8 digit hexadecimal number.

- If you do not specify a *hex\_value* number, the compiler initializes the value of each byte of automatic storage to zero.
- To initialize each byte of storage to a specific value, specify 1 or 2 digits for the *hex\_value*. If you specify only 1 digit, the compiler pads the *hex\_value* on the left with a zero.
- To initialize each word of storage to a specific value, specify 3 to 8 digits for the *hex\_value*. If you specify more than 2 but fewer than 8 digits, the compiler pads the *hex\_value* on the left with zeros.
- In the case of word initialization, if automatic variables are not a multiple of 4 bytes in length, the *hex\_value* may be truncated on the left to fit. For example, if you specify 5 digits for the *hex\_value* and an automatic variable is only 1 byte long, the compiler truncates the 3 digits on the left-hand side of the *hex\_value* and assigns the two right-hand digits to the variable.
- You can specify alphabetic digits as either upper- or lower-case.

## Usage

This option helps you to locate variables that are referenced before being defined. For example, by using both the **-qinitauto** option to initialize **REAL** variables with a signaling NAN value and the **-qfltrap** option, it is possible to identify references to uninitialized **REAL** variables at run time.

Setting *hex\_value* to zero ensures that all automatic variables are cleared before being used. Some programs assume that variables are initialized to zero and do not work when they are not. Other programs may work if they are not optimized but fail when they are optimized. Typically, setting all the variables to all zero bytes prevents such runtime errors. It is better to locate the variables that require zeroing and insert code in your program to do so than to rely on this option to do it for you. Using this option will generally zero more things than necessary and may result in slower programs.

To locate and fix these errors, set the bytes to a value other than zero, which will consistently reproduce incorrect results. This method is especially valuable in cases where adding debugging statements or loading the program into a symbolic debugger makes the error go away.

Setting the *hex\_value* to **FF** (255) gives **REAL** and **COMPLEX** variables an initial value of “negative not a number”, or quiet NAN. Any operations on these variables will also result in quiet NAN values, making it clear that an uninitialized variable has been used in a calculation.

This option can help you to debug programs with uninitialized variables in subprograms. For example, you can use it to initialize **REAL** variables with a signaling NAN value. You can initialize 8-byte **REAL** variables to double-precision signaling NAN values by specifying an 8-digit hexadecimal number, that, when repeated, has a double-precision signaling NAN value. For example, you could specify a number such as **7FBFFFFF**, that, when stored in a **REAL(4)** variable, has a single-precision signaling NAN value. The value **7FF7FFFF**, when stored in a **REAL(4)** variable, has a single-precision quiet NAN value. If the same number is stored twice in a **REAL(8)** variable (**7FF7FFFF7FF7FFFF**), it has a double-precision signaling NAN value.

## Restrictions

Equivalenced variables, structure components, and array elements are not initialized individually. Instead, the entire storage sequence is initialized collectively.

## Examples

The following example shows how to perform word initialization of automatic variables:

```
subroutine sub()
integer(4), automatic :: i4
character, automatic :: c
real(4), automatic :: r4
real(8), automatic :: r8
end subroutine
```

When you compile the code with the following option, the compiler performs word initialization, as the *hex\_value* is longer than 2 digits:

```
-qinitauto=0cf
```

The compiler initializes the variables as follows, padding the *hex\_value* with zeros in the cases of the *i4*, *r4*, and *r8* variables and truncating the first hexadecimal digit in the case of the *c* variable:

Variable	Value
i4	000000CF
c	CF
r4	000000CF
r8	000000CF000000CF

## Related information

- “-qfltrap” on page 132
- The **AUTOMATIC** attribute in the *XL Fortran Language Reference*

---

## -qinlglue

### Category

Object code control

### Purpose

When used with **-O2** or higher optimization, inlines glue code that optimizes external function calls in your application.

*Glue code or Program Linkage Table code*, generated by the linker, is used for passing control between two external functions. When **-qinlglue** is in effect, the optimizer inlines glue code for better performance. When **-qnoinlglue** is in effect, inlining of glue code is prevented.

### Syntax

►► -q inlglue  
noinlglue ►►

**@PROCESS:**

@PROCESS **INLGLUE** | NOINLGLUE

### Defaults

- **-qinlglue**

### Usage

Inlining glue code can cause the code size to grow. Specifying **-qcompact** overrides the **-qinlglue** setting to prevent code growth. If you want **-qinlglue** to be enabled, do not specify **-qcompact**.

Specifying **-qnoinlglue** or **-qcompact** can degrade performance; use these options with discretion.

### Related information

- “-qcompact” on page 108
- “-qtune” on page 250
- Inlining in the *XL Fortran Optimization and Programming Guide*
- Managing code size in the *XL Fortran Optimization and Programming Guide*

---

## -qinline

### Category

Optimization and tuning


### @PROCESS

None.

## Purpose

Attempts to inline procedures instead of generating calls to those procedures, for improved performance.

## Syntax



## Defaults

If **-qinline** is not specified, the default option is as follows:

- **-qnoinline** at the **-00** or **-qnoopt** optimization level
- **-qinline=noauto:level=5** at the **-02** optimization level
- **-qinline=auto:level=5** at the **-02 -qipa, -03** or higher optimization level

If **-qinline** is specified without any suboptions, the default option is **-qinline=auto:level=5**.

## Parameters

### auto | noauto

Enables or disables automatic inlining. When option **-qinline=auto** is in effect, all the functions are considered for inlining by the compiler. When option **-qinline=noauto** is in effect, the compiler only inlines small functions that it deems appropriate for inlining.

### level=number

Indicates the relative degree of inlining. The values for *number* must be integers from 0 to 10 inclusive. The default value for *number* is 5. The greater the value of *number*, the more aggressive inlining the compiler conducts.

### procedure\_name

If the *procedure\_name* is specified after the **-qinline+** option, the named procedure must be inlined. If the *procedure\_name* is specified after the **-qinline-** option, the named procedure must not be inlined.

## Usage

You can specify **-qinline** with any optimization level of **-02, -03, -04, or -05** to enable inlining of procedures.

When **-qinline** is in effect, the compiler determines whether inlining a specific procedure can improve performance. That is, whether a procedure is appropriate for inlining is subject to two factors: limits on the number of inlined calls and the

amount of code size increase as a result. Therefore, enabling inlining a procedure does not guarantee that procedure will be inlined.

Because inlining does not always improve runtime performance, you need to test the effects of this option on your code. Do not attempt to inline recursive or mutually recursive procedures.

You can use the **-qinline+<procedure\_name>** or **-qinline-<procedure\_name>** option to specify the procedures that must be inlined or must not be inlined.

Specifying **-qnoinline** disables all inlining, including that achieved by the high-level optimizer with the **-qipa** option.

If you specify the **-g** option to generate debugging information, the inlining effect of **-qinline** might be suppressed.

If you specify the **-qcompact** option to avoid optimizations that increase code size, the inlining effect of **-qinline** might be suppressed.

By default, **-qinline** affects only internal or module procedures. To turn on inline expansion for calls to procedures in different scopes, you must also use the **-qipa** option.

Conflicting **@PROCESS** directives or compilation options that are applied to different compilation units can impact inlining effectiveness. For example, if you specify inlining for a procedure, some **@PROCESS** compiler directives can be rendered ineffective. See the *XL Fortran Optimization and Programming Guide* for more information about inlining and IPA.

If you specify inlining for a procedure, the following **@PROCESS** compiler directives are only effective if they come before the first compilation unit in the file: **ALIAS**, **ALIGN**, **ATTR**, **COMPACT**, **DBG**, **EXTCHK**, **EXTNAME**, **FLOAT**, **FLTTRAP**, **HALT**, **IEEE**, **LIST**, **MAXMEM**, **OBJECT**, **OPTIMIZE**, **PHSINFO**, **SPELLSIZE**, **STRICT**, and **XREF**.

## Examples

### Example 1

To compile `myprogram.f` so that no procedures are inlined, use the following command:

```
xlf myprogram.f -O2 -qnoinline
```

If you want to enable automatic inlining, you use the `auto` suboption:

```
-O2 -qinline=auto
```

You can specify an inlining level 6 - 10 to achieve more aggressive automatic inlining. For example:

```
-O2 -qinline=auto:level=7
```

If automatic inlining is already enabled by default and you want to specify an inlining level of 7, you enter:

```
-O2 -qinline=level=7
```

### Example 2

Assuming `myprogram.f` contains the `salary`, `taxes`, `expenses`, and `benefits` procedures, you can use the following command to compile `myprogram.f` to inline these procedures:

```
xlf myprogram.f -O2 -qinline+salary:taxes:expenses:benefits
```

If you do not want the procedures `salary`, `taxes`, `expenses`, and `benefits` to be inlined, use the following command to compile `myprogram.f`:

```
xlf myprogram.f -O2 -qinline-salary:taxes:expenses:benefits
```

You can also disable automatic inlining and specify certain procedures to be inlined with the **-qinline+** option. For example:

```
-O2 -qinline=noauto -qinline+salary:taxes:benefits
```

In this case, the procedures `salary`, `taxes`, and `benefits` are inlined. Besides, procedures that are declared with the `inline` specifier are also inlined. No other procedures are inlined.

You cannot mix the **+** and **-** suboptions with each other or with other **-qinline** suboptions. For example, the following options are invalid suboption combinations:

```
-qinline+increase-decrease // invalid  
-qinline=level=5+increase // invalid
```

However, you can use **-qinline** separately. For example:

```
-qinline+increase -qinline-decrease -qinline=noauto:level=5
```

## Related information

- “-g” on page 73
- “-qipa” on page 163
- "Interprocedural analysis" in the *XL Fortran Optimization and Programming Guide*
- The `-qinline` inlining option in the *XL Fortran Optimization and Programming Guide*

---

## -qintlog

### Category

Floating-point and integer control

### Purpose

Specifies that you can mix integer and logical data entities in expressions and statements.

### Syntax

```
►► -q nointlog  
intlog ◀◀
```

**@PROCESS:**

@PROCESS INTLOG | NOINTLOG

### Defaults

-qnointlog

## Usage

When **-qintlog** is specified, logical operators that you specify with integer operands act upon those integers in a bit-wise manner, and integer operators treat the contents of logical operands as integers.

The following operations do not allow the use of logical variables:

- **ASSIGN** statement variables
- Assigned **GOTO** variables
- **DO** loop index variables
- Implied-**DO** loop index variables in **DATA** statements
- Implied-**DO** loop index variables in either input and output or array constructors
- Index variables in **FORALL** constructs

You can also use the intrinsic functions **IAND**, **IOR**, **IEOR**, and **NOT** to perform bitwise logical operations.

The **MOVE\_ALLOC** intrinsic function cannot take one integer and one logical argument.

## Examples

```
INTEGER I, MASK, LOW_ORDER_BYTE, TWOS_COMPLEMENT
I = 32767
MASK = 255
! Find the low-order byte of an integer.
LOW_ORDER_BYTE = I .AND. MASK
! Find the twos complement of an integer.
TWOS_COMPLEMENT = (.NOT. I) + 1
END
```

## Related information

- **-qport=clogicals** option.

---

## -qintsize

### Category

Floating-point and integer control

### Purpose

Sets the size of default **INTEGER** and **LOGICAL** data entities that have no length or kind specified.

This option is not intended as a general-purpose method for increasing the sizes of data entities. Its use is limited to maintaining compatibility with code that is written for other systems.

### Syntax

►► -q-intsize = 

4
2
8

 ►►

@PROCESS:

@PROCESS INTSIZE(*bytes*)

## Defaults

**-qintsize=4**

## Parameters

*bytes*

Allowed sizes are 2, 4, or 8.

## Usage

This option is intended to allow you to port programs unchanged from systems that have different default sizes for data. For example, you might need **-qintsize=2** for programs that are written for a 16-bit microprocessor or **-qintsize=8** for programs that are written for a CRAY computer.

You might need to add **PARAMETER** statements to give explicit lengths to constants that you pass as arguments.

The specified size<sup>1</sup> applies to these data entities:

- **INTEGER** and **LOGICAL** specification statements with no length or kind specified.
- **FUNCTION** statements with no length or kind specified.
- Intrinsic functions that accept or return default **INTEGER** or **LOGICAL** arguments or return values unless you specify a length or kind in an **INTRINSIC** statement. Any specified length or kind must agree with the default size of the return value.
- Variables that are implicit integers or logicals.
- Integer and logical literal constants with no kind specified. If the value is too large to be represented by the number of bytes that you specified, the compiler chooses a size that is large enough. The range for 2-byte integers is  $-(2^{*15})$  to  $2^{*15}-1$ , for 4-byte integers is  $-(2^{*31})$  to  $2^{*31}-1$ , and for 8-byte integers is  $-(2^{*63})$  to  $2^{*63}-1$ .
- Typeless constants in integer or logical contexts.
- In addition to types **INTEGER** and **LOGICAL**, **-qintsize** also works for **vector(integer)**. Specifying **-qintsize=2** is equivalent to specifying **vector(integer\*2)**. Similarly, specifying **-qintsize=4** is equivalent to specifying **vector(integer\*4)**. Specifying **-qintsize=8** is equivalent to **vector(integer\*8)**.

## Examples

In the following example, note how variables, literal constants, intrinsics, arithmetic operators, and input/output operations all handle the changed default integer size.

```
@PROCESS INTSIZE(8)
PROGRAM INTSIZETEST
INTEGER I
I = -9223372036854775807    ! I is big enough to hold this constant.
J = ABS(I)                 ! So is implicit integer J.
```

---

1. In Fortran 90/95 terminology, these values are referred to as *kind type parameters*.

```

IF (I .NE. J) THEN
  PRINT *, I, '.NE.', J
END IF
END

```

The following example only works with the default size for integers:

```

CALL SUB(17)
END

SUBROUTINE SUB(I)
  INTEGER(4) I           ! But INTSIZE may change "17"
                        ! to INTEGER(2) or INTEGER(8).
...
END

```

If you change the default value, you must either declare the variable I as **INTEGER** instead of **INTEGER(4)** or give a length to the actual argument, as follows:

```

@PROCESS INTSIZE(8)
  INTEGER(4) X
  PARAMETER(X=17)
  CALL SUB(X)           ! Use a parameter with the right length, or
  CALL SUB(17_4)       ! use a constant with the right kind.
END

```

### Related information

- “-qrealsize” on page 213
- *Type declaration: type parameters and specifiers in the XL Fortran Language Reference*

## -qipa

### Category

Optimization and tuning

### @PROCESS

None.

### Purpose

Enables or customizes a class of optimizations known as interprocedural analysis (IPA).

IPA is a two-step process: the first step, which takes place during compilation, consists of performing an initial analysis and storing interprocedural analysis information in the object file. The second step, which takes place during linking, and causes a complete recompilation of the entire application, applies the optimizations to the entire program.


You can use **-qipa** during the compilation step, the link step, or both. If you compile and link in a single compiler invocation, only the link-time suboptions are relevant. If you compile and link in separate compiler invocations, only the compile-time suboptions are relevant during the compile step, and only the link-time suboptions are relevant during the link step.

## Syntax

### -qipa compile-time syntax



### -qipa link-time syntax



## Defaults

- `-qnoipa`

## Parameters

You can specify the following parameters during a separate compile step only:

### object | noobject

Specifies whether to include standard object code in the output object files.

Specifying **noobject** can substantially reduce overall compile time by not generating object code during the first IPA phase. Note that if you specify **-S** with **noobject**, **noobject** will be ignored.

If compiling and linking are performed in the same step and you do not specify the **-S** or any listing option, **-qipa=noobject** is implied.

Specifying **-qipa** with no suboptions on the compile step is equivalent to **-qipa=object**.

You can specify the following parameters during a combined compile and link in the same compiler invocation, or during a separate link step only:

#### **exits**

Specifies names of procedures which represent program exits. Program exits are calls which can never return and can never call any procedure which has been compiled with IPA pass 1. The compiler can optimize calls to these procedures (for example, by eliminating save/restore sequences), because the calls never return to the program. These procedures must not call any other parts of the program that are compiled with **-qipa**.

#### **isolated**

Specifies a comma-separated list of procedures that are not compiled with **-qipa**. Procedures that you specify as *isolated* or procedures within their call chains cannot refer directly to any global variable.

#### **level**

Specifies the optimization level for interprocedural analysis. Valid suboptions are one of the following suboptions:

- 0** Performs only minimal interprocedural analysis and optimization.
- 1** Enables inlining, limited alias analysis, and limited call-site tailoring.
- 2** Performs full interprocedural data flow and alias analysis.

If you do not specify a level, the default is 1.

To generate data reorganization information, specify the optimization level **-qipa=level=2** or **-05** together with **-qreport**. During the IPA link phase, the data reorganization messages for program variable data are written to the data reorganization section of the listing file. Reorganizations include common block splitting, array splitting, array transposing, memory allocation merging, array interleaving, and array coalescing.

#### **list**

Specifies that a listing file be generated during the link phase. The listing file contains information about transformations and analyses performed by IPA, as well as an optional object listing for each partition.

If you do not specify a *list\_file\_name*, the listing file name defaults to a.lst. If you specify **-qipa=list** together with any other option that generates a listing file, IPA generates an a.lst file that overwrites any existing a.lst file. If you have a source file named a.f, the IPA listing will overwrite the regular compiler listing a.lst. You can use the **-qipa=list=list\_file\_name** suboption to specify an alternative listing file name.

Additional suboptions are one of the following suboptions:

- short** Requests less information in the listing file. Generates the Object File Map, Source File Map and Global Symbols Map sections of the listing.

**long** Requests more information in the listing file. Generates all of the sections generated by the **short** suboption, plus the Object Resolution Warnings, Object Reference Map, Inliner Report and Partition Map sections.

**lowfreq**

Specifies procedures that are likely to be called infrequently. These are typically error handling, trace, or initialization procedures. The compiler may be able to make other parts of the program run faster by doing less optimization for calls to these procedures.

**missing**

Specifies the interprocedural behavior of procedures that are not compiled with **-qipa** and are not explicitly named in an **unknown**, **safe**, **isolated**, or **pure** suboption.

Valid suboptions are one of the following suboptions:

**safe** Specifies that the missing procedures do not indirectly call a visible (not missing) function either through direct call or through a procedure pointer.

**isolated**

Specifies that the missing procedures do not directly reference global variables accessible to visible procedures. Procedures bound from shared libraries are assumed to be *isolated*.

**pure** Specifies that the missing procedures are *safe* and *isolated* and do not indirectly alter storage accessible to visible procedures. *pure* procedures also have no observable internal state.

**unknown**

Specifies that the missing procedures are not known to be *safe*, *isolated*, or *pure*. This suboption greatly restricts the amount of interprocedural optimization for calls to missing procedures.

The default is to assume **unknown**.

**partition**

Specifies the size of each program partition created by IPA during pass 2. Valid suboptions are one of the following suboptions:

- **small**
- **medium**
- **large**

Larger partitions contain more procedures, which result in better interprocedural analysis but require more storage to optimize. Reduce the partition size if compilation takes too long because of paging.

**pure**

Specifies *pure* procedures that are not compiled with **-qipa**. Any procedure specified as *pure* must be *isolated* and *safe*, and must not alter the internal state or have side-effects, defined as potentially altering any data visible to the caller.

**safe**

Specifies *safe* procedures that are not compiled with **-qipa** and do not call any other part of the program. Safe procedures can modify global variables and dummy arguments, but may not call procedures compiled with **-qipa**.

### **stdexits | nostdexits**

Specifies that certain predefined routines can be optimized as with the **exits** suboption. The procedures are: `abort`, `exit`, `_exit`, and `_assert`.

### **threads | nothreads**

Runs portions of the IPA optimization process during pass 2 in parallel threads, which can speed up the compilation process on multi-processor systems. Valid suboptions for the **threads** suboption are as follows:

#### **auto | noauto**

When **auto** is in effect, the compiler selects a number of threads heuristically based on machine load. When **noauto** is in effect, the compiler creates one thread per machine processor.

#### *number*

Instructs the compiler to use a specific number of threads. *number* can be any integer value in the range of 1 to 32 767. However, *number* is effectively limited to the number of processors available on your system.

Specifying **threads** with no suboptions implies **-qipa=threads=auto**.

### **unknown**

Specifies *unknown* procedures that are not compiled with **-qipa**. Any procedure specified as *unknown* can make calls to other parts of the program compiled with **-qipa**, and modify global variables and dummy arguments.

### *file\_name*

Gives the name of a file which contains suboption information in a special format.

The file format is shown as follows:

```
# ... comment
attribute{, attribute} = name{, name}

missing = attribute{, attribute}
exits = name{, name}
lowfreq = name{, name}
list [ = file-name | short | long ]
level = 0 | 1 | 2
partition = small | medium | large
```

where *attribute* is one of:

- exits
- lowfreq
- unknown
- safe
- isolated
- pure

## **Usage**

Specifying **-qipa** automatically sets the optimization level to **-O2**. For additional performance benefits, you can also specify the **-qinline** option. The **-qipa** option extends the area that is examined during optimization and inlining from a single procedure to multiple procedures (possibly in different source files) and the linkage between them.

If any object file used in linking with **-qipa** was created with the **-qipa=noobject** option, any file containing an entry point (the main program for an executable program, or an exported function for a library) must be compiled with **-qipa**.

You can link objects created with different releases of the compiler, but you must ensure that you use a linker that is at least at the same release level as the newer of the compilers used to create the objects being linked.

Some symbols which are clearly referenced or set in the source code may be optimized away by IPA, and may be lost to **debug** or **nm** outputs. Using IPA together with the **-g** compiler will usually result in non-steppable output.

Note that if you specify **-qipa** with **-#**, the compiler does not display linker information subsequent to the IPA link step.

For recommended procedures for using **-qipa**, see "Optimizing your applications" in the *XL Fortran Optimization and Programming Guide*.

## Examples

The following example shows how you might compile a set of files with interprocedural analysis:

```
xlf -c *.f -qipa
xlf -o product *.o -qipa
```

Here is how you might compile the same set of files, improving the optimization of the second compilation, and the speed of the first compile step. Assume that there exist a set of routines, `user_trace1`, `user_trace2`, and `user_trace3`, which are rarely executed, and the routine `user_abort` that exits the program:

```
xlf95 -c *.f -qipa=noobject
xlf95 -c *.o -qipa=lowfreq=user_trace[123]:exit=user_abort
```

## Related information

- `-qinline`
- "`-qlibmpi`" on page 172
- "`-qpdf1`, `-qpdf2`" on page 195
- "`-S`" on page 272
- Correct settings for environment variables

---

## **-qkeepparm**

### Category

Error checking and debugging

### @PROCESS

None.

### Purpose

When used with **-O2** or higher optimization, specifies whether procedure parameters are stored on the stack.

A procedure usually stores its incoming parameters on the stack at the entry point. However, when you compile code with optimization options enabled, the optimizer may remove these parameters from the stack if it sees an optimizing advantage in doing so.

## Syntax

►► -q 

nokeepparm
keepparm

 ►►

## Defaults

-qnokeepparm

## Usage

When **-qkeepparm** is in effect, parameters are stored on the stack even when optimization is enabled. This option then provides access to the values of incoming parameters to tools, such as debuggers, simply by preserving those values on the stack. However, this may negatively impact execution performance.

When **-qnokeepparm** is in effect, parameters are removed from the stack if this provides an optimization advantage.

---

## -qlanglvl

### Category

Language element control

### Purpose

Determines which language standard (or superset, or subset of a standard) to consult for nonconformance. It identifies nonconforming source code and also options that allow such nonconformances.

## Syntax

►► -q-`langlvl` = 

extended
77std
90std
90pure
95std
95pure
2003std
2003pure
2008std
2008pure
ts
tspure

 ►►

@PROCESS:

@PROCESS `LANGLVL` ({*suboption*})

## Defaults

-qlanglvl=extended

## Parameters

**77std** Accepts the language that the ANSI FORTRAN 77 standard specifies and reports anything else using language-level messages.

**90std** Accepts the language that the ISO Fortran 90 standard specifies and reports anything else using language-level messages.

### **90pure**

The same as **90std** except that it also issues language-level messages for any obsolescent Fortran 90 features used.

**95std** Accepts the language that the ISO Fortran 95 standard specifies and reports anything else using language-level messages.

### **95pure**

The same as **95std** except that it also issues language-level messages for any obsolescent Fortran 95 features used.

### **2003std**

Accepts the language that the ISO Fortran 2003 standard specifies and reports anything else using language-level messages.

### **2003pure**

The same as **2003std** except that it also issues language-level messages for any obsolescent Fortran 2003 features used.

### **2008std**

Accepts the language that the ISO Fortran 2003 standard specifies and all Fortran 2008 features supported by XL Fortran, and reports anything else using language-level messages.

### **2008pure**

The same as **2008std** except that it also issues language-level messages for any obsolescent Fortran 2008 features used.

**ts** Accepts the language that a standard technical specification specifies, and reports anything else using language-level messages. A technical specification is an addition to the current Fortran standard. This option checks for the language that the ISO Fortran 2003 standard specifies, all supported ISO Fortran 2008 standard features, and all supported ISO TS 29113 features.

**tspure** The same as **ts** except that it also issues language-level messages for any obsolescent Technical Specification features used.

### **extended**

Accepts the full Fortran 2003 language standard, all Fortran 2008 features supported by XL Fortran, and all extensions, effectively turning off language-level checking.

## Usage

When a **-qlanglvl** setting is specified, the compiler issues a message with severity code **L** if it detects syntax that is not allowed by the language level that you specified.

The **-qflag** option can override the **-qlanglvl** option.

The **langlvl** runtime option, which is described in “Setting runtime options” on page 34, helps to locate runtime extensions that cannot be checked for at compile time.

## Examples

The following example contains source code that conforms to a mixture of Fortran standards:

```
!-----  
! in free source form  
program tt  
  integer :: a(100,100), b(100), i  
  real :: x, y  
  ...  
  goto (10, 20, 30), i  
10 continue  
  pause 'waiting for input'  
  
20 continue  
  y= gamma(x)  
  
30 continue  
  b = maxloc(a, dim=1, mask=a .lt 0)  
  
end program  
!-----
```

The following table shows examples of how some **-qlanglvl** suboptions affect this sample program:

<b>-qlanglvl Suboption Specified</b>	<b>Result</b>	<b>Reason</b>
<b>95pure</b>	Flags <b>PAUSE</b> statement Flags computed <b>GOTO</b> statement Flags <b>GAMMA</b> intrinsic	Deleted feature in Fortran 95 Obsolescent feature in Fortran 95 Extension to Fortran 95
<b>95std</b>	Flags <b>PAUSE</b> statement Flags <b>GAMMA</b> intrinsic	Deleted feature in Fortran 95 Extension to Fortran 95
<b>extended</b>	No errors flagged	

### Related information

- “-qflag” on page 126
- “-qhalt” on page 140
- “-qsaa” on page 218
- **langlvl** runtime option in “Setting runtime options” on page 34

---

## **-qlibansi**

### **Category**

Optimization and tuning

### **@PROCESS**

None.

## Purpose

Assumes that all functions with the name of an ANSI C library function are, in fact, the library functions and not a user function with different semantics.

## Syntax

►► -q no libansi libansi ►►

## Defaults

-qnolibansi

## Usage

This option will allow the optimizer to generate better code because it will know about the behavior of a given function, such as whether or not it has any side effects.

**Note:** Do not use this option if your application contains your own version of a library function that is incompatible with the standard one.

## Related information

See “-qipa” on page 163.

---

## -qlibmpi

### Category

“Optimization and tuning” on page 57

### @PROCESS

None

## Purpose

Asserts that all functions with Message Passing Interface (MPI) names are in fact MPI functions and not a user function with different semantics.

## Syntax

►► -q no libmpi libmpi ►►

## Defaults

-qnolibmpi

## Usage

MPI is a library interface specification for message passing. It addresses the message-passing parallel programming model in which data is moved from the address space of one process to another through cooperative operations. For details about MPI, see the Message Passing Interface Forum.

**-q1ibmpi** allows the compiler to generate better code because it knows about the behavior of a given function, such as whether or not it has any side effects.

When you use **-q1ibmpi**, the compiler assumes that all functions with the name of an MPI library function are in fact MPI functions. **-qnolibmpi** makes no such assumptions.

**Note:** You cannot use this option if your application contains your own version of the library function that is incompatible with the standard one.

## Examples

To compile `myprogram.f`, enter the following command:

```
xlf -O5 myprogram.f -q1ibmpi
```

## Related information

- Message Passing Interface Forum
- “-qipa” on page 163

---

## -qlinedebug

### Category

Error checking and debugging

### Purpose

Generates only line number and source file name information for a debugger.

When **-qlinedebug** is in effect, the compiler produces minimal debugging information, so the resulting object size is smaller than that produced by the **-g** debugging option. You can use the debugger to step through the source code, but you will not be able to see or query variable information. The traceback table, if generated, will include line numbers.

**-qlinedebug** is equivalent to **-g1**.

### Syntax

```
└─no┬linedebug┘  
└─q┬linedebug┘
```

**@PROCESS:**

@PROCESS LINEDEBUG | NOLINEDEBUG

### Defaults

-qnolinedebug

## Usage

When **-q`l`indebug** is in effect, function inlining is disabled.

As with all debugging information, the output of **-q`l`indebug** might be incomplete or misleading if the code is optimized.

The **-g** option overrides the **-q`l`indebug** option. If you specify **-g** with **-q`no`l`i`ndebug** on the command line, **-q`no`l`i`ndebug** is ignored and a warning is issued.

## Related information

- “-g” on page 73
- “-O” on page 83

---

## -q`l`ist

### Category


Listings, messages, and compiler information

### Purpose

Produces a compiler listing file that includes object and constant area sections.

When **-q`l`ist** is in effect, a listing file is generated with a `.lst` suffix for each source file named on the command line.

### Syntax



@PROCESS:

@PROCESS LIST[[[NO]OFFSET]] | NOLIST

### Defaults

-qnolist

### Parameters

**offset** | **nooffset**

When **-q`l`ist=offset** is in effect, the listing will show the offset from the start of the procedure rather than the offset from the start of code generation. This suboption allows any program reading the `.lst` file to add the value of the PDEF and the line in question, and come up with the same value whether **offset** or **nooffset** is specified.

The **offset** suboption is relevant only if there are multiple procedures in a compilation unit. For example, this may occur if nested procedures are used in a program.

Specifying **-q`l`ist** with no suboption is equivalent to **-q`l`ist=nooffset**.

## Usage

You can use the object listing to help understand the performance characteristics of the generated code and to diagnose execution problems.

If you specify **-qipa** and want to generate the IPA listing file, use the **-qipa=list=filename** suboption to get an alternative listing.

The **-qnoprint** compiler option overrides this option.

## Related information

- “Listings, messages, and compiler information” on page 56
- “Object section” on page 296
- “-qnoprint” on page 188
- “-S” on page 272
- *Program units and procedures* in the *XL Fortran Language Reference*

---

## -qlistfmt

### Category

Listings, messages, and compiler information


### @PROCESS

None.

### Purpose

Creates an XML or HTML report to assist with finding optimization opportunities.

### Syntax



### Defaults

This option is off by default. If no `contentSelectionList` options are selected in their positive form, all available report information is produced. For example, specifying **-qlistfmt=xml** is equivalent to **-qlistfmt=xml=all**.

### Parameters

The following list describes **-qlistfmt** parameters:

#### xml | html

Indicates that the report should be generated in XML or HTML format. If an XML report has been generated before, you can convert the report to the

HTML format using the **genhtml** command. For more information about this command, see “genhtml command” on page 177.

#### *contentSelectionList*

The following suboptions provide a filter to limit the type and quantity of information in the report:

##### **data** | **nodata**

Produces data reorganization information.

##### **inlines** | **noinlines**

Produces inlining information.

##### **pdf** | **nopdf**

Produce profile-directed feedback information.

##### **transforms** | **notransforms**

Produces loop transformation information.

##### **all**

Produces all available report information.

##### **none**

Does not produce a report.

#### **filename**

Specifies the name of the report file. One file is produced during the compile phase, and one file is produced during the IPA link phase. If no filename is specified, a file with the suffix `.xml` or `.html` is generated in a way that is consistent with the rules of name generation for the given platform. For example, if compiling `foo.f` the generated XML files are `foo.xml` from the compile step and `a.xml` from the link step.

**Note:** If you compile and link in one step and use this suboption to specify a file name for the report, the information from the IPA link step will overwrite the information generated during the compile step.

The same will be true if you compile multiple files using the `filename` suboption. The compiler creates an report for each file so the report of the last file compiled will overwrite the previous reports. For example,

```
xlf -qlistfmt=xml=all:filename=abc.xml -O3 myfile1.f myfile2.f myfile3.f
```

will result in only one report, `abc.xml` based on the compilation of the last file `myfile3.f`

#### **stylesheet**

Specifies the name of an existing XML stylesheet for which an `xml-stylesheet` directive is embedded in the resulting report. The default behavior is to not include a stylesheet. The stylesheet shipped with XL Fortran is `x1style.xml`. This stylesheet renders the XML to an easily read format when viewed using a browser that supports XSLT.

To view the XML report created with the **stylesheet** suboption, you must place the actual stylesheet (`x1style.xml`) and the XML message catalogue (`XMLMessages-locale.xml` where *locale* refers to the locale set on the compilation machine) in the path specified by the **stylesheet** suboption. For example, if `a.xml` is generated with **stylesheet=x1style.xml**, `x1style.xml` and `XMLMessages-locale.xml` must be in the same directory as `a.xml`, before you can properly view `a.xml` with a browser. The message catalogs and stylesheet are installed in the `/opt/ibm/xlf/15.1.1/listings/` directory.

## version

Specifies the major version of the content that is emitted. If you have written a tool that requires a certain version of this report, you should specify the version. IBM XL Fortran for Linux, V15.1.1 creates reports at XML v1.1, so if you have written a tool to consume these reports, specify `version=v1`.

## Usage

The information produced in the report by the `-qlistfmt` option depends on which optimization options are used to compile the program.

- When used with an option that enables inlining such as `-qinline`, the report shows which functions were inlined and why others were not inlined.
- When used with an option that enables loop unrolling, the report contains a summary of how program loops are optimized. The report also includes diagnostic information to show why specific loops cannot be vectorized. For `-qlistfmt` to generate information about loop transformations, you must also specify at least one of the following options:
  - `-qsimd=auto`
  - `-qsmp`
  - `-O5`
  - `-qipa=level=2`
- When used with an option that enables parallel transformations, the report contains information about parallel transformations. For `-qlistfmt` to generate information about parallel transformations or parallel performance messages, you must also specify at least one of the following options:
  - `-qsmp`
  - `-O5`
  - `-qipa=level=2`
- When used with the option that enables profiling, `-qpdf`, the report contains information about call and block counts and cache misses.
- When used with an option that produces data reorganizations such as `-qipa=level=2`, the report contains information about those reorganizations.

If no *contentSelectionList* options are selected in their positive form, all available report information is produced.

## Examples

If you want to compile `myprogram.f` to produce an XML report that shows how loops are optimized, enter:

```
xlf -qhot -O3 -qlistfmt=xml=transforms myprogram.f
```

If you want to compile `myprogram.f` to produce an XML report that shows which functions are inlined, enter:

```
xlf -qinline -qlistfmt=xml=inlines myprogram.f
```

## genhtml command

To view the HTML version of an XML report that has already been generated, you can use the `genhtml` tool.

Use the following command to view the existing XML report in HTML format. This command generates the HTML content to standard output.

```
genhtml xml_file
```

Use the following command to generate the HTML content into a defined HTML file. You can use a web browser to view the generated HTML file.

```
genhtml xml_file > target_html_file
```

**Note:** The suffix of the HTML file name must be compliant with the static HTML page standard, for example, `.html` or `.htm`. Otherwise, the web browser might not be able to open the file.

### Related information

- “-qreport” on page 216
- “Using compiler reports to diagnose optimization opportunities” in the *XL Fortran Optimization and Programming Guide*

---

## -qlistopt

### Category

Listings, messages, and compiler information

### Purpose

Produces a compiler listing file that includes all options in effect at the time of compiler invocation.

When **listopt** is in effect, a listing file is generated with a `.lst` suffix for each source file named on the command line. The listing shows options in effect as set by the compiler defaults, the configuration file, and command line settings.

### Syntax

```
►► -q no listopt  
listopt
```

**@PROCESS:**

@PROCESS LISTOPT | NOLISTOPT

### Defaults

-qnolistopt

### Usage

You can use the option listing during debugging to check whether a problem occurs under a particular combination of compiler options, or during performance testing to record the optimization options in effect for a particular compilation.

Options that are always displayed in the listing include:

- All “on/off” options that are on by default: for example, **-qobject**
- All “on/off” options that are explicitly turned off through the configuration file, command-line options, or **@PROCESS** directives
- All options that take arbitrary numeric arguments (typically sizes)
- All options that have multiple suboptions

The **-qnoprint** compiler option overrides this option.

### Related information

- “Listings, messages, and compiler information” on page 56
- “Options section” on page 291

---

## -qlog4

### Category

Portability and migration

### Purpose

Specifies whether the result of a logical operation with logical operands is a **LOGICAL(4)** or is a **LOGICAL** with the maximum length of the operands.

You can use this option to port code that was originally written for the IBM VS FORTRAN compiler.

### Syntax

►► — -q no1og4  
log4 —►►

**@PROCESS:**

@PROCESS LOG4 | NOLOG4

### Defaults

**-qno1og4**, which makes the result depend on the lengths of the operands.

### Usage

Specifying **-q1og4** makes the result a **LOGICAL(4)**.

If you use **-qintsize** to change the default size of logicals, **-q1og4** is ignored.

---

## -qmakedep

### Category

Output control

**@PROCESS**

None.

### Purpose

Produces a dependency output file containing targets suitable for inclusion in a description file for the **make** command.

The dependency output file is named with a **.d** suffix.

**-qmakedep** is the long form of **-MMD**.

## Syntax

► `-qmakedep` \_gcc\_ ►

## Defaults

Not applicable.

## Parameters

### gcc

The format of the generated **make** rule to match the GCC format: the dependency output file includes a single target that lists all of the source file's dependencies.

If you specify **-qmakedep** with no suboption, the dependency output file specifies a separate rule for each of the source file's dependencies.

## Usage

The **make** command uses dependency information to determine the compilation order of the files. The **make** command also uses the dependency information to determine the minimum set of files that must be recompiled when a file is changed.

XL Fortran recognizes the following types of source file dependencies:

- Dependencies on the files that are included through C preprocessor `#include` directives.
- Dependencies on the files that are included through the Fortran `INCLUDE` directive.
- Dependencies on the module symbol files in the files that use or extend one or more Fortran modules.
- ▶ F2008 Dependencies on the submodule symbol files in the files that extend one or more Fortran submodules. ◀ F2008

For each source file that is named on the command line, a dependency output file is generated with the same name as the object file but with a `.d` suffix. Dependency output files are not created for any other types of input files. If you use the **-o** option to rename the object file, the name of the dependency output file is based on the name specified in the **-o** option. For more information, see the Examples section.

The dependency output files generated by these options are not **make** description files; they must be linked before they can be used with the **make** command. For more information about this command, see your operating system documentation.

If the **-qfullpath** option is also specified, the absolute path names of the source and include files are recorded in the dependency output file.

You can also use **-qmakedep** with the following options:

**-MF** *file\_path*

Sets the name of the dependency output file, where *file\_path* is the full or partial path or file name for the dependency output file. For more information, see “-MF” on page 79.

**-MT** *target*

Specifies the target name of the object file in the **make** rule in the generated dependency file. For more information, see “-MT” on page 81.

## Examples

**Example 1:** To compile `mysource.f` and create a dependency output file named `mysource.d`, enter:

```
xlf -c -qmakedep mysource.f
```

**Example 2:** To compile `source.f` and create an object file named `object.o` and a dependency output file named `object.d`, enter:

```
xlf -c -qmakedep source.f -o object.o
```

**Example 3:** If you have the following files in the current working directory:

- `options.h`
- `constants.h`
- `n.F`
- `m.f`

The `options.h` file contains following code:

```
@PROCESS free(f90)
```

The `constants.h` file contains following code:

```
real(4), parameter :: pi = 3.14
```

The `n.F` file contains following code:

```
#include "options.h"
module n
contains
  subroutine my_print(x)
    real, value :: x
    print *, x
  end subroutine
end module
```

The `m.f` file contains following code:

```
#include "options.h"
module m
  use n
contains
  subroutine sub
    implicit none
    include 'constants.h'
    call my_print(pi)
  end subroutine
end module
```

To compile `n.F` and create a dependency output file named `n.d` in the `./dependencies` directory, enter:

```
xlf -c n.F -qmakedep -MF./dependencies -o n_obj.o
```



The severity levels (from lowest to highest) are listed as follows. See “Error severity” on page 283 for more information about these levels.

- i** Informational messages.
- l** Language-level messages, such as those produced under the **-qlanglvl** option.
- w** Warning messages.
- e** Error messages.
- s** Severe error messages.

## Usage

If the **-qmaxerr** option does not specify the severity level, it uses the severity that is in effect by the **-qhalt** option; otherwise, the severity level is specified by either **-qmaxerr** or **-qhalt** that appears last.

When **-qflag** is specified, **-qmaxerr** counts the messages that are allowed by the **-qflag** option.

When **-qsuppress** is specified, **-qmaxerr** does not count the messages that are censored by the **-qsuppress** option.

## Examples

To stop compilation of `myprogram.f` when 5 error messages are encountered, enter the following command:

```
xlf myprogram.f -qmaxerr=5:e
```

To stop compilation of `myprogram.f` when 5 severe errors are encountered, enter the following command:

```
xlf myprogram.f -qmaxerr=5
```

To stop compilation of `myprogram.f` when 5 language level messages are encountered, enter the following command:

```
xlf myprogram.f -qmaxerr=5:l
```

or:

```
xlf myprogram.f -qmaxerr=5 -qhalt=1
```

## Related information

- “-qflag” on page 126
- “-qhalt” on page 140
- “-qsuppress” on page 245
- “Error severity” on page 283

---

## **-qmaxmem**

### Category

Optimization and tuning

## Purpose

Limits the amount of memory that the compiler allocates while performing specific, memory-intensive optimizations to the specified number of kilobytes.

## Syntax

► — `-q—maxmem=Kbytes` —►

### @PROCESS:

@PROCESS MAXMEM(*Kbytes*)

## Defaults

- `maxmem=8192` when `-02` is in effect.
- `maxmem=-1` when `-03` or higher optimization is in effect.

## Parameters

### *Kbytes*

The number of kilobytes worth of memory to be used by optimizations. The limit is the amount of memory for specific optimizations, and not for the compiler as a whole. Tables required during the entire compilation process are not affected by or included in this limit.

A value of `-1` permits each optimization to take as much memory as it needs without checking for limits.

## Usage

If the specified amount of memory is insufficient for the compiler to compute a particular optimization, the compiler issues a message and reduces the degree of optimization.

This option has no effect except in combination with the `-0` option.

When compiling with `-02`, you only need to increase the limit if a compile-time message instructs you to do so. When compiling with `-03`, you might need to establish a limit if compilation stops because the machine runs out of storage; start with a value of 8192 or higher, and decrease it if the compilation continues to require too much storage.

### Note:

1. Reduced optimization does not necessarily mean that the resulting program will be slower. It only means that the compiler cannot finish looking for opportunities to improve performance.
2. Increasing the limit does not necessarily mean that the resulting program will be faster. It only means that the compiler is better able to find opportunities to improve performance if they exist.
3. Setting a large limit has no negative effect when compiling source files for which the compiler does not need to use so much memory during optimization.
4. As an alternative to raising the memory limit, you can sometimes move the most complicated calculations into procedures that are then small enough to be fully analyzed.
5. Not all memory-intensive compilation stages can be limited.

6. Only the optimizations done for **-02** and **-03** can be limited; **-04** and **-05** optimizations cannot be limited.
7. The **-04** and **-05** optimizations may also use a file in the `/tmp` directory. This is not limited by the `-qmaxmem` setting.
8. Some optimizations back off automatically before they exceed the maximum available address space, but not before they exceed the paging space available at that time, which depends on machine workload.

## Restrictions

Depending on the source file being compiled, the size of subprograms in the source code, the machine configuration, and the workload on the system, setting the limit too high might fill up the paging space. In particular, a value of **-1** can fill up the storage of even a well-equipped machine.

## Related information

- “-O” on page 83
- *Optimizing your applications* in the *XL Fortran Optimization and Programming Guide*

## -qmbcs

### Category

Language element control

### Purpose

Indicates to the compiler whether character literal constants, Hollerith constants, H edit descriptors, and character string edit descriptors can contain Multibyte Character Set (MBCS) or Unicode characters.

This option is intended for applications that must deal with data in a multibyte language, such as Japanese.

### Syntax

```

▶▶ -q nombcs  
mbsc

```

**@PROCESS:**

@PROCESS MBCS | NOMBCS

### Defaults

-qnombcs

### Usage

Each byte of a multibyte character is counted as a column.

To process the multibyte data correctly at run time, set the locale (through the **LANG** environment variable or a call to the **libc setlocale** routine) to the same value as during compilation.

To read or write Unicode data, set the locale value to **UNIVERSAL** at run time. If you do not set the locale, you might not be able to interchange data with Unicode-enabled applications.

---

## **-qmixed**

### **Category**

Input control

### **Purpose**

This is the long form of the “-U” on page 274 option.

### **Syntax**

►► -q nomixed  
mixed \_\_\_\_\_►►

**@PROCESS:**

@PROCESS MIXED | NOMIXED

### **Defaults**

-qnomixed

---

## **-qmkshrobj**

### **Category**

Output control

### **@PROCESS**

None.

### **Purpose**

Creates a shared object from generated object files.

Use this option, together with the related options described later in this topic, instead of calling the linker directly to create a shared object. The advantage of using this option is that it is compatible with **-qipa** link-time optimizations (such as those performed at **-05**).

### **Syntax**

►► -q—mkshrobj \_\_\_\_\_►►

### **Defaults**

By default, the output object is linked with the runtime libraries and startup routines to create an executable file.

## Usage

The compiler automatically exports all global symbols from the shared object unless you specify which symbols to export by using the **--version-script** linker option.  Symbols that have the hidden or internal visibility attribute are not exported. 

Specifying **-qmkshrobj** implies **-qp ic**.

You can also use the following related options with **-qmkshrobj**:

**-o** *shared\_file*

The name of the file that holds the shared file information. The default is `a.out`.

**-e** *name*

Sets the entry name for the shared executable to *name*.

**Note:** Options **-qmkshrobj** and **-qstaticlink** are incompatible and cannot be specified together.

For detailed information about using **-qmkshrobj** to create shared libraries, see “Compiling and linking a library” on page 21.

## Examples

To construct the shared library `big_lib.so` from three smaller object files, enter the following command:

```
xlf -qmkshrobj -o big_lib.so lib_a.o lib_b.o lib_c.o
```

## Related information

- “-e” on page 70
- “-qipa” on page 163
- “-o” on page 85
- “-qp ic” on page 205

---

## -qmoddir

### Category

Output control

### @PROCESS

None.

### Purpose

Specifies the location for any module (`.mod`) or  submodule  (`.smod`) files that the compiler writes.

### Syntax

 `-q—moddir—=—directory` 

## Defaults

Not applicable.

## Usage

If you do not specify **-qmoddir**, the `.mod` or `.smod` files are placed in the current directory.

To read the `.mod` or `.smod` files from this directory when compiling files that reference the modules or `F2008` submodules `F2008`, use the **-I** option.

## Related information

- “XL Fortran output files” on page 24
- “-I” on page 76
- *Modules* section in the *XL Fortran Language Reference*.
- `F2008` *Submodules* section in the *XL Fortran Language Reference*. `F2008`

---

## -qnoprint

### Category

Listings, messages, and compiler information

### @PROCESS

None.

### Purpose

Prevents the compiler from creating the listing file, regardless of the settings of other listing options.

### Syntax

▶▶ -q—noprint —————▶▶

## Defaults

Not applicable.

## Usage

Specifying **-qnoprint** on the command line enables you to put other listing options in a configuration file or on **@PROCESS** directives and still prevent the listing file from being created.

A listing file is usually created when you specify any of the following options: **-qattr**, **-qlist**, **-qlistopt**, **-qphsinfo**, **-qsource**, **-qreport**, or **-qxref**. **-qnoprint** prevents the listing file from being created by changing its name to `/dev/null`, a device that discards any data that is written to it.

## Related information

- “Listings, messages, and compiler information” on page 56

---

## -qnullterm

### Category

Language element control

### Purpose

Appends a null character to each character constant expression that is passed as a dummy argument, making it more convenient to pass strings to C functions.

This option allows you to pass strings to C functions without having to add a null character to each string argument.

### Syntax

► — -q nonullterm  
nullterm —►

@PROCESS:

@PROCESS NULLTERM | NONULLTERM

### Defaults

-qnonullterm

### Usage

This option affects arguments that are composed of any of the following objects:

- Basic character constants
- Concatenations of multiple character constants
- Named constants of type character
- Hollerith constants
- Binary, octal, or hexadecimal typeless constants when an interface block is available
- Any character expression composed entirely of these objects.

The result values from the **CHAR** and **ACHAR** intrinsic functions also have a null character added to them if the arguments to the intrinsic function are constant expressions.

### Rules

This option does not change the length of the dummy argument, which is defined by the additional length argument that is passed as part of the XL Fortran calling convention.

### Restrictions

This option affects those arguments that are passed with or without the **%REF** built-in function, but it does not affect those that are passed by value. This option does not affect character expressions in input and output statements.

## Examples

Here are two calls to the same C function; one with, and one without the option:

```
@PROCESS NONULLTERM
SUBROUTINE CALL_C_1
  CHARACTER*9, PARAMETER :: HOME = "/home/luc"
! Call the libc routine mkdir() to create some directories.
  CALL mkdir ("/home/luc/testfiles\0", %val(448))
! Call the libc routine unlink() to remove a file in the home directory.
  CALL unlink (HOME // "/.hushlogin" // CHAR(0))
END SUBROUTINE

@PROCESS NULLTERM
SUBROUTINE CALL_C_2
  CHARACTER*9, PARAMETER :: HOME = "/home/luc"
! With the option, there is no need to worry about the trailing null
! for each string argument.
  CALL mkdir ("/home/luc/testfiles", %val(448))
  CALL unlink (HOME // "/.hushlogin")
END SUBROUTINE
!
```

## Related information

See *Passing character types between languages* in the *XL Fortran Optimization and Programming Guide*.

---

## -qobject


### Category

Error checking and debugging

### Purpose

Specifies whether to produce an object file or to stop immediately after checking the syntax of the source files.

### Syntax

►► -q  ◀◀

@PROCESS:

@PROCESS OBJect | NOOBJect

### Defaults

-qobject

### Usage

When debugging a large program that takes a long time to compile, you might want to use the **-qnoobject** option. It allows you to quickly check the syntax of a program without incurring the overhead of code generation. The **.lst** file is still produced, so you can get diagnostic information to begin debugging.

After fixing any program errors, you can change back to the default (**-qobject**) to test whether the program works correctly. If it does not work correctly, compile with the **-g** option for interactive debugging.

The **-qhalt** option can override the **-qobject** option, and **-qnoobject** can override **-qhalt**.

### Related information

- “Listings, messages, and compiler information” on page 56
- “Object section” on page 296
- The compiler phases in the *Getting Started with XL Fortran*.
- “-qhalt” on page 140

---

## -qonetrip

### Category

Language element control

### Purpose

This is the long form of the **-1** option.

### Syntax

►► -q onetrip noonetrip ◀◀

@PROCESS:

@PROCESS ONETRIP | NOONETRIP

### Defaults

-qnoonetrip

---

## -qoptfile

### Category

Compiler customization

### @PROCESS directive

None.

### Purpose

Specifies a response file that contains a list of additional command line options to be used for the compilation. Response files typically have the **.rsp** suffix.

### Syntax

►► -q-optfile=filename ◀◀

## Defaults

None.

## Parameters

*filename*

Specifies the name of the response file that contains a list of additional command line options. *filename* can contain a relative path or absolute path, or it can contain no path. It is a plain text file with one or more command line options per line.

## Usage

The format of the response file follows these rules:

- Specify the options you want to include in the file with the same syntax as on the command line. The response file is a whitespace-separated list of options. The following special characters indicate whitespace: \n, \v, \t. (All of these characters have the same effect.)
- A character string between a pair of single or double quotation marks are passed to the compiler as one option.
- You can include comments in the response file. Comment lines start with the # character and continue to the end of the line. The compiler ignores comments and empty lines.

When processed, the compiler removes the **-qoptfile** option from the command line, and sequentially inserts the options included in the file before the other subsequent options that you specify.

The **-qoptfile** option is also valid within a response file. The files that contain another response file are processed in a depth-first manner. The compiler avoids infinite loops by detecting and ignoring cycles in response file inclusion.

If **-qoptfile** and **-qsaveopt** are specified on the same command line, the original command line is used for **-qsaveopt**. A new line for each response file is included representing the contents of each response file. The options contained in the file are saved to the compiled object file.

## Example 1

This is an example of specifying a response file.

```
$ cat options.file
# To perform optimization at -O3 level, and high-order
# loop analysis and transformations during optimization
-O3 -qhot
# To indicate that the input source program is in fixed source form
-qfixed

$ xlf95 -qlist -qoptfile=options.file -qipa test.f
```

The preceding example is equivalent to the following invocation:

```
$ xlf95 -qlist -O3 -qhot -qfixed -qipa test.f
```

## Example 2

This is an example of specifying a response file that contains **-qoptfile** with a cycle.

```

$ cat options.file1
# To perform optimization at -O3 level, and high-order
# loop analysis and transformations during optimization
-O3 -qhot
# To include the -qoptfile option in the same response file
-qoptfile=options.file1
# To indicate that the input source program is in fixed source form
-qfixed
# To indicate that the source code is in free source form
-qfree

$ xlf95 -qlist -qoptfile=options.file1 -qipa test.f

```

The preceding example is equivalent to the following invocation:

```
$ xlf95 -qlist -O3 -qhot -qfixed -qfree -qipa test.f
```

### Example 3

This is an example of specifying a response file that contains **-qoptfile** without a cycle.

```

$ cat options.file1
-O3 -qhot
-qoptfile=options.file2

$ cat options.file2
-qfixed

$ xlf95 -qoptfile=options.file1 test.f

```

The preceding example is equivalent to the following invocation:

```
$ xlf95 -O3 -qhot -qfixed test.f
```

### Example 4

This is an example of specifying **-qsaveopt** and **-qoptfile** on the same command line.

```

$ cat options.file3
-O3
-qassert=contiguous

$ xlf95 -qsaveopt -qipa -qoptfile=options.file3 test.f -c

$ what test.o
test.o:
opt f xlf95 -qsaveopt -qipa -qoptfile=options.file3 test.f -c
optfile options.file3 -O3 -qassert=contiguous

```

### Related information

- “-qsaveopt” on page 220

---

## -qoptimize

### Purpose

This is the long form of the **-O** option.

### Syntax



**@PROCESS:**

@PROCESS OPTimize[(level)] | NOOPTimize

### Defaults

-qnooptimize

## -qpath

### Category


Compiler customization

### Purpose

Specifies substitute path names for XL Fortran components such as the assembler, C preprocessor, and linker.

You can use this option if you want to keep multiple levels of some or all of the XL Fortran components and have the option of specifying which one you want to use.

### Syntax



### Defaults

By default, the compiler uses the paths for compiler components defined in the configuration file.

### Parameters

*directory\_path*

The path to the directory where the compiler components are located. It must be an existing directory. It can be relative or absolute.

The following table shows the correspondence between **-qpath** parameters and the component names:

Parameter	Description	Component name
a	The assembler	as

Parameter	Description	Component name
b	The low-level optimizer	xlfcode
c	The compiler front end	xlfentry
d	The disassembler	dis
F	The C preprocessor	cpp
h	The array language optimizer	xlfhot
I (uppercase i)	The high-level optimizer, compile step	ipa
l (lowercase L)	The linker	ld
z	The binder	bolt

## Usage

The **-qpath** option overrides the **-F**, **-t**, and **-B** options.

## Examples

To compile `myprogram.f` using a substitute compiler front end and linker from `/fix/FE/` and the remaining compiler components from default locations, enter:

```
xlf myprogram.f -qpath=c:/fix/FE
```

To compile `myprogram.f` using a substitute compiler front end from `/fix/FE`, a substitute linker from the current directory, and the remaining compiler components from default locations, enter:

```
xlf95 myprogram.f -qpath=c:/fix/FE -qpath=l:.
```

## Related information

- “-B” on page 67
- “-t” on page 272

---

## -qpdf1, -qpdf2

### Category

Optimization and tuning

### @PROCESS


None.

### Purpose

Tunes optimizations through *profile-directed feedback* (PDF), where results from sample program execution are used to improve optimization near conditional branches and in frequently executed code sections.

Optimizes an application for a typical usage scenario based on an analysis of how often branches are taken and blocks of code are run.

## Syntax



## Defaults

`-qnopdf1, -qnopdf2`

## Parameters

### defname

Reverts a PDF file to its default file name.

### exename

Specifies the name of the generated PDF file according to the output file name specified by the `-o` option. For example, you can use `-qpdf1=exename -o func func.f` to generate a PDF file called `.func_pdf`.

### level=0 | 1 | 2

Specifies different levels of profiling information to be generated by the resulting application. The following table shows the type of profiling information supported on each level. The plus sign (+) indicates that the profiling type is supported.

Table 18. Profiling type supported on each `-qpdf1` level

Profiling type	Level		
	0	1	2
Block-counter profiling	+	+	+
Call-counter profiling	+	+	+
Single-pass profiling	+	+	
Value profiling		+	+
Multiple-pass profiling			+
Cache-miss profiling			+

`-qpdf1=level=1` is the default level. It is equivalent to `-qpdf1`. Higher PDF levels profile more optimization opportunities but have a larger overhead.

### Notes:

- Only one application compiled with the `-qpdf1=level=2` option can be run at a time on a particular computer.

- Cache-miss profiling information has several levels. If you want to gather different levels of cache-miss profiling information, set the PDF\_PM\_EVENT environment variable to L1MISS, L2MISS, or L3MISS (if applicable) accordingly. Only one level of cache-miss profiling information can be instrumented at a time. L2 cache-miss is the default level.
- If you want to bind your application to the specified processor for cache-miss profiling, set the PDF\_BIND\_PROCESSOR environment variable. Processor 0 is set by default.

**pdfname=** *file\_path*

Specifies the directories and names for the PDF files and any existing PDF map files. By default, if the PDFDIR environment variable is set, the compiler places the PDF and PDF map files in the directory specified by PDFDIR. Otherwise, if the PDFDIR environment variable is not set, the compiler places these files in the current working directory. If the PDFDIR environment variable is set but the specified directory does not exist, the compiler issues a warning message. The name of the PDF map file follows the name of the PDF file if the **-qpdf1=unique** option is not specified. For example, if you specify the **-qpdf1=pdfname=/home/joe/func** option, the generated PDF file is called func, and the PDF map file is called func\_map. Both of the files are placed in the /home/joe directory. You can use the **pdfname** suboption to do simultaneous runs of multiple executable applications by using the same directory. It is especially useful when tuning with PDF process on dynamic libraries.

**unique** | nounique

You can use the **-qpdf1=unique** option to avoid locking a single PDF file when multiple processes are writing to the same PDF file in the PDF training step. This option specifies whether a unique PDF file is created for each process during run time. The PDF file name is *<pdf\_file\_name>.<pid>*. *<pdf\_file\_name>* is *.\_pdf* by default or specified by other **-qpdf1** suboptions, which include **pdfname**, **exename**, and **defname**. *<pid>* is the ID of running process in the PDF training step. For example, if you specify the **-qpdf1=unique:pdfname=abc** option, and there are two processes for PDF training with the IDs 12345678 and 87654321, two PDF files abc.12345678 and abc.87654321 are generated.

**Note:**

- When **-qpdf1=unique** is specified, only one PDF map file is generated. The default name of the PDF map file is *.\_pdf\_map*.
- When **-qpdf1=unique** is specified, multiple PDF files with process IDs as suffixes are generated. You must use the **mergepdf** program to merge all these PDF files into one after the PDF training step.

## Usage

The PDF process consists of the following three steps:

1. Compile your program with the **-qpdf1** option and a minimum optimization level of **-O2**. A PDF map file named *.\_pdf\_map* by default and a resulting application are generated.
2. Run the resulting application with a typical data set. Profiling information is written to a PDF file named *.\_pdf* by default. This step is called the PDF training step.
3. Recompile and link or relink the program with the **-qpdf2** option and the optimization level used for the **-qpdf1** option. The **-qpdf2** process fine-tunes the optimizations according to the profiling information collected when the resulting application is run.

### Notes:

- The **showpdf** utility uses the PDF map file to display part of the profiling information in text or XML format. For details, see "Viewing profiling information with showpdf" in the *XL Fortran Optimization and Programming Guide*. If you do not need to view the profiling information, specify the **-qnoshowpdf** option during the **-qpdf1** phase so that the PDF map file is not generated. For details of **-qnoshowpdf**, see **-qshowpdf** in the *XL Fortran Compiler Reference*.
- When option **-04**, **-05**, or any level of option **-qipa** is in effect, and you specify the **-qpdf1** or **-qpdf2** option at the link step but not at the compile step, the compiler issues a warning message. The message indicates that you must recompile your program to get all the profiling information.
- When the **-qpdf1=pdfname** option is used during the **-qpdf1** phase, you must use the **-qpdf2=pdfname** option during the **-qpdf2** phase for the compiler to recognize the correct PDF file. This rule also applies to the **-qpdf[1|2]=exename** option.

The compiler issues an information message with a number in the range of 0 - 100 during the **-qpdf2** phase. If you have not changed your program between the **-qpdf1** and **-qpdf2** phases, the number is 100, which means that all the profiling information can be used to optimize the program. If the number is 0, it means that the profiling information is completely outdated, and the compiler cannot take advantage of any information. When the number is less than 100, you can choose to recompile your program with the **-qpdf1** option and regenerate the profiling information.

### Single-pass profiling

Single-pass profiling is supported on level 0 and 1 of the **-qpdf1** phase. If you recompile your program and use either of the **-qpdf1=level=0** or **-qpdf1=level=1** option, the compiler removes the existing PDF file and the possible existing PDF map file before generating a new application.

### Multiple-pass profiling

Multiple-pass profiling is supported on level 2 of the **-qpdf1** phase. After compiling a program with the **-qpdf1=level=2** option when you train the resulting application, you can recompile your program with the **-qpdf1=level=2** option. The profile information gathered previously is used to guide further instrumentation. When you train the resulting application again, the profiling information is written to a new profile file named `._pdf.1` by default. If you repeat this compiling and PDF training several times, the PDF files are generated up to five times (`._pdf.1` to `._pdf.5`). If the compiler detects that all the PDF files names have been used, it issues a warning message and overwrites the last PDF file `._pdf.5`. If the compiler cannot read any PDF files when compiling a program with the **-qpdf1=level=2** option, it issues a warning message to indicate that PDF files are not found. You can get initial profiling information by using the **-qpdf1=level=0** or **-qpdf1=level=1** option, and then use the **-qpdf1=level=2** option for more profiling information.

### Notes:

- If you have not specified the **-qnoshowpdf** option, PDF map files that correspond to the PDF files are also generated, with the default names `._pdf_map`, `._pdf.1_map`, and so on up to `._pdf.5_map`.
- If you use the **-qpdf2=pdfname** option to specify a PDF file, specify a file name that does not end with a numeric suffix from `.1` to `.5`. Otherwise, the compiler

looks for wrong files. For example, if you specify the `-qpdf2=pdfname=func.2` option during the `-qpdf2` phase, the compiler looks for the PDF files named (func.2, func.2.1, func.2.2, func.2.3), which might not exist. If you specify the `-qpdf2=pdfname=func` option without the numeric suffix, the compiler looks for (func, func.1, func.2, func.3).

### Other related options

You can use the following option with the `-qpdf1` option:

#### `-qprefetch`

When you run the `-qprefetch=assistthread` option to generate data prefetching assist threads, the compiler uses the delinquent load information to perform analysis and generate them. The delinquent load information can be gathered from dynamic profiling using the `-qpdf1=level=2` option. For more information, see `-qprefetch`.


#### `-qshowpdf`

Provides additional information to the profile file. See “`-qshowpdf`” on page 223 for more information.

For recommended procedures of using PDF, see "Profile-directed feedback" in the *XL Fortran Optimization and Programming Guide*.

The following utility programs, found in `/opt/ibm/xlf/15.1.1/bin/`, are available for managing the directory to which profiling information is written:

#### `cleanpdf`



Removes all PDF files or the specified PDF files, including PDF files with process ID suffixes. Removing profiling information reduces runtime overhead if you change the program and then go through the PDF process again.

*pdfdir* Specifies the directory that contains the PDF files to be removed. If *pdfdir* is not specified, the directory is set by the PDFDIR environment variable; if PDFDIR is not set, the directory is the current directory.

#### `-f pdfname`

Specifies the name of the PDF file to be removed. When specified, files with the naming convention *pdfname.<multiple\_pass\_profiling\_times>*, if applicable, are also removed. *<multiple\_pass\_profiling\_times>* is a numeric suffix from 1 to 5.

If `-f pdfname` is not specified, `._pdf` and files with the naming convention `._pdf.<multiple_pass_profiling_times>`, if applicable, are removed.

#### `-u`

Removes the PDF file that is specified by *pdfname* and files with the following naming convention when applicable:


- *pdfname.<pid>*, where *<pid>* is the ID of running process in the PDF training step
- *pdfname.<multiple\_pass\_profiling\_times>.<pid>*

If `-f pdfname` is not specified, removes `._pdf` and files with the following naming convention when applicable:

- `._pdf.<pid>`
- `._pdf.<multiple_pass_profiling_times>.<pid>`

Run **cleanpdf** only when you finish the PDF process for a particular application. Otherwise, if you want to resume by using PDF process with that application, you must compile all of the files again with **-qpdf1**.

## mergepdf



Merges two or more PDF files into a single PDF file.

### **-r** *scaling*

Specifies the scaling ratio for the PDF file. This value must be greater than zero and can be either an integer or a floating-point value. If not specified, a ratio of 1.0 is assumed.

*input* Specifies the name of a PDF input file, or a directory that contains PDF files.


### **-o** *output*

Specifies the name of the PDF output file, or a directory to which the merged output is written.

**-n** If specified, PDF files are not normalized. If not specified, **mergepdf** normalizes files based on an internally calculated ratio before applying any user-defined scaling factor.

**-v** Specifies verbose mode, and causes internal and user-specified scaling ratios to be displayed to standard output.

## resetpdf



Same as **cleanpdf**.

## showpdf

Displays part of the profiling information written to PDF and PDF map files. To use this command, you must first compile your program and use the **-qpdf1** option. See "Viewing profiling information with showpdf" in the *XL Fortran Optimization and Programming Guide* for more information.

## Predefined macros

None.

## Examples

The following example uses the **-qpdf1=level=0** option to reduce possible runtime instrumentation overhead:

```
#Compile all the files with -qpdf1=level=0
xlf -qpdf1=level=0 -O3 file1.f file2.f file3.f
```

```
#Run with one set of input data
./a.out < sample.data
```

```
#Recompile all the files with -qpdf2
xlf -qpdf2 -O3 file1.f file2.f file3.f
```

```
#If the sample data is typical, the program
#can now run faster than without the PDF process
```

The following example uses the **-qpdf1=level=1** option:

```
#Compile all the files with -qpdf1
xlf -qpdf1 -O3 file1.f file2.f file3.f
```

```
#Run with one set of input data
./a.out < sample.data
```

```
#Recompile all the files with -qpdf2
xlf -qpdf2 -O3 file1.f file2.f file3.f
```

```
#If the sample data is typical, the program
#can now run faster than without the PDF process
```

The following example uses the **-qpdf1=level=2** option to gather cache-miss profiling information:

```
#Compile all the files with -qpdf1=level=2
xlf -qpdf1=level=2 -O3 file1.f file2.f file3.f
```

```
#Set PM_EVENT=L2MISS to gather L2 cache-miss profiling
#information
export PDF_PM_EVENT=L2MISS
```

```
#Run with one set of input data
./a.out < sample.data
```

```
#Recompile all the files with -qpdf2
xlf -qpdf2 -O3 file1.f file2.f file3.f
```

```
#If the sample data is typical, the program
#can now run faster than without the PDF process
```

The following example uses the **-qpdf1=level=2** option with multiple runs to gather cache-miss profiling information at different cache levels:

```
#Compile all the files with -qpdf1=level=2
xlf -qpdf1=level=2 -O3 file1.f file2.f file3.f
```

```
#Set PM_EVENT=L1MISS to gather L1 cache-miss profiling
#information
export PDF_PM_EVENT=L1MISS
```

```
#Run with one set of input data
./a.out < sample.data
```

```
#Set PM_EVENT=L2MISS to gather L2 cache-miss profiling
#information
export PDF_PM_EVENT=L2MISS
```

```
#Run with one set of input data
./a.out < sample.data
```

```
#Recompile all the files with -qpdf2
xlf -qpdf2 -O3 file1.f file2.f file3.f
```

```
#If the sample data is typical, the program
#can now run faster than without the PDF process
```

The following example demonstrates the process of multiple-pass profiling:

```
#Compile all the files with -qpdf1=level=2. The static profiling
#information is recorded in a file named ._pdf_map by default
xlf -qpdf1=level=2 -O3 file1.f file2.f file3.
```

```
#Run with one set of input data, the profiling information
#is recorded in a file named ._pdf by default
./a.out < sample.data
```

```
#Recompile all the files with -qpdf1=level=2 again
#The compiler reads the previous profiling information, refines
#instrumentation, and generates a new instrumented
#executable. The static profiling information
#is recorded in ._pdf.1_map
xlf -qpdf1=level=2 -O3 file1.f file2.f file3.f
```

```
#Run it again, the profiling information is recorded in
#._pdf.1
./a.out < sample.data
```

```
#Recompile all the files with -qpdf2
xlf -qpdf2 -O3 file1.f file2.f file3.f
```

```
#If the sample data is typical, the program
#can now run faster than without the PDF process
```

The following example demonstrates the use of the PDF\_BIND\_PROCESSOR environment variable:

```
#Compile all the files with -qpdf1=level=1
xlf -qpdf1=level=1 -O3 file1.f file2.f file3.
```

```
#Set PDF_BIND_PROCESSOR environment variable so that
#all processes for this executable are run on Processor 1
export PDF_BIND_PROCESSOR=1
```

```
#Run executable with sample input data
./a.out < sample.data
```

```
#Recompile all the files with -qpdf2
xlf -qpdf2 -O3 file1.f file2.f file3.f
```

```
#If the sample data is typical, the program
#can now run faster than without the PDF process
```

The following example demonstrates the use of the **-qpdf[1|2]=exename** option:

```
#Compile all the files with -qpdf1=exename
xlf -qpdf1=exename -O3 -o final file1.f file2.f file3.f
```

```
#Run executable with sample input data
./final < typical.data
```

```
#List the content of the directory
>ls -lrta
```

```
-rw-r--r-- 1 user staff 50 Dec 05 13:18 file1.f
-rw-r--r-- 1 user staff 50 Dec 05 13:18 file2.f
-rw-r--r-- 1 user staff 50 Dec 05 13:18 file3.f
-rwxr-xr-x 1 user staff 12243 Dec 05 17:00 final
-rwxr-Sr-- 1 user staff 762 Dec 05 17:03 .final_pdf
```

```
#Recompile all the files with -qpdf2=exename
```

```
xlf -qpdf2=exename -O3 -o final file1.f file2.f file3.f
```

```
#The program is now optimized using PDF information
```

The following example demonstrates the use of the **-qpdf[1|2]=pdfname** option:

```
#Compile all the files with -qpdf1=pdfname. The static profiling
```

```
#information is recorded in a file named final_map
```

```
xlf -qpdf1=pdfname=final -O3 file1.f file2.f file3.f
```

```
#Run executable with sample input data. The profiling
```

```
#information is recorded in a file named final
```

```
./a.out < typical.data
```

```
#List the content of the directory
```

```
>ls -lrta
```

```
-rw-r--r-- 1 user staff 50 Dec 05 13:18 file1.f
```

```
-rw-r--r-- 1 user staff 50 Dec 05 13:18 file2.f
```

```
-rw-r--r-- 1 user staff 50 Dec 05 13:18 file3.f
```

```
-rwxr-xr-x 1 user staff 12243 Dec 05 18:30 a.out
```

```
-rwxr-Sr-- 1 user staff 762 Dec 05 18:32 final
```

```
#Recompile all the files with -qpdf2=pdfname
```

```
xlf -qpdf2=pdfname=final -O3 file1.f file2.f file3.f
```

```
#The program is now optimized using PDF information
```

## Related information

- “-qshowpdf” on page 223
- “-qipa” on page 163
- -qprefetch
- “-qreport” on page 216
- “XL Fortran input files” on page 23
- “XL Fortran output files” on page 24
- “Profile-directed feedback” in the *XL Fortran Optimization and Programming Guide*
- Correct settings for environment variables

---

## -qphsinfo

### Category

Listings, messages, and compiler information

### Purpose

Reports the time taken in each compilation phase to standard output.

### Syntax

```
►► -q nophsinfo  
phsinfo ◀◀
```

@PROCESS:

@PROCESS PHSINFO | NOPHSINFO

### Defaults

-qnophsinfo

## Usage

The output takes the form *number1/number2* for each phase where *number1* represents the CPU time used by the compiler and *number2* represents the total of the compile time and the time that the CPU spends handling system calls.

The time reported by `-qphsinfo` is in seconds.

## Examples

To compile `app.f`, which consists of 3 compilation units, and report the time taken for each phase of the compilation, enter:

```
xlf90 app.f -qphsinfo
```

The output looks similar to:

```
FORTRAN phase 1 ftphas1      TIME = 0.000 / 0.000
** m_module   === End of Compilation 1 ===
FORTRAN phase 1 ftphas1      TIME = 0.000 / 0.000
** testassign === End of Compilation 2 ===
FORTRAN phase 1 ftphas1      TIME = 0.000 / 0.010
** dataassign  === End of Compilation 3 ===
HOT           - Phase Ends; 0.000/ 0.000
HOT           - Phase Ends; 0.000/ 0.000
HOT           - Phase Ends; 0.000/ 0.000
W-TRANS       - Phase Ends; 0.000/ 0.010
OPTIMIZ       - Phase Ends; 0.000/ 0.000
REGALLO       - Phase Ends; 0.000/ 0.000
AS            - Phase Ends; 0.000/ 0.000
W-TRANS       - Phase Ends; 0.000/ 0.000
OPTIMIZ       - Phase Ends; 0.000/ 0.000
REGALLO       - Phase Ends; 0.000/ 0.000
AS            - Phase Ends; 0.000/ 0.000
W-TRANS       - Phase Ends; 0.000/ 0.000
OPTIMIZ       - Phase Ends; 0.000/ 0.000
REGALLO       - Phase Ends; 0.000/ 0.000
AS            - Phase Ends; 0.000/ 0.000
1501-510     Compilation successful for file app.f.
```

Each phase is invoked three times, corresponding to each compilation unit. FORTRAN represents front-end parsing and semantic analysis, HOT loop transformations, W-TRANS intermediate language translation, OPTIMIZ high-level optimization, REGALLO register allocation and low-level optimization, and AS final assembly.

Compile `app.f` at the `-O4` optimization level with `-qphsinfo` specified:

```
xlf90 myprogram.f -qphsinfo -O4
```

The following output results:

```
FORTRAN phase 1 ftphas1      TIME = 0.010 / 0.020
** m_module   === End of Compilation 1 ===
FORTRAN phase 1 ftphas1      TIME = 0.000 / 0.000
** testassign === End of Compilation 2 ===
FORTRAN phase 1 ftphas1      TIME = 0.000 / 0.000
** dataassign  === End of Compilation 3 ===
HOT           - Phase Ends; 0.000/ 0.000
HOT           - Phase Ends; 0.000/ 0.000
HOT           - Phase Ends; 0.000/ 0.000
IPA           - Phase Ends; 0.080/ 0.100
1501-510     Compilation successful for file app.f.
IPA           - Phase Ends; 0.050/ 0.070
```

W-TRANS	- Phase Ends;	0.010/	0.030
OPTIMIZ	- Phase Ends;	0.020/	0.020
REGALLO	- Phase Ends;	0.040/	0.040
AS	- Phase Ends;	0.000/	0.000

Note that during the IPA (interprocedural analysis) optimization phases, the program has resulted in one compilation unit; that is, all procedures have been inlined.

## Related information

.The compiler phases in the *Getting Started with XL Fortran*

## -qpic

### Category

Object code control

### @PROCESS

None.

### Purpose

Generates position-independent code required for use in shared libraries.

### Syntax

►► -q nopic  
pic ◀◀

### Defaults

- -qnopic

### Usage

When **-qpic** is in effect, the compiler generates two instructions for each TOC access to enlarge the accessing range. This helps avoid TOC overflow conditions when the Table of Contents is larger than 64 Kb.

You must specify **-qpic** when you build shared libraries.

When you specify **-qpic -qtls**, thread local storage (TLS) symbols are not affected by **-qpic**.

You can use different TOC access options for different compilation units in an application.

## -qport

### Category

Portability and migration



parameters. The result will be of the same basic type as the argument, but with the larger kind type parameter value.

### **nullarg** | **nonnullarg**

For an external or internal procedure reference, specifying **nullarg** causes the compiler to treat an empty argument, which is delimited by a left parenthesis and a comma, two commas, or a comma and a right parenthesis, as a null argument. This suboption has no effect if the argument list is empty.

Examples of empty arguments are:

```
call foo(,,z)
```

```
call foo(x,,z)
```

```
call foo(x,y,)
```

The following program includes a null argument.

#### **Fortran program:**

```
program nularg
real(4) res/0.0/
integer(4) rc
integer(4), external :: add
rc = add(%val(2), res, 3.14, 2.18,) ! The last argument is a
                                ! null argument.

if (rc == 0) then
print *, "res = ", res
else
print *, "number of arguments is invalid."
endif
end program
```

#### **C program:**

```
int add(int a, float *res, float *b, float *c, float *d)
{
    int ret = 0;
    if (a == 2)
        *res = *b + *c;
    else if (a == 3)
        *res = (*b + *c + *d);
    else
        ret = 1;
    return (ret);
}
```

### **sce** | **nosce**

By default, the compiler performs short circuit evaluation in selected logical expressions using XL Fortran rules. Specifying **sce** allows the compiler to use non-XL Fortran rules. The compiler will perform short circuit evaluation if the current rules allow it.

### **typestmt** | **notypestmt**

The `TYPE` statement, which behaves in a manner similar to the `PRINT` statement, is supported whenever **typestmt** is specified.

### **typlessarg** | **notyplessarg**

Converts all typeless constants to default integers if the constants are actual arguments to an intrinsic procedure whose associated dummy arguments are of integer type. Typeless actual arguments associated with dummy arguments of noninteger type remain unaffected by this option.

Using this option may cause some intrinsic procedures to become mismatched in kinds. Specify `-qxlf77=intarg` to convert the kind to that of the longest argument.

### Related information

- “-qintlog” on page 160
- “-qxlf77” on page 259
- See the section on the *INT* and *MOD* intrinsic functions in the *XL Fortran Language Reference* for further information.

---

## -qposition

### Category

Language element control

### Purpose

Positions the file pointer at the end of the file when data is written after an **OPEN** statement with no **POSITION=** specifier and the corresponding **STATUS=** value (**OLD** or **UNKNOWN**) is specified.

The position becomes **APPEND** when the first I/O operation moves the file pointer if that I/O operation is a **WRITE** or **PRINT** statement. If it is a **BACKSPACE**, **ENDFILE**, **READ**, or **REWIND** statement instead, the position is **REWIND**.

### Syntax

►► -qposition=  
└─ appendold ─┘  
└─ appendunknown ┘

### @PROCESS:

@PROCESS POSITION({APPENDOLD | APPENDUNKNOWN} ...)

### Defaults

The default setting depends on the I/O specifiers in the **OPEN** statement and on the compiler invocation command:

- **-qposition=appendold** for the **xlf** and **xlf\_r** commands when they are used to compile the **.f**, **.F**, **.f77**, or **.F77** files
- **-qposition=appendold** for the **f77** and **fort77** commands
- The defined Fortran 90, Fortran 95, Fortran 2003, and Fortran 2008 behaviors for the **xlf90**, **f90**, **xlf90\_r**, **xlf95**, **f95**, **xlf95\_r**, **xlf2003**, **f2003**, **xlf2003\_r**, **xlf2008**, **f2008**, and **xlf2008\_r** commands

### Examples

In the following example, **OPEN** statements that do not specify a **POSITION=** specifier, but specify **STATUS='old'** will open the file as if **POSITION='append'** was specified.

```
xlf95 -qposition=appendold opens_old_files.f
```

In the following example, **OPEN** statements that do not specify a **POSITION=** specifier, but specify **STATUS='unknown'** will open the file as if **POSITION='append'** was specified.

```
xlf95 -qposition=appendunknown opens_unknown_files.f
```

In the following example, **OPEN** statements that do not specify a **POSITION=** specifier, but specify either **STATUS='old'** or **STATUS='unknown'** will open the file as if **POSITION='append'** was specified.

```
xlf95 -qposition=appendold:appendunknown opens_many_files.f
```

### Related information

- *File positioning* in the *XL Fortran Optimization and Programming Guide*
- *OPEN statement* in the *XL Fortran Language Reference*

---

## -qppsuborigarg

### Category

Input control

### @PROCESS

None.

### Purpose

Instructs the C preprocessor to substitute original macro arguments before further macro expansion.

### Syntax

```
→ -WF, -q [nopppsuborigarg] pppsuborigarg →
```

### Defaults

- -WF, -qnopppsuborigarg

### Usage

**-qppsuborigarg** is a C preprocessor option, and must therefore be specified using the **-WF** option.

### Examples

Consider the following sample code, x.F:

```
#define PRINT_COMP(a) PRINT_4(SPLIT_COMP(a))
#define SPLIT_COMP(a) "Real:", real(a), "Imag:", imag(a)
#define PRINT_4(list) PRINT_LIST(list)
#define PRINT_LIST(list) print *, list
```

```
complex a
a = (3.5, -3.5)
PRINT_COMP(a)
end
```

If this code is compiled with **-qnoppsuborigarg**, the C preprocessor reports an error because the parameter "list" in the function-like macro PRINT\_4 is the expanded substitution text of the macro SPLIT\_COMP(a). The C preprocessor therefore complains because PRINT\_LIST is being called with four arguments but only expects one.

```
> xlf95 x.F -d
"x.F", line 8.1: 1506-215 (E) Too many arguments specified for macro PRINT_LIST.
** _main    === End of Compilation 1 ===
1501-510  Compilation successful for file x.F.
> cat Fx.f
```

```
complex a
a = (3.5, -3.5)
print *, "Real:"
end
```

When the code is compiled with **-qppsuborigarg**, the C preprocessor uses the text "SPLIT\_COMP(a)" rather than the expanded substitution text of SPLIT\_COMP(a) as the argument to the function-like macro PRINT\_LIST. Only after the macro PRINT\_LIST has been expanded, does the C preprocessor proceed to expand the macro "SPLIT\_COMP(a)". As a result, the macro PRINT\_LIST only receives the expected single argument "SPLIT\_COMP(a)" rather than the four arguments.

```
> xlf95 x.F -d -WF,-qppsuborigarg
** _main    === End of Compilation 1 ===
1501-510  Compilation successful for file x.F.
> cat Fx.f
```

```
complex a
a = (3.5, -3.5)
print *, "Real:", real(a), "Imag:", imag(a)
end
```

### Related information

- "-W" on page 276
- "-qfpp" on page 131
- "Passing Fortran files through the C preprocessor" on page 29

---

## -qprefetch

### Category

Optimization and tuning

### @PROCESS


None.

### Purpose

Inserts prefetch instructions automatically where there are opportunities to improve code performance.

When **-qprefetch** is in effect, the compiler may insert prefetch instructions in compiled code. When **-qnoprefetch** is in effect, prefetch instructions are not inserted in compiled code.

## Syntax



## Defaults

**-qprefetch=noassistthread:noaggressive:dscr=0**

## Parameters

### **assistthread** | **noassistthread**

When you work with applications that generate a high cache-miss rate, you can use **-qprefetch=assistthread** to exploit assist threads for data prefetching. This suboption guides the compiler to exploit assist threads at optimization level **-O3 -qhot** or higher. If you do not specify **-qprefetch=assistthread**, **-qprefetch=noassistthread** is implied.

#### **CMP**

For systems based on the chip multi-processor architecture (CMP), you can use **-qprefetch=assistthread=cmp**.

#### **SMT**

For systems based on the simultaneous multi-threading architecture (SMT), you can use **-qprefetch=assistthread=smt**.

**Note:** If you do not specify either CMP or SMT, the compiler uses the default setting based on your system architecture.

### **aggressive** | **noaggressive**

This suboption guides the compiler to generate aggressive data prefetching at optimization level **-O3** or higher. If you do not specify **aggressive**, **-qprefetch=noaggressive** is implied.

#### **dscr**

You can specify a value for the dscr suboption to improve the runtime performance of your applications. The compiler sets the Data Stream Control Register (DSCR) to the specified dscr value to control the hardware prefetch engine. The value is valid only when **-qarch=pwr8** is in effect and the optimization level is **-O2** or greater. The default value of dscr is 0.

#### **value**

The value that you specify for dscr must be 0 or greater, and representable as a 64-bit unsigned integer. Otherwise, the compiler issues a warning message and sets dscr to 0. The compiler accepts both decimal and hexadecimal numbers, and a hexadecimal number requires the prefix of 0x. The value range depends on your system architecture. See the product information about the POWER Architecture for details. If you specify multiple dscr values, the last one takes effect.

## Usage

The **-qnoprefetch** option does not prevent built-in functions such as **\_\_prefetch\_by\_stream** from generating prefetch instructions.

When you run **-qprefetch=assistthread**, the compiler uses the delinquent load information to perform analysis and generates prefetching assist threads. The delinquent load information can either be provided through the built-in **\_\_mem\_delay** function (const void \*delinquent\_load\_address, const unsigned int delay\_cycles), or gathered from dynamic profiling using **-qpdf1=level=2**.

When you use **-qpdf** to call **-qprefetch=assistthread**, you must use the traditional two-step PDF invocation:

1. Run **-qpdf1=level=2**
2. Run **-qpdf2 -qprefetch=assistthread**

## Example

```
DO i = 1, 1000

!IBM* MEM_DELAY(x(i), 10)
x(i) = x(i) + 1

END DO
```

## Examples

```
DO I = 1, 1000

!IBM* MEM_DELAY(X(I), 10)

X(I) = X(I) + 1

END DO
```

## Related information

- **-qarch**
- “-qhot” on page 143
- “-qpdf1, -qpdf2” on page 195
- “-qreport” on page 216
- **MEM\_DELAY** section in the *XL Fortran Language Reference*

---

## -qqcount

### Category

Language element control

### Purpose

Accepts the **Q** character-count edit descriptor (**Q**) as well as the extended-precision **Q** edit descriptor (**Qw.d**).

### Syntax

→ -q noqcount qcount →

@PROCESS:

@PROCESS QCOUNT | NOQCOUNT

## Defaults

With **-qnoqcount**, all **Q** edit descriptors are interpreted as the extended-precision **Q** edit descriptor.

## Usage

The compiler interprets a **Q** edit descriptor as one or the other depending on its syntax and issues a warning if it cannot determine which one is specified.

## Related information

- *Q (Character Count) Editing* in the *XL Fortran Language Reference*

---

## -qrealsize

### Category

Floating-point and integer control

### Purpose

Sets the default size of **REAL**, **DOUBLE PRECISION**, **COMPLEX**, and **DOUBLE COMPLEX** values.

This option is intended for maintaining compatibility with code that is written for other systems. You may find it a useful alternative to the **-qautodbl** option in some situations.

### Syntax

►► -qrealsize=4  
8►►

@PROCESS:

@PROCESS REALSIZE(*bytes*)

### Defaults

**-qrealsize=4**

### Parameters

The allowed values for *bytes* are:

- 4
- 8

### Usage

This option is intended to allow you to port programs unchanged from systems that have different default sizes for data. For example, you might need **-qrealsize=8** for programs that are written for a CRAY computer.

**-qautodbl** is related to **-qreal size**, although you cannot combine these options. When the **-qautodbl** option turns on automatic doubling, padding, or both, the **-qreal size** option has no effect.

Setting **-qreal size** to 8 overrides the setting of the **-qdp c** option.

In addition to type **REAL**, **-qreal size** also works for type **vector(real)**.

## Results

The option affects the sizes<sup>2</sup> of constants, variables, derived type components, and functions (which include intrinsic functions) for which no kind type parameter is specified. Objects that are declared with a kind type parameter or length, such as **REAL(4)** or **COMPLEX\*16**, are not affected.

This option determines the sizes of affected objects as follows:

Data Object	REALSIZE(4) in Effect	REALSIZE(8) in Effect
1.2	REAL(4)	REAL(8)
1.2e0	REAL(4)	REAL(8)
1.2d0	REAL(8)	REAL(16)
1.2q0	REAL(16)	REAL(16)
REAL	REAL(4)	REAL(8)
DOUBLE PRECISION	REAL(8)	REAL(16)
COMPLEX	COMPLEX(4)	COMPLEX(8)
DOUBLE COMPLEX	COMPLEX(8)	COMPLEX(16)

Similar rules apply to intrinsic functions:

- If an intrinsic function has no type declaration, its argument and return types may be changed by the **-qreal size** setting.
- Any type declaration for an intrinsic function must agree with the default size of the return value.

## Examples

This example shows how changing the **-qreal size** setting transforms some typical entities:

```
@PROCESS REALSIZE(8)
      REAL R                ! treated as a real(8)
      REAL(8) R8           ! treated as a real(8)
      VECTOR(REAL)        ! treated as a vector(real(8))
      VECTOR(REAL(4))     ! treated as a vector(real(4))
      DOUBLE PRECISION DP ! treated as a real(16)
      DOUBLE COMPLEX DC   ! treated as a complex(16)
      COMPLEX(4) C        ! treated as a complex(4)
      PRINT *,DSIN(DP)    ! treated as qsin(real(16))
! Note: we cannot get dsin(r8) because dsin is being treated as qsin.
END
```

Specifying **-qreal size=8** affects intrinsic functions, such as **DABS**, as follows:

```
INTRINSIC DABS           ! Argument and return type become REAL(16).
DOUBLE PRECISION DABS   ! OK, because DOUBLE PRECISION = REAL(16)
                        ! with -qreal size=8 in effect.
REAL(16) DABS           ! OK, the declaration agrees with the option setting.
REAL(8) DABS            ! The declaration does not agree with the option
                        ! setting and is ignored.
```

2. In Fortran 90/95 terminology, these values are referred to as *kind type parameters*.

## Related information

- “-qintsize” on page 161 is a similar option that affects integer and logical objects.
- “-qautodbl” on page 97
- *Type declaration: type parameters and specifiers* in the *XL Fortran Language Reference*

---

## -qrecur

### Purpose

Specifies whether external subprograms may be called recursively.

**Not recommended.**

### Syntax

►► -q norecur  
recur ◀◀

**@PROCESS:**

@PROCESS RECUR | **NORECUR**

### Defaults

-qnorecur

### Usage

For new programs, use the **RECURSIVE** keyword, which provides a standards-conforming way of using recursive procedures.

If you specify the **-qrecur** option, the compiler must assume that any procedure could be recursive. Code generation for recursive procedures may be less efficient. With the **RECURSIVE** keyword, you can specify exactly which procedures are recursive.

When you use the following commands to compile programs that contain recursive calls, specify **-qnosave** to make the default storage class automatic:

- For `.f`, `.F`, `.f77` and `.F77` files: **xlf** and **xlf\_r**
- For any source files: **f77** and **fort77**

### Examples

! The following RECUR recursive function:

```
@process recur
function factorial (n)
integer factorial
if (n .eq. 0) then
    factorial = 1
else
    factorial = n * factorial (n-1)
end if
end function factorial
```

! can be rewritten to use F90/F95 RECURSIVE/RESULT features:

```
recursive function factorial (n) result (res)
integer res
```

```

if (n .eq. 0) then
  res = 1
else
  res = n * factorial (n-1)
end if
end function factorial

```

## -qreport

### Category

Listings, messages, and compiler information

### Purpose


Produces listing files that show how sections of code have been optimized.

A listing file is generated with a .lst suffix for each source file named on the command line. When used with an option that enables automatic parallelization or vectorization, the listing file shows a pseudo-Fortran code listing and a summary of how program loops are parallelized or optimized. The report also includes diagnostic information to show why specific loops could not be parallelized or vectorized. For instance, when **-qreport** is used with **-qsimd=auto**, messages are provided to identify non-stride-one references that can prevent loop vectorization.

The compiler also reports the number of streams created for a given loop, which include both load and store streams. This information is included in the Loop Transformation section of the listing file. You can use this information to understand your application code and to tune your code for better performance. For example, you can distribute a loop which has more streams than the number supported by the underlying architecture. The POWER8 processors support both load and store stream prefetch.

### Syntax

#### Option:



#### @PROCESS:

@PROCESS REPORT[({SMPLIST |HOTLIST}...)] | NOREPORT

### Defaults

-qnoreport

### Parameters

smplist | hotlist

When **-qreport=smplist** is in effect, produces a pseudo-Fortran listing that shows how the program is parallelized. This listing is produced before loop and other optimizations are performed. It includes messages that

point out places in the program that can be modified to be more efficient. This report is only produced if **-qsmp** is in effect.

When **-qreport=hotlist** is in effect, produces a pseudo-Fortran listing that shows how loops are transformed, to assist you in tuning the performance of all loops. This report is only produced if **-qhot** is in effect.

In addition, if you specify the **-qreport=hotlist** option when **-qsmp** is in effect, a pseudo-Fortran listing will be produced that shows the calls to the SMP runtime library and the procedures created for parallel constructs.

Specifying **-qreport** with no suboptions is equivalent to **-qreport=hotlist**.

## Usage

For **-qreport** to generate a loop transformation listing, you must also specify one of the following options on the command line:

- **-qipa=level=2**

For **-qreport** to generate PDF information in the listing, you must specify the following option in the command line:

- **-qpdf2 -qreport**

For **-qreport** to generate a parallel transformation listing or parallel performance messages, you must also specify one of the following options on the command line:

- **-qsmp**
- **-05**
- **-qipa=level=2**

To generate data reorganization information, specify the optimization level **-qipa=level=2** or **-05** together with **-qreport**. Reorganizations include common block splitting, array splitting, array transposing, memory allocation merging, array interleaving, and array coalescing.

To generate information about data prefetch insertion locations, use the optimization level of **-qhot**, or any other option that implies **-qhot** together with **-qreport**. This information appears in the LOOP TRANSFORMATION SECTION of the listing file. In addition, when you use **-qprefetch=assistthread** to generate prefetching assist threads, the message: Assist thread for data prefetching was generated also appears in the LOOP TRANSFORMATION SECTION of the listing file.

To generate a list of aggressive loop transformations and parallelization performed on loop nests in the LOOP TRANSFORMATION SECTION of the listing file, use the optimization level of **-qhot=level=2** and **-qsmp** together with **-qreport**.

The pseudo-Fortran code listing is not intended to be compilable. Do not include any of the pseudo-Fortran code in your program, and do not explicitly call any of the internal routines whose names may appear in the pseudo-Fortran code listing.

## Examples

To compile `myprogram.f` so the compiler listing includes a report showing how loops are optimized, enter:

```
xlf -qhot -03 -qreport myprogram.f
```

To compile `myprogram.f` so the compiler listing also includes a report showing how parallelized loops are transformed, enter:

```
xlf_r -qhot -qsmp -qreport=smp1ist myprogram.f
```

### Related information

- “-qhot” on page 143
- “-qsimd” on page 225
- “-qipa” on page 163
- “-qsmp” on page 227

---

## -qsaa

### Category

Language element control

### Purpose

Checks for conformance to the SAA FORTRAN language definition. It identifies nonconforming source code and also options that allow such nonconformances.

### Syntax

►► -q nosaa  
saa ◀◀

@PROCESS:

@PROCESS SAA | NOSAA

### Defaults

-qnosaa

### Usage

The **-qflag** option can override this option.

Use the **-qlanglvl** option to check your code for conformance to international standards.

### Results

Warnings have a prefix of **(L)**, indicating a problem with the language level.

### Related information

- “-qflag” on page 126
- “-qlanglvl” on page 169

---

## -qsave


### Category

Language element control

## Purpose

Specifies the default storage class for local variables.

## Syntax



**@PROCESS:**

```
@PROCESS SAVE[({ALL | DEFAULTINIT})] | NOSAVE
```

## Defaults

The default for this option depends on the invocation command used:

- When **xlf** is used to compile the **.f**, **.F**, **.f77**, or **.F77** files, the default is **-qsave=all**.
- For the **f77** and **fort77** invocation commands, the default is **-qsave=all**.
- For all the other invocation commands, the default is **-qnosave**.

## Parameters

The **-qsave** suboptions include:

### all

The default storage class is **STATIC**.

### **defaultinit**

The default storage class is **STATIC** for variables of derived type that have default initialization specified, and **AUTOMATIC** otherwise.

The **all** and **defaultinit** suboptions are mutually exclusive.

## Usage

The **-qnosave** option sets the default storage class to **AUTOMATIC**. This usage is usually necessary for multithreaded applications and subprograms that are compiled with the **-qrecur** option.

You can specify the **-qsave** option to duplicate the behavior of FORTRAN 77 programs. The **xlf**, **f77** and **fort77** commands have **-qsave** listed as a default option in the configuration file to preserve the previous behavior. The default configuration file path is `/opt/ibm/xlf/15.1.1/etc/xlf.cfg`

## Examples

The following example illustrates the impact of the **-qsave** option on derived data type:

```
PROGRAM P
  CALL SUB
  CALL SUB
END PROGRAM P

SUBROUTINE SUB
```

```

LOGICAL, SAVE :: FIRST_TIME = .TRUE.
STRUCTURE /S/
  INTEGER I/17/
END STRUCTURE
RECORD /S/ LOCAL_STRUCT
INTEGER LOCAL_VAR

IF (FIRST_TIME) THEN
  LOCAL_STRUCT.I = 13
  LOCAL_VAR = 19
  FIRST_TIME = .FALSE.
ELSE
  ! Prints " 13" if compiled with -qsave or -qsave=all
  ! Prints " 13" if compiled with -qsave=defaultinit
  ! Prints " 17" if compiled with -qnosave
  PRINT *, LOCAL_STRUCT
  ! Prints " 19" if compiled with -qsave or -qsave=all
  ! Value of LOCAL_VAR is undefined otherwise
  PRINT *, LOCAL_VAR
END IF
END SUBROUTINE SUB

```

### Related information

- “-qrecur” on page 215
- See *Storage classes for variables* in the *XL Fortran Language Reference* for information on how this option affects the storage class of variables.

## -qsaveopt

### Category

Object code control

### @PROCESS

None.

### Purpose

Saves the command-line options used for compiling a source file, the user's configuration file name and the options specified in the configuration files, the version and level of each compiler component invoked during compilation, and other information to the corresponding object file.

### Syntax

```

▶▶ -q nosaveopt
saveopt

```

### Defaults

-qnosaveopt

## Usage

This option has effect only when compiling to an object (.o) file (that is, using the `-c` option). Though each object might contain multiple compilation units, only one copy of the command-line options is saved. Compiler options specified with `@PROCESS` directives are ignored.

Command-line compiler options information is copied as a string into the object file, using the following format:

```

▶▶ @(#) opt { f | c | C } invocation options

```

```

▶▶ @(#) cfg config_file_options_list

```

```

▶▶ @(#) env env_var_definition

```

where:

- f** Signifies a Fortran language compilation.
- c** Signifies a C language compilation.
- C** Signifies a C++ language compilation.

*invocation*

Shows the command used for the compilation, for example, `xlfc`.

*options* The list of command line options specified on the command line, with individual options separated by space.

*config\_file\_options\_list*

The list of options specified by the **options** attribute in all configuration files that take effect in the compilation, separated by space.

*env\_var\_definition*

The environment variables that are used by the compiler. Currently only **XLF\_USR\_CONFIG** is listed.

**Note:** You can always use this option, but the corresponding information is only generated when the environment variable **XLF\_USR\_CONFIG** is set.

For more information about the environment variable **XLF\_USR\_CONFIG**, see **XLF\_USR\_CONFIG**.

**Note:** The string of the command-line options is truncated after 64k bytes.

Compiler version and release information, as well as the version and level of each component invoked during compilation, are also saved to the object file in the format:

```

▶▶ @(#) version Version :--VV.RR.MMMM.LLLL
      component_name Version :--VV.RR--(product_name)--Level--:--YYMMDD--:--component_level_ID

```

where:

- V** Represents the version.
- R** Represents the release.
- M** Represents the modification.
- L** Represents the level.

*component\_name*

Specifies the components that were invoked for this compilation, such as the low-level optimizer.

*product\_name*

Indicates the product to which the component belongs (for example, C/C++ or Fortran).

*YYMMDD*

Represents the year, month, and date of the installed update. If the update installed is at the base level, the level is displayed as BASE.

*component\_level\_ID*

Represents the ID associated with the level of the installed component.

If you want to simply output this information to standard output without writing it to the object file, use the **-qversion** option.

## Examples

Compile `t.f` with the following command:

```
xlf t.f -c -qsaveopt -qhot
```

Issuing the **strings -a** command on the resulting `t.o` object file produces information similar to the following:

```
opt f /opt/ibm/xlf/15.1.1/bin/xlf t.f -c -qsaveopt -qhot
cfg -qnozerosize -qsave -qalias=intptr -qposition=appendold
-qxlf90=noautodealloc:nosignedzero:oldpad
-qxlf77=intarg:intxor:persistent:noleadzero:gedit77:noblankpad:oldboz:softeof
-qxlf2003=nopolymorphic:nobozlitargs:nostopexcept:novolatile:noautorealloc:oldnaninf -bh:4
version IBM XL Fortran for Linux, V15.1.1
version Version: 15.01.0000.0000
version Driver Version: 15.01(Fortran) Level: 121020 ID: _acSDAheyEeK928eKYVtYGg
version Fortran Front End and Run Time Version: 15.01(Fortran) Level: 121020 ID: _01piYhmQEeK928eKYVtYGg
version Fortran Transformer Version: 15.01(Fortran) Level: 121021 ID: _gYSYgRpREeK928eKYVtYGg
version High-Level Optimizer Version: 13.01(C/C++) and 15.01(Fortran) Level: 121021 ID: _d6ZMohn3EeK928eKYVtYGg
version Low-Level Optimizer Version: 13.01(C/C++) and 15.01(Fortran) Level: 121020 ID: _H1wQwhn3EeK928eKYVtYGg
```

In the first line, `t.f` identifies the source used as Fortran, `bin/xlf` shows the invocation command used, and `-qhot -qsaveopt` shows the compilation options. The second line, which starts with `cfg`, shows the compiler options added by the configuration file.

The remaining lines list each compiler component invoked during compilation, and its version and level. Components that are shared by multiple products may show more than one version number. Level numbers shown may change depending on the updates you have installed on your system.

## Related information

- “-qversion” on page 255
- COMPILER\_OPTIONS

---

## -qsck

### Category

Language element control

### @PROCESS

None.



application in text or XML format. For details of the **showpdf** utility, see "Viewing profiling information with showpdf" in the *XL Fortran Optimization and Programming Guide*.

If you do not need to view the profiling information, specify the **-qnoshowpdf** option during the **-qpdf1** phase so that the PDF map file is not generated. This can reduce your compile time.

---

## **-qsigtrap**

### **Category**

Error checking and debugging

### **@PROCESS**

None.

### **Purpose**

Sets up the specified trap handler to catch **SIGTRAP** and **SIGFPE** exceptions when compiling a file that contains a main program.

This option enables you to install a handler for **SIGTRAP** or **SIGFPE** signals without calling the **SIGNAL** subprogram in the program.

### **Syntax**

►► -q-sigtrap =trap\_handler ◀◀

### **Defaults**

Not applicable.

### **Usage**

If you specify the **-qsigtrap** option without a handler name, the `xl__trce` trap handler is used by default. To use a different trap handler, specify its name with the **-qsigtrap** option.

If you specify a different handler, ensure that the object module that contains it is linked with the program. To show more detailed information in the tracebacks generated by the trap handlers provided by XL Fortran (such as `xl__trce`), specify the **-qlinedebug** or **-g** option.

### **Related information**

- "XL Fortran runtime exceptions" on page 45 describes the possible causes of exceptions.
- *Detecting and trapping floating-point exceptions in the XL Fortran Optimization and Programming Guide* describes a number of methods for dealing with exceptions that result from floating-point computations.
- *Installing an exception handler in the XL Fortran Optimization and Programming Guide* lists the exception handlers that XL Fortran supplies.

---

## -qsimd

### Category

Optimization and tuning

### @PROCESS

None.

### Purpose

Controls whether the compiler can automatically take advantage of vector instructions for processors that support them.

These instructions can offer higher performance when used with algorithmic-intensive tasks such as multimedia applications.

### Syntax

►► `-qsimd=`

auto
noauto

 ►►

### Defaults

Whether `-qsimd` is specified or not, `-qsimd=auto` is implied at the `-O3` or higher optimization level; `-qsimd=noauto` is implied at the `-O2` or lower optimization level.

### Usage

The `-qsimd=auto` option enables automatic generation of vector instructions for processors that support them.

The `-qsimd=auto` option controls the autosimdization, which was performed by the deprecated `-qhot=simd` option. If you specify `-qhot=simd`, the compiler ignores it and does not issue any warning message.

When `-qsimd=auto` is in effect, the compiler converts certain operations that are performed in a loop on successive elements of an array into vector instructions. These instructions calculate several results at one time, which is faster than calculating each result sequentially. Applying this option is useful for applications with significant image processing demands.

The `-qsimd=noauto` option disables the conversion of loop array operations into vector instructions. Finer control can be achieved by using `-qstrict=ieefp`, `-qstrict=operationprecision`, and `-qstrict=vectorprecision`. For details, see “`-qstrict`” on page 238.

**Note:** Using vector instructions to calculate several results at one time might delay or even miss detection of floating-point exceptions on some architectures. If detecting exceptions is important, do not use `-qsimd=auto`.

### Rules

The following rules apply when you use the `-qsimd` option:



## Defaults

The default, `-qnosmallstack`, implies that all suboptions are off.

## Usage

Using this option may adversely affect program performance; it should be used only for programs that allocate large amounts of data on the stack.

`-qsmallstack` with no suboptions enables only the general small stack transformations.

`-qnosmallstack` only disables the general small stack transformations. To disable `dynlenonheap` transformations, specify `-qsmallstack=nodynlenonheap` as well.

`-qsmallstack=dynlenonheap` enables the dynamic length variable allocation and general small stack transformations.

To enable only the `dynlenonheap` transformations, specify `-qsmallstack=dynlenonheap -qnosmallstack`.

When both `-qsmallstack` and `-qstacktemp` options are used, the `-qstacktemp` setting will be used to allocate applicable temporary variables if it is set to a non-zero value, even if this setting conflicts with that of `-qsmallstack`. The `-qsmallstack` setting will continue to apply transformations not affected by `-qstacktemp`.

## Related information

- “`-qstacktemp`” on page 234

---

## **-qsmp**

### Category

Optimization and tuning


### @PROCESS

None.

### Purpose

Enables parallelization of program code.

### Syntax



## Defaults

**-qnosmp.** Code is produced for a uniprocessor machine.

## Parameters

### **auto** | **noauto**

Enables or disables automatic parallelization and optimization of program code. By default, the compiler tries to parallelize explicitly coded DO loops and those that are generated by the compiler for processing arrays. When **noauto** is in effect, only program code explicitly parallelized with OpenMP directives is optimized. **noauto** is implied if you specify **-qsmp=omp** or **-qsmp=noopt**.

### **omp** | **noomp**

Enforces or relaxes strict compliance with the OpenMP standard. When **noomp** is in effect, **auto** is implied. When **omp** is in effect, **noauto** is implied and only OpenMP parallelization directives are recognized. The compiler issues warning messages if your code contains any language constructs that do not conform to the OpenMP API.

**Note:** The **-qsmp=omp** option must be used to enable OpenMP parallelization.

Specifying **omp** also has the following effects:

- Automatic parallelization is disabled.
- All previously recognized directive triggers are ignored. The only recognized directive trigger is \$OMP. However, you can specify additional triggers on subsequent **-qdirective** options.
- The **-qcclines** compiler option is enabled.
- When the C preprocessor is invoked, the `_OPENMP` C preprocessor macro is defined based on the latest OpenMP API specification that XL Fortran supports. This macro is useful in supporting conditional compilation. See *Conditional Compilation* in the *XL Fortran Language Reference* for more information.

### **opt\_ | **noopt****

Enables or disables optimization of parallelized program code. When **noopt** is in effect, the compiler will do the smallest amount of optimization that is required to parallelize the code. This is useful for debugging because **-qsmp** enables the **-O2** and **-qhot** options by default, which may result in the movement of some variables into registers that are inaccessible to the debugger. However, if the **-qsmp=noopt** and **-g** options are specified, these variables will remain visible to the debugger.

### **rec\_locks | **norec\_locks****

Determines whether recursive locks are used to avoid problems associated with CRITICAL constructs. When **rec\_locks** is in effect, nested critical sections will not cause a deadlock; a thread can enter a CRITICAL construct from within the dynamic extent of another CRITICAL construct that has the same name. Note that the **rec\_locks** suboption specifies behavior for critical constructs that is inconsistent with the OpenMP API.

### **schedule**

Specifies the type of scheduling algorithms and, except in the case of **auto**, chunk size ( $n$ ) that are used for loops to which no other scheduling algorithm has been explicitly assigned in the source code. Suboptions of the **schedule** suboption are as follows:

#### **affinity[= $n$ ]**

The iterations of a loop are initially divided into  $n$  partitions, containing **ceiling**( $number\_of\_iterations / number\_of\_threads$ ) iterations. Each partition is initially assigned to a thread and is then further subdivided into chunks that each contain  $n$  iterations. If  $n$  is not specified, then the chunks consist of **ceiling**( $number\_of\_iterations\_left\_in\_partition / 2$ ) loop iterations.

When a thread becomes free, it takes the next chunk from its initially assigned partition. If there are no more chunks in that partition, then the thread takes the next available chunk from a partition initially assigned to another thread.

The work in a partition initially assigned to a sleeping thread will be completed by threads that are active.

The **affinity** scheduling type is not part of the OpenMP API specification.

**Note:** This suboption has been deprecated. You can use the **OMP\_SCHEDULE** environment variable with the **dynamic** clause for a similar functionality.

#### **auto**

Scheduling of the loop iterations is delegated to the compiler and runtime systems. The compiler and runtime system can choose any possible mapping of iterations to threads (including all possible valid schedule types) and these might be different in different loops. Do not specify chunk size ( $n$ ).

#### **dynamic[= $n$ ]**

The iterations of a loop are divided into chunks that contain  $n$  iterations each. If  $n$  is not specified, each chunk contains one iteration.

Active threads are assigned these chunks on a "first-come, first-do" basis. Chunks of the remaining work are assigned to available threads until all work has been assigned.

**guided[=*n*]**

The iterations of a loop are divided into progressively smaller chunks until a minimum chunk size of *n* loop iterations is reached. If *n* is not specified, the default value for *n* is 1 iteration.

Active threads are assigned chunks on a "first-come, first-do" basis. The first chunk contains **ceiling**(*number\_of\_iterations*/*number\_of\_threads*) iterations. Subsequent chunks consist of **ceiling**(*number\_of\_iterations\_left*/*number\_of\_threads*) iterations.

**runtime**

Specifies that the chunking algorithm will be determined at run time.

**static[=*n*]**

The iterations of a loop are divided into chunks containing *n* iterations each. Each thread is assigned chunks in a "round-robin" fashion. This is known as *block cyclic scheduling*. If the value of *n* is 1, then the scheduling type is specifically referred to as *cyclic scheduling*.

If *n* is not specified, the chunks will contain **floor**(*number\_of\_iterations*/*number\_of\_threads*) iterations. The first **remainder** (*number\_of\_iterations*/*number\_of\_threads*) chunks have one more iteration. Each thread is assigned a separate chunk. This is known as *block scheduling*.

If a thread is asleep and it has been assigned work, it will be awakened so that it may complete its work.

*n* Must be an integer of value 1 or greater.

Specifying **schedule** with no suboption is equivalent to **schedule=auto**.

For more information on chunking algorithms and SCHEDULE, refer to *Directives* in the *XL Fortran Language Reference*.

**stackcheck | nostackcheck**

Causes the compiler to check for stack overflow by slave threads at run time, and issue a warning if the remaining stack size is less than the number of bytes specified by the **stackcheck** option of the XLSMPOPTS environment variable. This suboption is intended for debugging purposes, and only takes effect when **XLSMPOPTS=stackcheck** is also set; see XLSMPOPTS in the *XL Fortran Optimization and Programming Guide* for more information.

**threshold[=*n*]**

When **-qsmp=auto** is in effect, controls the amount of automatic loop parallelization that occurs. The value of *n* represents the minimum amount of work required in a loop in order for it to be parallelized. Currently, the calculation of "work" is weighted heavily by the number of iterations in the loop. In general, the higher the value specified for *n*, the fewer loops are parallelized. Specifying a value of 0 instructs the compiler to parallelize all auto-parallelizable loops, whether or not it is profitable to do so. Specifying a value of 100 instructs the compiler to parallelize only those auto-parallelizable loops that it deems profitable. Specifying a value of greater than 100 will result in more loops being serialized.

*n* Must be a positive integer of 0 or greater.

If you specify **threshold** with no suboption, the program uses a default value of 100.

Specifying **-qsmp** without suboptions is equivalent to:

**-qsmp=auto:opt=noomp:norec\_locks:schedule=auto:nostackcheck:threshold=100**

## Usage

- Specifying the **omp** suboption always implies **noauto**. Specify **-qsmp=omp:auto** to apply automatic parallelization on OpenMP-compliant applications, as well.
- When **-qsmp** is in effect, the compiler recognizes all directives with the trigger constants **SMP\$**, **\$OMP**, and **IBMP**, unless you specify the **omp** suboption. If you specify **omp** and want the compiler to recognize directives specified with the other triggers, you can use the **-qdirective** option to do so.
- You should only use **-qsmp** with the **\_r**-suffixed invocation commands, to automatically link in all of the threadsafe components. You can use the **-qsmp** option with the non-**\_r**-suffixed invocation commands, but you are responsible for linking in the appropriate components. If you use the **-qsmp** option to compile any source file in a program, then you must specify the **-qsmp** option at link time as well, unless you link by using the **ld** command.
- If you use the **f77** or **fort77** command with the **-qsmp** option to compile programs, specify **-qnosave** to make the default storage class automatic, and specify **-qthreaded** to tell the compiler to generate threadsafe code.
- Object files generated with the **-qsmp=opt** option can be linked with object files generated with **-qsmp=noopt**. The visibility within the debugger of the variables in each object file will not be affected by linking.
- Specifying **-qsmp** implicitly sets **-O2**. The **-qsmp** option overrides **-qnooptimize**, but does not override **-O3**, **-O4**, or **-O5**. When debugging parallelized program code, you can disable optimization in parallelized program code by specifying **-qsmp=noopt**.
- The **-qsmp=noopt** suboption overrides performance optimization options anywhere on the command line unless **-qsmp** appears after **-qsmp=noopt**. For example, **-qsmp=noopt -O3** is equivalent to **-qsmp=noopt**, while **-qsmp=noopt -O3 -qsmp** is equivalent to **-qsmp -O3**.

## Examples

In the following example, you should specify **-qsmp=rec\_locks** to avoid a deadlock caused by critical constructs.

```
program t
  integer i, a, b

  a = 0
  b = 0
!smp$ parallel do
  do i=1, 10
!smp$ critical
    a = a + 1
!smp$ critical
    b = b + 1
!smp$ end critical
!smp$ end critical
  enddo
end
```

## Related information

- “-O” on page 83
- “-qthreaded” on page 249
- XLSMPOPTS environment variable and Parallelization directives in the *XL Fortran Optimization and Programming Guide*

---

## -qsource

### Category

Listings, messages, and compiler information

### Purpose

Produces a compiler listing file that includes the source section of the listing and provides additional source information when printing error messages.

### Syntax



**@PROCESS:**

@PROCESS SOURCE | NOSOURCE

### Defaults

-qnosource

### Usage

This option displays on the terminal each source line where the compiler detects a problem, which can be very useful in diagnosing program errors in the Fortran source files.

You can selectively print parts of the source code by using **SOURCE** and **NOSOURCE** in **@PROCESS** directives in the source files around those portions of the program you want to print. This is the only situation where the **@PROCESS** directive does not have to be before the first statement of a compilation unit.

### Examples

In the following example, the point at which the incorrect call is made is identified more clearly when the program is compiled with the **-qsource** option:

```
$ cat argument_mismatch.f
  subroutine mult(x,y)
    integer x,y
    print *,x*y
  end

  program wrong_args
  interface
    subroutine mult(a,b) ! Specify the interface for this
      integer a,b        ! subroutine so that calls to it
    end subroutine mult  ! can be checked.
  end interface
  real i,j
  i = 5.0
  j = 6.0
  call mult(i,j)
  end

$ xlf95 argument_mismatch.f
** mult === End of Compilation 1 ===
"argument_mismatch.f", line 15.20: 1513-061 (S) Actual argument attributes
```

```

do not match those specified by an accessible explicit interface.
** wrong_args === End of Compilation 2 ===
1501-511 Compilation failed for file argument_mismatch.f.
$ xlf95 -qsource argument_mismatch.f
** mult === End of Compilation 1 ===
    15 | call mult(i,j)
        | .....a...
a - "argument_mismatch.f", line 15.20: 1513-061 (S) Actual argument attributes do not match those
specified by an accessible explicit interface.
** wrong_args === End of Compilation 2 ===
1501-511 Compilation failed for file argument_mismatch.f.

```

### Related information

- “Listings, messages, and compiler information” on page 56
- “Source section” on page 292

## -qspillsize

### Category

Compiler customization

### Purpose

**-qspillsize** is the long form of **-NS**. See “-NS” on page 82.

### Syntax

►► `-qspillsize=bytes` ◀◀

**@PROCESS:**

@PROCESS SPILLSIZE(*bytes*)

### Defaults

Not applicable.

## -qstackprotect

### Category

“Object code control” on page 53

### @PROCESS

None.

### Purpose

Provides protection against malicious input data or programming errors that overwrite or corrupt the stack.

### Syntax

►► `-qstackprotect=all [size]` ◀◀

`nostackprotect` ————  
`stackprotect` ————  
`size` ————*N*

## Defaults

- `-qnostackprotect`

## Parameters

### `all`

`all` protects all procedures whether or not there are vulnerable objects. This option is not set by default.

### `size=N`

`size=N` protects all procedures containing automatic objects greater or equal to  $N$  bytes in size. The default size is 8 when `-qstackprotect` is enabled.

## Usage

`-qstackprotect` generates extra code to protect procedures with vulnerable objects against stack corruption. This option is disabled by default because it can cause performance degradation.

To generate code to protect all procedures with vulnerable objects:

```
xlf myprogram.f -qstackprotect=all
```

To generate code to protect procedures with objects of certain bytes:

```
xlf myprogram.f -qstackprotect=size=8
```

---

## `-qstacktemp`

### Category

Optimization and tuning

### Purpose

Determines where to allocate certain XL Fortran compiler temporaries at run time.

Applicable compiler temporaries are the set of temporary variables created by the compiler for its own use when it determines it can safely apply these. Most typically, the compiler creates these kinds of temporaries to hold copies of XL Fortran arrays for array language semantics, especially in conjunction with calls to intrinsic functions or user subprograms.

### Syntax

►► `-qstacktemp=`

0
-1
<i>value</i>

 ◀◀

**@PROCESS:**

`@PROCESS STACKTEMP={0 | -1 | value}`

### Defaults

`-qstacktemp=0`

## Parameters

The possible suboptions are:

- 0** Based on the target environment, the compiler determines whether it will allocate applicable temporaries on the heap or the stack. If this setting causes your program to run out of stack storage, try specifying a nonzero value instead, or try using the **-qsmallstack** option.
- 1** Allocates applicable temporaries on the stack. Generally, this is the best performing setting but uses the most amount of stack storage.
- value* Allocates applicable temporaries less than *value* on the stack and those greater than or equal to *value* on the heap. *value* is a positive integer. A value of 1 Mb has been shown to be a good compromise between stack storage and performance for many programs, but you may need to adjust this number based on your application's characteristics.

## Usage

If you have programs that make use of large arrays, you may need to use this option to help prevent stack space overflow when running them. For example, for SMP or OpenMP applications that are constrained by stack space, you can use this option to move some compiler temporaries onto the heap from the stack.

The compiler cannot detect whether or not the stack limits will be exceeded when an application runs. You will need to experiment with several settings before finding the one that works for your application. To override an existing setting, you must specify a new setting.

The **-qstacktemp** option can take precedence over the **-qsmallstack** option for certain compiler-generated temporaries.

## Related information

- “-qsmallstack” on page 226

---

## **-qstaticlink**

### Category

Linking

### @PROCESS directive


None.

### Purpose

Controls whether static or shared runtime libraries are linked into an application.

This option provides the ability to specify linking rules that are equivalent to those implied by the GNU options **-static**, **-static-libgcc**, **-shared**, and **-shared-libgcc**, used singly and in combination.

## Syntax



## Defaults

**-qnostaticlink**

## Parameters

### libgcc

- When you specify **libgcc** with **nostaticlink**, the compiler links the shared version of **libgcc**.
- When you specify **libgcc** with **staticlink**, the compiler links the static version of **libgcc**.

This suboption provides the equivalent ability that is enabled by the GNU options **-static-libgcc** and **-shared-libgcc**.

### xllibs

- When you specify **xllibs** with **-qnostaticlink**, the compiler links the shared version of the XL compiler libraries.
- When you specify **xllibs** with **-qstaticlink**, the compiler links the static version of the XL compiler libraries.

## Usage

When you specify **-qstaticlink** without suboptions, only static libraries are linked with the object file.

When you specify **-qnostaticlink** without suboptions, shared libraries are linked with the object file.

Conflicting compiler options are resolved as follows:

- If you first specify **-qnostaticlink** without suboptions and then specify **-qstaticlink** with or without suboptions, **-qnostaticlink** is overridden. For example, **-qnostaticlink -qstaticlink=xllibs** is equivalent to **-qstaticlink=xllibs**.
- If you specify **-qstaticlink** with or without suboptions followed by **-qnostaticlink** without suboptions, **-qnostaticlink** takes effect and shared libraries are linked. Otherwise, once you have specified **-qstaticlink** without suboptions, **-qstaticlink** takes effect and only static libraries are linked with the object file. See the following examples:

Table 19. Examples of conflicting compiler options and resolutions

Options combination	Compiler behavior
<b>-qstaticlink=libgcc</b> <b>-qnostaticlink</b>	Shared libraries are linked.

Table 19. Examples of conflicting compiler options and resolutions (continued)

Options combination	Compiler behavior
<code>-qstaticlink</code> <code>-qnostaticlink=libgcc</code>	All libraries are linked statically. The compiler issues the following warning message: (W) The options <code>-qnostaticlink=libgcc</code> and <code>-qstaticlink</code> are incompatible. Option <code>-qnostaticlink=libgcc</code> is ignored.
<code>-qstaticlink</code> <code>-qnostaticlink=libgcc</code> <code>:xllibs</code>	All libraries are linked statically. The compiler issues the following warning messages: (W) The options <code>-qnostaticlink=libgcc</code> and <code>-qstaticlink</code> are incompatible. Option <code>-qnostaticlink=libgcc</code> is ignored. (W) The options <code>-qnostaticlink=xllibs</code> and <code>-qstaticlink</code> are incompatible. Option <code>-qnostaticlink=xllibs</code> is ignored.
<code>-qstaticlink</code> <code>-qstaticlink=libgcc</code>	All libraries are linked statically.
<code>-qnostaticlink=libgcc</code> <code>-qstaticlink</code>	All libraries are linked statically.

**Note:** If a runtime library is linked in statically while its message catalog is not installed on the system, messages are issued with message numbers only, and no message text is shown.

**Important:** Any use of third-party libraries or products is subject to the provisions in their respective licenses. Using the `-qstaticlink` option can have significant legal consequences for the programs that you compile. It is strongly recommended that you seek legal advice before you use this option.

The following table shows the equivalent GNU and XL Fortran options for specifying linkage of shared and nonshared libraries.

Table 20. Option mappings: control of the GNU linker

GNU option	Meaning	XL Fortran option
<code>-shared</code>	Build a shared object.	<code>-qmkshrobj</code> <b>1</b>
<code>-static</code>	Build a static object and prevent linking with shared libraries. Every library that is linked to must be a static library.	<code>-qstaticlink</code>
<code>-shared-libgcc</code>	Link with the shared version of <code>libgcc</code> .	<code>-qnostaticlink=libgcc</code> <b>2</b>
<code>-static-libgcc</code>	Link with the static version of <code>libgcc</code> . You can still link your shared libraries.	<code>-qstaticlink=libgcc</code>
<b>Notes:</b> <b>1</b> Options <code>-qmkshrobj</code> and <code>-qstaticlink</code> are incompatible and cannot be specified together. <b>2</b> This is the default setting on SUSE Linux Enterprise Server (SLES) and Red Hat Enterprise Linux (RHEL).		

## Related information

- “`-qmkshrobj`” on page 186

# -qstrict

## Category

Optimization and tuning


## Purpose

Ensures that optimizations done by default at the **-O3** and higher optimization levels, and, optionally at **-O2**, do not alter certain program semantics mostly related to strict IEEE floating-point conformance.

This option is intended for situations where the changes in program execution in optimized programs produce different results from unoptimized programs.

**Note:** **-qstrict** affects the option default changes that are made by the optimization levels.

## Syntax



@PROCESS:

@PROCESS STRICT[(suboptions)] | NOSTRICT

## Defaults

- **-qstrict** or **-qstrict=all** is always in effect when the **-qnoopt** or **-O0** optimization level is in effect
- **-qstrict** or **-qstrict=all** is the default when the **-O2** or **-O** optimization level is in effect
- **-qnostrict** or **-qstrict=none** is the default when the **-O3** or higher optimization level is in effect

## Parameters

The **-qstrict** suboptions include the following:

**all** | **none**

**all** disables all semantics-changing transformations, including those controlled by the **ieee<sub>fp</sub>**, **order**, **library**, **constructcopy**, **precision**, and **exceptions** suboptions. **none** enables these transformations.

**precision** | **noprecision**

**precision** disables all transformations that are likely to affect floating-point precision, including those controlled by the **subnormals**, **operationprecision**, **vectorprecision**, **association**, **reductionorder**, and **library** suboptions. **noprecision** enables these transformations.

**exceptions** | **noexceptions**

**exceptions** disables all transformations likely to affect exceptions or be affected by them, including those controlled by the **nans**, **infinities**, **subnormals**, **guards**, **library**, and **constructcopy** suboptions. **noexceptions** enables these transformations.

**ieee<sub>fp</sub>** | **noieee<sub>fp</sub>**

**ieee<sub>fp</sub>** disables transformations that affect IEEE floating-point compliance, including those controlled by the **nans**, **infinities**, **subnormals**, **zerosigns**, **vectorprecision**, and **operationprecision** suboptions. **noieee<sub>fp</sub>** enables these transformations.

**nans** | **nonans**

**nans** disables transformations that may produce incorrect results in the presence of, or that may incorrectly produce IEEE floating-point NaN (not-a-number) values. **nonans** enables these transformations.

**infinities** | **noinfinities**

**infinities** disables transformations that may produce incorrect results in the presence of, or that may incorrectly produce floating-point infinities. **noinfinities** enables these transformations.

**subnormals** | **nosubnormals**

**subnormals** disables transformations that may produce incorrect results in the presence of, or that may incorrectly produce IEEE floating-point subnormals (formerly known as denorms). **nosubnormals** enables these transformations.

**zerosigns** | **nozerosigns**

**zerosigns** disables transformations that may affect or be affected by whether the sign of a floating-point zero is correct. **nozerosigns** enables these transformations.

**operationprecision | nooperationprecision**

**operationprecision** disables transformations that produce approximate results for individual floating-point operations. **nooperationprecision** enables these transformations.

**vectorprecision | novectorprecision**

**vectorprecision** disables vectorization in loops where it might produce different results in vectorized iterations than in nonvectorized residue iterations. **vectorprecision** ensures that every loop iteration of identical floating-point operations on identical data produces identical results.

**novectorprecision** enables vectorization even when different iterations might produce different results from the same inputs.

**order | noorder**

**order** disables all code reordering between multiple operations that may affect results or exceptions, including those controlled by the **association**, **reductionorder**, and **guards** suboptions. **noorder** enables code reordering.

**association | noassociation**

**association** disables reordering operations within an expression. **noassociation** enables reordering operations.

**reductionorder | noreductionorder**

**reductionorder** disables parallelizing floating-point reductions. **noreductionorder** enables parallelizing these reductions.

**guards | noguards**

**guards** disables moving operations past guards (that is, past **IF** statements, out of loops, or past subroutine or function calls that might end the program) which control whether the operation should be executed. **noguards** enables moving operations past guards.

**library | nolibrary**

**library** disables transformations that affect floating-point library functions; for example, transformations that replace floating-point library functions with other library functions or with constants. **nolibrary** enables these transformations.

**constructcopy | noconstructcopy**

**constructcopy** disables constructing arrays in place instead of using a temporary copy where an exception could occur. **noconstructcopy** enables constructing such arrays.

## Usage

The **all**, **precision**, **exceptions**, **ieeefp**, and **order** suboptions and their negative forms are group suboptions that affect multiple, individual suboptions. For many situations, the group suboptions will give sufficient granular control over transformations. Group suboptions act as if either the positive or the no form of every suboption of the group is specified. Where necessary, individual suboptions within a group (like **subnormals** or **operationprecision** within the **precision** group) provide control of specific transformations within that group.

With **-qnostrict** or **-qstrict=none** in effect, the following optimizations are turned on:

- Code that may cause an exception may be rearranged. The corresponding exception might happen at a different point in execution or might not occur at all. (The compiler still tries to minimize such situations.)

- Floating-point operations may not preserve the sign of a zero value. (To make certain that this sign is preserved, you also need to specify **-qfloat=rrm**, **-qfloat=nomaf**, or **-qfloat=strictnmaf**.)
- Floating-point expressions may be reassocated. For example,  $(2.0*3.1)*4.2$  might become  $2.0*(3.1*4.2)$  if that is faster, even though the result might not be identical.
- The optimization functions enabled by **-qfloat=rsqrt**. You can turn off the optimization functions by using the **-qstrict** option or **-qfloat=norsqrt**. With lower-level or no optimization specified, these optimization functions are turned off by default.

Specifying various suboptions of **-qstrict[=suboptions]** or **-qnostrict** combinations sets the following suboptions:

- **-qstrict** or **-qstrict=all** sets **-qfloat=norsqrt:rngchk**. **-qnostrict** or **-qstrict=none** sets **-qfloat=rsqrt:norngchk**.
- **-qstrict=infinities**, **-qstrict=operationprecision**, or **-qstrict=exceptions** sets **-qfloat=norsqrt**.
- **-qstrict=noinfinities:nooperationprecision:noexceptions** sets **-qfloat=rsqrt**.
- **-qstrict=nans**, **-qstrict=infinities**, **-qstrict=zerosigns**, or **-qstrict=exceptions** sets **-qfloat=rngchk**. Specifying all of **-qstrict=nonans:nozerosigns:noexceptions** or **-qstrict=noinfinities:nozerosigns:noexceptions**, or any group suboptions that imply all of them, sets **-qfloat=norngchk**.

**Note:** For details about the relationship between **-qstrict** suboptions and their **-qfloat** counterparts, see “-qfloat” on page 127.

To override any of these settings, specify the appropriate **-qfloat** suboptions after the **-qstrict** option on the command line.

## Examples

To compile `myprogram.f` so that the aggressive optimization of **-O3** are turned off, and division by the result of a square root is replaced by multiplying by the reciprocal (**-qfloat=rsqrt**), enter:

```
xlf myprogram.f -O3 -qstrict -qfloat=rsqrt
```

To enable all transformations except those affecting precision, specify:

```
xlf myprogram.f -qstrict=none:precision
```

To disable all transformations except those involving NaNs and infinities, specify:

```
xlf myprogram.f -qstrict=all:nonans:noinfinities
```

## Related information

- “-qsimd” on page 225
- “-qessl” on page 120
- “-qfloat” on page 127
- “-qhot” on page 143
- “-O” on page 83
- “-qxl90” on page 261

---

## -qstrictieemod

### Category

Floating-point and integer control

### Purpose

Specifies whether the compiler will adhere to the Fortran 2003 IEEE arithmetic rules for the `ieee_arithmetic` and `ieee_exceptions` intrinsic modules.

### Syntax

► — -q — strictieemod  
nostrictieemod — ►

@PROCESS:

@PROCESS STRICTIEEEMOD | NOSTRICTIEEEMOD

### Defaults

-qstrictieemod

### Usage

When you specify `-qstrictieemod`, the compiler adheres to the following rules:

- If there is an exception flag set on entry into a procedure that uses the IEEE intrinsic modules, the flag is set on exit. If a flag is clear on entry into a procedure that uses the IEEE intrinsic modules, the flag can be set on exit.
- If there is an exception flag set on entry into a procedure that uses the IEEE intrinsic modules, the flag clears on entry into the procedure and resets when returning from the procedure.
- When returning from a procedure that uses the IEEE intrinsic modules, the settings for halting mode and rounding mode return to the values they had at procedure entry.
- Calls to procedures that do not use the `ieee_arithmetic` or `ieee_exceptions` intrinsic modules from procedures that do use these modules, will not change the floating-point status except by setting exception flags.

Since the above rules can impact performance, specifying `-qnostrictieemod` will relax the rules on saving and restoring floating-point status. This prevents any associated impact on performance.

---

## -qstrict\_induction

### Category

Optimization and tuning

@PROCESS

None.

## Purpose

Prevents the compiler from performing induction (loop counter) variable optimizations. These optimizations may be unsafe (may alter the semantics of your program) when there are integer overflow operations involving the induction variables.

## Syntax

► — -q — nostrict\_induction  
strict\_induction — ►

## Defaults

-qnostrict\_induction

## Usage

You should avoid specifying **-qstrict\_induction** unless absolutely necessary, as it may cause performance degradation.

## Examples

Consider the following two examples:

### Example 1

```
integer(1) :: i, j           ! Variable i can hold a
j = 0                       ! maximum value of 127.

do i = 1, 200               ! Integer overflow occurs when 128th
  j = j + 1                 ! iteration of loop is attempted.
enddo
```

### Example 2

```
integer(1) :: i
i = 1_1                     ! Variable i can hold a maximum
                             ! value of 127.

100 continue
if (i == -127) goto 200     ! Go to label 200 once decimal overflow
  i = i + 1_1               ! occurs and i == -127.
  goto 100
200 continue
print *, i
end
```

If you compile these examples with the **-qstrict\_induction** option, the compiler does not perform induction variable optimizations, but the performance of the code may be affected. If you compile the examples with the **-qnostrict\_induction** option, the compiler may perform optimizations that may alter the semantics of the programs.

## Related information

- “-O” on page 83

---

## -qsuffix

### Category

Input control

### @PROCESS


None.

### Purpose

Specifies the source-file suffix on the command line.

This option saves time for the user by permitting files to be used as named with minimal makefile modifications. Only one setting is supported at any one time for any particular file type.

### Syntax



### Defaults

Not applicable.

### Parameters

**f=***suffix*

Where *suffix* represents the new *source-file-suffix*

**o=***suffix*

Where *suffix* represents the new *object-file-suffix*

**s=***suffix*

Where *suffix* represents the new *assembler-source-file-suffix*

**cpp=***suffix*

Where *suffix* represents the new *preprocessor-source-file-suffix*

### Rules

- The new suffix setting is case-sensitive.
- The new suffix can be of any length.

### Examples

For instance,

```
xlf a1.f2k a2.F2K -qsuffix=f=f2k:cpp=F2K
```

will cause these effects:

- The compiler is invoked for source files with a suffix of .f2k and .F2K.
- cpp is invoked for files with a suffix of .F2K.

---

## -qsuppress

### Category

Listings, messages, and compiler information


### @PROCESS

None.

### Purpose

Prevents specific informational or warning messages from being displayed or added to the listing file, if one is generated.

### Syntax



### Defaults

Not applicable.

### Parameters

`message_num[:message_num ...]`

Suppresses the display of a specific compiler message (*nnnn-mmm*) or a list of messages (*nnnn-mmm[:nnnn-mmm ...]*). To suppress a list of messages, separate each message number with a colon.

*nnnn-mmm* is the message number, where:

- *nnnn* must be a four-digit integer between 1500 and 1585; this is the range of XL Fortran message numbers.
- *mmm* must be any three-digit integer (with leading zeros if necessary).

#### `cmpmsg`

Suppresses the informational messages that report compilation progress and a successful completion.

This suboption has no effect on any error messages that are emitted.

### Usage

In some situations, users may receive an overwhelming number of compiler messages. In many cases, these compiler messages contain important information. However, some messages contain information that is either redundant or can be safely ignored. When multiple error or warning messages appear during compilation, it can be very difficult to distinguish which messages should be noted. By using `-qsuppress`, you can eliminate messages that do not interest you.

#### Note:

- The compiler tracks the message numbers specified with **-qsuppress**. If the compiler subsequently generates one of those messages, it will not be displayed or entered into the listing.
- Only compiler and driver messages can be suppressed. Linker or operating system message numbers will be ignored if specified with **-qsuppress**.
- To suppress IPA messages, enter **-qsuppress** before **-qipa** on the command line.
- The **-qhaltonmsg** option has precedence over **-qsuppress**. If you specify both **-qhaltonmsg** and **-qsuppress**, messages that **-qsuppress** suppresses are also printed and compilation stops.

## Examples

```
@process nullterm
  i = 1; j = 2;
  call printf("i=%d\n", %val(i));
  call printf("i=%d, j=%d\n", %val(i), %val(j));
end
```

Compiling this sample program would normally result in the following output:

```
"t.f", line 4.36: 1513-029 (W) The number of arguments to "printf" differ
from the number of arguments in a previous reference. You should use the
OPTIONAL attribute and an explicit interface to define a procedure with
optional arguments.
**_main   === End of Compilation 1 ===
1501-510  Compilation successful for file t.f.
```

When the program is compiled with **-qsuppress=1513-029**, the output is:

```
**_main   === End of Compilation 1 ===
1501-510  Compilation successful for file t.f.
```

## Related information

- “-qflag” on page 126
- “-qhaltonmsg” on page 142
- “-qmaxerr” on page 182

---

## -qswapomp

### Category

Portability and migration

### Purpose

Specifies that the compiler should recognize and substitute OpenMP routines in XL Fortran programs.

The OpenMP routines for Fortran and C have different interfaces. To support multi-language applications that use OpenMP routines, the compiler needs to recognize OpenMP routine names and substitute them with the XL Fortran versions of these routines, regardless of the existence of other implementations of such routines.

### Syntax

```
►► -q swapomp / noswapomp ◄◄
```

@PROCESS:

@PROCESS SWAPOMP | NOSWAPOMP

## Defaults

-qswapomp

## Usage

The compiler does not perform substitution of OpenMP routines when you specify the **-qnoswapomp** option.

The **-qswapomp** and **-qnoswapomp** options only affect Fortran subprograms that reference OpenMP routines that exist in the program.

## Rules

- If a call to an OpenMP routine resolves to a dummy procedure, module procedure, an internal procedure, a direct invocation of a procedure itself, or a statement function, the compiler will not perform the substitution.
- When you specify an OpenMP routine, the compiler substitutes the call to a different special routine depending upon the setting of the **-qintsize** option. In this manner, OpenMP routines are treated as generic intrinsic procedures.
- Unlike generic intrinsic procedures, if you specify an OpenMP routine in an **EXTERNAL** statement, the compiler will not treat the name as a user-defined external procedure. Instead, the compiler will still substitute the call to a special routine depending upon the setting of the **-qintsize** option.
- An OpenMP routine cannot be extended or redefined, unlike generic intrinsic procedures.

## Examples

In the following example, the OpenMP routines are declared in an **INTERFACE** statement.

```
@PROCESS SWAPOMP

      INTERFACE
      FUNCTION OMP_GET_THREAD_NUM()
        INTEGER OMP_GET_THREAD_NUM
      END FUNCTION OMP_GET_THREAD_NUM

      FUNCTION OMP_GET_NUM_THREADS()
        INTEGER OMP_GET_NUM_THREADS
      END FUNCTION OMP_GET_NUM_THREADS
      END INTERFACE

      IAM = OMP_GET_THREAD_NUM()
      NP = OMP_GET_NUM_THREADS()
      PRINT *, IAM, NP
      END
```

## Related information

See the *OpenMP execution environment, lock and timing routines* section in the *XL Fortran Optimization and Programming Guide*.

---

## **-qtbtable**

### **Category**

Object code control

### **@PROCESS**

None.

### **Purpose**

Controls the amount of debugging traceback information that is included in the object files.

### **Syntax**

►► -qtbtable=  
          full  
          none  
          small

### **Defaults**

Not applicable.

### **Parameters**

- full** The object code contains full traceback information. The program is debuggable, and if it stops because of a runtime exception, it produces a traceback listing that includes the names of all procedures in the call chain.
- none** The object code contains no traceback information at all. You cannot debug the program, because a debugger or other code-examination tool cannot unwind the program's stack at run time. If the program stops because of a runtime exception, it does not explain where the exception occurred.
- small** The object code contains traceback information but not the names of procedures or information about procedure parameters. You can debug the program, but some non-essential information is unavailable to the debugger. If the program stops because of a runtime exception, it explains where the exception occurred but reports machine addresses rather than procedure names.

### **Defaults**

- Code compiled with **-g** or without **-O** has full traceback information (**-qtbtable=full**).
- Code compiled with **-O** or higher optimization contains less traceback information (**-qtbtable=small**).

### **Usage**

This option is most suitable for programs that contain many long procedure names, such as the internal names constructed for module procedures. You may find it more applicable to C++ programs than to Fortran programs.

You can use this option to make your program smaller, at the cost of making it harder to debug. When you reach the production stage and want to produce a program that is as compact as possible, you can specify **-qbttable=none**. Otherwise, the usual defaults apply.

### Related information

- “-g” on page 73
- “-qcompact” on page 108
- “-O” on page 83
- *Debugging optimized code* in the *XL Fortran Optimization and Programming Guide*

---

## -qthreaded

### Category

Object code control

### @PROCESS

None.

### Purpose

Indicates to the compiler whether it must generate threadsafe code.

### Syntax

►► -q—threaded◄◄

### Defaults

**-qthreaded** is the default for the **xlf\_r**, **xlf90\_r**, **xlf95\_r**, **xlf2003\_r**, and **xlf2008\_r** commands.

### Usage

Specifying the **-qthreaded** option implies **-qdirective=ibmt**, and by default, the *trigger\_constant* **IBMT** is recognized.

The **-qthreaded** option does not imply the **-qnosave** option. The **-qnosave** option specifies a default storage class of automatic for user local variables. In general, both of these options need to be used to generate threadsafe code. Specifying these options ensures that variables and code created by the compiler are threadsafe; it does not guarantee the thread safety of user-written code.

The **-qthreaded** option does not imply the **-qxlf77=nopersistent** option. The **-qxlf77=nopersistent** option improves thread safety by preventing the compiler from saving the addresses of arguments to subprograms with the **ENTRY** statements in static storage.

---

## -qtimestamps

### Category

“Output control” on page 47


## @PROCESS

None.

### Purpose

Controls whether or not implicit time stamps are inserted into an object file.

### Syntax



### Defaults

`-qtimestamps`

### Usage

By default, the compiler inserts an implicit time stamp in an object file when it is created. In some cases, comparison tools may not process the information in such binaries properly. Controlling time stamp generation provides a way of avoiding such problems. To omit the time stamp, use the option **`-qnotimestamps`**.

This option does not affect time stamps inserted by pragmas and other explicit mechanisms.

---

## -qtune

### Category

Optimization and tuning


## @PROCESS

None.

### Purpose

Tunes instruction selection, scheduling, and other architecture-dependent performance enhancements to run best on a specific hardware architecture. Allows specification of a target SMT mode to direct optimizations for best performance in that mode.

### Syntax



## Defaults

`-qtune=balanced:balanced`

## Parameters for CPU suboptions

The following CPU suboptions allow you to specify a particular architecture for the compiler to target for best performance:

### **auto**

Optimizations are tuned for the platform on which the application is compiled.

### **balanced**

Optimizations are tuned across a selected range of recent hardware.

### **pwr8**

Optimizations are tuned for the POWER8 hardware platforms.

## Parameters for SMT suboptions

The following simultaneous multithreading (SMT) suboptions allow you to optionally specify an execution mode for the compiler to target for best performance.

### **balanced**

Optimizations are tuned for performance across various SMT modes for a selected range of recent hardware.

**st** Optimizations are tuned for single-threaded execution.

### **smt2**

Optimizations are tuned for SMT2 execution mode (two threads).

### **smt4**

Optimizations are tuned for SMT4 execution mode (four threads).

### **smt8**

Optimizations are tuned for SMT8 execution mode (eight threads).

## Usage

By arranging (scheduling) the generated machine instructions to take maximum advantage of hardware features such as cache size and pipelining, **-qtune** can improve performance. It only has an effect when used in combination with options that enable optimization.

Although changing the **-qtune** setting may affect the performance of the resulting executable, it has no effect on whether the executable can be executed correctly on a particular hardware platform.

## Examples

To specify that the executable program `testing` compiled from `myprogram.f` is to be optimized for a POWER8 hardware platform configured for the SMT4 mode, enter:

```
xlf -o testing myprogram.f -qtune=pwr8:smt4
```

## Related information

- “-qarch” on page 93

- "Optimizing your applications" in the *XL Fortran Optimization and Programming Guide*

## -qfmt

### Category

Portability and migration

### Purpose

Sets the byte order for I/O operations on unformatted data files.

### Syntax

►► -q-ufmt =le  
=be ◀◀

@PROCESS:

@PROCESS UFMT[({BE | LE})]

### Defaults

**-qfmt=le**

### Parameters

- be** Specifies that I/O operations on unformatted data files use the big-endian byte order. This option provides compatibility with big-endian data files at the cost of runtime performance. The byte order of non-character data and record markers is converted in real time during I/O operations.
- le** Specifies that I/O operations on unformatted data files use the little-endian byte order. This option provides better performance because data and record markers are read and written from unformatted data files without converting the byte order.

### Usage

You can specify the byte order for I/O operations in the following ways. If more than one is specified for the same unit and they conflict with each other, the one that comes first in the list takes precedence:

1. Runtime option **XLFRT\_OPTS=ufmt\_bigendian**
2. **CONVERT= char\_expr** in the **OPEN** statement, where *char\_expr* is **NATIVE**, **BIG\_ENDIAN**, or **LITTLE\_ENDIAN**
3. @PROCESS UFMT(BE) or @PROCESS UFMT(LE)
4. Compiler option **-qfmt=be** or **-qfmt=le**

### Related information

- "OPEN statement" in the *XL Fortran Language Reference*
- "The XLFRT\_OPTS environment variable" on page 34

---

## -qundef

### Category

Language element control

### Purpose

**-qundef** is the long form of the “-u” on page 274 option.

### Syntax

►► -q undef nundef ◀◀

@PROCESS:

@PROCESS UNDEF | NOUNDEF

### Defaults

-qnundef

---

## -qunroll

### Category

Optimization and tuning

### @PROCESS

None.

### Purpose

Specifies whether unrolling **DO** loops is allowed in a program. Unrolling is allowed on outer and inner **DO** loops.

### Syntax

►► -q nounroll unroll = n yes auto ◀◀

### Defaults

**-qunroll=auto**

### Parameters

auto The compiler performs basic loop unrolling.

yes The compiler looks for more opportunities to perform loop unrolling than that performed with **-qunroll=auto**. In general, this suboption has more

chances to increase compile time or program size than **-qunroll=auto** processing, but it may also improve your application's performance.

*n* Instructs the compiler to unroll loops by a factor of *n*. In other words, the body of a loop is replicated to create *n* copies and the number of iterations is reduced by a factor of  $1/n$ . The value of *n* must be a positive integer.

Specifying **-qunroll=1** disables loop unrolling, and is equivalent to specifying **-qnounroll**. If *n* is not specified and if **-qhot**, **-qsmp**, **-O4**, or **-O5** is specified, the optimizer determines an appropriate unrolling factor for each nested loop.

The compiler might limit unrolling to a number smaller than the value you specify for *n*. The reason is that the option affects all loops in source files to which it applies and large unrolling factors might significantly increase compile time without necessarily improving runtime performance. To specify a specific unrolling factor for particular loops, use the unroll directive in those loops.

If you decide to unroll a loop, specifying one of the above suboptions does not automatically guarantee that the compiler will perform the operation. Based on the performance benefit, the compiler will determine whether unrolling will be beneficial to the program. Experienced compiler users should be able to determine the benefit in advance.

## Usage

Specifying **-qunroll** with no suboptions is equivalent to **-qunroll=yes**.

The **-qnounroll** option prohibits unrolling unless you specify the **STREAM\_UNROLL**, **UNROLL**, or **UNROLL\_AND\_FUSE** directive for a particular loop. These directives always override the command line options.

## Examples

In the following example, the **UNROLL(2)** directive is used to tell the compiler that the body of the loop can be replicated so that the work of two iterations is performed in a single iteration. Instead of performing 1000 iterations, if the compiler unrolls the loop, it will only perform 500 iterations.

```
!IBM* UNROLL(2)
DO I = 1, 1000
  A(I) = I
END DO
```

If the compiler chooses to unroll the previous loop, the compiler translates the loop so that it is essentially equivalent to the following:

```
DO I = 1, 1000, 2
  A(I) = I
  A(I+1) = I + 1
END DO
```

## Related information

See the appropriate directive on unrolling loops in the *XL Fortran Language Reference*:

- **STREAM\_UNROLL**
- **UNROLL**
- **UNROLL\_AND\_FUSE**

## **-qunwind**

### **Category**

Optimization and tuning

### **Purpose**

Specifies that the compiler will preserve the default behavior for saves and restores to volatile registers during a procedure call.

### **Syntax**

►► -q unwind  
nounwind ►►

**@PROCESS:**

@PROCESS UNWIND | NOUNWIND

### **Defaults**

-qunwind

### **Usage**

If you specify **-qnounwind**, the compiler rearranges subprograms to minimize saves and restores to volatile registers. This rearrangement may make it impossible for the program or debuggers to walk through or "unwind" subprogram stack frame chains.

While code semantics are preserved, applications such as exception handlers that rely on the default behavior for saves and restores can produce undefined results. When using **-qnounwind** in conjunction with the **-g** compiler option, debugging information regarding exception handling when unwinding the program's stack can be inaccurate.

---

## **-qversion**

### **Category**

Listings, messages, and compiler information

**@PROCESS**

None.

### **Purpose**

Displays the version and release of the compiler being invoked.

## Syntax



## Defaults

-qnoverversion

## Parameters

### verbose

Displays information about the version, release, and level of each compiler component installed.

## Usage

When you specify **-qversion**, the compiler displays the version information and exits; compilation is stopped. If you want to save this information to the output object file, you can do so with the **-qsaveopt -c** options.

**-qversion** specified without the **verbose** suboption shows compiler information in the format:

```
product_nameVersion: VV.RR.MMMM.LLLL
```

where:

*V* Represents the version.  
*R* Represents the release.  
*M* Represents the modification.  
*L* Represents the level.

For more details, see Example 1.

**-qversion=verbose** shows component information in the following format:

```
component_name Version: VV.RR(product_name) Level: component_build_date ID:  
component_level_ID
```

where:

*component\_name*  
Specifies an installed component, such as the low-level optimizer.  
*component\_build\_date*  
Represents the build date of the installed component.  
*component\_level\_ID*  
Represents the ID associated with the level of the installed component.

For more details, see Example 2.

## Example 1

The output of specifying the **-qversion** option:

```
IBM XL Fortran for Linux, V15.1.1 (5765-J10; 5725-C75)  
Version: 15.01.0001.0000
```

## Example 2

The output of specifying the **-qversion=verbose** option:

```
IBM XL Fortran for Linux, V15.1.1 (5765-J10; 5725-C75)
Version: 15.01.0001.0000
Driver Version: 15.1.1(Fortran) Level: 141104 ID: _WxqVIGREeEST_ev78AmGSg
Fortran Front End and Run Time Version: 15.1.1(Fortran) Level: 141030
ID: _BzfPYmCREeST_ev78AmGSg
Fortran Transformer Version: 15.1.1(Fortran) Level: 141030
ID: _B6IC4mCREeST_ev78AmGSg
High-Level Optimizer Version: 13.1.1(C/C++) and 15.1.1(Fortran) Level: 141031
ID: _AbMzFGEbEeST_ev78AmGSg
Low-Level Optimizer Version: 13.1.1(C/C++) and 15.1.1(Fortran) Level: 141105
ID: _0oSNcWUyEeST_ev78AmGS
```

### Related information

- “-qsaveopt” on page 220
- COMPILER\_VERSION

---

## -qwarn64

### Category

Error checking and debugging

### @PROCESS

None.

### Purpose

Displays informational messages identifying statements that may cause problems with 32-bit to 64-bit migration.

This option aids in porting code from a 32-bit to a 64-bit environment by detecting the truncation of an 8-byte integer pointer to 4 bytes.

### Syntax

►► -q nowarn64  
warn64 ◀◀

### Defaults

**-qnowarn64**

### Usage

The compiler flags the following situations with informational messages:

- The assignment of a reference to the **LOC** intrinsic to an **INTEGER(4)** variable.
- The assignment between an **INTEGER(4)** variable or **INTEGER(4)** constant and an integer pointer.
- The specification of an integer pointer within a common block.
- The specification of an integer pointer within an equivalence statement.

You can use interface blocks for argument checking.

## Related information

- Chapter 7, “Using XL Fortran in a 64-bit environment,” on page 281

---

## **-qxflag=dvz**

### **Category**

Error checking and debugging

### **@PROCESS**

None.

### **Purpose**

Causes the compiler to generate code to detect floating-point divide-by-zero operations.

### **Syntax**

►► -qxflag= —dvz —————►►

### **Defaults**

Not applicable.

### **Usage**

This option takes effect at optimization levels of **-0** or higher.

With this option on, the extra code calls the external handler function `__xl_dzx` when the divisor is zero. The return value of this function is used as the result of the division. Users are required to provide the function to handle the divide-by-zero operations. Specifying **-qxflag=dvz** handles only single-precision (REAL\*4) and double-precision (REAL\*8) division.

The interface of the function is as follows:

```
real(8) function __xl_dzx(x, y, kind_type)
  real(8), value :: x, y
  integer, value :: kind_type
end function
```

where:

**x** is the dividend value

**y** is the divisor value

#### **kind\_type**

specifies the size of the actual arguments associated with **x** and **y**.

A **kind\_type** value equal to zero indicates that the actual arguments associated with **x** and **y** are of type REAL(8). A **kind\_type** value equal to one indicates that the actual arguments associated with **x** and **y** are of type REAL(4).

The division always executes before the handler routine is called. This means that any exception is posted and handled before the handler function is called.

### Related information

- *Implementation details of XL Fortran floating-point processing in the XL Fortran Optimization and Programming Guide*
- “-qfltrap” on page 132
- “Understanding XL Fortran error messages” on page 283

---

## -qxflag=oldtab

### Category

Portability and migration

### Purpose

Interprets a tab in columns 1 to 5 as a single character (for fixed source form programs).

### Syntax

►► — -qxflag— = —oldtab—————►►

#### @PROCESS:

@PROCESS XFLAG(OLDTAB)

### Defaults

By default, the compiler allows 66 significant characters on a source line after column 6. A tab in columns 1 through 5 is interpreted as the appropriate number of blanks to move the column counter past column 6. This default is convenient for those who follow the earlier Fortran practice of including line numbers or other data in columns 73 through 80.

### Usage

If you specify the option **-qxflag=oldtab**, the source statement still starts immediately after the tab, but the tab character is treated as a single character for counting columns. This setting allows up to 71 characters of input, depending on where the tab character occurs.

---

## -qxlf77

### Category

Language element control


### Purpose

Provides compatibility with FORTRAN 77 aspects of language semantics and I/O data format that have changed.

Most of these changes are required by the Fortran 90 standard.

## Syntax

### Option:



### @PROCESS:

@PROCESS XLF77(*settings*)

## Defaults

By default, the compiler uses settings that apply to Fortran 90, Fortran 95, Fortran 2003, Fortran 2008, and the most recent compiler version in all cases.

The default suboptions are: **blankpad**, **nogedit77**, **nointarg**, **nointxor**, **leadzero**, **nooldboz**, **nopersistent**, and **nosofteof**.

These defaults are only used by the **xlf90**, **f90**, **xlf90\_r**, **xlf95**, **f95**, **xlf95\_r**, **xlf2003**, **f2003**, **xlf2003\_r**, **xlf2008**, **f2008**, and **xlf2008\_r** commands, which you should use to compile new programs.

## Parameters

To get various aspects of XL Fortran Version 2 behavior, select the nondefault choice for one or more of the following suboptions. The descriptions explain what happens when you specify the nondefault choices.

### blankpad | noblankpad

For internal, direct-access, and stream-access files, uses a default setting equivalent to **pad='no'**. This setting produces conversion errors when reading from such a file if the format requires more characters than the record has. This suboption does not affect direct-access or stream-access files opened with a **pad=** specifier.

### gedit77 | nogedit77

Uses FORTRAN 77 semantics for the output of **REAL** objects with the **G** edit descriptor. Between FORTRAN 77 and Fortran 90, the representation of 0 for a list item in a formatted output statement changed, as did the rounding method, leading to different output for some combinations of values and **G** edit descriptors.

### intarg | nointarg

Converts all integer arguments of an intrinsic procedure to the kind of the

longest argument if they are of different kinds. Under Fortran 90/95 rules, some intrinsics (for example, **IBSET**) determine the result type based on the kind of the first argument; others (for example, **MIN** and **MAX**) require that all arguments be of the same kind.

**intxor** | **nointxor**

Treats **.XOR.** as a logical binary intrinsic operator. It has a precedence equivalent to the **.EQV.** and **.NEQV.** operators and can be extended with an operator interface. (Because the semantics of **.XOR.** are identical to those of **.NEQV.**, **.XOR.** does not appear in the Fortran 90 or Fortran 95 language standard.)

Otherwise, the **.XOR.** operator is only recognized as a defined operator. The intrinsic operation is not accessible, and the precedence depends on whether the operator is used in a unary or binary context.

**leadzero** | **noleadzero**

Produces a leading zero in real output under the **D**, **E**, **L**, **F**, and **Q** edit descriptors.

**oldboz** | **nooldboz**

Turns blanks into zeros for data read by **B**, **O**, and **Z** edit descriptors, regardless of the **BLANK=** specifier or any **BN** or **BZ** control edit descriptors. It also preserves leading zeros and truncation of too-long output, which is not part of the Fortran 90 or Fortran 95 standard.

**persistent** | **nopersistent**

Saves the addresses of arguments to subprograms with **ENTRY** statements in static storage. This is an implementation choice that has been changed for increased performance.

**softeof** | **nosofteof**

Allows **READ** and **WRITE** operations when a unit is positioned after its endfile record unless that position is the result of executing an **ENDFILE** statement. This suboption reproduces a FORTRAN 77 extension of earlier versions of XL Fortran that some existing programs rely on.

## Usage

If you only want to compile and run old programs unchanged, you can continue to use the appropriate invocation command and not concern yourself with this option.

You should only use this option if you are using existing source or data files with Fortran 90, Fortran 95, Fortran 2003, and Fortran 2008 and the **xlf90**, **f90**, **xlf90\_r**, **xlf95**, **f95**, **xlf95\_r**, **xlf2003**, **f2003**, **xlf2003\_r**, **xlf2008**, **f2008**, or **xlf2008\_r** command and find some incompatibility because of behavior or data format that has changed.

Eventually, you should be able to recreate the data files or modify the source files to remove the dependency on the old behavior.

---

## -qxlf90

### Category


Language element control

## Purpose

Provides compatibility with the Fortran 90 standard for certain aspects of the Fortran language.

## Syntax

**Option:**



**@PROCESS:**

@PROCESS XLF90(*settings*)

## Defaults

The default suboptions for **-qxlf90** depend on the invocation command that you specify.

For the **xlf95**, **f95**, **xlf95\_r**, **xlf2003**, **f2003**, **xlf2003\_r**, **xlf2008**, **f2008**, or **xlf2008\_r** command, the default suboptions are **signedzero**, **autodealloc**, and **nooldpad**.

For all other invocation commands, the defaults are **nosignedzero**, **noautodealloc** and **oldpad**.

## Parameters

### signedzero | nosignedzero

Determines how the **SIGN(A,B)** function handles signed real 0.0. If you specify the **-qxlf90=signedzero** compiler option, **SIGN(A,B)** returns  $-|A|$  when  $B=-0.0$ . This behavior conforms to the Fortran 95 standard and is consistent with the IEEE standard for binary floating-point arithmetic. Note that for the **REAL(16)** data type, XL Fortran never treats zero as negative zero.

This suboption also determines whether a minus sign is printed in the following cases:

- For a negative zero in formatted output. Again, note that for the **REAL(16)** data type, XL Fortran never treats zero as negative zero.
- For negative values that have an output form of zero (that is, where trailing non-zero digits are truncated from the output so that the resulting output looks like zero). Note that in this case, the **signedzero** suboption does affect the **REAL(16)** data type; non-zero negative values that have an output form of zero will be printed with a minus sign.

When using **-qxlf90=nosignedzero**, consider setting the **-qstrict=nozerosigns** option to improve performance.

### autodealloc | noautodealloc

Determines whether the compiler deallocates allocatable objects that are declared locally without either the **SAVE** or the **STATIC** attribute and have a status of currently allocated when the subprogram terminates. This

behavior conforms with the Fortran 95 standard. If you are certain that you are deallocating all local allocatable objects explicitly, you may want to turn off this suboption to avoid possible performance degradation.

#### **oldpad | nooldpad**

When the **PAD=specifier** is present in the **INQUIRE** statement, specifying **-qx1f90=nooldpad** returns UNDEFINED when there is no connection, or when the connection is for unformatted I/O. This behavior conforms with the Fortran 95 standard and above. Specifying **-qx1f90=oldpad** preserves the Fortran 90 behavior.

### **Examples**

Consider the following program:

```
PROGRAM TESTSIGN
REAL X, Y, Z
X=1.0
Y=-0.0
Z=SIGN(X,Y)
PRINT *,Z
END PROGRAM TESTSIGN
```

The output from this example depends on the invocation command and the **-qx1f90** suboption that you specify. For example:

Invocation Command/xlf2008 Suboption	Output
xlf2008	-1.0
xlf2008 -qxlf90=signedzero	-1.0
xlf2008 -qxlf90=nosignedzero	1.0
xlf2003	-1.0
xlf2003 -qxlf90=signedzero	-1.0
xlf2003 -qxlf90=nosignedzero	1.0
xlf95	-1.0
xlf95 -qxlf90=signedzero	-1.0
xlf95 -qxlf90=nosignedzero	1.0
xlf90	1.0
xlf	1.0

### **Related information**

- See the **SIGN** information in the *Intrinsic Procedures* section and the *Array concepts* section of the *XL Fortran Language Reference*.
- “-qstrict” on page 238

---

## **-qxlf2003**

### **Category**

Language element control


### **Purpose**

Provides the ability to use language features specific to the Fortran 2003 standard when compiling with compiler invocations that conform to earlier Fortran

standards, as well as the ability to disable these features when compiling with compiler invocations that conform to the Fortran 2003 standard.

## Syntax

**Option:**



**@PROCESS:**

@PROCESS XLF2003(*suboption,suboption,...*)

## Defaults

The default suboption depends on the invocation command that you specify.

The **f2003**, **xf2003**, or **xf2003\_r** command has the following defaults:

**autorealloc:bozlitargs:nodynamicacval:nooldnaninf:polymorphic:signdzerointr:  
stopexcept:volatile**

All other invocation commands have the following defaults:

**noautorealloc:nobozlitargs:nodynamicacval:oldnaninf:nopolymorphic:  
nosigndzerointr:nostopexcept:novolatile**

## Parameters

### autorealloc | noautorealloc

Controls whether the compiler automatically reallocates the left-hand-side (LHS) with the shape of the right-hand-side (RHS) when assigning into an allocatable variable. If the LHS variable was not allocated before the assignment, it is allocated automatically. The default is **autorealloc** for the **f2003**, **xf2003**, and **xf2003\_r** commands, and **noautorealloc** for all other commands. This suboption has no effect on reallocation when the values of length type parameters in the LHS and RHS differ.

### bozlitargs | nobozlitargs

The **bozlitargs** suboption ensures that the passing of boz-literal constants as arguments to the **INT**, **REAL**, **CMPLX**, or **DBLE** intrinsic function

conforms to the Fortran 2003 standard. The default is **bozlitargs** for the **f2003**, **xlF2003**, and **xlF2003\_r** commands. The **-qLanglvl=2003pure** or **-qLanglvl=2003std** option must be specified, as well. If **-qport=typ1ssarg** and **-qxlF2003=bozlitargs** are specified, passing boz-literal constants to the **CMPLX** intrinsic will yield non-standard results.

#### **dynamicacval** | **nodynamicacval**

When **dynamicacval** is in effect, the dynamic types of array constructor values are used to determine the type of the array constructors and you can use unlimited polymorphic entities in array constructors. When **nodynamicacval** is in effect, the declared types of array constructor values are used to determine the type of the array constructors and you cannot use unlimited polymorphic entities in array constructors.

**Note:** To make the **-qxlF2003=dynamicacval** option effective, you must also specify **-qxlF2003=polymorphic**.

#### **oldnaninf** | **nooldnaninf**

The **oldnaninf** suboption controls the formatting of the output of IEEE NaN and infinity exceptional values. This suboption has no effect on input. When **oldnaninf** is in effect, the compiler uses the XL Fortran V10.1 (and earlier) behavior for output. That is, INF for infinity, NAN for a quiet or signaling NaN.

When **nooldnaninf** is in effect, the compiler output for IEEE exceptional values is compliant with the Fortran 2003 standard. That is, Inf for infinity, NaN(Q) for a quiet NaN, and NaN(S) for a signaling NaN.

#### **polymorphic** | **nopolymorphic**

When **polymorphic** is in effect, the compiler allows polymorphic items in Fortran source files. You can specify the **CLASS** type specifier, the **SELECT TYPE** construct, and use polymorphic items in other Fortran statements. The use of the polymorphic argument also causes the compiler to produce runtime type information for each derived type definition.

When **nopolymorphic** is in effect, polymorphic items cannot be specified in Fortran source files and no runtime type information is generated.

#### **signdzerointr** | **nosigndzerointr**

When **signdzerointr** is in effect, the passing of signed zeros to the **SQRT**, **LOG**, and **ATAN2** intrinsic functions returns results consistent with the Fortran 2003 standard. The **-qxlF90=signedzero** option must be in effect, as well. For the **xlF**, **xlF\_r**, **f77**, **fort77**, **xlF90**, **xlF90\_r**, and **f90** invocations, specify both options to have the Fortran 2003 behavior.

The following example shows the use of this suboption:

```
! If the Test program is compiled with -qxlF2003=signdzerointr
! and -qxlF90=signedzero, then Fortran 2003 behavior is seen.
! Otherwise, this program will demonstrate Fortran 95 behavior.
```

Program Test

```
real a, b
complex j, l
a = -0.0
j = sqrt(cmplx(-1.0,a))
b = atan2(a,-1.0)
l = log(cmplx(-1.0,a))
print *, 'j=', j
print *, 'b=', b
print *, 'l=', l
end
```

```

! Fortran 95 output:

j= (-0.0000000000E+00,1.000000000)
b= 3.141592741
l= (0.0000000000E+00,3.141592741)

! Fortran 2003 output:

j= (0.0000000000E+00,-1.000000000)
b= -3.141592741
l= (0.0000000000E+00,-3.141592741)

```

### **stopexcept | nostopexcept**

When **stopexcept** is in effect, informational messages are displayed when IEEE floating-point exceptions are signaled by a **STOP** statement. Messages have the format:

```

STOP [stop-code]
(OVERFLOW, DIV-BY-ZERO, INVALID, UNDERFLOW, INEXACT)

```

where *stop-code* corresponds to the optional digit string or character constant specified in the **STOP** statement. OVERFLOW, DIV-BY-ZERO, INVALID, UNDERFLOW and INEXACT appear only if the corresponding flag is set.

The following example shows corresponding messages generated:

```

real :: r11, r12, r13, r14
logical :: l

r11 = 1.3
r12 = 0.0

r13 = r11 / r12 ! divide by zero

r14 = r13 ! to make sure r13 is actually used

r14 = log(-r11) ! invalid input for log

stop "The End"

end

```

#### **Output:**

```

STOP The End
(DIV-BY-ZERO, INVALID)

```

When **nostopexcept** is in effect, informational messages are suppressed.

### **volatile | novolatile**

When **volatile** is in effect, a nonvolatile entity that is use- or host-associated can be specified as VOLATILE in inner or local scope.

## **Usage**

If the application uses F2003 polymorphism, you must compile every unit with **polymorphic** specified. If the application does not use polymorphism, specify the **nopolymorphic** suboption; doing so may save compilation time and potentially improve runtime performance.

## Related information

See the following information in *XL Fortran Language Reference*:

- Polymorphic entities
- Array constructors

---

## -qxlf2008

### Category

Language element control

### Purpose

Provides the ability to use language features specific to the Fortran 2008 standard when compiling with compiler invocations that conform to earlier Fortran standards, as well as the ability to disable these features when compiling with compiler invocations that conform to the Fortran 2008 standard.

### Syntax

#### Option:

►► — -qxlf2008=checkpresence | nocheckpresence —►►

#### @PROCESS:

@PROCESS XLF2008(*suboption*,*suboption*,...)

### Defaults

The default suboption depends on the invocation command that you specify.

The **f2008**, **xlf2008**, and **xlf2008\_r** commands have the following default:

#### **checkpresence**

All other invocation commands have the following default:

#### **nocheckpresence**

### Parameters

#### checkpresence | nocheckpresence

When **checkpresence** is in effect, dummy argument presence is checked according to the Fortran 2008 standard. When **nocheckpresence** is in effect, dummy argument presence is checked according to previous Fortran standards. For more information about dummy argument presence, see Restrictions on optional dummy arguments not present.

**Note:** When **-qxlf2008=checkpresence** is in effect, the performance of your program is inhibited because of runtime checks for the allocation and association status of actual arguments. To avoid these performance impacts, consider using **-qxlf2008=nocheckpresence**.

---

## -qxlines

### Category

Input control

### Purpose

Specifies whether fixed source form lines with an X in column 1 are compiled or treated as comments.

This option is similar to the recognition of the character 'd' in column 1 as a conditional compilation (debug) character. The **-qxlines** option recognizes the character 'x' in column 1 as a conditional compilation character when this compiler option is enabled. The 'x' in column 1 is interpreted as a blank, and the line is handled as source code.

### Syntax

#### Option:

►► -q noxl  
lines ◀◀

#### @PROCESS:

@PROCESS XLINES | NOXLINES

### Defaults

-qnoxl

This option is set to **-qnoxl** by default, and lines with the character 'x' in column 1 in fixed source form are treated as comment lines.

While the **-qxlines** option is independent of **-D**, all rules for debug lines that apply to using 'd' as the conditional compilation character also apply to the conditional compilation character 'x'.

The **-qxlines** compiler option is only applicable to fixed source form.

### Usage

The conditional compilation characters 'x' and 'd' may be mixed both within a fixed source form program and within a continued source line. If a conditional compilation line is continued onto the next line, all the continuation lines must have 'x' or 'd' in column 1. If the initial line of a continued compilation statement is not a debugging line that begins with either 'x' or 'd' in column 1, subsequent continuation lines may be designated as debug lines as long as the statement is syntactically correct.

The OMP conditional compilation characters '!', 'C\$', and '\*\$' may be mixed with the conditional characters 'x' and 'd' both in fixed source form and within a continued source line. The rules for OMP conditional characters will still apply in this instance.

## Examples

An example of a base case of `-qxlines`:

```
C2345678901234567890
  program p
    i=3 ; j=4 ; k=5
X   print *,i,j
X   +      ,k
    end program p

<output>: 3 4 5 (if -qxlines is on)
          no output (if -qxlines is off)
```

In this example, conditional compilation characters 'x' and 'd' are mixed, with 'x' on the initial line:

```
C2345678901234567890
  program p
    i=3 ; j=4 ; k=5
X   print *,i,
D   +      j,
X   +      k
    end program p

<output>: 3 4 5 (if both -qxlines and -qdlines are on)
          3 5 (if only -qxlines is turned on)
```

Here, conditional compilation characters 'x' and 'd' are mixed, with 'd' on the initial line:

```
C2345678901234567890
  program p
    i=3 ; j=4 ; k=5
D   print *,i,
X   +      j,
D   +      k
    end program p

<output>: 3 4 5 (if both -qxlines and -qdlines are on)
          3 5 (if only -qdlines is turned on)
```

In this example, the initial line is not a debug line, but the continuation line is interpreted as such, since it has an 'x' in column 1:

```
C2345678901234567890
  program p
    i=3 ; j=4 ; k=5
    print *,i
X   +      ,j
X   +      ,k
    end program p

<output>: 3 4 5 (if -qxlines is on)
          3 (if -qxlines is off)
```

## Related information

- “-D” on page 69
- *Conditional compilation in the XL Fortran Language Reference*

---

## -qxref

### Category

Listings, messages, and compiler information

## Purpose

Produces a compiler listing that includes the cross-reference component of the attribute and cross-reference section of the listing.

## Syntax



@PROCESS:

@PROCESS XREF[(FULL)] | NOXREF

## Defaults

-qnoxref

## Usage

If you specify only **-qxref**, only identifiers that are used are reported. If you specify **-qxref=full**, the listing contains information about all identifiers that appear in the program, whether they are used or not.

If **-qxref** is specified after **-qxref=full**, the full cross-reference listing is still produced.

You can use the cross-reference listing during debugging to locate problems such as using a variable before defining it or entering the wrong name for a variable.

## Related information

- “Listings, messages, and compiler information” on page 56
- “Attribute and cross reference section” on page 295

---

## -qzerosize


### Category

“Language element control” on page 50

### Purpose

Determines whether checking for zero-sized character strings and arrays takes place in programs that might process such objects.

### Syntax



@PROCESS:

@PROCESS ZEROSIZE | NOZEROSIZE

## Defaults

The default setting depends on which command invokes the compiler:

- **-qzerosize** for the **f90**, **xlF90**, **xlF90\_r**, **f95**, **xlF95**, **xlF95\_r**, **f2003**, **xlF2003**, **xlF2003\_r**, **f2008**, **xlF2008**, and **xlF2008\_r** commands
- **-qnozerosize** for the **xlF** and **xlF\_r** commands when they are used to compile the **.f**, **.F**, **.f77**, or **.F77** files
- **-qnozerosize** for the **f77** and **fort77** commands

## Usage

Use **-qzerosize** for Fortran 90, Fortran 95, Fortran 2003, and Fortran 2008 programs that might process zero-sized character strings and arrays.

For FORTRAN 77 programs, where zero-sized objects are not allowed, or for Fortran 90 and Fortran 95 programs that do not use them, compiling with **-qnozerosize** can improve the performance of some array or character-string operations.

Runtime checking performed by the **-C** option takes slightly longer when **-qzerosize** is in effect.

---

**-r**

## Category

Object code control

## @PROCESS

None.

## Purpose

Produces a nonexecutable output file to use as an input file in another **ld** command call. This file may also contain unresolved symbols.

## Syntax

►► -r ◀◀

## Defaults

Not applicable.

## Usage

A file produced with this flag is expected to be used as an input file in another compiler invocation or **ld** command call.

## Predefined macros

None.

## Examples

To compile myprogram.f and myprog2.f into a single object file mytest.o, enter:  
xlf myprogram.f myprog2.f -r -o mytest.o

---

## -S

### Category

Output control

### @PROCESS

None.

### Purpose

Generates an assembler language file for each source file.

### Syntax

►► -S \_\_\_\_\_ ◀◀

### Rules

When this option is specified, the compiler produces the assembler source files as output instead of an object or an executable file.

### Restrictions

The generated assembler files do not include all the data that is included in a .o file by **-qipa** or **-g**.

### Examples

```
xlf95 -03 -qhot -S test.f           # Produces test.s
```

### Related information

The **-o** option can be used to specify a name for the resulting assembler source file.

---

## -t

### Category

Compiler customization

### @PROCESS

None.

### Purpose

Applies the prefix specified by the **-B** option to the designated components.

## Syntax



## Defaults

The default paths for all of the compiler components are defined in the compiler configuration file.

## Parameters

The following table shows the correspondence between `-t` parameters and the component names:

Parameter	Description	Component name
a	The assembler	as
b	The low-level optimizer	xlfcodes
c	The compiler front end	xlfcodes
d	The disassembler	dis
F	The C preprocessor	cpp
h	The array language optimizer	xlfcodes
I (uppercase i)	The high-level optimizer, compile step	ipa
l (lowercase L)	The linker	ld
z	The binder	bolt

## Usage

Use this option with the `-Bprefix` option.

## Examples

To compile `myprogram.f` so that the name `/u/newones/compilers/` is prefixed to the compiler and assembler program names, enter:

```
xlfcodes myprogram.f -B/u/newones/compilers/ -tca
```

## Related information

- “-B” on page 67

---

## **-U**

### **Category**

Input control

### **Purpose**

Makes the compiler sensitive to the case of letters in names.

### **Syntax**

▶▶ -U ◀◀

#### **@PROCESS:**

@PROCESS MIXED | NOMIXED

### **Defaults**

By default, the compiler interprets all names as if they were in lowercase. For example, `Abc` and `ABC` are both interpreted as `abc` and so refer to the same object.

### **Usage**

You can use this option when writing mixed-language programs, because Fortran names are all lowercase by default, while names in C and other languages may be mixed-case.

If `-U` is specified, case is significant in names. For example, the names `Abc` and `ABC` refer to different objects.

This option changes the link names used to resolve calls between compilation units. It also affects the names of modules and `▶ F2008` submodules `◀ F2008` and thus the names of their `.mod` and `.smod` files.

### **Restrictions**

The names of intrinsics must be all in lowercase when `-U` is in effect. Otherwise, the compiler may accept the names without errors, but the compiler considers them to be the names of external procedures, rather than intrinsics.

### **Related information**

This is the short form of `-qmixed`. See “`-qmixed`” on page 186.

---

## **-u**

### **Category**

Language element control

### **Purpose**

Specifies that no implicit typing of variable names is permitted.

It has the same effect as using the **IMPLICIT NONE** statement in each scope that allows implicit statements.

## Syntax

►► -u ◀◀

**@PROCESS:**

@PROCESS UNDEF | NUNDEF

## Defaults

**-qnundef**, which allows implicit typing.

## Related information

See **IMPLICIT** in the *XL Fortran Language Reference*.

This is the short form of **-qundef**. See “-qundef” on page 253.

---

## -v

### Category

Listings, messages, and compiler information

### @PROCESS

None.

### Purpose

Reports the progress of compilation, by naming the programs being invoked and the options being specified to each program.

### Syntax

►► -v ◀◀

### Defaults

Not applicable.

### Usage


For a particular compilation, examining the output that this option produces can help you determine:

- What files are involved
- What options are in effect for each step
- How far a compilation gets when it fails

### Related information

- “-#” on page 66 is similar to **-v**, but it does not actually execute any of the compilation steps.





## Parameters

*option*

Any option that is valid for the component to which it is being passed.

The following table shows the correspondence between **-W** parameters and the component names:

Parameter	Description	Component name
a	The assembler	as
b	The low-level optimizer	xlfcodes
c	The compiler front end	xlfcntry
d	The disassembler	dis
F	The C preprocessor	cpp
h	The array language optimizer	xlfcot
I (uppercase i)	The high-level optimizer, compile step	ipa
L	The high-level optimizer, link step	ipa
l (lowercase L)	The linker	ld
z	The binder	bolt

## Usage

In the string following the **-W** option, use a comma as the separator for each option, and do not include any spaces. If you need to include a character that is special to the shell in the option string, precede the character with a backslash. For example, if you use the **-W** option in the configuration file, you can use the escape sequence backslash comma (\,) to represent a comma in the parameter string.

You do not need the **-W** option to pass most options to the linker **ld**; unrecognized command-line options, except **-q** options, are passed to it automatically. Only linker options with the same letters as compiler options, such as **-v** or **-S**, strictly require **-W**.

## Examples

To compile the file `file.f` and pass the linker option **-berok** to the linker, enter the following command:

```
xlf -Wl,-berok file.f
```

To compile the file `uses_many_symbols.f` and the assembly file `produces_warnings.s` so that `produces_warnings.s` is assembled with the assembler option `-x` (issue warnings and produce cross-reference), and the object files are linked with the option `-s` (write list of object files and strip final executable file), issue the following command:

```
xlf -Wa,-x -Wl,-s produces_warnings.s uses_many_symbols.f
```

In the following example, the `\`, embeds a literal comma in the `-WF` string and causes three arguments, rather than four, to be supplied to the C preprocessor.

```
$ xlf -qfree=f90 '-WF,-Dint1=1,-Dint2=2,-Dlist=3\,4' a.F
$ cat a.F
print *, int1
print *, int2
print *, list
end
```

The output from the program will be:

```
$ ./a.out
1
2
3 4
```

## Related information

- “Passing command-line options to the `ld` or `as` command” on page 28

---

## -W

### Category

Listings, messages, and compiler information

### @PROCESS

None.

### Purpose

Suppresses informational, language-level and warning messages (equivalent to `-qflag=e:e`).

### Syntax

▶▶ — -W —————▶▶

### Defaults

Not applicable.

### Related information

- “`-qflag`” on page 126

---

## **-y**

### **Category**

Floating-point and integer control

### **Purpose**

Specifies the rounding mode for the compiler to use when evaluating constant floating-point expressions at compile time.

It is equivalent to the **-qieee** option.

### **Syntax**



**@PROCESS:**

@PROCESS IEEE(Near | Minus | Plus | Zero)

### **Defaults**

-yn

### **Parameters**

- n** Round to nearest.
- m** Round toward minus infinity.
- p** Round toward plus infinity.
- z** Round toward zero.

### **Usage**

If your program contains operations involving real(16) values, the rounding mode must be set to **-yn**, round-to-nearest.

### **Related information**

- “-O” on page 83
- “-qfloat” on page 127
- “-qieee” on page 146



---

## Chapter 7. Using XL Fortran in a 64-bit environment

The 64-bit environment addresses an 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 space.



---

## Chapter 8. Problem determination and debugging

This section describes some methods you can use for locating and fixing problems in compiling or executing your programs.

---

### Understanding XL Fortran error messages

Most information about potential or actual problems comes through messages from the compiler or application program. These messages are written to the standard error stream.

#### Error severity

Compilation errors can have the following severity levels (from the highest to the lowest):

- U** An unrecoverable error. Compilation failed because of an internal compile-time error.
- S** A severe error. Compilation failed due to one of the following:
  - An unrecoverable program error has been detected. Processing of the source file stops, and XL Fortran does not produce an object file. You can usually correct this error by fixing any program errors that were reported during compilation.
  - Conditions exist that the compiler could not correct. An object file is produced; however, you should not attempt to run the program.
  - An internal compiler table has overflowed. Processing of the program stops, and XL Fortran does not produce an object file.
  - An include file does not exist. Processing of the program stops, and XL Fortran does not produce an object file.
- E** An error that the compiler can correct. The program should run correctly.
- W** Warning message. It does not signify an error but may indicate some unexpected condition.
- L** Warning message that was generated by one of the compiler options that check for conformance to various language levels. It may indicate a language feature that you should avoid if you are concerned about portability.
- I** Informational message. It does not indicate any error, just something that you should be aware of to avoid unexpected behavior or to improve performance.

**Note:**

- The message levels **S** and **U** indicate a compilation failure.
- The message levels **I**, **L**, **W**, and **E** do not indicate a compilation failure.

By default, the compiler stops without producing output files if it encounters a severe error (severity **S**). However, you can control how the messages are generated by using the following options:

- You can make the compiler stop for less severe errors by specifying a different severity with the **-qhalt** option. For example, with **-qhalt=e**, the compiler stops

if it encounters any errors of severity E or higher severity. This technique can reduce the amount of compilation time that is needed to check the syntactic and semantic validity of a program.

- You can limit low-severity messages without stopping the compiler by using the **-qflag** option.
- You can stop compilation when a specified error message is generated by using the **-qhaltormsg** option.
- You can stop compilation when a specified number of errors of a specified minimum severity level is reached by using the **-qmaxerr** option.
- If you simply want to prevent specific messages from going to the output stream, see **-qsuppress**.

## Compiler return codes

The compiler return codes and their respective meanings are as follows:

- 0 The compiler did not encounter any errors severe enough to make it stop processing a compilation unit.
- 1 The compiler encountered an error of severity high enough to halt the compilation. Depending on the level of *halt\_severity*, the compiler might have continued processing the compilation units with errors.
- 40 An option error.
- 41 A configuration file error.
- 250 An out-of-memory error. The compiler cannot allocate any more memory for its use.
- 251 A signal received error. An unrecoverable error or interrupt signal is received.
- 252 A file-not-found error.
- 253 An input/output error. Cannot read or write files.
- 254 A fork error. Cannot create a new process.
- 255 An error while executing a process.

## Runtime return codes

If an XLF-compiled program ends abnormally, the return code to the operating system is 1.

If the program ends normally, the return code is 0 (by default) or is  $\text{MOD}(\text{digit\_string}, 256)$  if the program ends because of a **STOP** *digit\_string* statement.

## Format of XL Fortran diagnostic messages

In addition to the diagnostic message issued, the source line and a pointer to the position in the source line at which the error was detected are printed or displayed if you specify the **-qsource** compiler option. If **-qnosource** is in effect, the file name, the line number, and the column position of the error are displayed with the message.

The format of an XL Fortran diagnostic message is:

```

▶▶—15—cc—--nnn— ———┌—————┐—————▶
                       | (—severity_letter—) | message_text

```

where:

- 15 Indicates an XL Fortran message

*cc* Is the component number, as follows:

- 00 Indicates a code generation or optimization message
- 01 Indicates an XL Fortran common message
- 11-20 Indicates a Fortran-specific message
- 25 Indicates a runtime message from an XL Fortran application program
- 85 Indicates a loop-transformation message
- 86 Indicates an interprocedural analysis (IPA) message
- 87 Indicates a runtime message from the SMP library

*nnn* Is the message number

*severity\_letter*  
Indicates how serious the problem is, as described in the preceding section

*'message text'*  
Is the text describing the error

## Limiting the number of compile-time messages

If the compiler issues many low-severity (I or W) messages concerning problems you are aware of or do not care about, use the **-qflag** option or its short form **-w** to limit messages to high-severity ones:

```
# E, S, and U messages go in listing; U messages are displayed on screen.
xlf95 -qflag=e:u program.f
```

```
# E, S, and U messages go in listing and are displayed on screen.
```

```
xlf95 -w program.f
```

---

## Fixing installation or system environment problems

If individual users or all users on a particular machine have difficulty running the compiler, there may be a problem in the system environment. Here are some common problems and solutions:

*invocation\_command*: **not found**

Message:

```
xlf90: not found
xlf90_r: not found
xlf95: not found
xlf95_r: not found
xlf: not found
xlf_r: not found
xlf2003: not found
xlf2003_r: not found
xlf2008: not found
xlf2008_r: not found
f77: not found
fort77: not found
f90: not found
f95: not found
f2003: not found
f2008: not found
```

**System action:**

The shell cannot locate the command to execute the compiler.

**User response:**

Make sure that your **PATH** environment variable includes the directory `/opt/ibm/xlf/15.1.1/bin`. If the compiler is properly installed, the commands you need to execute it are in this directory.

**Could not load program *program*****Message:**

```
Could not load program program  
Error was: not enough space
```

**System action:**

The system cannot execute the compiler or an application program at all.

**User response:**

Set the storage limits for stack and data to “unlimited” for users who experience this problem. For example, you can set both your hard and soft limits with these **bash** commands:

```
ulimit -s unlimited  
ulimit -d unlimited
```

Or, you may find it more convenient to edit the file `/etc/security/limits.conf` to give all users unlimited stack and data segments (by entering -1 for these fields).

If the storage problem is in an XLF-compiled program, using the **-qsave** or **-qsmallstack** option might prevent the program from exceeding the stack limit.

**Explanation:**

The compiler allocates large internal data areas that may exceed the storage limits for a user. XLF-compiled programs place more data on the stack by default than in previous versions, also possibly exceeding the storage limit. Because it is difficult to determine precise values for the necessary limits, we recommend making them unlimited.

**Could not load library *library\_name*****Message:**

```
Could not load program program  
Could not load library library_name.so  
Error was: no such file or directory
```

**User response:**

Make sure the XL Fortran libraries are installed in `/opt/ibm/xlf/15.1.1/lib`, or set the **LD\_RUN\_PATH** environment variables to include the directory where `libxlf90.alibxlf90.so` is installed if it is in a different directory. See “Setting library search paths” on page 8 for details of this environment variable.

**A compilation fails with an I/O error.****System action:**

A compilation fails with an I/O error.

**User response:**

Increase the size of the /tmp filesystem, or set the environment variable **TMPDIR** to the path of a filesystem that has more free space.

**Explanation:**

The object file may have grown too large for the filesystem that holds it. The cause could be a very large compilation unit or initialization of all or part of a large array in a declaration.

**Too many individual makefiles and compilation scripts****System action:**

There are too many individual makefiles and compilation scripts to easily maintain or track.

**User response:**

Add stanzas to the configuration file, and create links to the compiler by using the names of these stanzas. By running the compiler with different command names, you can provide consistent groups of compiler options and other configuration settings to many users.

---

## Fixing compile-time problems

The following sections discuss common problems you might encounter while compiling and how to avoid them.

### Duplicating extensions from other systems

Some ported programs may cause compilation problems because they rely on extensions that exist on other systems. XL Fortran supports many extensions like these, but some require compiler options to turn them on. See “Portability and migration” on page 61 for a list of these options and *Porting programs to XL Fortran* in the *XL Fortran Optimization and Programming Guide* for a general discussion of porting.

### Isolating problems with individual compilation units

If you find that a particular compilation unit requires specific option settings to compile properly, you may find it more convenient to apply the settings in the source file through an **@PROCESS** directive. Depending on the arrangement of your files, this approach may be simpler than recompiling different files with different command-line options.

### Compiling with threadsafe commands

Threadsafe invocation commands like **xl\_f\_r** or **xl\_f90\_r**, for example, use different search paths and call different modules than the non threadsafe invocations. Your programs should account for the different usages. Programs that compile and run successfully for one environment may produce unexpected results when compiled and run for a different use. The configuration file, **xl\_f.cfg**, shows the paths, libraries, and so on for each invocation command. (See “Editing the default configuration file” on page 13 for an explanation of its contents.)

## Running out of machine resources

If the operating system runs low on resources (page space or disk space) while one of the compiler components is running, you should receive one of the following messages:

1501-229 Compilation ended because of lack of space.

1517-011 Compilation ended. No more system resources available.

1501-053 (S) Too much initialized data.

1501-511. Compilation failed for file [filename].

You may need to increase the system page space and recompile your program. See the man page information **man 8 mkswap swapon** for more information about page space.

If your program produces a large object file, for example, by initializing all or part of a large array, you may need to do one of the following:

- Increase the size of the filesystem that holds the /tmp directory.
- Set the **TMPDIR** environment variable to a filesystem with a lot of free space.
- For very large arrays, initialize the array at run time rather than statically (at compile time).

## Compilation failure because of internal alias table overflow

At an optimization level of **-O2** or higher, XL Fortran compiler computes alias information to ensure that the optimizer preserves the semantics of the application. Alias computation might be an expensive process in terms of space usage, depending on the object definitions and language constructs used in an application. In rare occasions, the internal alias table might overflow, which causes the compilation to fail. Such overflow might occur when a module defines or uses derived-type objects that contain a large number of ultimate components, and passes the objects as dummy arguments to a large number of procedures that are defined in the same module. Partitioning the set of module procedures and placing each partition in a separate module might resolve the overflow problem.

---

## Fixing link-time problems

After the XL Fortran compiler processes the source files, the linker links the resulting object files together. Any messages issued at this stage come from the **ld** command. A frequently encountered error and its solution are listed here for your convenience:

### Undefined or unresolved symbols detected

**Message:**

**filename.o(text+0x14): undefined reference to "p"**

**filename.o(text+0x14): relocation truncated to fit: R\_PPC\_REL24 p**

**System action:**

A program cannot be linked because of unresolved references.

**Explanation:**

Either needed object files or libraries are not being used during linking, there is an error in the specification of one or more external names, or there is an error in the specification of one or more procedure interfaces.

**User response:**

You may need to do one or more of the following actions:

- Compile again with the **-W1, -M** option to create a file that contains information about undefined symbols.
- Make sure that if you use the **-U** option, all intrinsic names are in lowercase.

---

## Fixing runtime problems

XL Fortran issues error messages during the running of a program in either of the following cases:

- XL Fortran detects an input/output error. “Setting runtime options” on page 34 explains how to control these kinds of messages.
- XL Fortran detects an exception error, and the default exception handler is installed (through the **-qsigtrap** option or a call to **SIGNAL**). To get a more descriptive message than Core dumped, you may need to run the program from within **gdb**.

The causes for runtime exceptions are listed in “XL Fortran runtime exceptions” on page 45.

You can investigate errors that occur during the execution of a program by using a symbolic debugger, such as **gdb**.

### Duplicating extensions from other systems

Some ported programs may not run correctly if they rely on extensions that are found on other systems. XL Fortran supports many such extensions, but you need to turn on compiler options to use some of them. See “Portability and migration” on page 61 for a list of these options and *Porting programs to XL Fortran* in the *XL Fortran Optimization and Programming Guide* for a general discussion of porting.

### Mismatched sizes or types for arguments

Arguments of different sizes or types might produce incorrect execution and results. To do the type-checking during the early stages of compilation, specify interface blocks for the procedures that are called within a program.

### Working around problems when optimizing

If you find that a program produces incorrect results when it is optimized and if you can isolate the problem to a particular variable, you might be able to work around the problem temporarily by declaring the variable as **VOLATILE**. This prevents some optimizations that affect the variable. (See **VOLATILE** in the *XL Fortran Language Reference*.) Because this is only a temporary solution, you should continue debugging your code until you resolve your problem, and then remove the **VOLATILE** keyword. If you are confident that the source code and program design are correct and you continue to have problems, contact your support organization to help resolve the problem.

### Input/Output errors

If the error detected is an input/output error and you have specified **IOSTAT** on the input/output statement in error, the **IOSTAT** variable is assigned a value according to *Conditions and IOSTAT values* in the *XL Fortran Language Reference*.

If you have installed the XL Fortran runtime message catalog on the system on which the program is executing, a message number and message text are issued to the terminal (standard error) for certain I/O errors. If you have specified **IOMSG** on the input/output statement, the **IOMSG** variable is assigned the error message text if an error is detected, or the content of **IOMSG** variable is not changed. If this catalog is not installed on the system, only the message number appears. Some of the settings in “Setting runtime options” on page 34 allow you to turn some of these error messages on and off.

If a program fails while writing a large data file, you may need to increase the maximum file size limit for your user ID. You can do this through a shell command, such as **ulimit** in **bash**.

## Tracebacks and core dumps

If a runtime exception occurs and an appropriate exception handler is installed, a message and a traceback listing are displayed. Depending on the handler, a core file might be produced as well. You can then use a debugger to examine the location of the exception.

To produce a traceback listing without ending the program, call the **xl\_\_trbk** procedure:

```
IF (X .GT. Y) THEN      ! X > Y indicates that something is wrong.
  PRINT *, 'Error - X should not be greater than Y'
  CALL XL__TRBK        ! Generate a traceback listing.
END IF
```

See *Installing an exception handler* in the *XL Fortran Optimization and Programming Guide* for instructions about exception handlers and “XL Fortran runtime exceptions” on page 45 for information about the causes of runtime exceptions.

---

## Debugging a Fortran program

You can use **gdb** and other symbolic debuggers to debug your programs. For instructions on using your chosen debugger, consult the online help within the debugger or its documentation.

Always specify the **-g** option when compiling programs for debugging.

**Note:** Debugging Fortran 2003 polymorphic objects and parameterized derived types is not supported in this release of XL Fortran.

### Related information:

- “Error checking and debugging” on page 54

---

## Chapter 9. Understanding XL Fortran compiler listings

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

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

- 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 precedes the section heading so that you can easily locate its beginning:

```
>>>> SECTION NAME <<<<<<
```

You can select which sections appear in the listing by specifying compiler options.

**Related information:** See “Listings, messages, and compiler information” on page 56.

---

### Header section

The listing file has a header section that contains the following items:

- A compiler identifier that consists of the following elements:
  - Compiler name
  - Version number
  - Release number
  - Modification number
  - Fix number
- Source file name
- Date of compilation
- Time of compilation

The header section is always present in a listing; it is the first line and appears only once. The following sections are repeated for each compilation unit when more than one compilation unit is present.

---

### Options section

The options section is always present in a listing. There is a separate section for each compilation unit. It indicates the specified options that are in effect for the compilation unit. This information is useful when you have conflicting options. If you specify the **-qlistopt** compiler option, this section lists the settings for all options.

---

## Source section

The source section contains the input source lines with a line number and, optionally, a file number. The file number indicates the source file (or include file) from which the source line originated. All main file source lines (those that are not from an include file) do not have the file number printed. Each include file has a file number associated with it, and source lines from include files have that file number printed. The file number appears on the left, the line number appears to its right, and the text of the source line is to the right of the line number. XL Fortran numbers lines relative to each file. The source lines and the numbers that are associated with them appear only if the **-qsource** compiler option is in effect. You can selectively print parts of the source by using the **@PROCESS** directives **SOURCE** and **NOSOURCE** throughout the program.

## Error messages

If the **-qsource** option is in effect, the error messages are interspersed with the source listing. The error messages that are generated during the compilation process contain the following elements:

- The source line
- A line of indicators that point to the columns that are in error
- The error message, which consists of the following elements:
  - The 4-digit component number
  - The number of the error message
  - The severity level of the message
  - The text that describes the error

For example:

```
          2 |      equivalence (i,j,i,j)
            |      .....a.b.
a - "t.f", line 2.24: 1514-117 (E) Same name appears more than once in an equivalence group.
b - "t.f", line 2.26: 1514-117 (E) Same name appears more than once in an equivalence group.
```

If the **-qnosource** option is in effect, the error messages are all that appear in the source section, and an error message contains:

- The file name in quotation marks
- The line number and column position of the error
- The error message, which consists of the following elements:
  - The 4-digit component number
  - The number of the error message
  - The severity level of the message
  - The text that describes the error

For example:

```
"doc.f", line 6.11: 1513-039 (S) Number of arguments is not
permitted for INTRINSIC function abs.
```

---

## PDF report section

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

### Loop iteration count

The most frequent loop iteration count and the average iteration count, for

a given set of input data, is calculated for most loops in a program. This information is only available when the program is compiled at optimization level -O5.

### **Block and call count**

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

### **Cache miss**

This section is printed in a single table. It reports the number of *Cache Misses* for certain functions, with additional information about the functions such as: Cache Level, Cache Miss Ratio, Line Number, File Name, and Memory Reference.

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

### **Relevance of profiling data**

This section includes the relevance between the source code during the `-qpdf1` phase and the profiling data. 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 levels are different between the `-qpdf1` and `-qpdf2` phases.

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

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

---

## **Transformation report section**

If the `-qreport` option is in effect, a transformation report listing shows how IBM XL Fortran for Linux, V15.1.1 optimized the program. This LOOP TRANSFORMATION section displays pseudo-Fortran code that corresponds to the original source code, so that you can see parallelization and loop transformations that the `-qhot` or `-qsmp` options have generated. This section of the report also shows information on the additional transformations and parallelization performed on loop nests if you compile with `-qsmp` and `-qhot=level=2`.

The compiler also reports the number of streams created for a given loop. You can use this information to understand your application code and to tune your code for better performance. For example, you can distribute a loop which has more streams than the number supported by the underlying architecture.

To generate information about where the compiler inserted data prefetch instructions, use the optimization level of **-qhot**, **-03**, **-04** or **-05** together with **-qreport**.

### Sample Report

The following report was created for the program **t.f** using the **xlf -qhot -qreport t.f**

command.

#### Program t.f:

```
integer a(100, 100)
integer i,j

do i = 1 , 100
  do j = 1, 100
    a(i,j) = j
  end do
end do
end
```

#### Transformation Report:

>>>> SOURCE SECTION <<<<<

\*\* \_main === End of Compilation 1 ===

>>>> LOOP TRANSFORMATION SECTION <<<<<

```
PROGRAM _main ()
4|   IF (.FALSE.) GOTO lab_9
   @CIV2 = 0
Id=1 DO @CIV2 = @CIV2, 24
5|   IF (.FALSE.) GOTO lab_11
   @LoopIV1 = 0
6|   @CSE0 = @CIV2 * 4
   @ICM0 = @CSE0 + 1
   @ICM1 = @CSE0 + 2
   @ICM2 = @CSE0 + 3
   @ICM3 = @CSE0 + 4
5|Id=2 DO @LoopIV1 = @LoopIV1, 99
   ! DIR_INDEPENDENT loopId = 0
   ! DIR_INDEPENDENT loopId = 0
6|   @CSE1 = @LoopIV1 + 1
   SHADV_M003_a(@CSE1,@ICM0) = @ICM0
   SHADV_M002_a(@CSE1,@ICM1) = @ICM1
   SHADV_M001_a(@CSE1,@ICM2) = @ICM2
   SHADV_M000_a(@CSE1,@ICM3) = @ICM3
7|   ENDDO
   lab_11
8|   ENDDO
   lab_9
4|   IF (.FALSE.) THEN
   @LoopIV0 = int((100 - MOD(100, int(4))))
Id=5 DO @LoopIV0 = @LoopIV0, 100
5|   IF (.FALSE.) GOTO lab_19
   @LoopIV1 = 0
6|   @ICM4 = @LoopIV0 + 1
```

```

5|Id=6          DO @LoopIV1 = @LoopIV1, 99
                ! DIR_INDEPENDENT loopId = 0
                ! DIR_INDEPENDENT loopId = 0
6|              a((@LoopIV1 + 1),@ICM4) = @ICM4
7|              ENDDO
                lab_19
8|              ENDDO
                ENDIF
9|              END PROGRAM _main

```

Source File	Source Line	Loop Id	Action / Information
0	4	1	Loop interchanging applied to loop nest.
0	4	1	Outer loop has been unrolled 4 time(s).

---

## Data reorganization report section

A summary of useful information about how program variable data gets reorganized by the compiler, *data reorganizations*.

To generate data reorganization information, specify the optimization level **-qipa=level=2** or **-05** together with **-qreport**. During the IPA link pass, the data reorganization messages for program variable data is produced to the data reorganization section of the listing file. Reorganization includes:

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

---

## Attribute and cross reference section

This section provides information about the entities that are used in the compilation unit. It is present if the **-qxref** or **-qattr** compiler option is in effect. Depending on the options in effect, this section contains all or part of the following information about the entities that are used in the compilation unit:

- Names of the entities
- Attributes of the entities (if **-qattr** is in effect). Attribute information may include any or all of the following details:
  - The class of the name
  - The type
  - The relative address of the name
  - Alignment
  - Dimensions
  - For an array, whether it is allocatable
  - Whether it is a pointer, target, or integer pointer
  - Whether it is a parameter
  - Whether it is volatile
  - For a dummy argument, its intent, whether it is value, and whether it is optional
  - For a module entity, whether it is private, public, or protected
- Coordinates to indicate where you have defined, referenced, or modified the entities. If you declared the entity, the coordinates are marked with a \$. If you initialized the entity, the coordinates are marked with a \*. If you both declared

and initialized the entity at the same place, the coordinates are marked with a &. If the entity is set, the coordinates are marked with a @. If the entity is referenced, the coordinates are not marked.

Class is one of the following:

- Automatic
- BSS (uninitialized static internal)
- Common
- Common block
- Construct name
- Controlled (for an allocatable object)
- Controlled automatic (for an automatic object)
- Defined assignment
- Defined operator
- Derived type definition
- Entry
- External subprogram
- Function
- Generic name
- Internal subprogram
- Intrinsic
- Module
- Module function
- Module subroutine
- Namelist
- Pointee
- Private component
- Program
- Reference argument
- Renames
- Static
- Submodule
- Subroutine
- Use associated
- Value parameter

If you specify the **full** suboption with **-qxref** or **-qattr**, XL Fortran reports all entities in the compilation unit. If you do not specify this suboption, only the entities you actually use appear.

---

## Object section

XL Fortran produces this section only when the **-qlist** compiler option is in effect. It contains the object code listing, which shows the source line number, the instruction offset in hexadecimal notation, the assembler mnemonic of the instruction, and the hexadecimal value of the instruction. On the right side, it also shows the cycle time of the instruction and the intermediate language of the compiler. Finally, the total number of machine instructions that are produced and the total cycle time (straight-line execution time) are displayed. There is a separate section for each compilation unit.

---

## **File table section**

This section contains a table that shows the file number and file name for each main source file and include file used. It also lists the line number of the main source file at which the include file is referenced. This section is always present. The table also includes the file creation date and time.

---

## **Compilation unit epilogue Section**

This is the last section of the listing for each compilation unit. It contains the diagnostics summary and indicates whether the unit was compiled successfully. This section is not present in the listing if the file contains only one compilation unit.

---

## **Compilation epilogue Section**

The compilation epilogue section occurs only once at the end of the listing. At completion of the compilation, XL Fortran presents a summary of the compilation: number of source records that were read, compilation start time, compilation end time, total compilation time, total CPU time, virtual CPU time, and a summary of diagnosed conditions. This section is always present in a listing.



---

## Chapter 10. XL Fortran technical information

This section contains details about XL Fortran that advanced programmers may need to diagnose unusual problems, run the compiler in a specialized environment, or do other things that a casual programmer is rarely concerned with.

---

### 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 routines in the runtime libraries are prefixed with an underscore (`_`), or `_xl`.

---

### The XL Fortran runtime environment

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

Unless you bind statically, you cannot run object code produced by the XL Fortran compiler without the XL Fortran Runtime Environment.

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.

---

## Technical details of the `-qfloat=hsflt` option

The `-qfloat=hsflt` option is unsafe for optimized programs that compute floating-point values that are outside the range of representation of single precision, not just outside the range of the result type. The range of representation includes both the precision and the exponent range.

Even when you follow the rules that are stated in the preceding paragraph and in “`-qfloat`” on page 127, programs that are sensitive to precision differences might not produce expected results. Because `-qfloat=hsflt` is not compliant with IEEE, programs will not always run as expected.

For example, in the following program, `X.EQ.Y` may be true or may be false:

```
REAL X, Y, A(2)
DOUBLE PRECISION Z
LOGICAL SAME

READ *, Z
X = Z
Y = Z
IF (X.EQ.Y) SAME = .TRUE.
! ...
! ... Calculations that do not change X or Y
! ...
CALL SUB(X)           ! X is stored in memory with truncated fraction.
IF (X.EQ.Y) THEN     ! Result might be different than before.
...

A(1) = Z
X = Z
A(2) = 1.             ! A(1) is stored in memory with truncated fraction.
IF (A(1).EQ.X) THEN ! Result might be different than expected.
...

```

If the value of `Z` has fractional bits that are outside the precision of a single-precision variable, these bits may be preserved in some cases and lost in others. This makes the exact results unpredictable when the double-precision value of `Z` is assigned to single-precision variables. For example, passing the variable as a dummy argument causes its value to be stored in memory with a fraction that is truncated rather than rounded.

---

## Implementation details for `-qautodbl` promotion and padding

The following sections provide additional details about how the `-qautodbl` option works, to allow you to predict what happens during promotion and padding.

### Terminology

The *storage relationship* between two data objects determines the relative starting addresses and the relative sizes of the objects. The `-qautodbl` option tries to preserve this relationship as much as possible.

Data objects can also have a *value relationship*, which determines how changes to one object affect another. For example, a program might store a value into one variable, and then read the value through a different storage-associated variable. With `-qautodbl` in effect, the representation of one or both variables might be different, so the value relationship is not always preserved.

An object that is affected by this option can be:

- *Promoted*, meaning that it is converted to a higher-precision data type. Usually, the resulting object is twice as large as it would be by default. Promotion applies to constants, variables, derived-type components, arrays, and functions (which include intrinsic functions) of the appropriate types.

**Note:** **BYTE**, **INTEGER**, **LOGICAL**, and **CHARACTER** objects are never promoted.

- *Padded*, meaning that the object keeps its original type but is followed by undefined storage space. Padding applies to **BYTE**, **INTEGER**, **LOGICAL**, and nonpromoted **REAL** and **COMPLEX** objects that may share storage space with promoted items. For safety, **POINTERS**, **TARGETS**, actual and dummy arguments, members of **COMMON** blocks, structures, pointer arrays, and pointer **COMPLEX** objects are always padded appropriately depending on the **-qautodbl** suboption. This is true whether or not they share storage with promoted objects.

Space added for padding ensures that the storage-sharing relationship that existed before conversion is maintained. For example, if array elements **I(20)** and **R(10)** start at the same address by default and if the elements of **R** are promoted and become twice as large, the elements of **I** are padded so that **I(20)** and **R(10)** still start at the same address.

Except for unformatted I/O statements, which read and write any padding that is present within structures, I/O statements do not process padding.

**Note:** The compiler does not pad **CHARACTER** objects.

## Examples of storage relationships for **-qautodbl** suboptions

The examples in this section illustrate storage-sharing relationships between the following types of entities:

- **REAL(4)**
- **REAL(8)**
- **REAL(16)**
- **COMPLEX(4)**
- **COMPLEX(8)**
- **COMPLEX(16)**
- **INTEGER(8)**
- **INTEGER(4)**
- **CHARACTER(16)**.

**Note:** In the diagrams, solid lines represent the actual data, and dashed lines represent padding.




Figure 5. Storage relationships without the `-qautodbl` option

The figure above illustrates the default storage-sharing relationship of the compiler.

```
@process autodbl(none)
  block data
    complex(4) x8      /(1.123456789e0,2.123456789e0)/
    real(16) r16(2)   /1.123q0,2.123q0/
    integer(8) i8(2)  /1000,2000/
    character*5 c(2)  /"abcde","12345"/
    common /named/ x8,r16,i8,c
  end

  subroutine s()
    complex(4) x8
    real(16) r16(2)
    integer(8) i8(2)
    character*5 c(2)
    common /named/ x8,r16,i8,c
    !      x8      = (1.123456e0,2.123456e0)      ! promotion did not occur
    !      r16(1) = 1.123q0                        ! no padding
    !      r16(2) = 2.123q0                        ! no padding
    !      i8(1)  = 1000                            ! no padding
    !      i8(2)  = 2000                            ! no padding
    !      c(1)   = "abcde"                        ! no padding
    !      c(2)   = "12345"                        ! no padding
  end subroutine s
```




Figure 6. Storage relationships with `-qautodbl=dbl`

```
@process autodbl(dbl)
  block data
    complex(4) x8
```

```

real(16) r16(2)    /1.123q0,2.123q0/
real(8)  r8
real(4)  r4        /1.123456789e0/
integer(8) i8(2)  /1000,2000/
character*5 c(2)  /"abcde","12345"/
equivalence (x8,r8)
common /named/ r16,i8,c,r4
! Storage relationship between r8 and x8 is preserved.
! Data values are NOT preserved between r8 and x8.
end

subroutine s()
  real(16) r16(2)
  real(8)  r4
  integer(8) i8(2)
  character*5 c(2)
  common /named/ r16,i8,c,r4
! r16(1) = 1.123q0                                ! no padding
! r16(2) = 2.123q0                                ! no padding
! r4 = 1.123456789d0                               ! promotion occurred
! i8(1) = 1000                                     ! no padding
! i8(2) = 2000                                     ! no padding
! c(1) = "abcde"                                   ! no padding
! c(2) = "12345"                                   ! no padding
end subroutine s

```




Figure 7. Storage relationships with `-qautobl=dbl4`

```

@process autodb1(db14)
  complex(8) x16    /(1.123456789d0,2.123456789d0)/
  complex(4) x8
  real(4)  r4(2)
  equivalence (x16,x8,r4)
! Storage relationship between r4 and x8 is preserved.
! Data values between r4 and x8 are preserved.
! x16 = (1.123456789d0,2.123456789d0)           ! promotion did not occur
! x8 = (1.123456789d0,2.123456789d0)           ! promotion occurred
! r4(1) = 1.123456789d0                          ! promotion occurred
! r4(2) = 2.123456789d0                          ! promotion occurred
end

```




Figure 8. Storage relationships with `-qautodbl=dbl8`

```

@process autodbl(dbl8)
  complex(8) x16  /(1.123456789123456789d0,2.123456789123456789d0)/
  complex(4) x8
  real(8) r8(2)
  equivalence (x16,x8,r8)
! Storage relationship between r8 and x16 is preserved.
! Data values between r8 and x16 are preserved.
! x16 = (1.123456789123456789q0,2.123456789123456789q0)
!
!                                     ! promotion occurred
! x8 = upper 8 bytes of r8(1)         ! promotion did not occur
! r8(1) = 1.123456789123456789q0     ! promotion occurred
! r8(2) = 2.123456789123456789q0     ! promotion occurred
end

```

## XL Fortran internal limits

Language Feature	Limit
Maximum number of iterations performed by <b>DO</b> loops with loop control with index variable of type <b>INTEGER(n)</b> for $n = 1, 2$ or $4$	$(2^{**}31)-1$
Maximum number of iterations performed by <b>DO</b> loops with loop control with index variable of type <b>INTEGER(8)</b>	$(2^{**}63)-1$
Maximum character format field width	$(2^{**}31)-1$
Maximum length of a format specification	$(2^{**}31)-1$
Maximum length of Hollerith and character constant edit descriptors	$(2^{**}31)-1$
Maximum length of a fixed source form statement	34 000
Maximum length of a free source form statement	34 000
Maximum number of continuation lines	n/a <b>1</b>
Maximum number of nested <b>INCLUDE</b> lines	64
Maximum number of nested interface blocks	1 024
Maximum number of statement numbers in a computed <b>GOTO</b>	999
Maximum number of times a format code can be repeated	$(2^{**}31)-1$
Allowable record numbers and record lengths for input/output files	The record number can be up to $(2^{**}63)-1$ , and the record length can be up to $(2^{**}63)-1$ bytes.  However, for unformatted sequential files, you must use the <b>uwidth=64</b> runtime option for the record length to be greater than $(2^{**}31)-1$ and up to $(2^{**}63)-1$ . If you use the default <b>uwidth=32</b> runtime option, the maximum length of a record in an unformatted sequential file is $(2^{**}31)-1$ bytes.
Allowable bound range of an array dimension	The bound of an array dimension can be positive, negative, or zero within the range $-(2^{**}63)$ to $2^{**}63-1$ .
Allowable external unit numbers	0 to $(2^{**}31)-1$ <b>2</b>
Maximum numeric format field width	2 000
Maximum number of concurrent open files	1 024 <b>3</b>

**1** You can have as many continuation lines as you need to create a statement with a maximum of 34 000 bytes.

**2** The value must be representable in an **INTEGER(4)** object, even if specified by an **INTEGER(8)** variable.

**3** In practice, this value is somewhat lower because of files that the runtime system may open, such as the preconnected units 0, 5, and 6.

---

## Glossary

This glossary defines terms that are commonly used in this document. It includes definitions developed by the American National Standards Institute (ANSI) and entries from the *IBM Terminology* website.

### A

#### **abstract interface**

An **ABSTRACT INTERFACE** consists of procedure characteristics and names of dummy arguments. Used to declare the interfaces for procedures and deferred bindings.

#### **abstract type**

A type that has the **ABSTRACT** attribute. A nonpolymorphic object cannot be declared to be of abstract type. A polymorphic object cannot be constructed or allocated to have a dynamic type that is abstract.

#### **active processor**

See *online processor*.

#### **actual argument**

An expression, variable, procedure, or alternate return specifier that is specified in a procedure reference.

**alias** A single piece of storage that can be accessed through more than a single name. Each name is an alias for that storage.

#### **alphanumeric character**

A letter or other symbol, excluding digits, used in a language. Usually the uppercase and lowercase letters A through Z plus other special symbols (such as \$ and \_) allowed by a particular language.

#### **alphanumeric**

Pertaining to a character set that contains letters, digits, and usually other characters, such as punctuation marks and mathematical symbols.

#### **American National Standard Code for Information Interchange**

See *ASCII*.

#### **argument**

An expression that is passed to a function or subroutine. See also *actual argument*, *dummy argument*.

#### **argument association**

The relationship between an actual argument and a dummy argument during the invocation of a procedure.

#### **arithmetic constant**

A constant of type integer, real, or complex.

#### **arithmetic expression**

One or more arithmetic operators and arithmetic primaries, the evaluation of which produces a numeric value. An arithmetic expression can be an unsigned arithmetic constant, the name of an arithmetic constant, or a reference to an arithmetic variable, function reference, or a combination of such primaries formed by using arithmetic operators and parentheses.

**arithmetic operator**

A symbol that directs the performance of an arithmetic operation. The intrinsic arithmetic operators are:

+	addition
-	subtraction
*	multiplication
/	division
**	exponentiation

**array** An entity that contains an ordered group of scalar data. All objects in an array have the same data type and type parameters.

**array declarator**

The part of a statement that describes an array used in a program unit. It indicates the name of the array, the number of dimensions it contains, and the size of each dimension.

**array element**

A single data item in an array, identified by the array name and one or more subscripts. See also *subscript*.

**array name**

The name of an ordered set of data items.

**array section**

A subobject that is an array and is not a structure component.

**ASCII** The standard code, using a coded character set consisting of 7-bit coded characters (8-bits including parity check), that is used for information interchange among data processing systems, data communication systems, and associated equipment. The ASCII set consists of control characters and graphic characters. See also *Unicode*.

**asynchronous**

Pertaining to events that are not synchronized in time or do not occur in regular or predictable time intervals.

**assignment statement**

An executable statement that defines or redefines a variable based on the result of expression evaluation.

**associate name**

The name by which a selector of a **SELECT TYPE** or **ASSOCIATE** construct is known within the construct.

**assumed-size array**

A dummy array whose size is assumed from the associated actual argument. Its last upper bound is specified by an asterisk.

**assumed-type object**

An entity declared with **TYPE(\*)**. An assumed-type object does not have a declared type, and its dynamic type and type parameters are assumed from its corresponding actual argument.

**attribute**

A property of a data object that may be specified in a type declaration statement, attribute specification statement, or through a default setting.

**automatic parallelization**

The process by which the compiler attempts to parallelize both explicitly coded **DO** loops and **DO** loops generated by the compiler for array language.

**B****base object**

An object that is designated by the leftmost *part\_name*.

**base type**

An extensible type that is not an extension of another type.

**binary constant**

A constant that is made of one or more binary digits (0 and 1).

**bind**

To relate an identifier to another object in a program; for example, to relate an identifier to a value, an address or another identifier, or to associate formal parameters and actual parameters.

**binding label**

A value of type default character that uniquely identifies how a variable, common block, subroutine, or function is known to a companion processor.

**blank common**

An unnamed common block.

**block data subprogram**

A subprogram headed by a **BLOCK DATA** statement and used to initialize variables in named common blocks.

**bounds\_remapping**

Allows a user to view a flat, rank-1 array as a multi-dimensional array.

**bss storage**

Uninitialized static storage.

**busy-wait**

The state in which a thread keeps executing in a tight loop looking for more work once it has completed all of its work and there is no new work to do.

**byte constant**

A named constant that is of type byte.

**byte type**

A data type representing a one-byte storage area that can be used wherever a **LOGICAL(1)**, **CHARACTER(1)**, or **INTEGER(1)** can be used.

**C****C descriptor**

A C structure of type `CFI_cdesc_t` that is defined in the `ISO_Fortran_binding.h` header file.

**character constant**

A string of one or more alphabetic characters enclosed in apostrophes or double quotation marks.

**character expression**

A character object, a character-valued function reference, or a sequence of them separated by the concatenation operator, with optional parentheses.

**character operator**

A symbol that represents an operation, such as concatenation (`//`), to be performed on character data.

**character set**

All the valid characters for a programming language or for a computer system.

**character string**

A sequence of consecutive characters.

**character substring**

A contiguous portion of a character string.

**character type**

A data type that consists of alphanumeric characters. See also *data type*.

**chunk** A subset of consecutive loop iterations.

**class** A set of types comprised of a base type and all types extended from it.

**collating sequence**

The sequence in which the characters are ordered for the purpose of sorting, merging, comparing, and processing indexed data sequentially.

**comment**

A language construct for the inclusion of text in a program that has no effect on the execution of the program.

**common block**

A storage area that may be referred to by a calling program and one or more subprograms.

**compile**

To translate a source program into an executable program (an object program).

**compiler comment directive**

A line in source code that is not a Fortran statement but is recognized and acted on by the compiler.

**compiler directive**

Source code that controls what XL Fortran does rather than what the user program does.

**complex constant**

An ordered pair of real or integer constants separated by a comma and enclosed in parentheses. The first constant of the pair is the real part of the complex number; the second is the imaginary part.

**complex number**

A number consisting of an ordered pair of real numbers, expressible in the form  $a+bi$ , where  $a$  and  $b$  are real numbers and  $i$  squared equals  $-1$ .

**complex type**

A data type that represents the values of complex numbers. The value is expressed as an ordered pair of real data items separated by a comma and enclosed in parentheses. The first item represents the real part of the complex number; the second represents the imaginary part.

**component**

A constituent of a derived type.

**component order**

The ordering of the components of a derived type that is used for intrinsic formatted input/output and for structure constructors.

**conform**

To adhere to a prevailing standard. An executable program conforms to the Fortran 95 Standard if it uses only those forms and relationships described therein and if the executable program has an interpretation according to the Fortran 95 Standard. A program unit conforms to the Fortran 95 Standard if it can be included in an executable program in a manner that allows the executable program to be standard-conforming. A processor conforms to the standard if it executes standard-conforming programs in a manner that fulfills the interpretations prescribed in the standard.

**connected unit**

In XL Fortran, a unit that is connected to a file in one of three ways: explicitly via the **OPEN** statement to a named file, implicitly, or by preconnection.

**constant**

A data object with a value that does not change. The four classes of constants specify numbers (arithmetic), truth values (logical), character data (character), and typeless data (hexadecimal, octal, and binary). See also *variable*.

**construct**

A sequence of statements starting with a **SELECT CASE**, **DO**, **IF**, or **WHERE** statement, for example, and ending with the corresponding terminal statement.

**contiguous**

An array is contiguous if it has array elements in order that are not separated by other data objects. A data object with multiple parts is contiguous if the parts in order are not separated by other data objects.

**continuation line**

A line that continues a statement beyond its initial line.

**control statement**

A statement that is used to alter the continuous sequential invocation of statements; a control statement may be a conditional statement, such as **IF**, or an imperative statement, such as **STOP**.

**D****data object**

A variable, constant, or subobject of a constant.

**data striping**

Spreading data across multiple storage devices so that I/O operations can be performed in parallel for better performance. Also known as *disk striping*.

**data transfer statement**

A **READ**, **WRITE**, or **PRINT** statement.

**data type**

The properties and internal representation that characterize data and functions. The intrinsic types are integer, real, complex, logical, and character. See also *intrinsic*.

**debug line**

Allowed only for fixed source form, a line containing source code that is to be used for debugging. Debug lines are defined by a D or X in column 1. The handling of debug lines is controlled by the **-qdlines** and **-qxlines** compiler options.

**decimal symbol**

The symbol that separates the whole and fractional parts of a real number.

**declared type**

The type that a data entity is declared to have. May differ from the type during execution (the dynamic type) for polymorphic data entities.

**default initialization**

The initialization of an object with a value specified as part of a derived type definition.

**deferred binding**

A binding with the **DEFERRED** attribute. A deferred binding can only appear in an abstract type definition.

**definable variable**

A variable whose value can be changed by the appearance of its name or designator on the left of an assignment statement.

**delimiters**

A pair of parentheses or slashes (or both) used to enclose syntactic lists.

**denormalized number**

An IEEE number with a very small absolute value and lowered precision. A denormalized number is represented by a zero exponent and a non-zero fraction.

**derived type**

A type whose data have components, each of which is either of intrinsic type or of another derived type.

**digit** A character that represents a nonnegative integer. For example, any of the numerals from 0 through 9.

**directive**

A type of comment that provides instructions and information to the compiler.

**disk striping**

See *data striping*.

**DO loop**

A range of statements invoked repetitively by a **DO** statement.

**DO variable**

A variable, specified in a **DO** statement, that is initialized or incremented prior to each occurrence of the statement or statements within a **DO** loop. It is used to control the number of times the statements within the range are executed.

**DOUBLE PRECISION constant**

A constant of type real with twice the precision of the default real precision.

**dummy argument**

An entity whose name appears in the parenthesized list following the procedure name in a **FUNCTION**, **SUBROUTINE**, **ENTRY**, or statement function statement.

**dynamic dimensioning**

The process of re-evaluating the bounds of an array each time the array is referenced.

**dynamic extent**

For a directive, the lexical extent of the directive and all subprograms called from within the lexical extent.

**dynamic type**

The type of a data entity during execution of a program. The dynamic type of a data entity that is not polymorphic is the same as its declared type.

**E****edit descriptor**

An abbreviated keyword that controls the formatting of integer, real, or complex data.

**effective item**

A scalar object resulting from expanding an input/output list.

**elemental**

Pertaining to an intrinsic operation, procedure or assignment that is applied independently to elements of an array or corresponding elements of a set of conformable arrays and scalars.

**embedded blank**

A blank that is surrounded by any other characters.

**entity** A general term for any of the following: a program unit, procedure, operator, interface block, common block, external unit, statement function, type, named variable, expression, component of a structure, named constant, statement label, construct, or namelist group.

**environment variable**

A variable that describes the operating environment of the process.

**epoch** The starting date used for time in POSIX. It is Jan 01 00:00:00 GMT 1970.

**executable program**

A program that can be executed as a self-contained procedure. It consists of a main program and, optionally, modules, submodules, subprograms and non-Fortran external procedures.

**executable statement**

A statement that causes an action to be taken by the program; for example, to perform a calculation, test conditions, or alter normal sequential execution.

**explicit initialization**

The initialization of an object with a value in a data statement initial value list, block data program unit, type declaration statement, or array constructor.

**explicit interface**

For a procedure referenced in a scoping unit, the property of being an internal procedure, module procedure, intrinsic procedure, external

procedure that has an interface block, recursive procedure reference in its own scoping unit, or dummy procedure that has an interface block.

**expression**

A sequence of operands, operators, and parentheses. It may be a variable, a constant, or a function reference, or it may represent a computation.

**extended-precision constant**

A processor approximation to the value of a real number that occupies 16 consecutive bytes of storage.

**extended type**

An extensible type that is an extension of another type. A type that is declared with the **EXTENDS** attribute.

**extensible type**

A type from which new types may be derived using the **EXTENDS** attribute. A nonsequence type that does not have the **BIND** attribute.

**extension type**

A base type is an extension type of itself only. An extended type is an extension type of itself and of all types for which its parent type is an extension.

**external file**

A sequence of records on an input/output device. See also *internal file*.

**external name**

The name of a common block, subroutine, or other global procedure, which the linker uses to resolve references from one compilation unit to another.

**external procedure**

A procedure that is defined by an external subprogram or by a means other than Fortran.

**F**

**field** An area in a record used to contain a particular category of data.

**file** A sequence of records. See also *external file*, *internal file*.

**file index**

See *i-node*.

**final subroutine**

A subroutine that is called automatically during finalization.

**finalizable**

A type that has final subroutines, or that has a finalizable component. An object of finalizable type.

**finalization**

The process of calling user-defined final subroutines immediately before destroying an object.

**floating-point number**

A real number represented by a pair of distinct numerals. The real number is the product of the fractional part, one of the numerals, and a value obtained by raising the implicit floating-point base to a power indicated by the second numeral.

**format**

A defined arrangement of such things as characters, fields, and lines, usually used for displays, printouts, or files.

To arrange such things as characters, fields, and lines.

**formatted data**

Data that is transferred between main storage and an input/output device according to a specified format. See also *list-directed* and *unformatted record*.

**function**

A procedure that returns the value of a single variable or an object and usually has a single exit. See also *intrinsic procedure*, *subprogram*.

**G****generic identifier**

A lexical token that appears in an **INTERFACE** statement and is associated with all the procedures in an interface block.

**H****hard limit**

A system resource limit that can only be raised or lowered by using root authority, or cannot be altered because it is inherent in the system or operating environments's implementation. See also *soft limit*.

**hexadecimal**

Pertaining to a system of numbers to the base sixteen; hexadecimal digits range from 0 (zero) through 9 (nine) and A (ten) through F (fifteen).

**hexadecimal constant**

A constant, usually starting with special characters, that contains only hexadecimal digits.

**high order transformations**

A type of optimization that restructures loops and array language.

**Hollerith constant**

A string of any characters capable of representation by XL Fortran and preceded with *nH*, where *n* is the number of characters in the string.

**host**

A main program or subprogram that contains an internal procedure is called the host of the internal procedure. A module or submodule that contains a module procedure is called the host of the module procedure. A module or submodule is called the host of its descendant submodules.

**host association**

The process by which an internal subprogram, module subprogram, derived-type definition, or submodule accesses the entities of its host.

**host instance**

An instance of the host procedure that supplies the host environment of the internal procedure.

**I****IPA**

Interprocedural analysis, a type of optimization that allows optimizations to be performed across procedure boundaries and across calls to procedures in separate source files.

**implicit interface**

A procedure referenced in a scoping unit other than its own is said to have an implicit interface if the procedure is an external procedure that does not

have an interface block, a dummy procedure that does not have an interface block, or a statement function.

**implied DO**

An indexing specification (similar to a **DO** statement, but without specifying the word **DO**) with a list of data elements, rather than a set of statements, as its range.

**infinity**

An IEEE number (positive or negative) created by overflow or division by zero. Infinity is represented by an exponent where all the bits are 1's, and a zero fraction.

**inherit**

To acquire from a parent. Type parameters, components, or procedure bindings of an extended type that are automatically acquired from its parent type without explicit declaration in the extended type are said to be inherited.

**inheritance association**

The relationship between the inherited components and the parent component in an extended type.

**i-node** The internal structure that describes the individual files in the operating system. There is at least one i-node for each file. An i-node contains the node, type, owner, and location of a file. A table of i-nodes is stored near the beginning of a file system. Also known as *file index*.

**input/output (I/O)**

Pertaining to either input or output, or both.

**input/output list**

A list of variables in an input or output statement specifying the data to be read or written. An output list can also contain a constant, an expression involving operators or function references, or an expression enclosed in parentheses.

**integer constant**

An optionally signed digit string that contains no decimal point.

**interface block**

A sequence of statements from an **INTERFACE** statement to its corresponding **END INTERFACE** statement.

**interface body**

A sequence of statements in an interface block from a **FUNCTION** or **SUBROUTINE** statement to its corresponding **END** statement.

**interference**

A situation in which two iterations within a **DO** loop have dependencies upon one another.

**internal file**

A sequence of records in internal storage. See also *external file*.

**interprocedural analysis**

See *IPA*.

**intrinsic**

Pertaining to types, operations, assignment statements, and procedures that are defined by Fortran language standards and can be used in any scoping unit without further definition or specification.

**intrinsic module**

A module that is provided by the compiler and is available to any program.

**intrinsic procedure**

A procedure that is provided by the compiler and is available to any program.

**K****keyword**

A statement keyword is a word that is part of the syntax of a statement (or directive) and may be used to identify the statement.

An argument keyword specifies the name of a dummy argument

**kind type parameter**

A parameter whose values label the available kinds of an intrinsic type or a derived-type parameter that is declared to have the **KIND** attribute.

**L****lexical extent**

All of the code that appears directly within a directive construct.

**lexical token**

A sequence of characters with an indivisible interpretation.

**link-edit**

To create a loadable computer program by means of a linker.

**linker** A program that resolves cross-references between separately compiled or assembled object modules and then assigns final addresses to create a single relocatable load module. If a single object module is linked, the linker simply makes it relocatable.

**list-directed**

A predefined input/output format that depends on the type, type parameters, and values of the entities in the data list.

**literal** A symbol or a quantity in a source program that is itself data, rather than a reference to data.

**literal constant**

A lexical token that directly represents a scalar value of intrinsic type.

**load balancing**

An optimization strategy that aims at evenly distributing the work load among processors.

**logical constant**

A constant with a value of either true or false (or T or F).

**logical operator**

A symbol that represents an operation on logical expressions:

.NOT.	(logical negation)
.AND.	(logical conjunction)
.OR.	(logical union)
.EQV.	(logical equivalence)
.NEQV.	(logical nonequivalence)
.XOR.	(logical exclusive disjunction)

**loop** A statement block that executes repeatedly.

**M**

**\_main** The default name given to a main program by the compiler if the main program was not named by the programmer.

**main program**

The first program unit to receive control when a program is run. See also *subprogram*.

**master thread**

The head process of a team of threads.

**module**

A program unit that contains or accesses definitions to be accessed by other program units.

**module procedure interface body**

An interface body whose initial statement contains the **MODULE** prefix specifier. A module procedure interface body specifies the interface of a separate module procedure.

**module subprogram**

A subprogram that is contained in a module or submodule, but is not an internal subprogram. A module subprogram is a function subprogram, a subroutine subprogram, or a separate module subprogram.

**mutex** A primitive object that provides mutual exclusion between threads. A mutex is used cooperatively between threads to ensure that only one of the cooperating threads is allowed to access shared data or run certain application code at a time.

**N**

**NaN (not-a-number)**

A symbolic entity encoded in floating-point format that does not correspond to a number. See also *quiet NaN*, *signaling NaN*.

**name** A lexical token consisting of a letter followed by up to 249 alphanumeric characters (letters, digits, and underscores). Note that in FORTRAN 77, this was called a symbolic name.

**named common**

A separate, named common block consisting of variables.

**namelist group name**

The first parameter in the NAMELIST statement that names a list of names to be used in READ, WRITE, and PRINT statements.

**negative zero**

An IEEE representation where the exponent and fraction are both zero, but the sign bit is 1. Negative zero is treated as equal to positive zero.

**nest** To incorporate a structure or structures of some kind into a structure of the same kind. For example, to nest one loop (the nested loop) within another loop (the nesting loop); to nest one subroutine (the nested subroutine) within another subroutine (the nesting subroutine).

**NEWUNIT value**

A negative number that is less than -2 and is unequal to the unit number of any currently connected file. It is a unit value that the runtime library assigns to the variable specified by the **NEWUNIT=** specifier.

**nonexecutable statement**

A statement that describes the characteristics of a program unit, data, editing information, or statement functions, but does not cause any action to be taken by the program.

**nonexisting file**

A file that does not physically exist on any accessible storage medium.

**normal**

A floating-point number that is not denormal, infinity, or NaN.

**not-a-number**

See *NaN*.

**numeric constant**

A constant that expresses an integer, real, complex, or byte number.

**numeric storage unit**

The space occupied by a nonpointer scalar object of type default integer, default real, or default logical.

**O**

**octal** Pertaining to a system of numbers to the base eight; the octal digits range from 0 (zero) through 7 (seven).

**octal constant**

A constant that is made of octal digits.

**one-trip DO-loop**

A **DO** loop that is executed at least once, if reached, even if the iteration count is equal to 0. (This type of loop is from FORTRAN 66.)

**online processor**

In a multiprocessor machine, a processor that has been activated (brought online). The number of online processors is less than or equal to the number of physical processors actually installed in the machine. Also known as *active processor*.

**operator**

A specification of a particular computation involving one or two operands.

**P**

**pad** To fill unused positions in a field or character string with dummy data, usually zeros or blanks.

**paging space**

Disk storage for information that is resident in virtual memory but is not currently being accessed.

**parent component**

The component of an entity of extended type that corresponds to its inherited portion.

**parent type**

The extensible type from which an extended type is derived.

**passed-object dummy argument**

The dummy argument of a type-bound procedure or procedure pointer component that becomes associated with the object through which the procedure was invoked.

**PDF** See *profile-directed feedback*.

**pointee array**

An explicit-shape or assumed-size array that is declared in an integer **POINTER** statement or other specification statement.

**pointer**

A variable that has the **POINTER** attribute. A pointer must not be referenced or defined unless it is pointer associated with a target. If it is an array, it does not have a shape unless it is pointer-associated.

**polymorphic**

Able to be of differing types during program execution. An object declared with the **CLASS** keyword is polymorphic.

**preconnected file**

A file that is connected to a unit at the beginning of execution of the executable program. Standard error, standard input, and standard output are preconnected files (units 0, 5 and 6, respectively).

**predefined convention**

The implied type and length specification of a data object, based on the initial character of its name when no explicit specification is given. The initial characters I through N imply type integer of length 4; the initial characters A through H, O through Z, \$, and \_ imply type real of length 4.

**present**

A dummy argument is present in an instance of a subprogram if it is associated with an actual argument and the actual argument is a dummy argument that is present in the invoking procedure or is not a dummy argument of the invoking procedure.

**primary**

The simplest form of an expression: an object, array constructor, structure constructor, function reference, or expression enclosed in parentheses.

**procedure**

A computation that may be invoked during program execution. It may be a function or a subroutine. It may be an intrinsic procedure, an external procedure, a module procedure, an internal procedure, a dummy procedure, or a statement function. A subprogram may define more than one procedure if it contains **ENTRY** statements.

**procedure binding**

See type-bound procedure.

**procedure pointer**

A procedure entity that has the **EXTERNAL** and **POINTER** attributes. It can be pointer associated with an external procedure, a module procedure, a dummy procedure or another procedure pointer.

**profile-directed feedback (PDF)**

A type of optimization that uses information collected during application execution to improve performance of conditional branches and in frequently executed sections of code.

**program state**

The values of user variables at certain points during the execution of a program.

**program unit**

A main program or subprogram.

**pure**

An attribute of a procedure that indicates there are no side effects.

## Q

### quiet NaN

A NaN (not-a-number) value that does not signal an exception. The intent of a quiet NaN is to propagate a NaN result through subsequent computations. See also *NaN*, *signaling NaN*.

## R

### random access

An access method in which records can be read from, written to, or removed from a file in any order. See also *sequential access*.

**rank** The number of dimensions of an array.

### real constant

A string of decimal digits that expresses a real number. A real constant must contain a decimal point, a decimal exponent, or both.

**record** A sequence of values that is treated as a whole within a file.

### relational expression

An expression that consists of an arithmetic or character expression, followed by a relational operator, followed by another arithmetic or character expression.

### relational operator

The words or symbols used to express a relational condition or a relational expression:

.GT.	greater than
.GE.	greater than or equal to
.LT.	less than
.LE.	less than or equal to
.EQ.	equal to
.NE.	not equal to

### result variable

The variable that returns the value of a function.

### return specifier

An argument specified for a statement, such as *CALL*, that indicates to which statement label control should return, depending on the action specified by the subroutine in the *RETURN* statement.

## S

**scalar** A single datum that is not an array.

Not having the property of being an array.

### scale factor

A number indicating the location of the decimal point in a real number (and, on input, if there is no exponent, the magnitude of the number).

**scope** That part of an executable program within which a lexical token has a single interpretation.

### scope attribute

That part of an executable program within which a lexical token has a single interpretation of a particular named property or entity.

### scoping unit

A derived-type definition.

A **BLOCK** construct (not including any nested **BLOCK** constructs, derived-type definitions, and interface bodies within it).

An interface body.

A program unit or subprogram, excluding derived-type definitions, **BLOCK** constructs, interface bodies, and subprograms contained within it.

**selector**

The object that is associated with the associate name in an **ASSOCIATE** construct.

**semantics**

The relationships of characters or groups of characters to their meanings, independent of the manner of their interpretation and use. See also *syntax*.

**separate module procedure**

A module procedure that is defined by a separate module subprogram or by a function or subroutine subprogram whose initial statement contains the **MODULE** prefix specifier. A separate module procedure defines a module procedure interface body.

**separate module subprogram**

A module subprogram whose initial statement contains the **MODULE PROCEDURE** statement. A separate module subprogram defines a module procedure interface body.

**sequential access**

An access method in which records are read from, written to, or removed from a file based on the logical order of the records in the file. See also *random access*.

**signaling NaN**

A NaN (not-a-number) value that signals an invalid operation exception whenever it appears as an operand. The intent of the signaling NaN is to catch program errors, such as using an uninitialized variable. See also *NaN*, *quiet NaN*.

**sleep** The state in which a thread completely suspends execution until another thread signals it that there is work to do.

**SMP** See *symmetric multiprocessing*.

**soft limit**

A system resource limit that is currently in effect for a process. The value of a soft limit can be raised or lowered by a process, without requiring root authority. The soft limit for a resource cannot be raised above the setting of the hard limit. See also *hard limit*.

**spill space**

The stack space reserved in each subprogram in case there are too many variables to hold in registers and the program needs temporary storage for register contents.

**specification statement**

A statement that provides information about the data used in the source program. The statement could also supply information to allocate data storage.

**stanza** A group of lines in a file that together have a common function or define a part of the system. Stanzas are usually separated by blank lines or colons, and each stanza has a name.

**statement**

A language construct that represents a step in a sequence of actions or a set of declarations. Statements fall into two broad classes: executable and nonexecutable.

**statement function**

A name, followed by a list of dummy arguments, that is equated with an intrinsic or derived-type expression, and that can be used as a substitute for the expression throughout the program.

**statement label**

A number made up of one to five digits that is used to identify a statement. Statement labels can be used to transfer control, to define the range of a **DO**, or to refer to a **FORMAT** statement.

**storage association**

The relationship between two storage sequences if a storage unit of one is the same as a storage unit of the other.

**structure**

A scalar data object of derived type.

**structure component**

The part of a data object of derived-type corresponding to a component of its type.

**submodule**

A program unit that extends a module or another submodule. A submodule accesses the definitions from its ancestor module or submodules through host association. It might contain definitions to be accessed by its descendant submodules through host association. It might also contain separate module procedures that define module procedure interface bodies declared in its ancestor module or submodules.

**subobject**

A portion of a named data object that may be referenced or defined independently of other portions. It can be an array element, array section, structure component, or substring.

**subprogram**

A function subprogram or a subroutine subprogram. Note that in FORTRAN 77, a block data program unit was called a subprogram. See also *main program*.

**subroutine**

A procedure that is invoked by a **CALL** statement or defined assignment statement.

**subscript**

A subscript quantity or set of subscript quantities enclosed in parentheses and used with an array name to identify a particular array element.

**substring**

A contiguous portion of a scalar character string. (Although an array section can specify a substring selector, the result is not a substring.)

**symmetric multiprocessing (SMP)**

A system in which functionally-identical multiple processors are used in parallel, providing simple and efficient load-balancing.

**synchronous**

Pertaining to an operation that occurs regularly or predictably with regard to the occurrence of a specified event in another process.

**syntax** The rules for the construction of a statement. See also *semantics*.

**T**

**target** A named data object specified to have the **TARGET** attribute, a data object created by an **ALLOCATE** statement for a pointer, or a subobject of such an object.

**thread** A stream of computer instructions that is in control of a process. A multithread process begins with one stream of instructions (one thread) and may later create other instruction streams to perform tasks.

**thread-visible variable**

A variable that can be accessed by more than one thread.

**time slice**

An interval of time on the processing unit allocated for use in performing a task. After the interval has expired, processing unit time is allocated to another task, so a task cannot monopolize processing unit time beyond a fixed limit.

**token** In a programming language, a character string, in a particular format, that has some defined significance.

**trigger constant**

A sequence of characters that identifies comment lines as compiler comment directives.

**Type-bound procedure**

A procedure binding in a type definition. The procedure may be referenced by the binding-name via any object of that dynamic type, as a defined operator, by defined assignment, or as part of the finalization process.

**type compatible**

All entities are type compatible with other entities of the same type. Unlimited polymorphic entities are type compatible with all entities; other polymorphic entities are type compatible with entities whose dynamic type is an extension type of the polymorphic entity's declared type.

**type declaration statement**

A statement that specifies the type, length, and attributes of an object or function. Objects can be assigned initial values.

**type parameter**

A parameter of a data type. **KIND** and **LEN** are the type parameters of intrinsic types. A type parameter of a derived type has either a **KIND** or a **LEN** attribute.

**Note:** The type parameters of a derived type are defined in the derived-type definition.

**U****unformatted record**

A record that is transmitted unchanged between internal and external storage.

**Unicode**

A universal character encoding standard that supports the interchange,

processing, and display of text that is written in any of the languages of the modern world. It also supports many classical and historical texts in a number of languages. The Unicode standard has a 16-bit international character set defined by ISO 10646. See also *ASCII*.

**unit** A means of referring to a file to use in input/output statements. A unit can be connected or not connected to a file. If connected, it refers to a file. The connection is symmetric: that is, if a unit is connected to a file, the file is connected to the unit.

**unsafe option**

Any option that could result in undesirable results if used in the incorrect context. Other options may result in very small variations from the default result, which is usually acceptable. Typically, using an unsafe option is an assertion that your code is not subject to the conditions that make the option unsafe.

**use association**

The association of names in different scoping units specified by a **USE** statement.

**V**

**variable**

A data object whose value can be defined and redefined during the execution of an executable program. It may be a named data object, array element, array section, structure component, or substring. Note that in FORTRAN 77, a variable was always scalar and named.

**X**

**XPG4** X/Open Common Applications Environment (CAE) Portability Guide Issue 4; a document which defines the interfaces of the X/Open Common Applications Environment that is a superset of POSIX.1-1990, POSIX.2-1992, and POSIX.2a-1992 containing extensions to POSIX standards from XPG3.

**Z**

**zero-length character**

A character object that has a length of 0 and is always defined.

**zero-sized array**

An array that has a lower bound that is greater than its corresponding upper bound. The array is always defined.



---

## Notices

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

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

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

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

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

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

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

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

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

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

Licensees of this program who 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:

Lab Director  
IBM Canada Ltd. Laboratory  
8200 Warden Avenue  
Markham, Ontario L6G 1C7  
Canada

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

#### COPYRIGHT LICENSE:

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

Each copy or any portion of these sample programs or any derivative work, must include a copyright notice as follows:

© (your company name) (year). Portions of this code are derived from IBM Corp. Sample Programs. © Copyright IBM Corp. 1998, 2014.

This software and documentation are based in part on the Fourth Berkeley Software Distribution under license from the Regents of the University of California. We acknowledge the following institution for its role in this product's development: the Electrical Engineering and Computer Sciences Department at the Berkeley campus.

---

## Trademarks and service marks

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

Adobe 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.

Windows is a trademark of Microsoft Corporation 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

`_OPENMP` C preprocessor macro 30

`-#` compiler option 66  
`-1` compiler option 67  
`-B` compiler option 67  
`-c` compiler option 69  
`-C` compiler option 68  
`-d` compiler option 70  
`-D` compiler option 69  
`-F` compiler option 71  
`-g` compiler option 73, 290  
`-I` compiler option 76  
`-k` compiler option 77  
`-l` compiler option 78  
`-L` compiler option 77  
`-NS` compiler option 82  
`-o` compiler option 85  
`-O` compiler option 83  
`-O2` compiler option 84  
`-O3` compiler option 84  
`-O4` compiler option 84  
`-O5` compiler option 84  
`-p` compiler option 86  
`-qalias` compiler option 87  
`-qalign` compiler option 90  
`-qarch` compiler option 29  
`-qassert` compiler option 95  
`-qattr` compiler option 96, 295  
`-qautodbl` compiler option 97, 300  
`-qbindcexname` compiler option 99  
`-qcache` compiler option 29, 100  
`-qclines` compiler option 103  
`-qcheck` compiler option 68, 103  
`-qci` compiler option 106  
`-qcompact` compiler option 108  
`-qcr` compiler option 108  
`-qctyp1ss` compiler option 109  
`-qdbg` compiler option 73, 110  
`-qddim` compiler option 112  
`-qdescriptor` compiler option 113  
`-qdirective` compiler option 114  
`-qdlines` compiler option 69, 116  
`-qdpc` compiler option 117  
`-qenum` compiler option 118  
`-qescape` compiler option 119  
`-qessl` compiler option 120  
`-qextern` compiler option 121  
`-qextname` compiler option 122  
`-qfdpr` compiler option 124  
`-qfixed` compiler option 125  
`-qflag` compiler option 126  
`-qfltrap` compiler option 132  
`-qfpp` option 131  
`-qfree` compiler option 134  
`-qfullpath` compiler option 135  
`-qfunctrace` 136  
`-qfunctrace_xlf_catch` 138  
`-qfunctrace_xlf_enter` 139  
`-qfunctrace_xlf_exit` 140  
`-qhalt` compiler option 140  
`-qhelp` compiler option 143  
`-qieee` compiler option 279  
`-qinit` compiler option 152  
`-qinitalloc` compiler option 152  
`-qinitauto` compiler option 154  
`-qinlgue` compiler option 157  
`-qintlog` compiler option 160  
`-qintsize` compiler option 161  
`-qkeepparm` compiler option 168  
`-qlanglvl` compiler option 169  
`-qlibansi` compiler option 171  
`-qlibmpi` 172  
`-qlinedebug` compiler option 173  
`-qlist` compiler option 174, 296  
`-qlistfmt` compiler option 175  
`-qlistopt` compiler option 178, 291  
`-qlog4` compiler option 179  
`-qmaxmem` compiler option 183  
`-qmbcs` compiler option 185  
`-qmixed` compiler option 186  
`-qmkshrobj` 21  
`-qmoddir` compiler option 187  
`-qnofunctrace` 136  
`-qnoprint` compiler option 188  
`-qnullterm` compiler option 189  
`-qobject` compiler option 190  
`-qonetric` compiler option 67, 191  
`-qoptimize` compiler option 83, 193  
`-qphsinfo` compiler option 203  
`-qpic` compiler option 205  
`-qport` compiler option 205  
`-qposition` compiler option 208  
`-qppsuborigarg` option 209  
`-qqcount` compiler option 212  
`-qrealize` compiler option 213  
`-qrecur` compiler option 215  
`-qreport` compiler option 216, 292, 293, 295  
`-qsaa` compiler option 218  
`-qsave` compiler option 218  
`-qsaveopt` compiler option 220  
`-qsc1k` compiler option 222  
`-qshowpdf` compiler option 223  
`-qsigtrap` compiler option 224  
`-qsmallstack` compiler option 226  
`-qsmmp` compiler option 227  
`-qsource` compiler option 232, 292  
`-qspillsize` compiler option 82, 233  
`-qstackprotect` compiler option 233  
`-qstacktemp` compiler option 234  
`-qstrict` compiler option 238  
`-qstrict_induction` compiler option 242  
`-qstrictieemod` compiler option 242  
`-qsuffix` compiler option 244  
`-qsuppress` compiler option 245  
`-qswapomp` compiler option 246  
`-qtbtable` compiler option 248  
`-qthreaded` compiler option 249  
`-qtimestamps` compiler option 249  
`-qtune` compiler option 29  
`-qundef` compiler option 253, 274  
`-qunroll` compiler option 253

`-qunwind` compiler option 255  
`-qversion` compiler option 255  
`-qwarn64` compiler option 257  
`-qxflag=dvz` compiler option 258  
`-qxflag=oldtab` compiler option 259  
`-qxf2003` compiler option 263  
`-qxf2008` compiler option 267  
`-qxf77` compiler option 259  
`-qxf90` compiler option 261  
`-qxlines` compiler option 268  
`-qxref` compiler option 269, 295  
`-qzerosize` compiler option 270  
`-S` compiler option 272  
`-u` compiler option 274  
`-U` compiler option 274  
`-v` compiler option 275  
`-V` compiler option 276  
`-w` compiler option 126, 278  
`-yn, -ym, -yp, -yz` compiler options 146, 279  
`/tmp` directory 9  
`.a` files 23  
`.cfg` files 23  
`.d` files 24  
`.f` and `.F` files 23  
`.lst` files 24  
`.mod` files 23, 24, 187  
`.o` files 23, 24  
`.s` files 23, 24  
`.so` files 23  
`.XOR` operator 260  
`@PROCESS` compiler directive 27  
`#if` and other `cpp` directives 31  
`#pragma nofunctrace` 136

## Numerics

1501-224, 1501-229, and 1517-011 error messages 288  
15xx identifiers for XL Fortran messages 284  
4K suboption of `-qalign` 90  
64-bit environment 281

## A

`a.out` file 24  
actual arguments  
  definition of 307  
addresses of arguments, saving 260  
aggressive array I/O 35  
`aggressive_array_io` runtime option 35  
`ALIAS @PROCESS` directive 87  
`ALIGN @PROCESS` directive 90  
alignment of `BIND(C)` derived types 90  
alignment of CSECTs and large arrays for data-stripped I/O 90  
allocatable arrays, automatic deallocation  
  with `-qxf90=autodealloc` 261  
alphabetic character, definition of 307

alphanumeric, definition of 307  
 alter program semantics 238  
 ANSI  
   checking conformance to the Fortran  
     90 standard 4, 40, 169  
   checking conformance to the Fortran  
     95 standard 4, 40, 169  
 appendold and appendunknown  
   suboptions of -qposition 208  
 architecture 93  
   -qarch compiler option 93  
   -qtune compiler option 250  
 archive files 23  
 argument addresses, saving 260  
 argument promotion (integer only) for  
   intrinsic procedures 260  
 arguments  
   definition of 307  
   passing null-terminated strings to C  
     functions 189  
 arrays  
   optimizing assignments 87  
   padding 143  
 arrays, initialization problems 288  
 aryovrlp suboption of -qalias 87  
 as attribute of configuration file 14  
 as command, passing command-line  
   options to 28  
 ASCII  
   definition of 308  
 asopt attribute of configuration file 15  
 assembler  
   source (.s) files 23, 24  
 ATTR @PROCESS directive 96  
 attribute section in compiler listing 295  
 AUTODBL @PROCESS directive 97  
 autodealloc suboption of -qxl90 261  
 autorealloc suboption, -qxl2003 264

## B

bash shell 7  
 basic example, described xii  
 big-endian I/O 43  
 BIND(C) derived types, alignment 90  
 BINDCEXTNAME @PROCESS  
   directive 99  
 blankpad suboption of -qxl77 260  
 bolt attribute of configuration file 14  
 bozlitargs suboption, -qxl2003 264  
 bss storage, alignment of arrays in 90  
 buffering runtime option  
   description 35  
   using with preconnected files 35

## C

C preprocessor (cpp) 29, 131, 209  
 carriage return character 108  
 catch routine 138  
 CCLINES @PROCESS 103  
 character constants and typeless  
   constants 109  
 CHECK @PROCESS directive 68, 103  
 checkpresence suboption, -qxl2008 267

chunk  
   definition of 310  
 CI @PROCESS directive 106  
 cleanpdf command 199  
 cnvrr runtime option 37  
 code attribute of configuration file 14  
 code generation for different systems 29  
 code optimization 6  
 command line options, summary 47  
 command line, specifying options on 26  
 COMPACT @PROCESS directive 108  
 compatibility  
   compatibility  
     options for compatibility 61  
 compilation order 22  
 compilation unit epilogue section in  
   compiler listing 297  
 compiler listings 291  
 compiler options  
   descriptions 65  
   scope and precedence 26  
   section in compiler listing 291  
   specifying in the source file 27  
   specifying on the command line 26  
 compiling  
   cancelling a compilation 23  
   description of how to compile a  
     program 17  
   Fortran 2003 programs 20  
   Fortran 2008 programs 20  
   problems 287  
   SMP programs 22  
 conditional compilation 29  
 configuration 9  
   custom configuration files 9  
   configuration file 23, 71  
 configuration file attributes 14  
 conflicting options  
   -C interferes with -qhot 105  
   -qautodbl overrides -qrealsize 99  
   -qdpc is overridden by -qautodbl and  
     -qrealsize 214  
   -qflag overrides -qlanglvl or  
     -qsa 127  
   -qhalt is overridden by  
     -qnoobject 191  
   -qhalt overrides -qobject 191  
   -qintsize overrides -qlog4 179  
   -qlanglvl is overridden by -qflag 170  
   -qlog4 is overridden by -qintsize 179  
   -qnoobject overrides -qhalt 141  
   -qobject is overridden by -qhalt 141  
   -qrealsize is overridden by  
     -qautodbl 99, 214  
   -qrealsize overrides -qdpc 214  
   -qsa is overridden by -qflag 218  
   @PROCESS overrides command-line  
     setting 26  
   command-line overrides configuration  
     file setting 26  
   specified more than once, last one  
     takes effect 26  
 conformance checking 4, 169, 218  
 control of implicit timestamps 249  
 control of transformations 238  
 conversion errors 37  
 core file 290

could not load program (error  
 message) 285  
 cpp attribute of configuration file 14  
 cpp command 29  
 cpp, cppoptions, and cppsuffix attributes  
   of configuration file 15  
 cppsuffix attribute of configuration  
   file 15  
 cpu\_time\_type runtime option 37  
 cross reference section in compiler  
   listing 295  
 crt\_64 attribute of configuration file 14  
 CSECTS, alignment of 90  
 csh shell 7  
 CTYPLSS @PROCESS directive 109  
 customizing configuration file (including  
   default compiler options) 13

## D

data limit 285  
 data reorganization report section in  
   compiler listing 295  
 data stripping  
   -qalign required for improved  
     performance 90  
 data types 92  
   -qaltivec compiler option 92  
 DBG @PROCESS directive 73, 110  
 dbl, dbl4, dbl8, suboptions of  
   -qautodbl 97  
 DDIM @PROCESS directive 112  
 debugger support  
   symbolic 6  
 debugging 283  
   using path names of original  
     files 135  
 default\_recl runtime option 38  
 defaultmsg attribute of configuration  
   file 16  
 defaults  
   customizing compiler defaults 13  
   search paths for include and .mod  
     files 76  
   search paths for libraries 8  
 dependency files 24  
 deps suboption of -qassert 95  
 descriptor data structure formats 113  
 diagnostics, compiler listings 291  
 DIRECTIVE @PROCESS directive 114  
 disassembly listing  
   from the -S compiler option 272  
 disk space, running out of 288  
 DLINES @PROCESS directive 69, 116  
 DPC @PROCESS directive 117  
 dummy argument  
   definition of 313  
 dynamic dimensioning of arrays 112  
 dynamic extent, definition of 313  
 dynamic library 21  
 dynamic linking 32

## E

E error severity 283  
edit descriptors (B, O, Z), differences between F77 and F90 260  
edit descriptors (G), difference between F77 and F90 260  
editing configuration file 13  
editing source files 17  
emacs text editor 17  
enable suboption of -qfltrap 132  
end-of-file, writing past 260  
enter routine 139  
ENTRY statements, compatibility with previous compiler versions 260  
environment problems 285  
environment variables  
  compile time 7  
  PDF\_BIND\_PROCESSOR 8  
  TMPDIR 9  
  LD\_LIBRARY\_PATH 8  
  LD\_RUN\_PATH 8  
  runtime 8  
  LD\_LIBRARY\_PATH 45  
  LD\_RUN\_PATH 45  
  TMPDIR 45  
  XLFRTOPTIONS 34  
  XLF\_USR\_CONFIG 45  
  XLFSCRATCH\_unit 9  
  XLFUNIT\_unit 9  
eof, writing past 260  
epilogue sections in compiler listing 297  
err\_recovery runtime option 38  
errloc runtime option 38  
error checking and debugging 54  
error messages 283  
  1501-229 288  
  1517-011 288  
  explanation of format 284  
  in compiler listing 292  
errereof runtime option 38  
errthrdnum runtime option 38  
errtrace runtime option 39  
ESCAPE @PROCESS directive 119  
exception handling 45  
  for floating point 132  
exclusive or operator 260  
executable files 24  
executing a program 33  
executing the compiler 17  
exit routine 140  
extensions, language 3  
external names  
  in the runtime environment 299  
EXTNAME @PROCESS directive 122

## F

faster array I/O 35  
file table section in compiler listing 297  
files  
  editing source 17  
  input 23  
  output 24  
  using suffixes other than .f for source files 15

FIPS FORTRAN standard, checking conformance to 4  
FIXED @PROCESS directive 125  
FLAG @PROCESS directive 126  
floating-point  
  exception handling 45  
  exceptions 132  
FLTRAP @PROCESS directive 132  
Fortran 2003  
  programs, compiling 20  
Fortran 2003 features 40  
Fortran 2003 iostat\_end behavior 39  
Fortran 2008  
  programs, compiling 20  
Fortran 2008 features 40  
Fortran 90  
  compiling programs written for 19  
Fortran 95  
  compiling programs written for 19  
FREE @PROCESS directive 134  
fsuffix attribute of configuration file 15  
full suboption of -qbttable 248  
FULLPATH @PROCESS directive 135  
function trace 136, 138, 139, 140  
functrace 136, 138, 139, 140

## G

G edit descriptor, difference between F77 and F90 260  
gcc\_libs\_64 attribute of configuration file 14  
gcc\_path\_64 attribute of configuration file 14  
gcr64 attribute of configuration file 14  
gedit77 suboption of -qxf77 260  
generating code for different systems 29

## H

HALT @PROCESS directive 140  
hardware, compiling for different types of 29  
header section in compiler listing 291  
hexint and nohexint suboptions of -qport 205  
high order transformation 143  
hot attribute of configuration file 14  
hsflt suboption of -qfloat 300

## I

I error severity 283  
i-node 41  
IEEE @PROCESS directive 146, 279  
IEEE infinity output 265  
IEEE NaN output 265  
implicit timestamps, control of 249  
include\_64 attribute of configuration file 16  
inexact suboption of -qfltrap 133  
informational message 283  
INIT @PROCESS directive 152  
initialize arrays, problems 288  
INLGLUE @PROCESS directive 157  
inlining 157

input files 23  
input/output  
  increasing throughput with data stripping 90  
  runtime behavior 34  
  when unit is positioned at end-of-file 260  
installation problems 285  
installing the compiler 7  
intarg suboption of -qxf77 260  
integer arguments of different kinds to intrinsic procedures 260  
internal limits for the compiler 305  
interprocedural analysis (IPA) 163  
INTLOG @PROCESS directive 160  
intptr suboption of -qalias 87  
intrinsic procedures accepting integer arguments of different kinds 260  
intrinths runtime option 39  
INTSIZE @PROCESS directive 161  
intxor suboption of -qxf77 260  
invalid suboption of -qfltrap 133  
invoking a program 33  
invoking the compiler 17  
iostat\_end runtime option 39  
ipa attribute of configuration file 14  
irand routine, naming restriction for 33  
ISO

  checking conformance to the Fortran 2003 standard 4, 169  
  checking conformance to the Fortran 2008 standard 4, 169  
  checking conformance to the Fortran 90 standard 4, 40, 169  
  checking conformance to the Fortran 95 standard 4, 40, 169  
itercnt suboption of -qassert 95

## K

kind type parameters 162, 214  
ksh shell 7

## L

L error severity 283  
LANGLVL @PROCESS directive 169  
langlvl runtime option 40  
language extensions 3  
language standards 3  
language support 3  
language-level error 283  
ld command, passing command-line options to 28  
LD\_LIBRARY\_PATH environment variable 8, 45  
LD\_RUN\_PATH environment variable 8, 45  
ldopt attributes of configuration file 14, 15  
leadzero suboption of -qxf77 260  
level of XL Fortran, determining 16  
lexical extent, definition of 317  
lib\*.so library files 23, 78  
libraries 23  
  default search paths 8

libraries (*continued*)  
 shared 299  
 library  
 shared (dynamic) 21  
 static 21  
 library path environment variable 285  
 libxlf90.t.so 19  
 libxlf90.so library 34  
 libxlsmp.so library 34  
 limit command 285  
 limits internal to the compiler 305  
 line feed character 108  
 LINEDEBUG @PROCESS directive 173  
 linking 31  
 dynamic 32  
 problems 288  
 static 32  
 LIST @PROCESS directive 174  
 listing files 24  
 listings, compiler 291  
 LISTOPT @PROCESS directive 178  
 LOG4 @PROCESS directive 179

## M

m suboption of -y 279  
 machines, compiling for different types 29, 93  
 macro expansion 29, 131, 209  
 macro, \_OPENMP C preprocessor 30  
 maf suboption of -qfloat 240  
 make command 66  
 makefiles  
 copying modified configuration files along with 14  
 malloc system routine 98  
 MAXMEM @PROCESS directive 183  
 MBCS @PROCESS directive 185  
 mclock routine, naming restrictions for 33  
 mcr64 attribute of configuration file 14  
 mergepdf 199  
 message suppression 245  
 messages  
 1501-053 error message 288  
 1501-229 error message 288  
 1517-011 error message 288  
 minus suboption of -qieee 146  
 MIXED @PROCESS directive 186, 274  
 mod and nomod suboptions of -qport 205  
 mod files 23, 24, 187  
 modules, effect on compilation order 22  
 mon.out file 23  
 mpi 172  
 MPI 172  
 multconn runtime option 41  
 multconnio runtime option 41

## N

n suboption of -y 279  
 name conflicts, avoiding 33  
 namelist runtime option 42

NaN values  
 specifying with -qinitalloc compiler option 152  
 specifying with -qinitauto compiler option 154  
 naninfoutput runtime option 42  
 nearest suboption of -qieee 146  
 nlwidth runtime option 42  
 nodblpad suboption of -qautodbl 97  
 nodeps suboption of -qassert 95  
 NOFUNCTRACE 136  
 none suboption of -qautodbl 97  
 none suboption of -qbttable 248  
 nooldnaninf suboption, -qxf2003 265  
 nooldpad suboption of -qxf90 261  
 null-terminated strings, passing to C functions  
 strings, passing to C functions 189  
 NULLTERM @PROCESS directive 189

## O

OBJECT @PROCESS directive 190  
 object files 23, 24  
 object output, implicit timestamps 249  
 oldboz suboption of -qxf77 260  
 oldpad suboption of -qxf90 261  
 ONETRIP @PROCESS directive 67, 191  
 optimization 6  
 loop optimization 143  
 OPTIMIZE @PROCESS directive 83, 193  
 options attribute of configuration file 15  
 options for performance optimization 57  
 options section in compiler listing 291  
 options that control linking 61  
 options that control listings and messages 56  
 osuffix attribute of configuration file 15  
 output files 24  
 overflow suboption of -qfltrap 133

## P

p suboption of -y 279  
 pad setting, changing for internal, direct-access and stream-access files 260  
 padding of data types with -qautodbl option 300  
 paging space  
 running out of 288  
 path name of source files, preserving with -qfullpath 135  
 PDF report section in compiler listing 292  
 PDF\_BIND\_PROCESSOR environment variable 8  
 PDF\_PM\_EVENT 8  
 PDF\_PM\_EVENT environment variables 8  
 PDFDIR 8  
 PDFDIR environment variables 8  
 performance of real operations, speeding up 98, 213  
 performance optimization options 57  
 persistent suboption of -qxf77 260

PHSINFO @PROCESS directive 203  
 platform, compiling for a specific type 93  
 plus suboption of -qieee 146  
 pointers (Fortran 90) and -qinit compiler option 152  
 polymorphic suboption of -qxf2003 263  
 PORT @PROCESS directive 205  
 POSITION @PROCESS directive 208  
 POSIX pthreads  
 API support 22  
 runtime libraries 34  
 PowerPC systems  
 compiling programs for 29  
 precision of real data types 98, 213  
 preprocessing Fortran source with the C preprocessor 29  
 problem determination 283  
 procedure trace 136, 138, 139, 140  
 prof command 24  
 profile-directed feedback (PDF) 195  
 -qpdf1 compiler option 195  
 -qpdf2 compiler option 195  
 profiling 86  
 -qpdf1 compiler option 195  
 -qpdf2 compiler option 195  
 profiling data files 24  
 Program Editor 17  
 promoting integer arguments to intrinsic procedures 260  
 promotion of data types with -qautodbl option 300  
 pteovrlp suboption of -qalias 87

## Q

QCOUNT @PROCESS directive 212  
 qdirectstorage compiler option 116  
 quiet NaN 155  
 quiet NaN suboption of -qfltrap 133

## R

rand routine, naming restriction for 33  
 random runtime option 43  
 READ statements past end-of-file 260  
 README file 7  
 REAL data types 98  
 REALSIZE @PROCESS directive 213  
 RECUR @PROCESS directive 215  
 recursion 215, 219  
 register flushing 168  
 resetpdf command 199  
 return code  
 from compiler 284  
 from Fortran programs 284  
 rpm command 16  
 rrm suboption of -qfloat 240  
 run time  
 exceptions 45  
 options 34  
 running a program 33  
 running the compiler 17  
 runtime  
 libraries 23  
 problems 289

runtime environment  
  external names in 299  
runtime options 34

## S

S error severity 283  
SAA @PROCESS directive 218  
SAA FORTRAN definition, checking  
  conformance to 4  
SAVE @PROCESS directive 218  
scratch\_vars runtime option 9, 43  
setrteopts service and utility  
  procedure 34  
severe error 283  
sh shell 7  
shared (dynamic) library 21  
shared libraries 299  
shared object files 23  
shared objects 186  
  -qmkshrobj 186  
shared-memory parallelism (SMP) 227  
showpdf 199  
SIGN intrinsic, effect of  
  -qxf90=signedzero on 261  
signal handling 45  
signedzero suboption of -qxf90 261  
SIGTRAP signal 45, 132  
small suboption of -qbttable 248  
SMP  
  programs, compiling 22  
  softeof suboption of -qxf77 260  
SOURCE @PROCESS directive 232  
source file options 27  
source files 23  
  allowing suffixes other than .f 15  
  preserving path names for  
  debugging 135  
  specifying options in 27  
source section in compiler listing 292  
source-code conformance checking 4  
source-level debugging support 6  
space problems 285  
SPILLSIZE @PROCESS directive 82, 233  
ssuffix attribute of configuration file 15  
stack  
  limit 285  
stackprotect  
  stackprotect 233  
standards, language 3  
static library 21  
static linking 32  
static storage, alignment of arrays in 90  
std suboption of -qalias 87  
storage limits 285  
storage relationship between data  
  objects 300  
storage-associated arrays, performance  
  implications of 87  
STRICT @PROCESS directive 238  
strictieemod @PROCESS directive 242  
suffix, allowing other than .f on source  
  files 15  
suffixes for source files 242, 244  
summary of command line options 47  
SWAPOMP @PROCESS directive 246  
symbolic debugger support 6

system problems 285

## T

target machine, compiling for 93  
temporary arrays, reducing 87  
temporary file directory 9  
text editors 17  
threads, controlling 39  
throughput for I/O, increasing with data  
  striping 90  
times routine, naming restriction for 33  
TMPDIR environment variable 45, 288  
  compile time 9  
trace 136, 138, 139, 140  
Trace/breakpoint trap 45  
traceback listing 224, 290  
transformation report section in compiler  
  listing 293  
transformations, control of 238  
trigger\_constant  
  IBM\* 114  
  IBMT 249  
  setting values 114  
trigraphs 31  
tuning 250  
  -qarch compiler option 250  
  -qtune compiler option 250  
typeless constants and character  
  constants 109  
tpeymt and notpeymt suboptions of  
  -qport 205

## U

U error severity 283  
ufmt\_bigendian runtime option 43  
ulimit command 285  
UNDEF @PROCESS directive 253, 274  
underflow suboption of -qflttrap 133  
unformatted data files, little-endian  
  I/O 43  
Unicode data 186  
unit\_vars runtime option 9, 44  
UNIVERSAL setting for locale 186  
unrecoverable error 283  
unrolling DO LOOPS 253  
UNWIND @PROCESS directive 255  
use attribute of configuration file 14  
UTF-8 encoding for Unicode data 186  
uwidth runtime option 44

## V

value relationships between data  
  objects 300  
vector data types 92  
  -qaltivec compiler option 92  
vector processing 225  
  -qaltivec compiler option 92  
vi text editor 17

## W

W error severity 283  
warning error 283  
what command 16  
WRITE statements past end-of-file 260

## X

XFLAG(OLDTAB) @PROCESS  
  directive 259  
xl\_trbk library procedure 290  
xl\_trce exception handler 224  
xlf attribute of configuration file 14  
xlf\_r command  
  for compiling SMP programs 22  
XLF\_USR\_CONFIG environment  
  variable 45  
xlf.cfg configuration file 13  
xlf.cfg.nn configuration file 71  
xlf.cfgconfiguration file 71  
XLF2003 @PROCESS directive 263  
XLF2008 @PROCESS directive 267  
XLF77 @PROCESS directive 259  
XLF90 @PROCESS directive 261  
xlf90\_r command  
  for compiling SMP programs 22  
xlf95\_r command  
  for compiling SMP programs 22  
XLFRTTEOPTS environment variable 34  
XLFSCRATCH\_unit environment  
  variable 9, 43  
XLFUNIT\_unit environment variable 9,  
  44  
XLINES @PROCESS 268  
XOR 260  
XREF @PROCESS directive 269  
xrf\_messages runtime option 44

## Z

z suboption of -y 279  
zero suboption of -qieee 146  
zerodivide suboption of -qflttrap 133  
zeros (leading), in output 260  
ZEROSIZE @PROCESS directive 270







Product Number: 5765-J10; 5725-C75

Printed in USA

SC27-6610-00

