

IBM XL C/C++ for AIX, V12.1



# Compiler Reference

*Version 12.1*



IBM XL C/C++ for AIX, V12.1



# Compiler Reference

*Version 12.1*

**Note**

Before using this information and the product it supports, read the information in "Notices" on page 619.

**First edition**

This edition applies to IBM XL C/C++ for AIX, V12.1 (Program 5765-J02; 5725-C72) and to all subsequent releases and modifications until otherwise indicated in new editions. Make sure you are using the correct edition for the level of the product.

© Copyright IBM Corporation 1996, 2012.

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

---

# Contents

<b>About this information</b> . . . . .	<b>ix</b>
Who should read this information . . . . .	ix
How to use this information . . . . .	ix
How this information is organized . . . . .	ix
Conventions . . . . .	x
Related information . . . . .	xiii
IBM XL C/C++ information . . . . .	xiii
Standards and specifications. . . . .	xv
Other IBM information . . . . .	xv
Other information . . . . .	xvi
Technical support . . . . .	xvi
How to send your comments . . . . .	xvi

## Chapter 1. Compiling and linking applications . . . . . 1

Invoking the compiler . . . . .	1
Command-line syntax . . . . .	2
Types of input files . . . . .	3
Types of output files. . . . .	4
Specifying compiler options . . . . .	5
Specifying compiler options on the command line	6
Specifying compiler options in a configuration file	7
Specifying compiler options in program source files	8
Resolving conflicting compiler options. . . . .	8
Specifying compiler options for architecture-specific compilation . . . . .	9
Reusing GNU C/C++ compiler options with gxc and gxc++ . . . . .	11
gxc or gxc++ syntax . . . . .	11
Preprocessing. . . . .	12
Directory search sequence for include files . . . . .	13
Linking. . . . .	14
Order of linking . . . . .	14
Redistributable libraries . . . . .	15
Compatibility with earlier versions . . . . .	16
Compiler messages and listings. . . . .	17
Compiler messages . . . . .	17
Compiler return codes. . . . .	19
Compiler listings . . . . .	20
Message catalog errors. . . . .	22
Paging space errors during compilation . . . . .	23

## Chapter 2. Configuring compiler defaults . . . . . 25

Setting environment variables . . . . .	25
Compile-time and link-time environment variables . . . . .	26
Runtime environment variables. . . . .	27
Environment variables for parallel processing . . . . .	27
Using custom compiler configuration files . . . . .	40
Creating custom configuration files . . . . .	40
Configuring the gxc or gxc++ option mapping . . . . .	43

## Chapter 3. Tracking and reporting compiler usage . . . . . 47

Understanding utilization tracking and reporting. . . . .	47
Overview . . . . .	47
Four usage scenarios . . . . .	48
Preparing to use this feature. . . . .	56
Time synchronization . . . . .	56
License types and user information . . . . .	56
Central configuration . . . . .	57
Concurrent user considerations. . . . .	57
Usage file considerations . . . . .	58
Regular utilization checking. . . . .	60
Testing utilization tracking . . . . .	60
Configuring utilization tracking . . . . .	61
Editing utilization tracking configuration file entries . . . . .	62
Understanding the utilization reporting tool . . . . .	65
Utilization reporting tool command-line options	66
Generating usage reports . . . . .	69
Understanding usage reports . . . . .	70
Pruning usage files . . . . .	72
Diagnostic messages from utilization tracking and reporting . . . . .	73

## Chapter 4. Compiler options reference 75

Summary of compiler options by functional category . . . . .	75
Output control . . . . .	75
Input control . . . . .	76
Language element control . . . . .	77
Template control (C++ only). . . . .	79
Floating-point and integer control . . . . .	80
Object code control. . . . .	80
Error checking and debugging . . . . .	83
Listings, messages, and compiler information . . . . .	85
Optimization and tuning . . . . .	87
Linking. . . . .	91
Portability and migration. . . . .	91
Compiler customization . . . . .	92
Deprecated options. . . . .	93
Individual option descriptions . . . . .	94
-+ (plus sign) (C++ only) . . . . .	95
-# (pound sign) . . . . .	96
-q32, -q64 . . . . .	96
-qaggrcopy . . . . .	97
-qalias . . . . .	98
-qalign . . . . .	100
-qalignrulefor (C++ only) . . . . .	103
-qalloca, -ma (C only) . . . . .	104
-qaltivec . . . . .	105
-qarch . . . . .	106
-qasm . . . . .	110
-qasm_as . . . . .	112
-qassert . . . . .	113
-qattr . . . . .	114

-b	115	-qkeepparam	210
-B	116	-qkeyword	211
-qbitfields	117	-l	213
-bmaxdata	117	-L	214
-brtl	118	-qlanglvl	215
-c	119	-qlargepage	238
-C, -C!	120	-qdbl128, -qlongdouble	239
-qcache	121	-qlib	240
-qchars	123	-qlibansi	241
-qcheck	125	-qlibmpi	241
-qcinc (C++ only)	126	-qlinedebug	242
-qcompact	127	-qlist	243
-qconcurrentupdate (C only)	128	-qlistfmt	245
-qplusplus (C only)	129	-qlistopt	248
-qcrtr	129	-qlonglit	249
-qc_stdinc (C only)	130	-qlonglong	249
-qcpp_stdinc (C++ only)	131	-ma (C only)	251
-D	132	-qmacpstr	251
-qdataimported, -qdatalocal, -qtocdata	134	-qmakedep, -M	252
-qdbgfmt	135	-qmaxerr	254
-qdbxextra (C only)	136	-qmaxmem	256
-qdfp	137	-qmbcs, -qdbcs	257
-qdigraph	138	-MF	258
-qdirectstorage	139	-qminimaltoc	259
-qdollar	140	-qmkshrobj	260
-qdpcl	140	-qnamemangling (C++ only)	262
-e	141	-o	268
-E	142	-O, -qoptimize	269
-qeh (C++ only)	144	-qobjmodel (C++ only)	272
-qenum	144	-qoldpassbyvalue (C++ only)	274
-qexpfile	148	-qoptdebug	275
-qextchk	149	-qoptfile	277
-f	150	-p, -pg, -qprofile	279
-F	150	-P	280
-qfdpr	152	-qpath	281
-qflag	153	-qpdf1, -qpdf2	283
-qfloat	154	-qphsinfo	290
-qfltrap	159	-qpics	291
-qformat	163	-qppline	293
-qfullpath	164	-qprefetch	293
-qfuncsect	165	-qprint	296
-qfunctrace	166	-qpriority (C++ only)	297
-g	169	-qprocimported, -qproclcal, -qprocunknown	298
-G	172	-qproto (C only)	300
-qgenproto (C only)	172	-r	301
-qhalt	174	-qreport	302
-qhaltmsg	175	-qreserved_reg	304
-qheapdebug	177	-qro	305
-qhot	178	-qroconst	306
-I	181	-qroptr	307
-qidirfirst	182	-qrtti (C++ only)	308
-qignerno	183	-s	309
-qignprag	184	-S	310
-qinclude	185	-qsaveopt	311
-qinfo	187	-qshowinc	313
-qinitauto	193	-qshowmacros	315
-qinlglue	196	-qshowpdf	316
-qinline	197	-qsimd	316
-qipa	199	-qskipsrc	318
-qisolated_call	206	-qsmallstack	319
-qkeepinlines (C++ only)	209	-qsmp	320

-qsource . . . . .	324	Error checking and debugging. . . . .	385
-qsourcectype. . . . .	325	Listings, messages and compiler information	386
-qspeculateabsolutes . . . . .	326	Optimization and tuning . . . . .	386
-qspill . . . . .	327	Object code control . . . . .	387
-qsrcmsg (C only) . . . . .	328	Portability and migration . . . . .	388
-qstackprotect . . . . .	329	Deprecated directives. . . . .	388
-qstaticinline (C++ only). . . . .	330	Individual pragma descriptions . . . . .	389
-qstatsym. . . . .	331	#pragma align . . . . .	389
-qstdinc . . . . .	332	#pragma alloca (C only). . . . .	389
-qstrict . . . . .	333	#pragma block_loop . . . . .	390
-qstrict_induction . . . . .	337	#pragma chars . . . . .	393
-qsuppress . . . . .	338	#pragma comment. . . . .	393
-qsymtab (C only) . . . . .	340	#pragma define, #pragma instantiate (C++ only)	394
-qsyntaxonly (C only) . . . . .	341	#pragma disjoint . . . . .	395
-t . . . . .	342	#pragma do_not_instantiate (C++ only). . . . .	397
-qtabsize . . . . .	343	#pragma enum . . . . .	398
-qtbtable . . . . .	344	#pragma execution_frequency . . . . .	398
-qtempinc (C++ only). . . . .	345	#pragma expected_value . . . . .	400
-qtemplatedepth (C++ only) . . . . .	346	#pragma fini (C only) . . . . .	401
-qtemplaterecompile (C++ only) . . . . .	347	#pragma hashome (C++ only) . . . . .	401
-qtemplateregistry (C++ only) . . . . .	348	#pragma ibm iterations . . . . .	403
-qtempmax (C++ only) . . . . .	349	#pragma ibm max_iterations . . . . .	404
-qthreaded . . . . .	350	#pragma ibm min_iterations . . . . .	405
-qtimestamps . . . . .	351	#pragma ibm snapshot . . . . .	406
-qtls . . . . .	352	#pragma implementation (C++ only) . . . . .	406
-qtmplinst (C++ only) . . . . .	353	#pragma info . . . . .	407
-qtmplparse (C++ only) . . . . .	354	#pragma init (C only) . . . . .	407
-qtocdata . . . . .	355	#pragma ishome (C++ only) . . . . .	408
-qtocmerge . . . . .	355	#pragma isolated_call . . . . .	409
-qtrigraph . . . . .	356	#pragma langlvl (C only) . . . . .	409
-qtune . . . . .	357	#pragma leaves. . . . .	409
-qtwolink (C++ only) . . . . .	359	#pragma loopid . . . . .	410
-U . . . . .	361	#pragma map . . . . .	411
-qunique . . . . .	362	#pragma mc_func . . . . .	413
-qunroll . . . . .	363	#pragma namemangling (C++ only) . . . . .	415
-qunwind. . . . .	365	#pragma namemanglingrule (C++ only) . . . . .	415
-qupconv (C only). . . . .	366	#pragma nofunctrace . . . . .	418
-qutf . . . . .	367	#pragma nosimd . . . . .	418
-v, -V . . . . .	368	#pragma novector . . . . .	418
-qvecnvolf. . . . .	369	#pragma object_model (C++ only) . . . . .	419
-qversion . . . . .	370	#pragma operator_new (C++ only) . . . . .	419
-w . . . . .	371	#pragma options . . . . .	420
-W . . . . .	372	#pragma option_override . . . . .	422
-qwarn0x (C++0x) . . . . .	374	#pragma pack . . . . .	424
-qwarn64 . . . . .	376	#pragma pass_by_value (C++ only) . . . . .	427
-qweakexp . . . . .	377	#pragma priority (C++ only) . . . . .	427
-qweaksymbol . . . . .	378	#pragma reachable . . . . .	427
-qxcall. . . . .	378	#pragma reg_killed_by . . . . .	428
-qxref . . . . .	379	#pragma report (C++ only). . . . .	430
-y . . . . .	381	#pragma simd_level . . . . .	431
-Z . . . . .	382	#pragma STDC cx_limited_range. . . . .	432
		#pragma stream_unroll . . . . .	433
		#pragma strings . . . . .	434
		#pragma unroll. . . . .	434
		#pragma unrollandfuse . . . . .	435
		#pragma weak . . . . .	436
		Pragma directives for parallel processing . . . . .	439
<b>Chapter 5. Compiler pragmas</b>		<b>Chapter 6. Compiler predefined</b>	
<b>reference . . . . .</b>	<b>383</b>	<b>macros . . . . .</b>	<b>463</b>
Pragma directive syntax . . . . .	383	General macros. . . . .	463
Scope of pragma directives . . . . .	384		
Summary of compiler pragmas by functional			
category . . . . .	384		
Language element control . . . . .	385		
C++ template pragmas . . . . .	385		
Floating-point and integer control . . . . .	385		

Macros indicating the XL C/C++ compiler product	464
Macros related to the platform	465
Macros related to compiler features	465
Macros related to compiler option settings	465
Macros related to architecture settings	469
Macros related to language levels	470

## Chapter 7. Compiler built-in functions 479

Fixed-point built-in functions	479
Absolute value functions	480
Assert functions	480
Bit permutation functions	480
Comparison functions	481
Count zero functions	481
Division functions	481
Load functions	483
Multiply functions	483
Population count functions	484
Rotate functions	485
Store functions	486
Trap functions	487
Binary floating-point built-in functions	488
Absolute value functions	488
Add functions	488
Conversion functions	489
FPSCR functions	491
Multiply functions	493
Multiply-add/subtract functions	494
Reciprocal estimate functions	495
Rounding functions	495
Select functions	497
Square root functions	497
Software division functions	497
Store functions	498
Decimal floating-point built-in functions	499
Absolute value functions	499
Coefficient functions	500
Comparison functions	501
Conversion functions	502
Exponent functions	507
NaN functions	508
Register transfer functions	509
Rounding functions	510
Test functions	512
Miscellaneous functions	517
Synchronization and atomic built-in functions	518
Check lock functions	518
Clear lock functions	519
Compare and swap functions	520
Fetch functions	521
Load functions	523
Store functions	523
Synchronization functions	524
Cache-related built-in functions	525
Data cache functions	525
Prefetch built-in functions	527
Block-related built-in functions	534
__bcopy	534
__bzero	534
Vector built-in functions	535
vec_abs	535

vec_add	536
vec_all_eq	536
vec_all_ge	537
vec_all_gt	538
vec_all_le	539
vec_all_lt	540
vec_all_nan	541
vec_all_ne	542
vec_all_nge	543
vec_all_ngt	544
vec_all_nle	544
vec_all_nlt	545
vec_all_numeric	545
vec_and	546
vec_andc	547
vec_any_eq	548
vec_any_ge	549
vec_any_gt	550
vec_any_le	551
vec_any_lt	552
vec_any_nan	553
vec_any_ne	554
vec_any_nge	555
vec_any_ngt	556
vec_any_nle	556
vec_any_nlt	557
vec_any_numeric	557
vec_ceil	558
vec_cmpeq	558
vec_cmpge	559
vec_cmpgt	560
vec_cmple	560
vec_cmplt	561
vec_cpsgn	562
vec_ctd	562
vec_ctf	562
vec_cts	563
vec_ctsl	563
vec_ctu	564
vec_ctul	564
vec_cvf	565
vec_div	565
vec_extract	566
vec_floor	567
vec_insert	567
vec_madd	568
vec_max	568
vec_mergeh	569
vec_mergel	570
vec_min	570
vec_msub	571
vec_mul	572
vec_nabs	573
vec_neg	573
vec_nmadd	574
vec_nmsub	574
vec_nor	575
vec_or	576
vec_permi	577
vec_promote	578
vec_re	578



vec_round . . . . .	579	omp_get_max_active_levels. . . . .	611
vec_roundc . . . . .	579	omp_set_max_active_levels . . . . .	611
vec_roundm . . . . .	580	omp_get_schedule . . . . .	611
vec_roundp . . . . .	580	omp_set_schedule . . . . .	612
vec_roundz . . . . .	581	omp_get_thread_limit . . . . .	612
vec_rsqrte . . . . .	581	omp_get_level . . . . .	613
vec_sel . . . . .	582	omp_get_ancestor_thread_num . . . . .	613
vec_sl . . . . .	583	omp_get_team_size . . . . .	613
vec_sldw . . . . .	583	omp_get_active_level. . . . .	613
vec_splat . . . . .	584	omp_get_num_threads . . . . .	613
vec_splats . . . . .	585	omp_set_num_threads . . . . .	614
vec_sqrt . . . . .	586	omp_get_max_threads . . . . .	614
vec_sr . . . . .	586	omp_get_thread_num . . . . .	614
vec_sra . . . . .	587	omp_get_num_procs . . . . .	615
vec_sub . . . . .	587	omp_in_final . . . . .	615
vec_trunc. . . . .	588	omp_in_parallel . . . . .	615
vec_xld2 . . . . .	588	omp_set_dynamic . . . . .	615
vec_xlds . . . . .	589	omp_get_dynamic. . . . .	615
vec_xlw4 . . . . .	590	omp_set_nested . . . . .	616
vec_xor . . . . .	591	omp_get_nested . . . . .	616
vec_xstd2. . . . .	592	omp_init_lock, omp_init_nest_lock . . . . .	616
vec_xstw4 . . . . .	593	omp_destroy_lock, omp_destroy_nest_lock . . . . .	616
GCC atomic memory access built-in functions . . . . .	594	omp_set_lock, omp_set_nest_lock. . . . .	617
Atomic lock, release, and synchronize functions . . . . .	595	omp_unset_lock, omp_unset_nest_lock . . . . .	617
Atomic fetch and operation functions . . . . .	596	omp_test_lock, omp_test_nest_lock . . . . .	617
Atomic operation and fetch functions . . . . .	599	omp_get_wtime . . . . .	617
Atomic compare and swap functions . . . . .	602	omp_get_wtick . . . . .	618
Miscellaneous built-in functions . . . . .	603	<b>Notices . . . . .</b>	<b>619</b>
Optimization-related functions . . . . .	603	Trademarks and service marks . . . . .	621
Move to/from register functions . . . . .	604	<b>Index . . . . .</b>	<b>623</b>
Memory-related functions . . . . .	607		
Built-in functions for parallel processing . . . . .	609		
IBM SMP built-in functions. . . . .	609		

**Chapter 8. OpenMP runtime functions  
for parallel processing . . . . . 611**



---

## About this information

This information is a reference for the IBM® XL C/C++ for AIX®, V12.1 compiler. Although it provides information on compiling and linking applications written in C and C++, it is primarily intended as a reference for compiler command-line options, pragma directives, predefined macros, built-in functions, environment variables, and error messages and return codes.

---

## Who should read this information

This information is for experienced C or C++ developers who have some familiarity with the XL C/C++ compilers or other command-line compilers on UNIX operating systems. It assumes thorough knowledge of the C or C++ programming language, and basic knowledge of operating system commands. Although this information is intended as a reference guide, programmers new to XL C/C++ can still find information in it on the capabilities and features unique to the XL C/C++ compiler.

---

## How to use this information

Unless indicated otherwise, all of the text in this reference pertains to both C and C++ languages. Where there are differences between languages, these are indicated through qualifying text and icons, as described in “Conventions” on page x.

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

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

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

---

## How this information is organized

Chapter 1, “Compiling and linking applications,” on page 1 discusses topics related to compilation tasks, including invoking the compiler, preprocessor, and linker; types of input and output files; different methods for setting include file path names and directory search sequences; different methods for specifying compiler options and resolving conflicting compiler options; how to reuse GNU C/C++ compiler options through the use of the compiler utilities `gxlc` and `gxlc++`; and compiler listings and messages.

Chapter 2, “Configuring compiler defaults,” on page 25 discusses topics related to setting up default compilation settings, including setting environment variables, customizing the configuration file, and customizing the `gxc` and `gxc++` option mappings.

Chapter 3, “Tracking and reporting compiler usage,” on page 47 discusses topics related to tracking compiler utilization. This chapter provides information that helps you to detect whether compiler utilization exceeds your floating user license entitlements.

Chapter 4, “Compiler options reference,” on page 75 begins with a summary of options according to functional category, which allows you to look up and link to options by function; and includes individual descriptions of each compiler option sorted alphabetically.

Chapter 5, “Compiler pragmas reference,” on page 383 begins with a summary of pragma directives according to functional category, which allows you to look up and link to pragmas by function; and includes individual descriptions of pragmas sorted alphabetically, including OpenMP and SMP directives.

Chapter 6, “Compiler predefined macros,” on page 463 provides a list of compiler macros according to category.

Chapter 7, “Compiler built-in functions,” on page 479 contains individual descriptions of XL C/C++ built-in functions for Power® architectures, categorized by their functionality.

---

## Conventions

### Typographical conventions

The following table explains the typographical conventions used in the IBM XL C/C++ for AIX, V12.1 information.










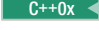
*Table 1. Typographical conventions*

Typeface	Indicates	Example
<b>bold</b>	Lowercase commands, executable names, compiler options, and directives.	The compiler provides basic invocation commands, <code>xlc</code> and <code>xlc</code> ( <code>xlc++</code> ), along with several other compiler invocation commands to support various C/C++ language levels and compilation environments.
<i>italics</i>	Parameters or variables whose actual names or values are to be supplied by the user. Italics are also used to introduce new terms.	Make sure that you update the <i>size</i> parameter if you return more than the <i>size</i> requested.
<u>underlining</u>	The default setting of a parameter of a compiler option or directive.	<code>nomaf</code>   <u><code>maf</code></u>
monospace	Programming keywords and library functions, compiler builtins, examples of program code, command strings, or user-defined names.	To compile and optimize <code>myprogram.c</code> , enter: <code>xlc myprogram.c -O3</code> .

## Qualifying elements (icons)





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

Table 2. Qualifying elements

Qualifier/Icon	Meaning
C only, or C only begins   C only ends	The text describes a feature that is supported in the C language only; or describes behavior that is specific to the C language.
C++ only, or C++ only begins   C++ only ends	The text describes a feature that is supported in the C++ language only; or describes behavior that is specific to the C++ language.
IBM extension begins   IBM extension ends	The text describes a feature that is an IBM extension to the standard language specifications.
C1X, or C1X begins   C1X ends	The text describes a feature that is introduced into standard C as part of C1X.
C++0x, or C++0x begins   C++0x ends	The text describes a feature that is introduced into standard C++ as part of C++0x.

## Syntax diagrams

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

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

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


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



- Optional items are shown below the main path:



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



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



- An arrow returning to the left above the main line (a repeat arrow) indicates that you can make more than one choice from the stacked items or repeat an item. The separator character, if it is other than a blank, is also indicated:




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



- Keywords are shown in nonitalic letters and should be entered exactly as shown.
- Variables are shown in italicized lowercase letters. They represent user-supplied names or values.
- If punctuation marks, parentheses, arithmetic operators, or other such symbols are shown, you must enter them as part of the syntax.

### Sample syntax diagram

The following syntax diagram example shows the syntax for the `#pragma comment` directive.



**Notes:**

- 1 This is the start of the syntax diagram.
- 2 The symbol # must appear first.
- 3 The keyword pragma must appear following the # symbol.
- 4 The name of the pragma comment must appear following the keyword pragma.
- 5 An opening parenthesis must be present.
- 6 The comment type must be entered only as one of the types indicated: compiler, date, timestamp, copyright, or user.
- 7 A comma must appear between the comment type copyright or user, and an optional character string.
- 8 A character string must follow the comma. The character string must be enclosed in double quotation marks.
- 9 A closing parenthesis is required.
- 10 This is the end of the syntax diagram.

The following examples of the **#pragma comment** directive are syntactically correct according to the diagram shown above:

```
#pragma comment(date)
#pragma comment(user)
#pragma comment(copyright,"This text will appear in the module")
```

**Examples in this information**

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

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

**Related information**

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

**IBM XL C/C++ information**

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

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

Man pages are provided for the compiler invocations and all command-line utilities provided with the product. Instructions for installing and accessing the man pages are provided in the *IBM XL C/C++ for AIX, V12.1 Installation Guide*.

- Information center

The information center of searchable HTML files can be launched on a network and accessed remotely or locally. Instructions for installing and accessing the online information center are provided in the *IBM XL C/C++ for AIX, V12.1 Installation Guide*.

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

- PDF documents

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

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

Table 3. XL C/C++ PDF files

Document title	PDF file name	Description
<i>IBM XL C/C++ for AIX, V12.1 Installation Guide, GC14-7329-00</i>	install.pdf	Contains information for installing XL C/C++ and configuring your environment for basic compilation and program execution.
<i>Getting Started with IBM XL C/C++ for AIX, V12.1, SC14-7328-00</i>	getstart.pdf	Contains an introduction to the XL C/C++ product, with information on setting up and configuring your environment, compiling and linking programs, and troubleshooting compilation errors.
<i>IBM XL C/C++ for AIX, V12.1 Compiler Reference, SC14-7330-00</i>	compiler.pdf	Contains information about the various compiler options, pragmas, macros, environment variables, and built-in functions, including those used for parallel processing.
<i>IBM XL C/C++ for AIX, V12.1 Language Reference, SC14-7331-00</i>	langref.pdf	Contains information about the C and C++ programming languages, as supported by IBM, including language extensions for portability and conformance to nonproprietary standards.
<i>IBM XL C/C++ for AIX, V12.1 Optimization and Programming Guide, SC14-7332-00</i>	proguide.pdf	Contains information on advanced programming topics, such as application porting, interlanguage calls with Fortran code, library development, application optimization and parallelization, and the XL C/C++ high-performance libraries.
<i>Standard C++ Library Reference, SC14-7333-00</i>	stdlib.pdf	Contains reference information about the standard C++ runtime libraries and headers.
<i>C/C++ Legacy Class Libraries Reference, SC09-7652-00</i>	legacy.pdf	Contains reference information about the USL I/O Stream Library and the Complex Mathematics Library.

To read a PDF file, use the Adobe Reader. If you do not have the Adobe Reader, you can download it (subject to license terms) from the Adobe website at <http://www.adobe.com>.

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



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

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

## Standards and specifications

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

- *Information Technology - Programming languages - C, ISO/IEC 9899:1990*, also known as C89.
- *Information Technology - Programming languages - C, ISO/IEC 9899:1999*, also known as C99.
- *Information Technology - Programming languages - C++, ISO/IEC 14882:1998*, also known as C++98.
- *Information Technology - Programming languages - C++, ISO/IEC 14882:2003(E)*, also known as *Standard C++*.
- *Information Technology - Programming languages - Extensions for the programming language C to support new character data types, ISO/IEC DTR 19769*. This draft technical report has been accepted by the C standards committee, and is available at <http://www.open-std.org/JTC1/SC22/WG14/www/docs/n1040.pdf>.
- *Draft Technical Report on C++ Library Extensions, ISO/IEC DTR 19768*. This draft technical report has been submitted to the C++ standards committee, and is available at <http://www.open-std.org/JTC1/SC22/WG21/docs/papers/2005/n1836.pdf>.
- *AltiVec Technology Programming Interface Manual*, Motorola Inc. This specification for vector data types, to support vector processing technology, is available at [http://www.freescale.com/files/32bit/doc/ref\\_manual/ALTIVECPIM.pdf](http://www.freescale.com/files/32bit/doc/ref_manual/ALTIVECPIM.pdf).
- *Information Technology - Programming Languages - Extension for the programming language C to support decimal floating-point arithmetic, ISO/IEC WDTR 24732*. This draft technical report has been submitted to the C standards committee, and is available at <http://www.open-std.org/JTC1/SC22/WG14/www/docs/n1176.pdf>.
- *Decimal Types for C++: Draft 4* <http://www.open-std.org/jtc1/sc22/wg21/docs/papers/2006/n1977.html>
- *ANSI/IEEE Standard for Binary Floating-Point Arithmetic, ANSI/IEEE Std 754-1985*.
- *OpenMP Application Program Interface Version 3.1*, available at <http://www.openmp.org>

## Other IBM information

- *Parallel Environment for AIX: Operation and Use*
- The IBM Systems Information Center, at <http://publib.boulder.ibm.com/infocenter/systems/index.jsp?topic=/com.ibm.aix.doc/doc/base/aixparent.htm> is a resource for AIX information.

You can find the following books for your specific AIX system:

- *AIX Commands Reference, Volumes 1 - 6*
- *Technical Reference: Base Operating System and Extensions, Volumes 1 & 2*
- *AIX National Language Support Guide and Reference*

- *AIX General Programming Concepts: Writing and Debugging Programs*
- *AIX Assembler Language Reference*
- *ESSL for AIX V5.1/ESSL for Linux on POWER V5.1 Guide and Reference* available at the Engineering and Scientific Subroutine Library (ESSL) and Parallel ESSL web page.

## Other information

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

---

## Technical support

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

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

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

---

## How to send your comments

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

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

---

# Chapter 1. Compiling and linking applications

By default, when you call the XL C/C++ compiler, all of the following phases of translation are performed:

- preprocessing of program source
- compiling and assembling into object files
- linking into an executable

These different translation phases are actually performed by separate executables, which are referred to as compiler *components*. However, you can use compiler options to perform only certain phases, such as preprocessing, or assembling. You can then reinvoke the compiler to resume processing of the intermediate output to a final executable.

The following sections describe how to invoke the XL C/C++ compiler to preprocess, compile and link source files and libraries:

- “Invoking the compiler”
- “Types of input files” on page 3
- “Types of output files” on page 4
- “Specifying compiler options” on page 5
- “Reusing GNU C/C++ compiler options with `gxc` and `gxc++`” on page 11
- “Preprocessing” on page 12
- “Linking” on page 14
- “Compiler messages and listings” on page 17

---

## Invoking the compiler

Different forms of the XL C/C++ compiler invocation commands support various levels of the C and C++ languages. In most cases, you should use the `xlc` command to compile your C source files, and the `xlc++` command to compile C++ source files. Use `xlc++` to link if you have both C and C++ object files.

You can use other forms of the command if your particular environment requires it. Table 4 lists the different basic commands, with the special versions of each basic command. Special commands are described in Table 5 on page 2.

**Note:** For each invocation command, the compiler configuration file defines default option settings and, in some cases, macros; for information about the defaults implied by a particular invocation, see the `/etc/vac.cfg` for your system..

Table 4. Compiler invocations

Basic invocations	Description	Equivalent special invocations
<code>xlc</code>	Invokes the compiler for C source files. This command supports all of the ISO C99 standard features, and most IBM language extensions. This invocation is recommended for all applications.	<code>xlc_r</code> , <code>xlc_r7</code> , <code>xlc128</code> , <code>xlc128_r</code> , <code>xlc128_r4</code> , <code>xlc128_r7</code>
<code>c99</code>	Invokes the compiler for C source files. This command supports all ISO C99 language features, but does not support IBM language extensions. Use this invocation for strict conformance to the C99 standard.	<code>c99_r</code> , <code>c99_r4</code> , <code>c99_r7</code> , <code>c99_128</code> , <code>c99_128_r</code> , <code>c99_128_r4</code> , <code>c99_128_r7</code>

Table 4. Compiler invocations (continued)


Basic invocations	Description	Equivalent special invocations
c89	Invokes the compiler for C source files. This command supports all ANSI C89 language features, but does not support IBM language extensions. Use this invocation for strict conformance to the C89 standard.	c89_r, c89_r4, c89_r7, c89_128, c89_128_r, c89_128_r4, c89_128_r7
cc	Invokes the compiler for C source files. This command supports pre-ANSI C, and many common language extensions. You can use this command to compile legacy code that does not conform to standard C.	cc_r, cc_r4, cc_r7, cc128, cc128_r, cc128_r4, cc128_r7
gxc	Invokes the compiler for C source files. This command accepts many common GNU C options, maps them to their XL C option equivalents, and then invokes <code>xlc</code> . For more information, refer to “Reusing GNU C/C++ compiler options with <code>gxc</code> and <code>gxc++</code> ” on page 11.	
xlc++, xlc	Invokes the compiler for C++ source files. If any of your source files are C++, you must use this invocation to link with the correct runtime libraries.  Files with <code>.c</code> suffixes, assuming you have not used the <code>--</code> compiler option, are compiled as C language source code.	xlc++_r, xlc++_r4, xlc++_r7, xlc++128, xlc++128_r, xlc++128_r4, xlc++128_r7, xlc_r, xlc_r4, xlc_r7, xlc128, xlc128_r, xlc128_r4, xlc128_r7
xlc++core, xlccore	Invokes the compiler as described above for <code>xlc++</code> and <code>xlc</code> , but links only to the core of the runtime library. Use this invocation if you want to link your application to a runtime library other than that supplied with XL C++.	xlc++core_r, xlc++core_r7, xlc++core128, xlc++core128_r, xlc++core128_r4, xlc++core128_r7, xlccore_r, xlccore_r7, xlc128core, xlc128core_r, xlc128core_r7
gxc++, gxlc	Invokes the compiler for C++ files. This command accepts many common GNU C/C++ options, maps them to their XL C/C++ option equivalents, and then invokes <code>xlc++</code> . For more information, refer to “Reusing GNU C/C++ compiler options with <code>gxc</code> and <code>gxc++</code> ” on page 11.	

Table 5. Suffixes for special invocations

128-suffixed invocations	All <b>128</b> -suffixed invocation commands are functionally similar to their corresponding base compiler invocations. They specify the <code>-qldbl128</code> option, which increases the length of long double types in your program from 64 to 128 bits. They also link with the 128-bit versions of the C and C++ runtime libraries.
_r-suffixed invocations	All <code>_r</code> -suffixed invocations allow for threadsafe compilation and you can use them to link the programs that use multi-threading. Use these commands if you want to create threaded applications.  The <code>_r7</code> invocations are provided to help migrate programs based on Posix Draft 7 to Posix Draft 10. The <code>_r4</code> invocations should be used for DCE threaded applications. For more information about DCE, see What is DCE? in CICS® Transaction Server for z/OS® Information Center.

## Command-line syntax

You invoke the compiler using the following syntax, where *invocation* can be replaced with any valid XL C/C++ invocation command listed in Table 4 on page 1:



The parameters of the compiler invocation command can be the names of input files, compiler options, and linker options.

Your program can consist of several input files. All of these source files can be compiled at once using only one invocation of the compiler. Although more than one source file can be compiled using a single invocation of the compiler, you can specify only one set of compiler options on the command line per invocation. Each distinct set of command-line compiler options that you want to specify requires a separate invocation.

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

By default, the invocation command calls *both* the compiler and the linker. It passes linker options to the linker. Consequently, the invocation commands also accept all linker options. To compile without linking, use the `-c` compiler option. The `-c` option stops the compiler after compilation is completed and produces as output, an object file `file_name.o` for each `file_name.nnn` input source file, unless you use the `-o` option to specify a different object file name. The linker is not invoked. You can link the object files later using the same invocation command, specifying the object files without the `-c` option.

#### Related information

- “Types of input files”

---

## Types of input files

The compiler processes the source files in the order in which they are displayed. If the compiler cannot find a specified source file, it produces an error message and the compiler proceeds to the next specified file. However, the linker does not run and temporary object files are removed.

By default, the compiler preprocesses and compiles all the specified source files. Although you usually want to use this default, you can use the compiler to preprocess the source file without compiling; see “Preprocessing” on page 12 for details.

You can input the following types of files to the XL C/C++ compiler:

#### C and C++ source files

These are files containing C or C++ source code.

To use the C compiler to compile a C language source file, the source file must have a `.c` (lowercase c) suffix, unless you compile with the `-qsource=c` option.

To use the C++ compiler, the source file must have a `.C` (uppercase C), `.cc`, `.cp`, `.cpp`, `.cxx`, or `.c++` suffix, unless you compile with the `-+` or `-qsource=c++` option.

### Preprocessed source files

Preprocessed source files have a `.i` suffix, for example, `file_name.i`. The compiler sends the preprocessed source file, `file_name.i`, to the compiler where it is preprocessed again in the same way as a `.c` or `.C` file. Preprocessed files are useful for checking macros and preprocessor directives.

### Object files

Object files must have a `.o` suffix, for example, `file_name.o`. Object files, library files, and unstripped executable files serve as input to the linker. After compilation, the linker links all of the specified object files to create an executable file.

### Assembler files

Assembler files must have a `.s` suffix, for example, `file_name.s`, unless you compile with the `-qsource=assembler` option. Assembler files are assembled to create an object file.

### Unpreprocessed assembler files

Unpreprocessed assembler files must have a `.S` suffix, for example, `file_name.S`, unless you compile with the `-qsource=assembler-with-cpp` option. The compiler compiles all source files with a `.S` extension as if they are assembler language source files that need preprocessing.

### Shared library files

Shared library files generally have a `.a` suffix, for example, `file_name.a`, but they can also have a `.so` suffix, for example, `file_name.so`.

### Unstripped executable files

Extended Common Object File Format (XCOFF) files that have not been stripped with the operating system `strip` command can be used as input to the compiler. See the `strip` command in the *AIX Commands Reference* and the description of `a.out` file format in the *AIX Files Reference* for more information.

### Related information

- Options summary by functional category: Input control

---

## Types of output files

You can specify the following types of output files when invoking the XL C/C++ compiler:

### Executable files

By default, executable files are named `a.out`. To name the executable file something else, use the `-o file_name` option with the invocation command. This option creates an executable file with the name you specify as `file_name`. The name you specify can be a relative or absolute path name for the executable file.

The format of the `a.out` file is described in the *AIX Files Reference*.

### Object files

If you specify the `-c` option, an output object file, `file_name.o`, is produced for each input file. The linker is not invoked, and the object files are placed in your current directory. All processing stops at the completion of the compilation. The compiler gives object files a `.o` suffix, for example, `file_name.o`, unless you specify the `-o file_name` option, giving a different suffix or no suffix at all.

You can link the object files later into a single executable file by invoking the compiler.

#### Shared library files

If you specify the **-qmkshrobj** option, the compiler generates a single shared library file for all input files. The compiler names the output file `shr.o`, unless you specify the **-o** *file\_name* option, and give the file a `.so` suffix.

#### Assembler files

If you specify the **-S** option, an assembler file, *file\_name.s*, is produced for each input file.

You can then assemble the assembler files into object files and link the object files by reinvoking the compiler.

#### Preprocessed source files

If you specify the **-P** option, a preprocessed source file, *file\_name.i*, is produced for each input file.

You can then compile the preprocessed files into object files and link the object files by reinvoking the compiler.

#### Listing files

If you specify any of the listing-related options, such as **-qlist** or **-qsource**, a compiler listing file, *file\_name.lst*, is produced for each input file. The listing file is placed in your current directory.

#### Target files

If you specify the **-M** or **-qmakedep** option, a target file suitable for inclusion in a makefile, *file\_name.u* is produced for each input file.

#### Related information

- Options summary by functional category: Output control

---

## Specifying compiler options

Compiler options perform a wide variety of functions, such as setting compiler characteristics, describing the object code and compiler output to be produced, and performing some preprocessor functions. You can specify compiler options in one or more of the following ways:

- On the command line
- In a custom configuration file, which is a file with a `.cfg` extension
- In your source program
- As system environment variables
- In a makefile

The compiler assumes default settings for most compiler options not explicitly set by you in the ways listed above.

When specifying compiler options, it is possible for option conflicts and incompatibilities to occur. The XL C/C++ compiler resolves most of these conflicts and incompatibilities in a consistent fashion, as follows:

In most cases, the compiler uses the following order in resolving conflicting or incompatible options:

1. Pragma statements in source code override compiler options specified on the command line.

2. Compiler options specified on the command line override compiler options specified as environment variables or in a configuration file. If conflicting or incompatible compiler options are specified in the same command line compiler invocation, the subsequent option in the invocation takes precedence.
3. Compiler options specified as environment variables override compiler options specified in a configuration file.
4. Compiler options specified in a configuration file, command line or source program override compiler default settings.

Option conflicts that do not follow this priority sequence are described in “Resolving conflicting compiler options” on page 8.


## Specifying compiler options on the command line

Most options specified on the command line override both the default settings of the option and options set in the configuration file. Similarly, most options specified on the command line are in turn overridden by pragma directives, which provide you a means of setting compiler options right in the source file. Options that do not follow this scheme are listed in “Resolving conflicting compiler options” on page 8.

There are two kinds of command-line options:

- `-qoption_keyword` (compiler-specific)
- Flag options

### -q options



Command-line options in the `-qoption_keyword` format are similar to on and off switches. For *most* `-q` options, if a given option is specified more than once, the last appearance of that option on the command line is the one used by the compiler. For example, `-qsource` turns on the source option to produce a compiler listing, and `-qnosource` turns off the source option so no source listing is produced. For example:

```
xlc -qnosource MyFirstProg.c -qsource MyNewProg.c
```

would produce a source listing for both `MyNewProg.c` and `MyFirstProg.c` because the last **source** option specified (`-qsource`) takes precedence.

You can have multiple `-qoption_keyword` instances in the same command line, but they must be separated by blanks. Option keywords can appear in either uppercase or lowercase, but you must specify the `-q` in lowercase. You can specify any `-qoption_keyword` before or after the file name. For example:

```
xlc -qLIST -qfloat=nomaf file.c
xlc file.c -qxref -qsource
```

You can also abbreviate many compiler options. For example, specifying `-qopt` is equivalent to specifying `-qoptimize` on the command line.



Some options have suboptions. You specify these with an equal sign following the *-qoption*. If the option permits more than one suboption, a colon (:) must separate each suboption from the next. For example:

```
xlc -qflag=w:e -qattr=full file.c
```

compiles the C source file `file.c` using the option **-qflag** to specify the severity level of messages to be reported. The **-qflag** suboption **w** (warning) sets the minimum level of severity to be reported on the listing, and suboption **e** (error) sets the minimum level of severity to be reported on the terminal. The **-qattr** with suboption **full** will produce an attribute listing of all identifiers in the program.

## Flag options

XL C/C++ supports a number of common conventional flag options used on UNIX systems. Lowercase flags are different from their corresponding uppercase flags. For example, **-c** and **-C** are two different compiler options: **-c** specifies that the compiler should only preprocess and compile and not invoke the linker, while **-C** can be used with **-P** or **-E** to specify that user comments should be preserved.

XL C/C++ also supports flags directed to other programming tools and utilities (for example, the `ld` command). The compiler passes on those flags directed to **ld** at link time.

Some flag options have arguments that form part of the flag. For example:

```
xlc stem.c -F/home/tools/test3/new.cfg:xlc
```

where `new.cfg` is a custom configuration file.

You can specify flags that do not take arguments in one string. For example:

```
xlc -0cv file.c
```

has the same effect as:

```
xlc -0 -c -v file.c
```

and compiles the C source file `file.c` with optimization (**-O**) and reports on compiler progress (**-v**), but does not invoke the linker (**-c**).

A flag option that takes arguments can be specified as part of a single string, but you can only use one flag that takes arguments, and it must be the last option specified. For example, you can use the **-o** flag (to specify a name for the executable file) together with other flags, only if the **-o** option and its argument are specified last. For example:

```
xlc -0vo test test.c
```

has the same effect as:

```
xlc -0 -v -otest test.c
```

Most flag options are a single letter, but some are two letters. Note that specifying **-pg** (extended profiling) is not the same as specifying **-p -g** (**-p** for profiling, and **-g** for generating debug information). Take care not to specify two or more options in a single string if there is another option that uses that letter combination.

## Specifying compiler options in a configuration file

The default configuration file (`/etc/vac.cfg`) defines values and compiler options for the compiler. The compiler refers to this file when compiling C or C++ programs. The configuration file is a plain text file. You can edit this file, or create

an additional customized configuration file to support specific compilation requirements. For more information, see “Using custom compiler configuration files” on page 40.

## Specifying compiler options in program source files

You can specify compiler options within your program source by using pragma directives. A pragma is an implementation-defined instruction to the compiler. For those options that have equivalent pragma directives, you can have several ways to specify the syntax of the pragmas:

- Using **#pragma options *option\_name*** syntax — You can use command-line options with the **#pragma options** syntax, which takes the same name as the option, and suboptions with a syntax identical to that of the option. For example, if the command-line option is:

```
-qhalt=w
```

The pragma form is:

```
#pragma options halt=w
```

The descriptions for each individual option indicates whether this form of the pragma is supported. For details, see “#pragma options” on page 420.

- Using **#pragma *name*** syntax — Some options also have corresponding pragma directives that use a pragma-specific syntax, which may include additional or slightly different suboptions. Throughout the section “Individual option descriptions” on page 94, each option description indicates whether this form of the pragma is supported, and the syntax is provided.
- Using the standard C99 `_Pragma` operator — For options that support either forms of the pragma directives listed above, you can also use the C99 `_Pragma` operator syntax in both C and C++.

Complete details on pragma syntax are provided in “Pragma directive syntax” on page 383.

Other pragmas do not have equivalent command-line options; these are described in detail throughout Chapter 5, “Compiler pragmas reference,” on page 383.

Options specified with pragma directives in program source files override all other option settings, except other pragma directives. The effect of specifying the same pragma directive more than once varies. See the description for each pragma for specific information.

Pragma settings can carry over into included files. To avoid potential unwanted side effects from pragma settings, you should consider resetting pragma settings at the point in your program source where the pragma-defined behavior is no longer required. Some pragma options offer **reset** or **pop** suboptions to help you do this. These suboptions are listed in the detailed descriptions of the pragmas to which they apply.

## Resolving conflicting compiler options

In general, if more than one variation of the same option is specified (with the exception of **-qxref** and **-qattr**), the compiler uses the setting of the last one specified. Compiler options specified on the command line must appear in the order you want the compiler to process them.

Two exceptions to the rules of conflicting options are the `-Idirectory` and `-Ldirectory` options, which have cumulative effects when they are specified more than once.

In most cases, the compiler uses the following order in resolving conflicting or incompatible options:

1. Pragma statements in source code override compiler options specified on the command line.
2. Compiler options specified on the command line override compiler options specified as environment variables or in a configuration file. If conflicting or incompatible compiler options are specified on the command line, the option appearing later on the command line takes precedence.
3. Compiler options specified as environment variables override compiler options specified in a configuration file.
4. Compiler options specified in a configuration file override compiler default settings.

Not all option conflicts are resolved using the preceding rules. The following table summarizes exceptions and how the compiler handles conflicts between them. Rules for resolving conflicts between compiler mode and architecture-specific options are discussed in “Specifying compiler options for architecture-specific compilation.”

Option	Conflicting options	Resolution
<code>-qalias=allptrs</code>	<code>-qalias=noansi</code>	<code>-qalias=noansi</code>
<code>-qalias=typeptr</code>	<code>-qalias=noansi</code>	<code>-qalias=noansi</code>
<code>-qhalt</code>	Multiple severities specified by <code>-qhalt</code>	Lowest severity specified
<code>-qnoprint</code>	<code>-qxref</code> , <code>-qattr</code> , <code>-qsource</code> , <code>-qlistopt</code> , <code>-qlist</code>	<code>-qnoprint</code>
<code>-qfloat=rsqrt</code>	<code>-qnoignerrno</code>	Last option specified
<code>-qxref</code>	<code>-qxref=full</code>	<code>-qxref=full</code>
<code>-qattr</code>	<code>-qattr=full</code>	<code>-qattr=full</code>
<code>-qfloat=hsflt</code>	<code>-qfloat=spnans</code>	<code>-qfloat=hsflt</code>
<code>-qfloat=hssngl</code>	<code>-qfloat=spnans</code>	<code>-qfloat=hssngl</code>
<code>-E</code>	<code>-P</code> , <code>-o</code> , <code>-S</code>	<code>-E</code>
<code>-P</code>	<code>-c</code> , <code>-o</code> , <code>-S</code>	<code>-P</code>
<code>-#</code>	<code>-v</code>	<code>-#</code>
<code>-F</code>	<code>-B</code> , <code>-t</code> , <code>-W</code> , <code>-qpath</code>	<code>-B</code> , <code>-t</code> , <code>-W</code> , <code>-qpath</code>
<code>-qpath</code>	<code>-B</code> , <code>-t</code>	<code>-qpath</code>
<code>-S</code>	<code>-c</code>	<code>-S</code>
<code>-qnostdinc</code>	<code>-qc_stdinc</code> , <code>-qcpp_stdinc</code>	<code>-qnostdinc</code>

## Specifying compiler options for architecture-specific compilation

You can use the `-q32`, `-q64`, `-qarch`, and `-qtune` compiler options to optimize the output of the compiler to suit:

- The broadest possible selection of target processors
- A range of processors within a given processor architecture family
- A single specific processor

Generally speaking, the options do the following:

- **-q32** selects 32-bit execution mode.
- **-q64** selects 64-bit execution mode.
- **-qarch** selects the general family processor architecture for which instruction code should be generated. Certain **-qarch** settings produce code that will run *only* on systems that support *all* of the instructions generated by the compiler in response to a chosen **-qarch** setting.
- **-qtune** selects the specific processor for which compiler output is optimized. Some **-qtune** settings can also be specified as **-qarch** options, in which case they do not also need to be specified as a **-qtune** option. The **-qtune** option influences only the performance of the code when running on a particular system but does not determine where the code will run.

The compiler evaluates compiler options in the following order, with the last allowable one found determining the compiler mode:

1. Internal default (32-bit mode)
2. OBJECT\_MODE environment variable setting
3. Configuration file settings
4. Command line compiler options (**-q32**, **-q64**, **-qarch**, **-qtune**)
5. Source file statements (**#pragma options tune=***suboption*)

The compilation mode actually used by the compiler depends on a combination of the settings of the **-q32**, **-q64**, **-qarch**, and **-qtune** compiler options, subject to the following conditions:

- *Compiler mode* is set according to the last-found instance of the **-q32** or **-q64** compiler options. If neither of these compiler options is set, the compiler mode is set by the value of the OBJECT\_MODE environment variable. If the OBJECT\_MODE environment variable is also not set, the compiler assumes 32-bit compilation mode.
- *Architecture target* is set according to the last-found instance of the **-qarch** compiler option, provided that the specified **-qarch** setting is compatible with the *compiler mode* setting. If the **-qarch** option is not set, the compiler sets **-qarch** to the appropriate default based on the effective compiler mode setting. See “-qarch” on page 106 for details.
- Tuning of the architecture target is set according to the last-found instance of the **-qtune** compiler option, provided that the **-qtune** setting is compatible with the *architecture target* and *compiler mode* settings. If the **-qtune** option is not set, the compiler assumes a default **-qtune** setting according to the **-qarch** setting in use. If **-qarch** is not specified, the compiler sets **-qtune** to the appropriate default based on the effective **-qarch** as selected by default based on the effective compiler mode setting.

Allowable combinations of these options are found in “-qtune” on page 357.

The following list describes possible option conflicts and compiler resolution of these conflicts:

- **-q32** or **-q64** setting is incompatible with user-selected **-qarch** option.

**Resolution:** **-q32** or **-q64** setting overrides the **-qarch** option; compiler issues a warning message, sets **-qarch** to its default setting, and sets the **-qtune** option accordingly to its default value.

- **-qarch** option is incompatible with user-selected **-qtune** option.

**Resolution:** Compiler issues a warning message, and sets **-qtune** to the **-qarch** setting's default **-qtune** value.

- Selected **-qarch** or **-qtune** options are not known to the compiler.

**Resolution:** Compiler issues a warning message, sets **-qarch** and **-qtune** to their default settings. The compiler mode (32-bit or 64-bit) is determined by the **OBJECT\_MODE** environment variable or **-q32/-q64** compiler settings.

#### Related information

- “-qarch” on page 106
- “-qtune” on page 357
- “-q32, -q64” on page 96

---

## Reusing GNU C/C++ compiler options with **gxc** and **gxc++**


Each of the **gxc** and **gxc++** utilities accepts GNU C or C++ compiler options and translates them into comparable XL C/C++ options. Both utilities use the XL C/C++ options to create an **xlc** or **xlc++** invocation command, which they use to invoke the compiler. These utilities are provided to facilitate the reuse of make files created for applications previously developed with GNU C/C++. However, to fully exploit the capabilities of XL C/C++, it is recommended that you use the XL C/C++ invocation commands and their associated options.

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

The **gxc** and **gxc++** option mappings are modifiable. For information on adding to or editing the **gxc** or **gxc++** configuration file, see “Configuring the **gxc** or **gxc++** option mapping” on page 43.

### **gxc** or **gxc++** syntax

The following diagram shows the **gxc** or **gxc++** syntax:



where:

*filename*

Is the name of the file to be compiled.

- v** Verifies the command that is used to invoke XL C/C++. The utility displays the XL C/C++ invocation command that it has created, before using it to invoke the compiler.
- vv** Runs a simulation. The utility displays the XL C/C++ invocation command that it has created, but does not invoke the compiler.
- Wx, xlc\_or\_xlc++\_options** Sends the given XL C/C++ options directly to the **xlc** or **xlc++** invocation command. The utility adds the given options to the XL C/C++ invocation it is creating, without attempting to translate them. Use this option with

known XL C/C++ options to improve the performance of the utility. Multiple `xl_c_or_xl_cpp_options` are delimited by a comma.

`-gcc_or_gplusplus_options`

The GNU C/C++ options that are translated to XL C/C++ options. The utility emits a warning for any option it cannot translate. The GNU C/C++ options that are currently recognized by `gxl_c` or `gxl_cpp` are in the configuration file `gxl_c.cfg`. Multiple `-gcc_or_gplusplus_options` are delimited by the space character.

## Examples

To use the GCC `-fstrict-aliasing` option to compile the C version of the Hello World program, you can use:

```
gxl_c -fstrict-aliasing hello.c
```

which translates into:

```
xl_c -qalias=ansi hello.c
```

This command is then used to invoke the XL C compiler.

## Related information

- “Configuring the `gxl_c` or `gxl_cpp` option mapping” on page 43

---

## Preprocessing

Preprocessing manipulates the text of a source file, usually as a first phase of translation that is initiated by a compiler invocation. Common tasks accomplished by preprocessing are macro substitution, testing for conditional compilation directives, and file inclusion.

You can invoke the preprocessor separately to process text without compiling. The output is an intermediate file, which can be input for subsequent translation. Preprocessing without compilation can be useful as a debugging aid because it provides a way to see the result of include directives, conditional compilation directives, and complex macro expansions.

The following table lists the options that direct the operation of the preprocessor.

Option	Description
<code>“-E”</code> on page 142	Preprocesses the source files and writes the output to standard output. By default, <code>#line</code> directives are generated.
<code>“-P”</code> on page 280	Preprocesses the source files and creates an intermediary file with a <code>.i</code> file name suffix for each source file. By default, <code>#line</code> directives are not generated.
<code>“-qppline”</code> on page 293	Toggles on and off the generation of <code>#line</code> directives for the <code>-E</code> and <code>-P</code> options.
<code>“-C, -C!”</code> on page 120	Preserves comments in preprocessed output.
<code>“-D”</code> on page 132	Defines a macro name from the command line, as if in a <code>#define</code> directive.
<code>“-U”</code> on page 361	Undefines a macro name defined by the compiler or by the <code>-D</code> option.

Option	Description
"-qshowmacros" on page 315	Emits macro definitions to preprocessed output.

## Directory search sequence for include files

The XL C/C++ compiler supports the following types of include files:





- Header files supplied by the compiler (referred to throughout this document as *XL C/C++ headers*)
- Header files mandated by the C and C++ standards (referred to throughout this document as *system headers*)
- Header files supplied by the operating system (also referred to throughout this document as *system headers*)
- User-defined header files

You can use any of the following methods to include any type of header file:

- Use the standard `#include <file_name>` preprocessor directive in the including source file.
- Use the standard `#include "file_name"` preprocessor directive in the including source file.
- Use the **-qinclude** compiler option.

If you specify the header file using a full (absolute) path name, you can use these methods interchangeably, regardless of the type of header file you want to include. However, if you specify the header file using a *relative* path name, the compiler uses a different directory search order for locating the file depending on the method used to include the file.

Furthermore, the **-qidirfirst** and **-qstdinc** compiler options can affect this search order. The following summarizes the search order used by the compiler to locate header files depending on the mechanism used to include the files and on the compiler options that are in effect:

1. Header files included with **-qinclude** only: The compiler searches the current (working) directory from which the compiler is invoked.<sup>1</sup>
2. Header files included with **-qinclude** or `#include "file_name"`: The compiler searches the directory in which the including file is located.<sup>1</sup>
3. All header files: The compiler searches each directory specified by the **-I** compiler option, in the order that it displays on the command line.
4. All header files:  The compiler searches the standard directory for the XL C headers. The default directory for these headers is specified in the compiler configuration file. This is normally `/usr/vac/include/`, but the search path can be changed with the **-qc\_stdinc** compiler option.  The compiler searches the standard directory for the XL C++ headers. The default directory for these headers is specified in the compiler configuration file. This is normally `/usr/vacpp/include/`, but the search path can be changed with the **-qcpp\_stdinc** compiler option.<sup>2</sup>
5. All header files:  The compiler searches the standard directory for the system headers. The default directory for these headers is specified in the compiler configuration file. This is normally `/usr/include/`. But the search path can be changed with **-qc\_stdinc**.  The compiler searches the standard directory for the system headers. The default directory for these headers is

specified in the compiler configuration file. This is normally `/usr/include/`. But the search path can be changed with `-qcpp_stdinc2`

**Note:**

1. If the `-qidirfirst` compiler option is in effect, step 3 is performed before steps 1 and 2.
2. If the `-qnostdinc` compiler option is in effect, steps 4 and 5 are omitted.

**Related information**

- “-I” on page 181
- “-qc\_stdinc (C only)” on page 130
- “-qcpp\_stdinc (C++ only)” on page 131
- “-qidirfirst” on page 182
- “-qinclude” on page 185
- “-qstdinc” on page 332

---

## Linking

The linker links specified object files to create one executable file. Invoking the compiler with one of the invocation commands automatically calls the linker unless you specify one of the following compiler options: `-E`, `-P`, `-c`, `-S`, `-qsyntaxonly` or `-#`.

**Input files**

Object files, unstripped executable files, and library files serve as input to the linker. Object files must have a `.o` suffix, for example, `filename.o`. Static library file names have an `.a` suffix, for example, `filename.a`. Dynamic library file names typically have a `.so` suffix, for example, `filename.so`.

**Output files**

The linker generates an *executable file* and places it in your current directory. The default name for an executable file is `a.out`. To name the executable file explicitly, use the `-o file_name` option with the compiler invocation command, where *file\_name* is the name you want to give to the executable file. For example, to compile `myfile.c` and generate an executable file called `myfile`, enter:

```
xlc myfile.c -o myfile
```

If you use the `-qmkshrobj` option to create a shared library, the default name of the shared object created is `shr.o`. You can use the `-o` option to rename the file and give it a `.so` suffix.

You can invoke the linker explicitly with the `ld` command. However, the compiler invocation commands set several linker options, and link some standard files into the executable output by default. In most cases, it is better to use one of the compiler invocation commands to link your object files. For a complete list of options available for linking, see “Linking” on page 91.

**Related information**

“-qmkshrobj” on page 260

## Order of linking

The compiler links libraries in the following order:

1. System startup libraries



2. User .o files and libraries
3. XL C/C++ libraries
4. C++ standard libraries
5. C standard libraries

#### Related information

- “Linking” on page 91
- “Redistributable libraries”
- `ld` in the *AIX Commands Reference, Volume 5: s through u*

## Redistributable libraries

If you build your application using XL C/C++, it might use one or more of the following redistributable libraries. If you ship the application, ensure that the users of the application have the filesets containing the libraries. To make sure the required libraries are available to users, you must do one of the following:

- You can ship the filesets that contain the redistributable libraries with the application. The filesets are stored under the `runtime/` directory on the installation CD.
- The user can download the filesets that contain the redistributable libraries from the XL C/C++ support website at:

<http://www.ibm.com/software/awdtools/xlcpp/aix/support/>

For information about the licensing requirements related to the distribution of these filesets see the `LicAgree.pdf` file on the CD.

Table 6. Redistributable libraries

Fileset	Libraries (and default installation path)	Description
xlC.rte	<code>/usr/lpp/xlC/lib/libibmcls.a</code> <code>/usr/lpp/xlC/lib/libibmuis.a</code>	XL C++ runtime libraries
xlC.aix50.rte	<code>/usr/lpp/xlC/lib/aix53/libC.a</code> <code>/usr/lpp/xlC/lib/aix53/libC128.a</code> <code>/usr/lpp/xlC/lib/profiled/aix53/libC.a</code> <code>/usr/lpp/xlC/lib/profiled/aix53/libC128.a</code>	XL C++ runtime environment for AIX 5.3 libraries
xlC.msg.en_US.rte	<code>/usr/lib/nls/msg/en_US/ibmcl.cat</code>	XL C++ runtime messages (English)
xlC.msg.ja_JP.rte	<code>/usr/lib/nls/msg/ja_JP/ibmcl.cat</code>	XL C++ runtime messages (Japanese, IBM-eucJP)
xlC.msg.Ja_JP.rte	<code>/usr/lib/nls/msg/Ja_JP/ibmcl.cat</code>	XL C++ runtime messages (Japanese, IBM-943)
xlsmprte	<code>/usr/include/omp.h</code> <code>/usr/lpp/xlsmpr/default_msg/smprt.cat</code>	SMP runtime environment
xlsmprte.aix53.rte	<code>/usr/lpp/xlsmpr/aix53/libx1omp_ser.a</code> <code>/usr/lpp/xlsmpr/aix53/libx1smpr.a</code> <code>/usr/lpp/xlsmpr/aix53/libx1smprdebug.a</code>	SMP runtime environment for AIX
xlsmprte.msg.en_US.rte	<code>/usr/lib/nls/msg/en_US/smprt.cat</code>	SMP runtime messages (English, ISO8859-1)
xlsmprte.msg.EN_US.rte	<code>/usr/lib/nls/msg/EN_US/smprt.cat</code>	SMP runtime messages (English, UTF-8)
xlsmprte.msg.ja_JP.rte	<code>/usr/lib/nls/msg/ja_JP/smprt.cat</code>	SMP runtime messages (Japanese, IBM-eucJP)
xlsmprte.msg.Ja_JP.rte	<code>/usr/lib/nls/msg/Ja_JP/smprt.cat</code>	SMP runtime messages (Japanese, IBM-943)
xlsmprte.msg.JA_JP.rte	<code>/usr/lib/nls/msg/JA_JP/smprt.cat</code>	SMP runtime messages (Japanese, UTF-8)

Table 6. Redistributable libraries (continued)

Fileset	Libraries (and default installation path)	Description
xlsmp.msg.zh_CN.rte	/usr/lib/nls/msg/zh_CN/smp.rtc	SMP runtime messages (Chinese, IBM-eucCN)
xlsmp.msg.ZH_CN.rte	/usr/lib/nls/msg/ZH_CN/smp.rtc	SMP runtime messages (Chinese, UTF-8)
xlsmp.msg.Zh_CN.rte	/usr/lib/nls/msg/Zh_CN/smp.rtc	SMP runtime messages (Chinese, GBK)
vac.aix53.lib	/usr/vac/lib/aix53/libx1.a /usr/vac/lib/aix53/libxlopt.a	XL C libraries for AIX
vacpp.cmp.rte	/usr/vacpp/lib/libC.a /usr/vacpp/lib/libC_r.a /usr/vacpp/lib/libC128.a /usr/vacpp/lib/libC128_r.a /usr/vacpp/lib/profiled/libC.a /usr/vacpp/lib/profiled/libC_r.a /usr/vacpp/lib/profiled/libC128.a /usr/vacpp/lib/profiled/libC128_r.a	XL C++ compiler application runtime libraries
vacpp.cmp.tools	/usr/vacpp/bin/c++filt /usr/vacpp/bin/linkx1C /usr/vacpp/bin/makeC++SharedLib /usr/vacpp/bin/makeC++SharedLib_r /usr/vacpp/bin/makeC++SharedLib128 /usr/vacpp/bin/makeC++SharedLib_r7 /usr/vacpp/exe/aix53/munch	XL C++ utilities
vacpp.memdbg.aix53.rte	/usr/vacpp/lib/aix53/libhC.a /usr/vacpp/lib/aix53/libhC_r.a /usr/vacpp/lib/profiled/aix53/libhC.a /usr/vacpp/lib/profiled/aix53/libhC_r.a	User heap/memory debug toolkit for AIX
vacpp.memdbg.rte	/usr/vacpp/lib/libhm.a /usr/vacpp/lib/libhm_r.a /usr/vacpp/lib/libhmd.a /usr/vacpp/lib/libhmd_r.a /usr/vacpp/lib/libhmu.a /usr/vacpp/lib/libhmu_r.a /usr/vacpp/lib/libhu.a /usr/vacpp/lib/libhu_r.a /usr/vacpp/lib/profiled/libhm.a /usr/vacpp/lib/profiled/libhm_r.a /usr/vacpp/lib/profiled/libhmd.a /usr/vacpp/lib/profiled/libhmd_r.a /usr/vacpp/lib/profiled/libhmu.a /usr/vacpp/lib/profiled/libhmu_r.a /usr/vacpp/lib/profiled/libhu.a /usr/vacpp/lib/profiled/libhu_r.a	User heap/memory debug toolkit

## Compatibility with earlier versions

This section describes issues about compatibility with earlier versions and their workarounds.

### Compiler option compatibility issues

In IBM XL C/C++ for AIX, V12.1, the implementation of the threadprivate data, that is, OpenMP threadprivate variable, has been improved. The operating system thread local storage is used instead of the runtime implementation. The new implementation might improve performance on some applications.

If you plan to mix the object files .o that you have compiled with levels prior to 11.1 with the object files that you compiled with IBM XL C/C++ for AIX, V12.1,

and the same OpenMP threadprivate variables are referenced in both old and new object files, different implementations might cause incompatibility issues. A link error, a compile time error or other undefined behaviors might occur. To support compatibility with earlier versions, you can use the `-qsmp=noostls` suboption to switch back to the old implementation. You can recompile the entire program with the default suboption `-qsmp=ostls` to get the benefit of the new implementation.

If you are not sure whether the object files you have compiled with levels prior to 11.1 contain any old implementation, you can use the `nm` command to determine whether you need to use the `-qsmp=noostls` suboption. The following code is an example that shows how to use the `nm` command:

```
> nm oldfiles.o
...
._xlGetThStorageBlock U          -
._xlGetThValue          U          -
...
```

In the preceding example, if `_xlGetThStorageBlock` or `_xlGetThValue` is found, this means the object files contain old implementation. In this case, you must use `-qsmp=noostls`; otherwise, use the default suboption `-qsmp=ostls`.

---


## Compiler messages and listings

The following sections discuss the various methods of reporting provided by the compiler after compilation.

- “Compiler messages”
- “Compiler return codes” on page 19
- “Compiler listings” on page 20
- “Message catalog errors” on page 22
- “Paging space errors during compilation” on page 23

### Compiler messages

When the compiler encounters a programming error while compiling a C or C++ source program, it issues a diagnostic message to the standard error device, or to a listing file if you compile with the `-qsource` option. These diagnostic messages are specific to the C or C++ language.

 If you specify the compiler option `-qsrcmsg` and the error is applicable to a particular line of code, the reconstructed source line or partial source line is included with the error message. A reconstructed source line is a preprocessed source line that has all the macros expanded.

You can control the diagnostic messages issued, according to their severity, using either the `-qflag` option or the `-w` option. To get additional informational messages about potential problems in your program, use the `-qinfo` option.

#### Related information

- “-qsource” on page 324
- “-qsrcmsg (C only)” on page 328
- “-qflag” on page 153
- “-w” on page 371
- “-qinfo” on page 187

## Compiler message format

Diagnostic messages have the following format:

*"file", line line\_number.column\_number: 15dd-number (severity) text.*

where:

*file*

Is the name of the C or C++ source file with the error.

*line\_number*

Is the source code line number where the error was found.

*column\_number*

Is the source code column number where the error was found.

**15** Is the compiler product identifier.

*dd* Is a two-digit code indicating the compiler component that issued the message.  
*dd* can have the following values:

- 00** - code generating or optimizing message
- 01** - compiler services message
- 05** - message specific to the C compiler
- 06** - message specific to the C compiler
- 40** - message specific to the C++ compiler
- 47** - message specific to the C++ linkage helper
- 86** - message specific to interprocedural analysis (IPA)

*number*


Is the message number.

*severity*

Is a letter representing the severity of the error. See "Message severity levels and compiler response" for a description of these.

*text*

Is a message describing the error.

 If you compile with **-qsrcmsg**, diagnostic messages have the following format:

*x - 15dd-*nnn*(severity) text.*



where *x* is a letter referring to a finger in the finger line.

## Message severity levels and compiler response

The XL C/C++ compiler uses a multilevel classification scheme for diagnostic messages. Each level of severity is associated with a compiler response. The following table provides a key to the abbreviations for the severity levels and the associated default compiler response. You can adjust the default compiler response by using any of the following options:

- **-qhalt** halts the compilation phase at a lower severity level than the default
- **-qmaxerr** halts the compilation phase as soon as a specific number of errors at a specific severity level is reached
- **-qhaltonmsg** halts the compilation phase as soon as a specific error is encountered

Table 7. Compiler message severity levels

Letter	Severity	Compiler response
I	Informational	Compilation continues and object code is generated. The message reports conditions found during compilation.
W	Warning	Compilation continues and object code is generated. The message reports valid but possibly unintended conditions.
 E	Error	Compilation continues and object code is generated. Error conditions exist that the compiler can correct, but the program might not produce the expected results.
S	Severe error	Compilation continues, but object code is not generated. Error conditions exist that the compiler cannot correct: <ul style="list-style-type: none"> <li>• If the message indicates a resource limit (for example, file system full or paging space full), provide additional resources and recompile.</li> <li>• If the message indicates that different compiler options are needed, recompile using them.</li> <li>• Check for and correct any other errors reported prior to the severe error.</li> <li>• If the message indicates an internal compile-time error, the message should be reported to your IBM service representative.</li> </ul>
 U	Unrecoverable error	The compiler halts. An internal compile-time error has occurred. The message should be reported to your IBM service representative.

#### Related information

- “-qhalt” on page 174
- “-qmaxerr” on page 254
- “-qhaltmsg” on page 175
- Options summary by functional category: Listings and messages

## Compiler return codes

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

- No messages are issued.
- The highest severity level of all errors diagnosed is less than the setting of the **-qhalt** compiler option, and the number of errors did not reach the limit set by the **-qmaxerr** compiler option.
- No message specified by the **-qhaltmsg** compiler option is issued.

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

Return code	Error type
1	Any error with a severity level higher than the setting of the <b>-qhalt</b> compiler option has been detected.
40	An option error or an unrecoverable error has been detected.
41	A configuration file error has been detected.
249	A no-files-specified error has been detected.
250	An out-of-memory error has been detected. The compiler cannot allocate any more memory for its use.

251	A signal-received error has been detected. That is, an unrecoverable error or interrupt signal has occurred.
252	A file-not-found error has been detected.
253	An input/output error has been detected: files cannot be read or written to.
254	A fork error has been detected. A new process cannot be created.
255	An error has been detected while the process was running.

**Note:** Return codes can also be displayed for runtime errors. For example, a runtime return code of 99 indicates that a static initialization has failed.

### gxc and gxc++ return codes

Like other invocation commands, gxc and gxc++ return output, such as listings, diagnostic messages related to the compilation, warnings related to unsuccessful translation of GNU options, and return codes. If gxc or gxc++ cannot successfully call the compiler, it sets the return code to one of the following values:

**40** A gxc or gxc++ option error or unrecoverable error has been detected.

**255** An error has been detected while the process was running.

## Compiler listings

A listing is a compiler output file (with a .lst suffix) that contains information about a particular compilation. As a debugging aid, a compiler listing is useful for determining what has gone wrong in a compilation. For example, any diagnostic messages emitted during compilation are written to the listing.

To produce a listing, you can compile with any of the following options, which provide different types of information:

- -qsource
- -qlistopt
- -qattr
- -qxref
- -qlist
- -qreport

When any of these options is in effect, a listing file *filename.lst* is saved in the current directory for every input file named in the compilation.

Listing information is organized in sections. A listing contains a header section and a combination of other sections, depending on other options in effect. The contents of these sections are described as follows.

#### Header section

Lists the compiler name, version, release, the source file name, and the date and time of the compilation.

#### Source section

If you use the **-qsource** option, lists the input source code with line numbers. If there is an error at a line, the associated error message is displayed after the source line. Lines containing macros have additional lines showing the macro expansion. By default, this section only lists the main source file. Use the **-qshowinc** option to expand all header files as well.

### Options section

Lists the non-default options that were in effect during the compilation. To list all options in effect, specify the `-qlistopt` option.

### Attribute and cross-reference listing section

If you use the `-qattr` or `-qxref` options, provides information about the variables used in the compilation unit, such as type, storage duration, scope, and where they are defined and referenced. Each of these options provides different information about the identifiers used in the compilation.

### File table section

Lists the file name and number for each main source file and include file. Each file is associated with a file number, starting with the main source file, which is assigned file number 0. For each file, the listing shows from which file and line the file was included. If the `-qshowinc` option is also in effect, each source line in the source section has a file number to indicate which file the line came from.

### PDF report section

The following information is included in this section when you use the `-qreport` option with the `-qpdf2` option:

#### Loop iteration count

The most frequent loop iteration count and the average iteration count, for a given set of input data, are 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 option `-qpdf1=level=2` to get this report. You can also select the level of cache to profile using the environment variable `PDF_PM_EVENT` during run time.

#### Relevance of profiling data

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

#### Missing profiling data

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

### Outdated profiling data

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

### Transformation report section

If the **-qreport** option is in effect, this section displays pseudo code that corresponds to the original source code, so that you can see parallelization and loop transformations that the **-qhot** or **-qsmp** option has generated. This section of the report also shows additional loop transformation and parallelization information about loop nests if you compile with **-qsmp** and **-qhot=level=2**.

This section also reports the number of streams created for a given loop and the location of data prefetch instructions inserted by the compiler. To generate information about data prefetch insertion locations, use the optimization level of **-qhot**, **-O3 -qhot**, **-O4** or **-O5** together with **-qreport**.

### Data reorganization section

Displays data reorganization messages for program variable data during the IPA link pass when **-qreport** is used with **-qipa=level=2** or **-O5**. Reorganization information includes:

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

### Compilation epilogue section

Displays a summary of the diagnostic messages by severity level, the number of source lines read, and whether the compilation was successful.

### Object section

If you use the **-qlist** option, lists the object code generated by the compiler. This section is useful for diagnosing execution-time problems, if you suspect the program is not performing as expected due to code generation error.

### Related information

- Summary of command line options: Listings and messages

## Message catalog errors

Before the compiler can compile your program, the message catalogs must be installed and the environment variables *LANG* and *NLSPATH* must be set to a language for which the message catalog has been installed.

If you see the following message during compilation, the appropriate message catalog cannot be opened:

```
Error occurred while initializing the message system in
file: message_file
```

where *message\_file* is the name of the message catalog that the compiler cannot open. This message is issued in English only.



You must then verify that the message catalogs and the environment variables are in place and correct. If the message catalog or environment variables are not correct, compilation can continue, but diagnostic messages are suppressed and the following message is issued instead:

```
No message text for message_number
```

where *message\_number* is the compiler internal message number. This message is issued in English only.

To determine which message catalogs are installed on your system, assuming that you have installed the compiler to the default location, you can list all of the file names for the catalogs by the following command:

```
ls /usr/lib/nls/msg/$LANG/*.cat
```

where *LANG* is the environment variable on your system that specifies the system locale.

The compiler calls the default message catalogs in `/usr/vacpp/exe/default_msg/` when:

- The message catalogs for the locale specified by *LANG* cannot be found.
- The locale has never been changed from the default, *C*.

For more information about the *NLSPATH* and *LANG* environment variables, see your operating system documentation.

## Paging space errors during compilation

If the operating system runs low on paging space during a compilation, the compiler issues one of the following messages:

```
1501-229 Compilation ended due to lack of space.  
1501-224 fatal error in ../exe/xlCcode: signal 9 received.
```

If lack of paging space causes other compiler programs to fail, the following message is displayed:

```
Killed.
```

To minimize paging-space problems, do any of the following and recompile your program:

- Reduce the size of your program by splitting it into two or more source files
- Compile your program without optimization
- Reduce the number of processes competing for system paging space
- Increase the system paging space

To check the current paging-space settings enter the command: **lsp**s -a or use the AIX System Management Interface Tool (SMIT) command **smit pgsp**.

For more information about paging space and how to allocate it, see your operating system documentation.



---

## Chapter 2. Configuring compiler defaults

When you compile an application with XL C/C++, the compiler uses default settings that are determined in a number of ways:

- Internally defined settings. These settings are predefined by the compiler and you cannot change them.
- Settings defined by system environment variables. Certain environment variables are required by the compiler; others are optional. You might have already set some of the basic environment variables during the installation process (for more information, see the XL C/C++ Installation Guide). “Setting environment variables” provides a complete list of the required and optional environment variables you can set or reset after installing the compiler, including those used for parallel processing.
- Settings defined in the compiler configuration file, `vac.cfg`. The compiler requires many settings that are determined by its configuration file. Normally, the configuration file is automatically generated during the installation procedure. (For more information, see the XL C/C++ Installation Guide). However, you can customize this file after installation, to specify additional compiler options, default option settings, library search paths, and other settings. Information on customizing the configuration file is provided in “Using custom compiler configuration files” on page 40.
- Settings defined by the GCC options configuration file. If you are using the `gxc` or `gxc++` utility to map GCC options, the default option mappings are defined in the `/etc/gxc.cfg` file. You can customize this file to suit your requirements; for more information, see “Configuring the `gxc` or `gxc++` option mapping” on page 43.

---

### Setting environment variables

To set environment variables in Bourne, Korn, and BASH shells, use the following commands:

```
variable=value  
export variable
```

where *variable* is the name of the environment variable, and *value* is the value you assign to the variable.

To set environment variables in the C shell, use the following command:

```
setenv variable value
```

where *variable* is the name of the environment variable, and *value* is the value you assign to the variable.

To set the variables so that all users have access to them, in Bourne, Korn, and BASH shells, add the commands to the file `/etc/profile`. To set them for a specific user only, add the commands to the file `.profile` in the user's home directory. In C shell, add the commands to the file `/etc/csh.cshrc`. To set them for a specific user only, add the commands to the file `.cshrc` in the user's home directory. The environment variables are set each time the user logs in.

The following sections discuss the environment variables you can set for XL C/C++ and applications you have compiled with it:

- “Compile-time and link-time environment variables”
- “Runtime environment variables” on page 27

## Compile-time and link-time environment variables

The following environment variables are used by the compiler when you are compiling and linking your code. Many are built into the AIX operating system. With the exception of *LANG* and *NLSPATH*, which must be set if you are using a locale other than the default *en\_US*, all of these variables are optional.

### LANG

Specifies the locale for your operating system. The default locale used by the compiler for messages and help files is United States English, *en\_US*, but the compiler supports other locales. For a list of these, see *National language support* in the *XL C/C++ Installation Guide*. For more information on setting the *LANG* environment variable to use an alternate locale, see your operating system documentation.

### LIBPATH

Specifies an alternate directory search path for dynamically linked libraries at application run time. If shared libraries required by your application have been moved to an alternate directory that was not specified at link time, and you do not want to relink the executable, you can set this environment variable to allow the dynamic linker to locate them at run time. For more information about this environment variable, see your operating system documentation.

### NLSPATH

Specifies the directory search path for finding the compiler message and help files. You only need to set this environment variable if the national language to be used for the compiler message and help files is not English. For information on setting the *NLSPATH*, see *Enabling the XL C/C++ error messages* in the *XL C/C++ Installation Guide*.

### OBJECT\_MODE

Optionally specifies the bit mode for compilation to either 32 or 64 bits. This is equivalent to the **-q32** and **-q64** compiler options. Set the *OBJECT\_MODE* environment variable to a value of 32 for 32-bit compilation mode, or 64 for 64-bit compilation mode. If unspecified, the default compilation mode is 32 bits. See also “-q32, -q64” on page 96 for more information.

**PATH** Specifies the directory search path for the executable files of the compiler. Executables are in */usr/vac/bin/* and */usr/vacpp/bin/* if installed to the default location.

### TMPDIR

Optionally specifies the directory in which temporary files are created during compilation. The default location, */tmp/*, may be inadequate at high levels of optimization, where paging and temporary files can require significant amounts of disk space, so you can use this environment variable to specify an alternate directory.

### XLC\_USR\_CONFIG

Specifies 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 40.

## Runtime environment variables

The following environment variables are used by the system loader or by your application when it is executed. All of these variables are optional.

### LIBPATH

Specifies an alternate directory search path for dynamically linked libraries at application run time. If shared libraries required by your application have been moved to an alternate directory that was not specified at link time, and you do not want to relink the executable, you can set this environment variable to allow the dynamic linker to locate them at run time. For more information about this environment variable, see your operating system documentation.

### MALLOCALIGN=16

Specifies that dynamic memory allocations return 16-byte aligned addresses. See also “-qipa” on page 199.

### PDFDIR

Optionally specifies the directory in which profiling information is saved when you run an application that you have compiled with the **-qpdf1** option. The default value is unset, and the compiler places the profile data file 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. When you recompile or relink your program with the **-qpdf2** option, the compiler uses the data saved in this directory to optimize the application. It is recommended that you set this variable to an absolute path if you use profile-directed feedback (PDF). See “-qpdf1, -qpdf2” on page 283 for more information.

### 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\_BIND\_PROCESSOR

If you want to bind your process to a particular processor, you can specify the PDF\_BIND\_PROCESSOR environment variable to bind the process tree from the executable to a different processor. Processor 0 is set by default.

### XL\_AR

To use your own archive files when generating a nonexecutable package with **-r -qipa=relink**, you can use the **ar** tool and set the **XL\_AR** environment variable to point to it. See **-qipa** for more information.

## Environment variables for parallel processing

The XLSMPOPTS environment variable sets options for program run time using loop parallelization. Suboptions for the XLSMPOPTS environment variables are discussed in detail in “XLSMPOPTS” on page 28.

If you are using OpenMP constructs for parallelization, you can also specify runtime options using the OMP environment variables, as discussed in “Environment variables for OpenMP” on page 33.

When runtime options specified by OMP and XLSMPOPTS environment variables conflict, OMP options will prevail.


**Note:** You must use a threadsafe compiler mode invocations when compiling parallelized program code.

#### Related information

- “Pragma directives for parallel processing” on page 439
- “Built-in functions for parallel processing” on page 609

## XLSMPOPTS

Runtime options affecting parallel processing can be specified with the XLSMPOPTS environment variable. This environment variable must be set before you run an application, and uses basic syntax of the form:



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 XLSMPOPTS option string contains imbedded blanks, you must enclose the entire option string in double quotation marks (").

For example, to have a program run time create 4 threads and use dynamic scheduling with chunk size of 5, you would set the XLSMPOPTS environment variable as shown below:

```
XLSMPOPTS=PARTHDS=4:SCHEDULE=DYNAMIC=5
```

The following are the available runtime option settings for the XLSMPOPTS environment variable:

Scheduling options are as follows:

#### **schedule**

Specifies the type of scheduling algorithms and chunk size ( $n$ ) that are used for loops to which no other scheduling algorithm has been explicitly assigned in the source code.

Work is assigned to threads in a different manner, depending on the scheduling type and chunk size used. Choosing chunking granularity is a tradeoff between overhead and load balancing. The syntax for this option is **schedule=***suboption*, where the suboptions are defined 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 does not appear in the OpenMP API standard.

**dynamic[=*n*]**

The iterations of a loop are divided into chunks containing *n* iterations each. If *n* is not specified, then the chunks consist of **ceiling**(*number\_of\_iterations*/*number\_of\_threads*) iterations.

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.

If a thread is asleep, its assigned work will be taken over by an active thread once that thread becomes available.

**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.

**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 **ceiling**(*number\_of\_iterations*/*number\_of\_threads*) iterations. Each thread is assigned one of these chunks. 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 integral assignment expression of value 1 or greater.

Specifying **schedule** with no suboption is equivalent to **schedule=runtime**.

Parallel environment options are as follows:

**parthds=*num***

Specifies the number of threads (*num*) requested, which is usually equivalent to the number of processors available on the system.

Some applications cannot use more threads than the maximum number of processors available. Other applications can experience significant performance improvements if they use more threads than there are processors. This option gives you full control over the number of user threads used to run your program.

The default value for *num* is the number of processors available on the system.

**usrthds=*num***

Specifies the maximum number of threads (*num*) that you expect your code will explicitly create if the code does explicit thread creation. The default value for *num* is 0.

**stack=num**

Specifies the largest amount of space in bytes (*num*) that a thread's stack needs. The default value for *num* is 4194304.

Set *num* so it is within the acceptable upper limit. *num* can be up to 256 MB for 32-bit mode, or up to the limit imposed by system resources for 64-bit mode. An application that exceeds the upper limit may cause a segmentation fault.

**stackcheck[=num]**

When the **-qsmp=stackcheck** is in effect, enables stack overflow checking for slave threads at runtime. *num* is the size of the stack in bytes; when the remaining stack size is less than this value, a runtime warning message is issued. If you do not specify a value for *num*, the default value is 4096 bytes. Note that this option only has an effect when the **-qsmp=stackcheck** has also been specified at compile time. See “-qsmp” on page 320 for more information.

**startproc=cpu\_id**

Enables thread binding and specifies the *cpu\_id* to which the first thread binds. If the value provided is outside the range of available processors, a warning message is issued and no threads are bound.

**procs=cpu\_id[,cpu\_id,...]**

Enables thread binding and specifies a list of *cpu\_id* to which the threads are bound. If the number of CPU IDs specified is less than the number of threads used by the program, the remaining threads are not bound.

**stride=num**

Specifies the increment used to determine the *cpu\_id* to which subsequent threads bind. *num* must be greater than or equal to 1. If the value provided causes a thread to bind to a CPU outside the range of available processors, a warning message is issued and no threads are bound.

**bind=SDL=n1,n2,n3**

Specifies different system detail levels to bind threads by using the Resource Set API. This suboption supports binding a thread to multiple logical processors.

*SDL* stands for System Detail Level and can be MCM, L2CACHE, PROC\_CORE, or PROC. If the *SDL* value is not specified, or an incorrect *SDL* value is specified, the SMP runtime issues an error message.

The list of three integers *n1,n2,n3* determines how to divide threads among resources (one of *SDL*s). *n1* is the starting *resource\_id*, *n2* is the number of requested resources, and *n3* is the stride, which specifies the increment used to determine the next *resource\_id* to bind. *n1,n2,n3* must all be specified; otherwise, the SMP runtime issues an error message and default binding rules apply.

When the number of resources specified in **bind** is greater than the number of threads, the extra resources are ignored.

When the number of threads *t* is greater than the number of resources *x*, *t* threads are divided among *x* resources according to the following formula:

The  $\text{ceil}(t/x)$  threads are bound to the first  $(t \bmod x)$  resources. The  $\text{floor}(t/x)$  threads will be bound to the remaining resources.

With the **XL SMP\_OPTS** environment variable being set as in the following example, a program runs with 16 threads. It binds threads to PROC 0, 2, 4, 6, 8, 10, 12, 14, 16, 18, 20, 22, 24, 26, 28, 30.



```
XLSMPOPTS="bind=PROC=0,16,2"
```

#### Notes:

- The **bind** suboption takes precedence over the **startproc/stride** and **procs** suboptions. However, **bindlist** takes precedence over **bind**.
- Resource Set can only be used by a user account with the **CAP\_NUMA\_ATTACH** and **CAP\_PROPAGATE** capabilities. These capabilities are set on a per-user basis by using the **chuser** command as follows:  

```
chuser "capabilities=CAP_PROPAGATE,CAP_NUMA_ATTACH" username
```
- If the *resource\_id* specified in **bind** is outside the range of 0 to **INT32\_MAX**, where **INT32\_MAX** is 2147483647 as defined in `stdint.h`, the SMP runtime issues an error message and default binding rules apply.
- The SMP runtime verifies that the *resource\_id* exists. If the *resource\_id* does not exist, a warning message is issued and the thread is left unbound.
- If you change the number of threads inside the program, for example, through **omp\_set\_num\_threads()** or **num\_threads** clause, the following situation occurs:
  - If the number of threads in the application is increased, rebinding takes place based on the environment variable settings.
  - If the number of threads is reduced after binding, the original binding remains.

**bindlist=SDL=i1,i2,...ix**

Specifies different system detail levels to bind threads by using the Resource Set API. This suboption supports binding a thread to multiple logical processors.

*SDL* stands for System Detail Level and can be **MCM**, **L2CACHE**, **PROC\_CORE**, or **PROC**. If the *SDL* value is not specified, or an incorrect *SDL* value is specified, the SMP runtime issues an error message.

The list of  $x$  integers  $i1,i2...ix$  enumerates the resources (one of *SDLs*) to be used during binding. When the number of integers in the list is greater than or equal to the number of threads, the position in the list determines the thread ID that will be bound to the resource.

When the number of resources specified in **bindlist** is greater than the number of threads, the extra resources are ignored.

When the number of threads  $t$  is greater than the number of resources  $x$ ,  $t$  threads will be divided among  $x$  resources according to the following formula:

The  $\text{ceil}(t/x)$  threads are bound to the first  $(t \bmod x)$  resources. The  $\text{floor}(t/x)$  threads will be bound to the remaining resources.

For example:

```
XLSMPOPTS="bindlist=MCM=0,1,2,3"
```

This example code shows that threads are bound to MCM 0,1,2,3. When the program runs with four threads, thread 0 is bound to MCM 0, thread 1 is bound to MCM 1, thread 2 is bound to MCM 2, and thread 3 is bound to MCM 3. When the program runs with six threads, threads 0 and 1 are

bound to MCM 0, threads 2 and 3 are bound to MCM 1, thread 4 is bound to MCM 2, and thread 5 is bound to MCM 3.

With the `XLSMPOPTS` environment variable being set as in the following example, a program runs with eight (or fewer) threads. It binds all even-numbered threads to L2CACHE 0 and all odd-numbered threads to L2CACHE 1.

```
XLSMPOPTS="bindlist=L2CACHE=0,1,0,1,0,1,0,1"
```

#### Notes:

- The `bindlist` suboption takes precedence over the `startproc/stride`, `procs`, and `bind` suboptions.
- Resource Set can only be used by a user account with the `CAP_NUMA_ATTACH` and `CAP_PROPAGATE` capabilities. These capabilities are set on a per-user basis by using the `chuser` command as follows:  

```
chuser "capabilities=CAP_PROPAGATE,CAP_NUMA_ATTACH" username
```
- The SMP runtime verifies that the thread ID specified for a resource is not less than 0 and greater than the available resources. Otherwise, the SMP runtime issues a warning message and the thread is left unbound.
- If you change the number of threads inside the program, for example, through `omp_set_num_threads()` or `num_threads` clause, the following situation occurs:
  - If the number of threads in the application is increased, rebinding takes place based on the environment variable settings.
  - If the number of threads is reduced after binding, the original binding remains.

Performance tuning options are as follows:

#### **spins=num**

Specifies the number of loop spins, or iterations, before a yield occurs.

When a thread completes its work, the thread continues executing in a tight loop looking for new work. One complete scan of the work queue is done during each busy-wait state. An extended busy-wait state can make a particular application highly responsive, but can also harm the overall responsiveness of the system unless the thread is given instructions to periodically scan for and yield to requests from other applications.

A complete busy-wait state for benchmarking purposes can be forced by setting both `spins` and `yields` to 0.

The default value for `num` is 100.

#### **yields=num**

Specifies the number of yields before a sleep occurs.

When a thread sleeps, it completely suspends execution until another thread signals that there is work to do. This provides better system utilization, but also adds extra system overhead for the application.

The default value for `num` is 100.

#### **delays=num**

Specifies a period of do-nothing delay time between each scan of the work queue. Each unit of delay is achieved by running a single no-memory-access delay loop.

The default value for *num* is 500.

Dynamic profiling options are as follows:

**profilefreq=*n***

Specifies the frequency with which a loop should be revisited by the dynamic profiler to determine its appropriateness for parallel or serial execution. The runtime library uses dynamic profiling to dynamically tune the performance of automatically parallelized loops. Dynamic profiling gathers information about loop running times to determine if the loop should be run sequentially or in parallel the next time through. Threshold running times are set by the **parthreshold** and **seqthreshold** dynamic profiling options, described below.

The allowed values for this option are the numbers from 0 to 32. If *num* is 0, all profiling is turned off, and overheads that occur because of profiling will not occur. If *num* is greater than 0, running time of the loop is monitored once every *num* times through the loop. The default for *num* is 16. Values of *num* exceeding 32 are changed to 32.

It is important to note that dynamic profiling is not applicable to user-specified parallel loops.

**parthreshold=*num***

Specifies the time, in milliseconds, below which each loop must execute serially. If you set *num* to 0, every loop that has been parallelized by the compiler will execute in parallel. The default setting is 0.2 milliseconds, meaning that if a loop requires fewer than 0.2 milliseconds to execute in parallel, it should be serialized.

Typically, *num* is set to be equal to the parallelization overhead. If the computation in a parallelized loop is very small and the time taken to execute these loops is spent primarily in the setting up of parallelization, these loops should be executed sequentially for better performance.

**seqthreshold=*num***

Specifies the time, in milliseconds, beyond which a loop that was previously serialized by the dynamic profiler should revert to being a parallel loop. The default setting is 5 milliseconds, meaning that if a loop requires more than 5 milliseconds to execute serially, it should be parallelized.

**seqthreshold** acts as the reverse of **parthreshold**.

## Environment variables for OpenMP

OpenMP runtime options affecting parallel processing are set by specifying OMP environment variables. These environment variables use syntax of the form:

►►—*env\_variable*—=—*option\_and\_args*—◀◀

If an OMP environment variable is not explicitly set, its default setting is used.

For information about the OpenMP specification, see: <http://www.openmp.org>.

**OMP\_DYNAMIC:** The OMP\_DYNAMIC environment variable enables or disables dynamic adjustment of the number of threads available for running parallel regions.

If it is set to TRUE, the number of threads available for executing parallel regions can be adjusted at run time to make the best use of system resources. For more information, see the description for **profilefreq=*num*** in “XLSMPOPTS” on page 28.

If it is set to FALSE, dynamic adjustment is disabled.

The default setting is TRUE.

#### **OMP\_MAX\_ACTIVE\_LEVELS:**

Use OMP\_MAX\_ACTIVE\_LEVELS to set the *max-active-levels-var* internal control variable. This controls the maximum number of active nested parallel regions. In programs where nested parallelism is disabled, the initial value should be 1. In programs where nested parallelism is enabled, the initial value should be greater than 1. The function **omp\_get\_max\_active\_levels** can be used to retrieve this value at run time. The value for OMP\_MAX\_ACTIVE\_LEVELS is a positive integer. If a positive integer is not specified, the default value for *max-active-levels-var* is set by the runtime.

▶▶—OMP\_MAX\_ACTIVE\_LEVELS=*n*—————▶▶

**OMP\_NESTED:** The OMP\_NESTED=TRUE|FALSE environment variable enables or disables nested parallelism. Its setting can be overridden by calling the **omp\_set\_nested** runtime library function.

If nested parallelism is disabled, nested parallel regions are serialized and run in the current thread.

In the current implementation, nested parallel regions are always serialized. As a result, OMP\_SET\_NESTED does not have any effect, and **omp\_get\_nested** always returns 0. If **-qsmp=nested\_par** option is on (only in non-strict OMP mode), nested parallel regions might employ additional threads as available. However, no new team is created to run nested parallel regions.

The default value for OMP\_NESTED is FALSE.

**OMP\_NUM\_THREADS:** The OMP\_NUM\_THREADS environment variable specifies the number of threads to use for parallel regions.

The syntax of the environment variable is as follows:

▶▶—OMP\_NUM\_THREADS=*num\_list*—————▶▶

*num\_list*

A list of one or more positive integer values separated by commas.

If you do not set OMP\_NUM\_THREADS, the number of processors available is the default value to form a new team for the first encountered parallel construct. By default, any nested parallel constructs are run by one thread.

If *num\_list* contains a single value, dynamic adjustment of the number of threads is enabled (OMP\_DYNAMIC is set to true), and a parallel construct without a **num\_threads** clause is encountered, the value is the maximum number of threads that can be used to form a new team for the encountered parallel construct.

If *num\_list* contains a single value, dynamic adjustment of the number of threads is not enabled (OMP\_DYNAMIC is set to false), and a parallel construct without a **num\_threads** clause is encountered, the value is the exact number of threads that can be used to form a new team for the encountered parallel construct.

If *num\_list* contains multiple values, dynamic adjustment of the number of threads is enabled (**OMP\_DYNAMIC** is set to true), a parallel construct without a **num\_threads** clause is encountered, the first value is the maximum number of threads that can be used to form a new team for the encountered parallel construct. After the encountered construct is entered, the first value is removed and the remaining values form a new *num\_list*. The new *num\_list* is in turn used in the same way for any closely nested parallel constructs inside the encountered parallel construct.

If *num\_list* contains multiple values, dynamic adjustment of the number of threads is not enabled (**OMP\_DYNAMIC** is set to false), and a parallel construct without a **num\_threads** clause is encountered, the first value is the exact number of threads that can be used to form a new team for the encountered parallel construct. After the encountered construct is entered, the first value is removed and the remaining values form a new *num\_list*. The new *num\_list* is in turn used in the same way for any closely nested parallel constructs inside the encountered parallel construct.

**Note:** If the number of parallel regions is equal to or greater than the number of values in *num\_list*, the **omp\_get\_max\_threads** function returns the last value of *num\_list* in the parallel region.

If the number of threads requested exceeds the system resources available, the program stops.

The **omp\_set\_num\_threads** function sets the first value of *num\_list*. The **omp\_get\_max\_threads** function returns the first value of *num\_list*.

If you specify the number of threads for a given parallel region more than once with different settings, the compiler uses the following precedence order to determine which setting takes effect:

1. The number of threads set using the **num\_threads** clause takes precedence over that set using the **omp\_set\_num\_threads** function.
2. The number of threads set using the **omp\_set\_num\_threads** function takes precedence over that set using the **OMP\_NUM\_THREADS** environment variable.
3. The number of threads set using the **OMP\_NUM\_THREADS** environment variable takes precedence over that set using the **PARTHDS** suboption of the **XLSPMPOPTS** environment variable.

### Example

```
export OMP_NUM_THREADS=3,4,5
export OMP_DYNAMIC=false

// omp_get_max_threads() returns 3

#pragma omp parallel
{
// Three threads running the parallel region
// omp_get_max_threads() returns 4

    #pragma omp parallel if(0)
    {
// One thread running the parallel region
// omp_get_max_threads() returns 5

        #pragma omp parallel
        {
// Five threads running the parallel region
```

```

    // omp_get_max_threads() returns 5
    }
}
}

```

**OMP\_PROC\_BIND:** The OMP\_PROC\_BIND environment variable controls whether OpenMP threads can be moved between processors. The syntax is as follows:

► OMP\_PROC\_BIND= [TRUE] | [FALSE] ◄

By default, the OMP\_PROC\_BIND environment variable is not set. If you set OMP\_PROC\_BIND to TRUE, the threads are bound to processors. If you set OMP\_PROC\_BIND to FALSE, the threads can be moved between processors.

If you do not set OMP\_PROC\_BIND, but set the suboptions of XLSMPOPTS (**startproc/stride**, **procs**, **bind**, or **bindlist**), the threads are bound to processors according to the settings in the XLSMPOPTS environment variable.

If you set neither OMP\_PROC\_BIND nor the suboptions of XLSMPOPTS (**startproc/stride**, **procs**, **bind**, or **bindlist**), the threads are not bound to processors.

If you do not set OMP\_PROC\_BIND and the XLSMPOPTS setting (**startproc/stride**, **procs**, **bind**, or **bindlist**) is invalid, the threads are not bound to processors.

If you set OMP\_PROC\_BIND to TRUE and also set the suboptions of XLSMPOPTS (**startproc/stride**, **procs**, **bind**, or **bindlist**), the threads are bound to processors according to the settings in the XLSMPOPTS environment variable.

**Notes:**

- If **procs** is set and the number of CPU IDs specified is smaller than the number of threads used by the program, the remaining threads are not bound to processors.
- If XLSMPOPTS=**startproc** is used, the value specified by **startproc** is smaller than the number of CPUs, and the value specified by **stride** causes a thread to bind to a CPU outside the range of available processors, some of the threads are bound and some are not.

If you set OMP\_PROC\_BIND to TRUE, but do not set the XLSMPOPTS suboption (**startproc/stride**, **procs**, **bind**, or **bindlist**), the threads are bound to processors.

If you set OMP\_PROC\_BIND to TRUE and the XLSMPOPTS setting (**startproc/stride**, **procs**, **bind**, or **bindlist**) is invalid, the threads are bound to processors.

If you set OMP\_PROC\_BIND to FALSE and also set the suboptions of XLSMPOPTS (**startproc/stride**, **procs**, **bind**, or **bindlist**), the threads are not bound to processors.

If you set OMP\_PROC\_BIND to FALSE, but do not set the suboptions of XLSMPOPTS (**startproc/stride**, **procs**, **bind**, or **bindlist**), the threads are not bound to processors.

If you set `OMP_PROC_BIND` to `FALSE` and the `XLSMPOPTS` setting (`startproc/stride`, `procs`, `bind`, or `bindlist`) is invalid, the threads are not bound to processors.

The following table summarizes the previous thread binding rules:

*Table 8. Thread binding rule summary*

<b>OMP_PROC_BIND settings</b>	<b>XLSMPOPTS settings</b>	<b>Thread binding results</b>
OMP_PROC_BIND is not set	XLSMPOPTS is not set	Threads are not bound
OMP_PROC_BIND is not set	XLSMPOPTS is set ( <b>startproc/stride</b> , <b>procs</b> , <b>bind</b> , or <b>bindlist</b> )	Threads are bound according to the settings in XLSMPOPTS
OMP_PROC_BIND is not set	XLSMPOPTS setting is invalid	Threads are not bound
OMP_PROC_BIND=TRUE	XLSMPOPTS is not set	Threads are bound
OMP_PROC_BIND=TRUE	XLSMPOPTS is set ( <b>startproc/stride</b> , <b>procs</b> , <b>bind</b> , or <b>bindlist</b> )	Threads are bound according to the settings in XLSMPOPTS <b>Notes:</b> <ul style="list-style-type: none"> <li>• If <b>procs</b> is set and the number of CPU IDs specified is smaller than the number of threads used by the program, the remaining threads are not bound.</li> <li>• If XLSMPOPTS=<b>startproc</b> is used, the value specified by <b>startproc</b> is smaller than the number of CPUs, and the value specified by <b>stride</b> causes a thread to bind to a CPU outside the range of available processors, some of the threads are bound and some are not.</li> </ul>
OMP_PROC_BIND=TRUE	XLSMPOPTS setting is invalid	Threads are bound
OMP_PROC_BIND=FALSE	XLSMPOPTS is not set	Threads are not bound
OMP_PROC_BIND=FALSE	XLSMPOPTS is set ( <b>startproc/stride</b> , <b>procs</b> , <b>bind</b> , or <b>bindlist</b> )	Threads are not bound
OMP_PROC_BIND=FALSE	XLSMPOPTS setting is invalid	Threads are not bound

**Note:** The `OMP_PROC_BIND` environment variable provides a portable way to control whether OpenMP threads can be migrated. The **startproc/stride**, **procs**, **bind**, or **bindlist** suboption of the `XLSMPOPTS` environment variable, which is an IBM extension, provides a finer control to bind OpenMP threads to processors. If

portability of your application is important, use only the `OMP_PROC_BIND` environment variable to control thread binding.

**OMP\_SCHEDULE:** The `OMP_SCHEDULE` environment variable specifies the scheduling algorithm used for loops not explicitly assigned a scheduling algorithm with the `omp schedule` clause.

For example:

```
OMP_SCHEDULE="guided, 4"
```

Valid options for *algorithm* are:

- auto
- dynamic[, *n*]
- guided[, *n*]
- runtime
- static[, *n*]

If specifying a chunk size with *n*, the value of *n* must be a positive integer.

The default scheduling algorithm is **auto**.

**Related reference:**

“`omp_set_schedule`” on page 612

“`omp_get_schedule`” on page 611

**OMP\_STACKSIZE:**

The `OMP_STACKSIZE` environment variable indicates the stack size of threads created by the OpenMP run time. `OMP_STACKSIZE` sets the value of the *stacksize-var* internal control variable. `OMP_STACKSIZE` does not control the stack size of the master thread. The syntax is as follows:

```
▶▶—OMP_STACKSIZE=—size—————▶▶
```

By default, the size value is represented in Kilobytes. You can also use the suffixes B, K, M, or G if you want to indicate the size in Bytes, Kilobytes, Megabytes, or Gigabytes respectively. White space is allowed between and around the size value and the suffix. For example, these two examples both indicate a stack size of 10 Megabytes.

```
setenv OMP_STACKSIZE 10M
setenv OMP_STACKSIZE " 10 M "
```

If `OMP_STACKSIZE` is not set, the initial value of the *stacksize-var* internal control variable is set to the default value. The default value for 32-bit mode is 256M. For 64-bit mode, the default is up to the limit imposed by system resources. If the compiler cannot use the stack size specified or if `OMP_STACKSIZE` does not conform to the correct format, the compiler sets the environment variable to the default value. If the `STACK` suboption of the `XLSMPOPTS` environment variable and the `OMP_STACKSIZE` environment are specified, the `OMP_STACKSIZE` environment variable takes precedence.

**OMP\_THREAD\_LIMIT:**



The `OMP_THREAD_LIMIT` environment variable sets the number of OpenMP threads to use for the whole program. The syntax is as follows:

►► `OMP_THREAD_LIMIT=n` ◀◀

*n* The number of OpenMP threads to use for the whole program. It must be a positive scalar integer.

The value for `OMP_THREAD_LIMIT` is a positive integer.

If the `OMP_THREAD_LIMIT` environment variable is not set and the `OMP_NUM_THREADS` environment variable is set to a single value, the default value for `OMP_THREAD_LIMIT` is the value of `OMP_NUM_THREADS` or the number of available processors, whichever is greater.

If the `OMP_THREAD_LIMIT` environment variable is not set and the `OMP_NUM_THREADS` environment variable is set to a list, the default value for `OMP_THREAD_LIMIT` is the multiplication of all the numbers in the list or the number of available processors, whichever is greater.

If both the `OMP_THREAD_LIMIT` and `OMP_NUM_THREADS` environment variables are not set, the default value for `OMP_THREAD_LIMIT` is the number of available processors.

#### **OMP\_WAIT\_POLICY:**

The `OMP_WAIT_POLICY` environment variable gives hints to the compiler about the preferred behavior of waiting threads during program run time. The `OMP_WAIT_POLICY` environment variable sets the *wait-policy-var* internal control variable value.

The syntax is as follows:

►► `OMP_WAIT_POLICY=PASSIVE  
ACTIVE` ◀◀

The default value for `OMP_WAIT_POLICY` is `PASSIVE`.

Use `ACTIVE` if you want waiting threads to be mostly active. With `ACTIVE`, the thread consumes processor cycles while waiting, if possible.

Use `PASSIVE` if you want waiting threads to be mostly passive. That is, the preference is for the thread to not consume processor cycles while waiting. For example, you prefer waiting threads to sleep or to yield the processor to other threads.

**Note:** If the `OMP_WAIT_POLICY` environment variable is set and the `SPINS`, `YIELDS`, or `DELAYS` suboptions of the `XLSMPOPTS` environment variable are specified, `OMP_WAIT_POLICY` takes precedence.

---

## Using custom compiler configuration files

The XL C/C++ compiler generates a default configuration file `/etc/vac.cfg.nm` at installation time, where *nm* indicates which OS version the configuration file is for. 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 `xlc` compiler invocation command. This is to avoid forcing your users to specify this option on the command line for every compilation, because `-qlist` is automatically in effect every time the compiler is called with the `xlc` 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. Also, if you upgrade the operating system, you must change the symbolic links to the configuration file in `/etc/` to point to the correct version of the configuration file.
- 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 `XLC_USR_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 150
- “Compile-time and link-time environment variables” on page 26

## Creating custom configuration files

If you use the `XLC_USR_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 =DEFAULT
   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>
DEFAULT:
   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 `XLC_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 `xlC` 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.

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

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

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

DEFLT:
      options=<C>
```

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

Figure 3. Default configuration file vac.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 9. Assignment operators and attribute ordering

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

For example, assume that the XLC\_USR\_CONFIG environment variable is set to point to the custom user-defined configuration file at ~/userconfig2.

### Custom user-defined configuration file

~/userconfig2

```
xlc_prepend: use=xlc
              options--<B1>
xlc_replace: use=xlc
              options:=<B2>
xlc_append:  use=xlc
              options+=<B3>
```

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

### Default configuration file vac.cfg

```
xlc: 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 xlc uses B, D, and C
2. stanza xlc\_prepend uses B1, B, D, and C
3. stanza xlc\_replace uses B2
4. stanza xlc\_append uses B, D, C, and B3

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

```

xlc:
  use=xlc
  options--Isome_include_path
  options+=some options

```

Figure 4. Using additional assignment operations

## Examples of stanzas in custom configuration files

<pre> DEFLT: use=DEFLT        options = -g </pre>	<p>This example specifies that the <b>-g</b> option is to be used in all compilations.</p>
<pre> xlc: use=xlc options+=-qlist xlc_r: use=xlc_r options+=-qlist </pre>	<p>This example specifies that <b>-qlist</b> is to be used for any compilation called by the <b>xlc</b> and <b>xlc_r</b> commands. This <b>-qlist</b> specification overrides the default setting of <b>-qlist</b> specified in the system configuration file.</p>
<pre> DEFLT: use=DEFLT        libraries=-L/home/user/lib,-lmylib </pre>	<p>This example specifies that all compilations should link with <code>/home/user/lib/libmylib.a</code>.</p>

## Configuring the gxc or gxc++ option mapping

The **gxc** and **gxc++** utilities use the configuration file `/etc/gxc.cfg` to translate GNU C and C++ options to XL C/C++ options. Each entry in `gxc.cfg` describes how the utility should map a GNU C or C++ option to an XL C/C++ option and how to process it.

An entry consists of a string of flags for the processing instructions, a string for the GNU C/C++ option, and a string for the XL C/C++ option. The three fields must be separated by white space. If an entry contains only the first two fields and the XL C/C++ option string is omitted, the GNU C option in the second field will be recognized by **gxc** or **gxc++** and silently ignored.

The **#** character is used to insert comments in the configuration file. A comment can be placed on its own line, or at the end of an entry.

The following syntax is used for an entry in `gxc.cfg`:

```
abcd  "gcc_or_g++_option"  "xlc_or_xlc++_option"
```

where:

- a* Lets you disable the option by adding **no-** as a prefix. The value is either **y** for yes, or **n** for no. For example, if the flag is set to **y**, then **finline** can be disabled as **fno-inline**, and the entry is:

```
ynn*      "-finline"          "-qinline"
```

If given **-fno-inline**, then the utility will translate it to **-qnoinline**.

- b* Informs the utility that the XL C/C++ option has an associated value. The value is either **y** for yes, or **n** for no. For example, if option **-fmyvalue=n** maps to **-qmyvalue=n**, then the flag is set to **y**, and the entry is:

```
nyn*      "-fmyvalue"        "-qmyvalue"
```

The utility will then expect a value for these options.

- c Controls the processing of the options. The value can be any of the following:
  - n Tells the utility to process the option listed in the *gcc\_or\_g++\_option* field
  - i Tells the utility to ignore the option listed in the *gcc\_or\_g++\_option* field. The utility will generate a message that this has been done, and continue processing the given options.
  - e Tells the utility to halt processing if the option listed in the *gcc\_or\_g++\_option* field is encountered. The utility will also generate an error message.

For example, the GCC option `-I-` is not supported and must be ignored by `gxlc` or `gxlcpp`. In this case, the flag is set to `i`, and the entry is:

```
nni*      "-I-"
```

If the utility encounters this option as input, it will not process it and will generate a warning.

- d Lets `gxlc` or `gxlcpp` include or ignore an option based on the type of compiler. The value can be any of the following:
  - c Tells the utility to translate the option only for C.
  - x Tells the utility to translate the option only for C++.
  - \* Tells `gxlc` or `gxlcpp` to translate the option for C and C++.

For example, `-fwritable-strings` is supported by both compilers, and maps to `-qnor0`. The entry is:

```
nnn*      "-fwritable-strings"      "-qnor0"
```

*"gcc\_or\_g++\_option"*

Is a string representing a GNU C/C++ option. This field is required and must appear in double quotation marks.

*"xlc\_or\_xlcpp\_option"*

Is a string representing an XL C/C++ option. This field is optional, and, if present, must appear in double quotation marks. If left blank, the utility ignores the *gcc\_or\_g++\_option* in that entry.

It is possible to create an entry that will map a range of options. This is accomplished by using the asterisk (\*) as a wildcard. For example, the GCC `-D` option requires a user-defined name and can take an optional value. It is possible to have the following series of options:

```
-DCOUNT1=100
-DCOUNT2=200
-DCOUNT3=300
-DCOUNT4=400
```

Instead of creating an entry for each version of this option, the single entry is:

```
nnn*      "-D*"      "-D*"
```

where the asterisk will be replaced by any string following the `-D` option.

Conversely, you can use the asterisk to exclude a range of options. For example, if you want `gxlc` or `gxlcpp` to ignore all the `-std` options, then the entry would be:

```
nni*      "-std*"
```

When the asterisk is used in an option definition, option flags *a* and *b* are not applicable to these entries.

The character % is used with a GNU C/C++ option to signify that the option has associated parameters. This is used to insure that `gxl` or `gxl++` will ignore the parameters associated with an option that is ignored. For example, the `-isystem` option is not supported and uses a parameter. Both must be ignored by the application. In this case, the entry is:

```
nmi*      "-isystem %"
```

For a complete list of GNU C and C++ and XL C/C++ option mappings, see:

<http://www.ibm.com/support/docview.wss?uid=swg27022984>

#### **Related information**

- The GNU Compiler Collection online documentation at <http://gcc.gnu.org/onlinedocs/>





---

## Chapter 3. Tracking and reporting compiler usage

You can use the utilization tracking and reporting feature to record and analyze which users in your organization are using the compiler and the number of users using it concurrently. This information can help you determine whether your organization's use of the compiler exceeds your compiler license entitlements.

To use this feature, follow these steps:

1. Understand how the feature works. See “Understanding utilization tracking and reporting” for more information.
2. Investigate how the compiler is used in your organization, and decide how you track the compiler usage accordingly. See “Preparing to use this feature” on page 56 for more information.
3. Configure and enable utilization tracking. See “Configuring utilization tracking” on page 61 for more information.
4. Use the utilization reporting tool to generate usage reports or prune usage files. See “Generating usage reports” on page 69 or “Pruning usage files” on page 72 for more information.

---

### Understanding utilization tracking and reporting

The utilization tracking and reporting feature provides a mechanism for you to detect whether your organization's use of the compiler exceeds your compiler license entitlements. This section introduces the feature, describes how it works, and illustrates its typical usage scenarios.

#### Overview

When utilization tracking is enabled, all compiler invocations are recorded in a file. This file is called a usage file and it has the .cuf extension. You can then use the utilization reporting tool to generate a report from one or more of these usage files, and optionally prune the usage files.

You can use the utilization tracking and reporting feature in various ways based on how the compiler is used in your organization. “Four usage scenarios” on page 48 illustrates the typical usage scenarios of this feature.

The following sections introduce the configuration of the utilization tracking functionality and the usage of the utilization reporting tool.

#### Utilization tracking

A utilization tracking configuration file `urtxlc_cpp1201aix.cfg` is included in the default compiler installation. You can use this file to enable utilization tracking and control different aspects of the tracking.

A symlink `urt_client.cfg` is also included in the default compiler installation. It points to the location of the utilization tracking configuration file. If you want to put the utilization tracking configuration file in a different location, you can modify the symlink accordingly.

For more information, see “Configuring utilization tracking” on page 61.

**Note:** Utilization tracking is disabled by default.

## Utilization reporting tool

The utilization reporting tool generates compiler usage reports based on the information in the usage files. You can optionally prune the usage files with the tool. For more information, see “Generating usage reports” on page 69 and “Pruning usage files” on page 72.

## Four usage scenarios

This section describes four possible scenarios for managing the compiler usage, for recording the compiler usage information and for generating reports from this information.

The following scenarios describe some typical ways that your organization might be using the compiler and illustrates how you can use this feature to track compiler usage in each case.

**Note:** Actual usage is not limited to these scenarios.

“Scenario: One machine, one shared .cuf file”

“Scenario: One machine, multiple .cuf files” on page 50

“Scenario: Multiple machines, one shared .cuf file” on page 52

“Scenario: Multiple machines, multiple .cuf files” on page 54

### Scenario: One machine, one shared .cuf file

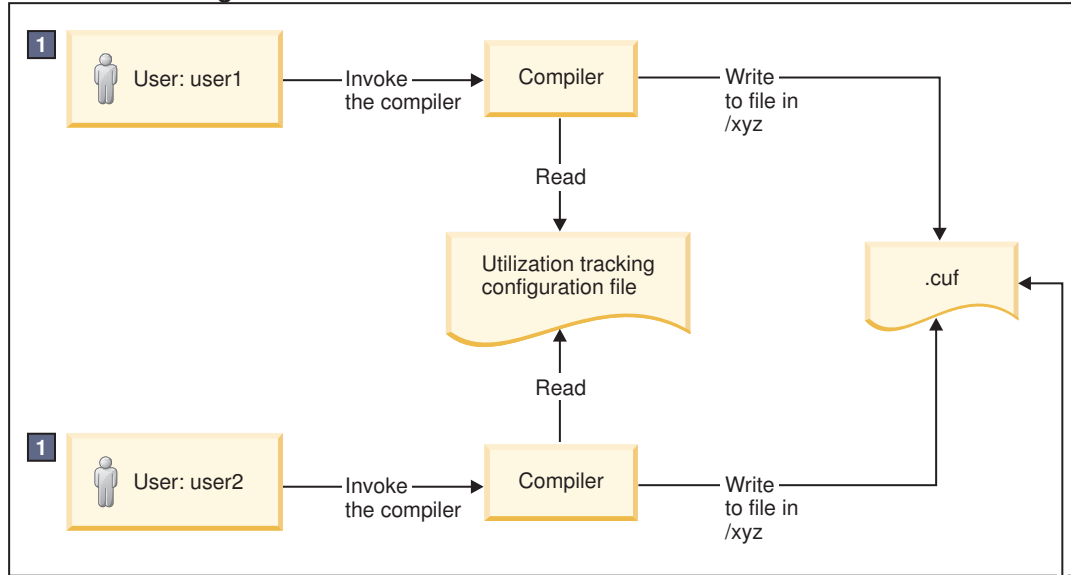
This scenario describes an environment where all the compilations are done on one machine and all users share one .cuf file.

The advantage of using the approach in this scenario is that it simplifies report generation and usage file pruning, because the utilization report tool only need to access one .cuf file. The disadvantage is that all compiler users need to compete for access to this file. Because the file might become large, it might have an impact on performance. Some setup work is also required to create the shared .cuf file and to give all compiler users write access. The “The number of usage files” on page 59 section provides detailed information about using a single usage file for all compiler users.

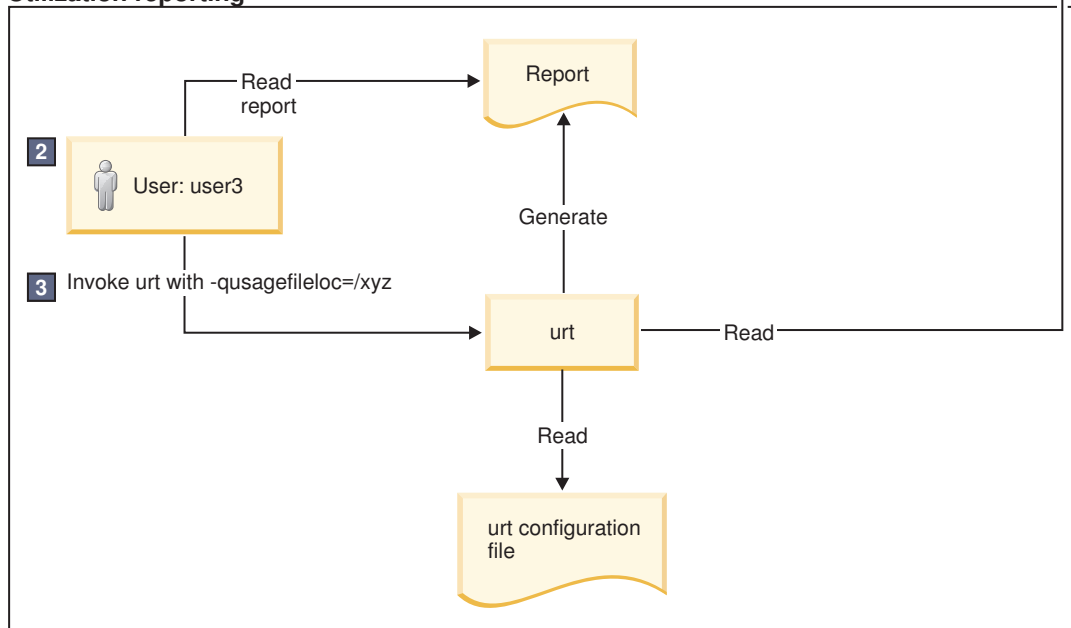
In this scenario, two compiler users run the compiler on the same machine and their utilization information is recorded in a shared .cuf file. The utilization tracking configuration file for the compiler is modified to point to the location of the .cuf file. When the compiler is invoked, it writes the utilization information to that file. You can then use the utilization reporting tool to retrieve the utilization information from the file and generate usage reports.

The following diagram illustrates this scenario.

## Utilization tracking



## Utilization reporting



1. Both user1 and user2 need write access to the .cuf file in /xyz.
2. user3 needs read access to the .cuf file in/xyz to generate the usage report, and write access to prune the .cuf file.
3. A cron job can be created to run **urt** automatically on a regular basis.

Figure 5. Compiler users use a single machine, with a shared .cuf file

The diagram reflects the following points:

1. user1 and user2 use the same utilization tracking configuration file, which manages the tracking functionality centrally. A common location /xyz is created to keep a shared .cuf file.
2. When user1 and user2 invoke the compiler, the utilization information is recorded in the .cuf file under the common directory /xyz.

3. user3 invokes `urt` with `-qusagefileloc=/xyz` to generate usage reports.

**Note:** Regular running of the utilization reporting tool can prevent these files from growing too big, because you can prune the usage files with this tool.

### **Scenario: One machine, multiple .cuf files**

This scenario describes an environment where all the compilations are done on one machine and all users have their own `.cuf` files.

The approach in this scenario has the following advantages:

- Compiler users do not have to compete for access to a single `.cuf` file, and this might result in better performance.
- You do not need to set up write access to a single common location for all compiler users. They already have write access to their own home directories.

However, using multiple `.cuf` files that are automatically created in each user's home directory might have the following issues:

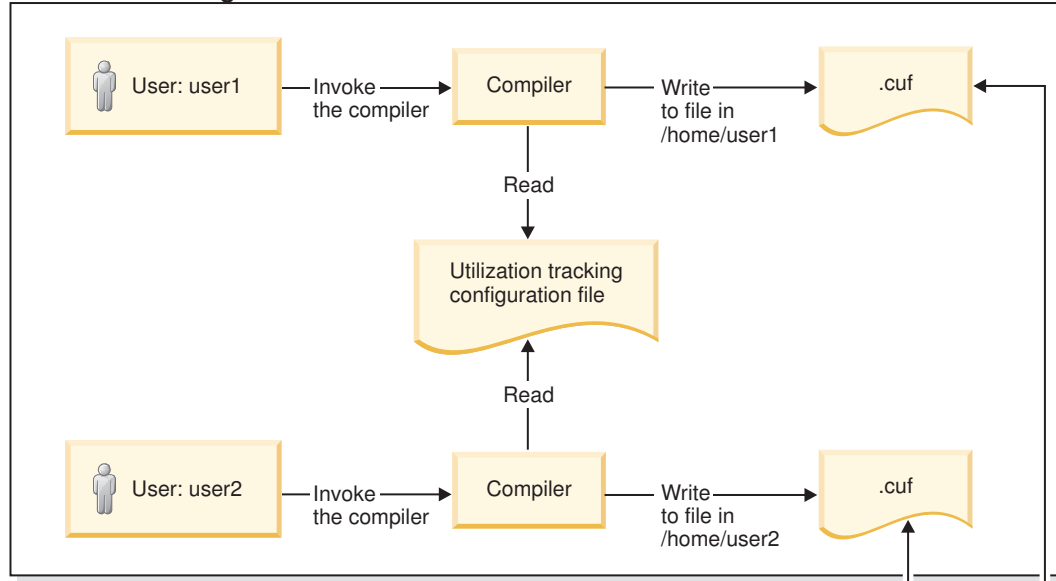
- Compiler users might not know that the file has been created or what it is when they see the file. In this case, they might delete the file.
- Some users' home directories might be on file systems that are mounted from a remote system. This causes utilization tracking to use a remote file, which might affect performance.
- Compiler users might not want `.cuf` files to take up space in their `/home` directories.

Instead of using each user's home directory, the `.cuf` files for each user can be created in a common location. The “Usage file location” on page 58 section provides detailed information about how to create these files in a common location.

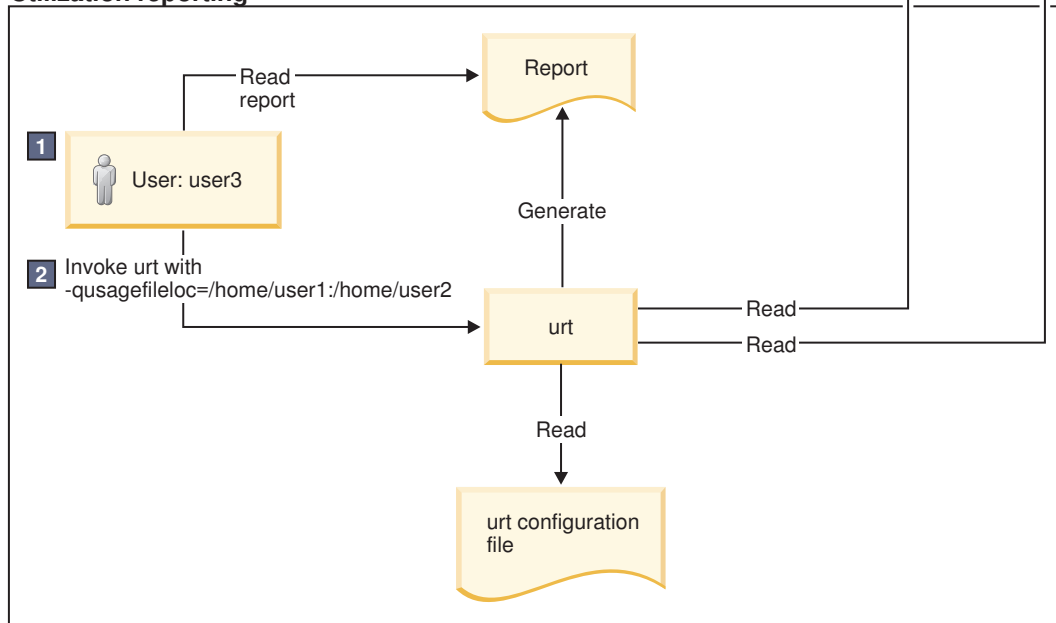
In this scenario, two compiler users run the compiler on the same machine and they have their own `.cuf` files. When the compiler is invoked, it automatically creates a `.cuf` file for each user and writes the utilization information to that file. You can then use the utilization reporting tool to retrieve the utilization information from the file and generate usage reports.

The following diagram illustrates this scenario.

## Utilization tracking



## Utilization reporting



1. user3 needs read access to .cuf files in /home/user1 and /home/user2 to generate the usage report, and write access to prune the usage files.
2. A cron job can be created to run **urt** automatically on a regular basis.

Figure 6. Compiler users use one machine, with separate .cuf files

This diagram reflects the following points:

1. user1 and user2 use the same utilization tracking configuration file, which manages the tracking functionality centrally.
2. When user1 and user2 invoke the compiler, the utilization information is recorded in the two .cuf files under their respective home directories, /home/user1 and /home/user2.
3. user3 invokes **urt** with `-qusagefileloc=/home/user1:/home/user2` to generate usage reports.

**Note:** If you need to find out which home directories contain usage files, you can invoke **urt** as follows:

```
urt -qusagefileloc=/home -qmaxsubdirs=1
```

In this case, **urt** looks for `.cuf` files in all users' home directories.

### **Scenario: Multiple machines, one shared .cuf file**

This scenario describes an environment where the compilations are done on multiple machines but all users share a single `.cuf` file.

The advantage of the approach in this scenario is that using one `.cuf` file can simplify the report generation and the usage file pruning process. The section “The number of usage files” on page 59 provides detailed information about using a single usage file for all compiler users. The `.cuf` file is already on the machine where the utilization reporting tool is installed. You do not need to copy the file to that machine or install the tool on multiple machines to prune the `.cuf` files.

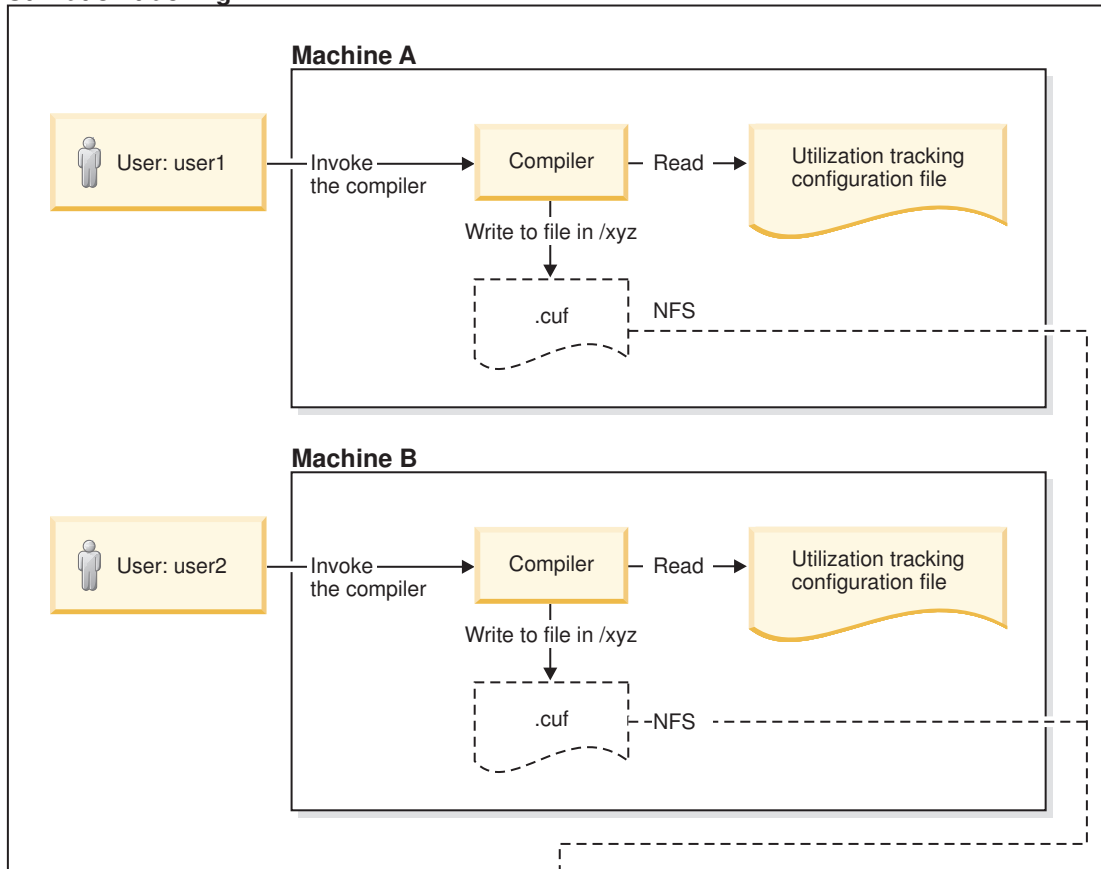
This approach has the following disadvantages:

- The compiler users must compete for access to one usage file. Because the file might become large, it might have an impact on performance.
- Some setup work is required to create the shared `.cuf` file and to give all compiler users write access on a network file system.
- The efficiency of the whole process depends on the speed and reliability of the network file system, because the compilers and the `.cuf` file are on different machines. For example, some file systems are better than others in supporting file locking, which is required for concurrent access by multiple users.

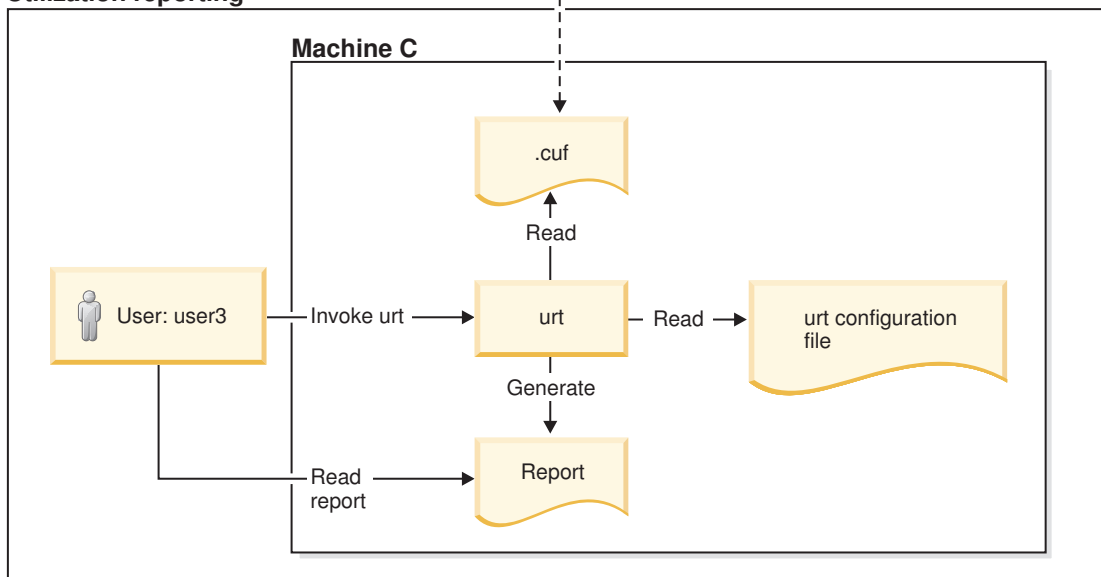
In this scenario, two compiler users run the compilers on separate machines and they use one shared `.cuf` file on a network file system, such as NFS, DFS, or AFS. When the compiler is invoked, it writes the utilization information to that file. You can then use the utilization reporting tool to retrieve the utilization information from the file and generate usage reports.

The following diagram illustrates this scenario.

## Utilization tracking



## Utilization reporting



1. On Machine A and Machine B, mount point /xyz is created to Machine C. All compiler utilization is recorded in the .cuf file, from which the usage report is generated.

Figure 7. Compiler users use multiple machines, with a shared .cuf file

This diagram reflects the following points:

1. Utilization tracking is configured respectively on Machine A and Machine B.

**Notes:**

- Although each machine has its own configuration file, the contents of these files must be the same.
  - Centrally managing the utilization tracking functionality can reduce your configuration effort and eliminate possible errors. The “Central configuration” on page 57 section provides detailed information about how you can use a common configuration file shared by compiler users using different machines.
2. A network file system is set up for the central management of the `.cuf` files. When user1 and user2 invoke the compilers from Machine A and Machine B, the utilization information of both compilers is written to the `.cuf` file on Machine C.
  3. user3 invokes `urt` to generate usage reports from the `.cuf` file on Machine C.

**Note:** You can use the utilization reporting tool to prune the usage files regularly to prevent them from growing too big.

**Scenario: Multiple machines, multiple `.cuf` files**

This scenario describes an environment where the compilations are done on multiple machines and all users have their own usage files.

In this scenario, two compiler users run the compilers on separate machines and they have their own `.cuf` files. When the compiler is invoked, it writes the utilization information to that file. You can then use the utilization reporting tool to retrieve the utilization information from the file and generate usage reports. This tool can be run on either of the machines on which the compiler is installed or on a different machine.

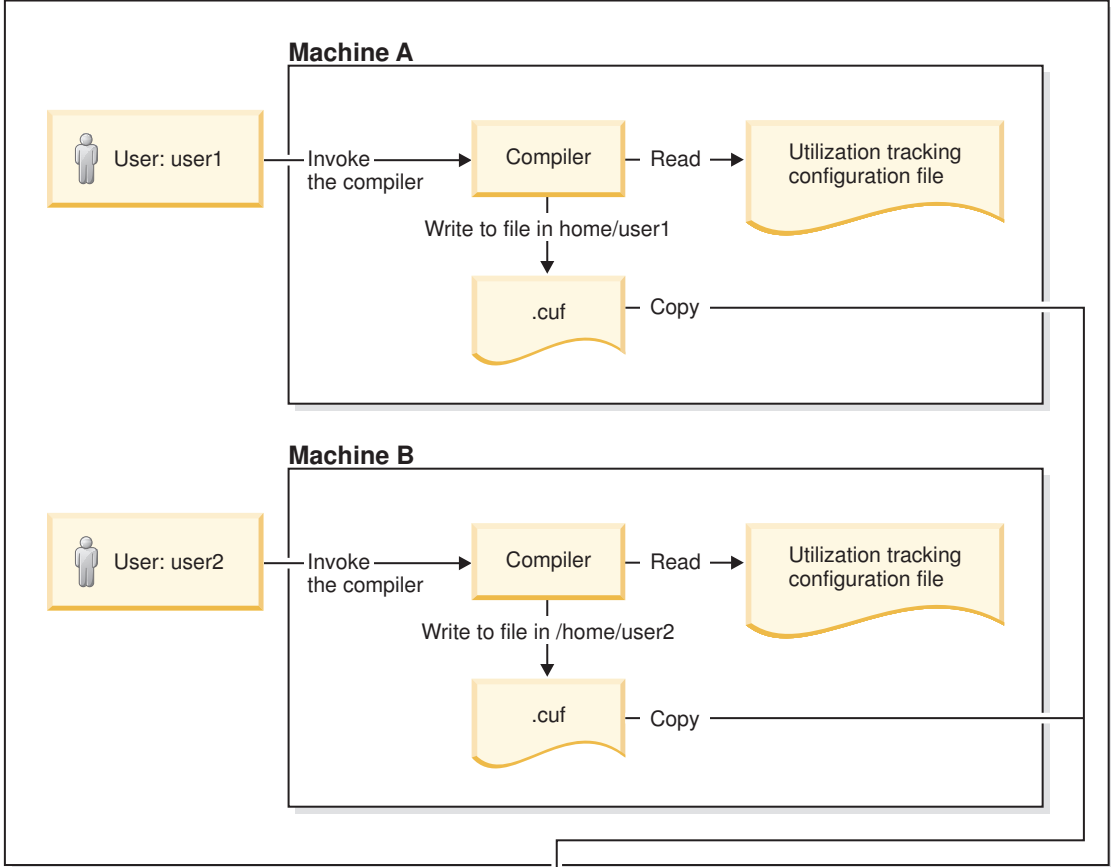
**Note:** The utilization reporting tool requires read access to all the `.cuf` files. You can use either of the following methods to make the files accessible in this example:

- Use a network file system, such as NFS, DFS, or AFS.
- Copy the files from their original locations to the machine where you plan to run the utilization reporting tool. You can use `ftp`, `rnp`, `rsync` or any other remote copy command to copy the files.

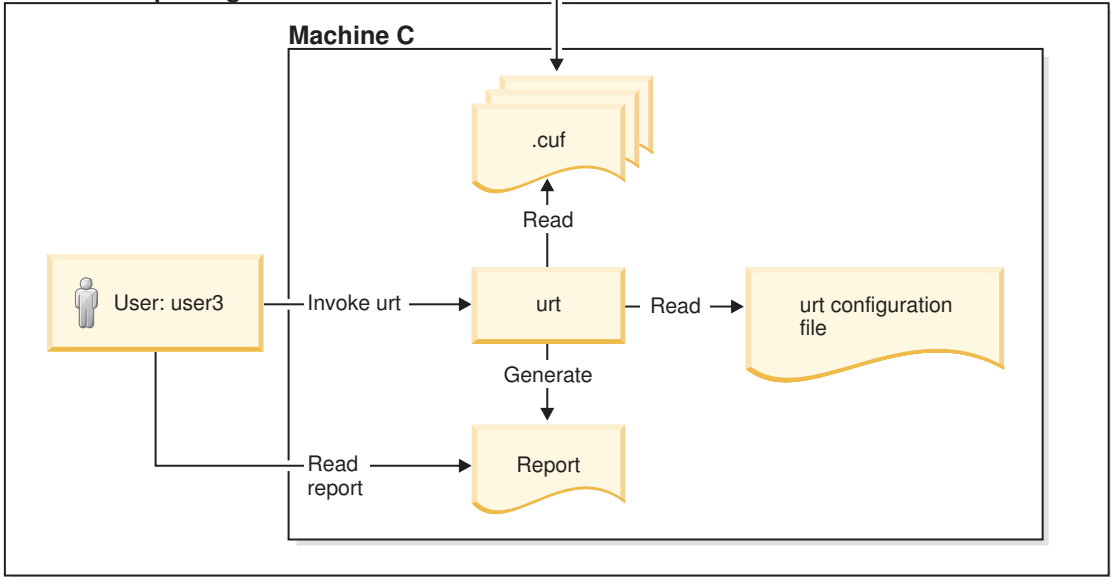
The following diagram illustrates this scenario.



**Utilization tracking**



**Utilization reporting**



1. user3 copies the .cuf files to Machine C. A cron job can be created to copy the files automatically on a regular basis.

Figure 8. Compiler users use multiple machines, with multiple .cuf files

This diagram reflects the following points:

1. Utilization tracking is configured respectively on Machine A and Machine B.

**Notes:**

- Although each machine has its own configuration file, the contents of these files must be the same.
  - Centrally managing the utilization tracking functionality can reduce your configuration effort and eliminate possible errors. The “Central configuration” on page 57 section provides detailed information about how you can use a common configuration file shared by compiler users using different machines.
2. When user1 and user2 invoke the compilers, the utilization information is recorded in the two .cuf files under their respective home directories, /home/user1 and /home/user2.

**Note:** These .cuf files can also be created in another common location, for example, /var/tmp. The “Usage file location” on page 58 section provides detailed information about how to create these files in a common location.

3. user3 copies the two .cuf files from Machine A and Machine B to Machine C.
4. user3 invokes **urt** to generate usage reports from the .cuf files on Machine C.

**Related information**

- “Preparing to use this feature”
- “Configuring utilization tracking” on page 61
- “Generating usage reports” on page 69
- “Pruning usage files” on page 72

---

## Preparing to use this feature

Before enabling utilization tracking within your organization, you must consider certain factors related to how the compiler is used in your organization.

The following sections describe those considerations in detail:

### Time synchronization

If you plan to track the utilization of the compiler on more than one machine, you must consider synchronizing the time across the machines.

The usage report generated by the utilization reporting tool lists the time when the compiler invocations start and end. The report also determines which invocations are concurrent. This information is much less reliable and useful if time is not synchronized across these machines.

If you are unable to synchronize time across different machines, you can use the **-qadjusttime** option to instruct the utilization reporting tool to adjust the times that have been recorded.

### License types and user information

Before you start to use this feature, you need the number and type of license entitlements for your organization.

The license and user information that you need is as follows:

- The number of Concurrent User licenses that you have for this compiler. This information is required for the **-qmaxconcurrentusers** entry in the utilization tracking configuration file.

- The users who have their own Authorized User license for this compiler. This information is used for the **-qexemptconcurrentusers** entry in the utilization tracking configuration file.
- The users who use the compiler with multiple accounts. This information is used for the **-qsameuser** option for the utilization reporting tool.

**Note:** It is not mandatory to specify the users who have their own Authorized User license and the users who use the compiler with multiple accounts, but it improves the accuracy of the usage reports generated by the utilization reporting tool. For detailed information, see “Concurrent user considerations.”

## Central configuration

Configuring utilization tracking the same for all compiler users is very important, because it can ensure the accuracy of your utilization tracking, and minimize your configuration and maintenance effort. You can achieve this by ensuring that all users use the same utilization tracking configuration file.

If you have only one installation of the compiler, you can directly edit the utilization tracking configuration file. Every compiler user can automatically use that configuration file.

If you have multiple installations of the compiler, you need to maintain a single utilization tracking config file and reference it from all installations. Any changes you make to the utilization tracking configuration file, including enabling or disabling utilization tracking, can automatically apply to all compiler installations when users invoke the compiler. In each installation, there is a symlink named `urt_client.cfg`, located in `usr/vacpp/urt`. Modify the symlink to point to this shared instance of the configuration file.

If the compiler is installed on multiple machines, the utilization tracking configuration file needs to be placed on a network file system, such as NFS, DFS, or AFS, to be used by the compiler on each machine.

**Note:** If it is not possible for you to use a single utilization tracking configuration file for all compiler users, you must ensure all utilization tracking configuration files for each compiler installation are consistent. Using different configurations for the same compiler is not supported.

## Concurrent user considerations

Invocations of the compiler are considered concurrent when their start time and end times overlap. This section provides the information about how the utilization reporting tool counts concurrent users and the ways to increase the accuracy of the usage reports.

When the utilization reporting tool counts concurrent users, it looks at the user account information that has been captured in the usage files. The account information consists of a user name, a user ID, and a host name. By default, each unique combination of this account information is considered and counted as a different user. However, invocations of the compiler by the following users must not be included in the count of concurrent users:

- Users who have their own Authorized User license are considered exempt users, because their use of the compiler does not consume any Concurrent User licences.

- Users who have multiple accounts. Because the accounts belong to the same user, invocations of the compiler while logged on using those accounts are counted as usage by a single user.

The utilization reporting tool can account for the above situations if you provide it with information regarding exempt users and users with multiple accounts. Here is how you can provide the information:

- Specify the **-qexemptconcurrentusers** entry in the utilization tracking configuration file. This entry specifies users with Authorized User licenses.
- Specify the **-qsameuser urt** command-line option. This option specifies users with multiple accounts.

**Notes:**

- When the number of concurrent users is adjusted with **-qexemptconcurrentusers** or **-qsameuser**, the utilization reporting tool generates a message to indicate that the concurrent usage information is adjusted.
- The number of concurrent users might be zero if all concurrent invocations are invoked by exempt users. The tool also generates a message with this information.

## Usage file considerations

Usage (.cuf) files are used to store compiler usage information. This section provides information that helps you decide how you want to generate and use these files.

### Usage file location

Usage files can be created in each user's home directory, or they can be created in a central location for all users.

With utilization tracking enabled, when a compiler user compiles a program, a .cuf file is automatically created in the user's home directory in case the file does not exist. This is convenient for testing the utilization tracking feature because users already have write access to their own home directories, which means no additional setup is required. However, this might have the following issues:

- Compiler users might not know that the file has been created or what it is when they see the file. In this case, they might delete the file.
- Some users' home directories might be on file systems that are mounted from a remote system. This causes utilization tracking to use a remote file, which might affect performance.
- Compiler users might not want usage files to take up space in the /home directory.

A good alternative is to set up a central location where the usage files can be created, and provide appropriate access to that location for both the compiler users and the utilization reporting tool users. This can be set up by using the other/world permissions or by using group permissions.

For example, if the central location is a directory named /var/tmp/track\_compiler\_use, you can modify the **-qusagefileloc** entry in the utilization tracking configuration file as follows:

```
-qusagefileloc=/var/tmp/track_compiler_use/$LOGNAME.cuf
```

This creates a .cuf file for each user in the specified location, such as user1.cuf or user2.cuf. It is easier to run the utilization reporting tool to generate the usage

report from the .cuf files in this central location. You only need to pass the path of the location, /var/tmp/track\_compiler\_use to the utilization reporting tool , and then the tool can read all the .cuf files in that location.

If the compiler users are running the compiler on more than one machine, you need to add `$HOSTNAME` to the `-qusagefileloc` entry to ensure that there are no collisions in the file names. For example, you can specify the `-qusagefileloc` entry as follows:

```
-qusagefileloc=/var/tmp/track_compiler_use/$HOSTNAME_$(LOGNAME).cuf
```

This creates a .cuf file for each user, and the name of that .cuf file also contains the name of the host on which the compiler is used, such as `host1_user1.cuf`.

## The number of usage files

You can use one usage file or separate usage files for different compiler users.

### Using separate usage files for different compiler users

The advantages of using separate usage files are as follows:

- It might provide better performance because compiler users access their own usage files instead of competing for access to a shared one and separate usage files are usually smaller.
- Usage file for a user can be automatically created when the user uses the compiler to compile a program. There is no need to explicitly create a usage file for each user beforehand. For more information, see “Usage file location” on page 58.
- When generating utilization reports, you usually include all compiler users. However, if there are circumstances in which you want to exclude some users, you can simply omit their usage files when you invoke the utilization reporting tool. For example, you might want to omit users who have their own Authorized User license.

The disadvantage is that you might have to maintain separate usage files for different users.

### Using a single usage file for all compiler users

The advantage of using a shared usage file for all users is that you only need to maintain a single file instead of multiple files. However, with a single usage file, you lose the flexibility and possible performance benefits of using multiple usage files, as described in the preceding subsection.

The compiler provides an empty usage file `urtstub.cuf` in the `usr/vac/urt` directory. You can create a usage file for all compiler users by copying the empty usage file to a directory where they all have write access. In this case, you need to change the `-qusagefileloc` entry in the utilization tracking configuration file to point to the location of the usage file.

### Usage files on multiple machines

If you use the compiler on multiple machines, you need to decide how to make the usage files available for the utilization reporting tool.

You can use various methods to make the usage files available for the utilization reporting tool to generate usage reports and prune the usage files. Choose one of the following approaches to manage usage files on multiple machines:

- Copy the usage files from the machines where the compiler is used to the machine where the utilization reporting tool is installed. You can use any remote copy command, for example, **ftp**, **rmp**, **scp**, and **rsync**. In this case, the usage files are being accessed locally by both the compiler, for utilization tracking, and by the utilization reporting tool, for generating the usage report. Accessing the files locally yields the best performance.
- Use a distributed file system to export the file system from the machines where the compiler is used, and mount those file systems on the machine where the utilization reporting tool is installed. When you run the utilization reporting tool, it can access the usage files remotely via the mounted file systems.
- You can also export the file system from the machine where the utilization reporting tool is installed, and mount that file system on each machine where the compiler is used, using it as the location of the usage files where the compiler is recording its utilization. In this approach, the compiler records utilization in a remote usage file, and the utilization reporting tool reads the usage file locally.

**Note:** If you find this degrades the performance of the compiler, consider using one of the first two approaches instead.

### Usage file size

You need to consider the fact that the size of the usage files might grow quickly, especially when you use a shared file for all compiler users. If the usage file gets too large, it might affect utilization tracking performance.

To keep the usage files from growing quickly, you can optionally prune the usage files when you generate usage reports. You can also prune the files regularly using **cron**.

For more information about how to prune files, see “Pruning usage files” on page 72.

## Regular utilization checking

You can run the utilization reporting tool on a regular basis to verify whether the usage of the compiler is compliant with the Concurrent User licenses you have purchased. You can create a **cron** job to do this automatically.

If the usage files need to be copied to the machine where the utilization reporting tool is running, you can also automate the copying task with a **cron** job.

Another reason for running the tool regularly is to prune the usage files to control the size of these files.

**Note:** To reduce contention for read and write access to the usage files, run the utilization reporting tool or copy the usage files when the compiler is not being used.

---

## Testing utilization tracking

Before you begin to track the compiler usage for all users in your organization, you can test the feature with a limited number of users or with a separate compiler installation. During this testing, you can try different configurations so as to decide the best setup for your organization.

## Testing with a limited number of users

To enable compiler utilization tracking for a limited number of users, you can use a separate utilization tracking configuration file and ask only these users to use the file. Other users of the same installation use the default utilization tracking configuration file in which utilization tracking is disabled, and their use of the compiler is therefore not recorded.

The default compiler configuration file, `vac.cfg.53` or `vac.cfg.61` contains two entries, `xlurt_cfg_path` and `xlurt_cfg_name`, which specify the location of the utilization tracking configuration file. You need to perform the following tasks to let the specified users use the separate utilization tracking configuration file:

1. Create a separate compiler configuration file or stanza, in which the `xlurt_cfg_path` and `xlurt_cfg_name` entries specify the location of the utilization tracking configuration file you want to use.
2. Ask these users to use the following compiler option or environment variable to instruct the compiler to use the separate compiler configuration file or stanza, which in turn allows them to use the separate utilization tracking configuration file.
  - The `-F` option
  - The `XLC_USR_CONFIG` environment variable

**Note:** This approach is only for testing the utilization tracking feature. Do not use it for tracking all compiler utilization in your organization unless you can ensure that all compiler invocations are done with the `-F` option or the `XLC_USR_CONFIG` environment variable set.

## Testing with a separate compiler installation

You can install a separate instance of the compiler for testing utilization tracking. In this case, you can directly modify the utilization tracking configuration file in that installation to enable and configure utilization tracking. The compiler users involved in the testing do not need to perform any task for the tracking.

When you are satisfied that you have found the best utilization tracking configuration for your organization, you can enable it for all compiler users in your organization.

### Related information

- “Configuring utilization tracking”

---

## Configuring utilization tracking

You can use the utilization tracking configuration file to enable and configure the utilization tracking functionality.

The default location of the configuration file is `usr/vacpp/urt` and its file name is `urtxlc_cpp1201aix.cfg`.

The compiler uses a symlink to specify the location of the utilization tracking configuration file. The symlink is also located in `usr/vacpp/urt` and its name is `urt_client.cfg`. In the following situations, you might need to change the symlink:

- If you want to use a utilization tracking configuration file in a different location, change the symlink to point to this location.

- If you have multiple installations of the same compiler, and you plan to use a single utilization tracking configuration file, change the symlink in each installation to point to that file. For more information, see “Central configuration” on page 57.

**Note:** Installing a PTF update does not overwrite the utilization tracking configuration file.

You can use the entries in the utilization tracking configuration file to configure how compiler usage is tracked. For details about the specific entries in that file and how they can be modified, see “Editing utilization tracking configuration file entries.”

## Editing utilization tracking configuration file entries

You can configure different aspects of utilization tracking by editing the entries in the utilization tracking configuration file.

The entries are divided into two categories.

1. The entries in the *Product information* category identify the compiler. Do not modify these entries.
2. The entries in the *Tracking configuration* category can be used to configure utilization tracking for this product. Changes to these entries take effect in the usage file upon the next compiler invocation. In this case, the compiler emits a message to indicate that the new configuration values have been saved in the usage file. When you generate a report from the usage file, the new values are used.

The following rules apply when you modify the entries:

- The following entries are written to the usage files whenever you change them, and they are used the next time the utilization reporting tool generates a report from the usage files. These configuration entries must be the same for all compiler users.
  - **-qmaxconcurrentusers**
  - **-qexemptconcurrentusers**
  - **-qqualhostname**
- If **-qqualhostname** is changed, you must discard any existing usage files and start tracking utilization again with new usage files. Otherwise some invocations are recorded with qualified host names and some are recorded with unqualified host names.

### Notes:

- The entries are not compiler options. They can only be used in the utilization tracking configuration file.
- If the **-qexemptconcurrentusers** entry is specified multiple times in the utilization tracking configuration file, all the specified instances are used. If other entries are specified multiple times, the value of the last one overrides previous ones.
- The compiler generates a message if you specify the above entries with different values for different users when using more than one utilization tracking configuration file. You must modify the entries to keep them consistent, or make sure all compiler users use a single utilization configuration file.



## Product information

**-qprodId**=*product\_identifier\_string*

Indicates the unique product identifier string.

**-qprodVer**=*product\_version*

Indicates the product version.

**-qprodRel**=*product\_release*

Indicates the product release.

**-qprodName**=*product\_name*

Indicates the product name.

**-qconcurrentusagescope**=**prod** | **ver** | **rel**

Specifies the level at which concurrent users are counted, and their numbers are limited. The suboptions are as follows:

- **prod** indicates the product level.
- **ver** indicates the version level.
- **rel** indicates the release level.

Default: **-qconcurrentusagescope=prod**

## Tracking configuration

**-qmaxconcurrentusers**=*number*

Specifies the maximum number of concurrent users. It is the number of Concurrent User license that you have purchased for the product. When the utilization reporting tool generates a report from the usage file, it determines whether your compiler usage in your organization has exceeded this maximum number of concurrent users.

**Note:** You must update this entry to reflect the actual number of Concurrent User licenses that you have purchased.

Default: 0

**-qexemptconcurrentusers** ="*user\_account\_info\_1* [| *user\_account\_info\_2* | ... | *user\_account\_info\_n*]"

Specify exempt users who have their own Authorized User license. Exempt users can have as many concurrent invocations of the compiler as they want, without affecting the number of Concurrent User licenses available in your organization. When the utilization reporting tool generates a usage report, it does not include such users in the count of concurrent users.

*user\_account\_info* can be any combination of the following items:

- *name(user\_name)*
- *uid(user\_ID)*
- *host(host\_name)*

Users whose information matches the specified criteria are considered exempt users. For example, to indicate that *user1@host1* and *user2@host1* are exempt users, you can specify this entry in either of the following forms:

- **-qexemptconcurrentusers**="name(user1)host(host1)"
- **-qexemptconcurrentusers**="name(user2)host(host1)"
- **-qexemptconcurrentusers**="name(user1)host(host1) | name(user2)host(host1)"

For *user\_name*, *user\_ID*, and *host\_name*, you can also use a list of user names, user IDs, or hostnames separated by a space within the parentheses. For example:

```
-qexemptconcurrentusers="name(user1 user2)host(host1)"
```

This is equivalent to the previous examples.

**Note:** Specifying this entry does not exempt users from compiler utilization tracking. It only exempts them from being counted as concurrent users. To optimize utilization tracking performance, the format of the specified value is not validated until the report is produced. For more information about counting concurrent users, see “Concurrent user considerations” on page 57.

#### **-qqualhostname | -qnoqualhostname**

Specifies whether host names that are captured in usage files and then listed in compiler usage reports are qualified with domain names.

If all compiler usage within your organization is on machines within a single domain, you can reduce the size of the usage files by using **-qnoqualhostname** to suppress domain name qualification.

Default: **-qqualhostname**, which means the host names are qualified with domain names.

#### **-qenabletracking | -qnoenabletracking**

Enables or disables utilization tracking.

Default: **-qnoenabletracking**, which means utilization tracking is disabled.

#### **-qusagefileloc=directory\_or\_file\_name**

Specifies the location of the usage file.

By default, a `.cuf` file is automatically created for each user in their home directory. You can set up a central location where the files for each user can be created. For more information, see “Usage file location” on page 58.

The following rules apply when you specify this entry:

- If a file name is specified, it must have the `.cuf` extension. If the file is a symlink, it must point to a file with the `.cuf` extension. If the specified file does not exist, a `.cuf` file is created, along with any parent directories that do not already exist.
- If a directory is specified, there must be exactly one file with the `.cuf` extension in the directory. A new file is not created in this case.
- The path of the specified directory can be a relative or an absolute path. Relative paths are relative to the compiler user's current working directory.

**Note:** If a compiler user cannot access the file, for example, because of insufficient permissions to use or create the file, the compiler generates a message and the compilation continues.

You can use the following variables for this option:

- `$HOME` for the user's home directory. This allows each user to have a `.cuf` file in their home directory or a subdirectory of their home directory.
- `$USER` or `$LOGNAME` for the user's login user name. You can use this variable to create a `.cuf` file for each user and include the user's login name in the name of the `.cuf` file or in the name of a parent directory.
- `$HOSTNAME` for the name of the host on which the compiler runs. This can be useful when you track utilization across different hosts.

#### **-qfileaccessmaxwait=number\_of\_milliseconds**

Specifies the maximum number of milliseconds to wait for accessing the usage file.

**Note:** This entry is used to account for unusual circumstances where the system is under extreme heavy load and there is a delay in accessing the usage file.

Default: 3000 milliseconds

**Notes:**


- These entries are not compiler options. They can only be used in the utilization tracking configuration file.
- If the **-qexemptconcurrentusers** entry is specified multiple times in the utilization tracking configuration file, all the specified instances are used. If other entries are specified multiple times, the value of the last one overrides previous ones.

---

## Understanding the utilization reporting tool

You can use the utilization reporting tool to generate compiler usage reports from the information in one or more usage files, and optionally prune the usage files when you generate the reports.

The tool is located in the `/opt/ibmurt/1.1/bin` directory. You can use the **urt** command to invoke it. The syntax of the **urt** command is as follows:



The generated report is displayed on the standard output. You can direct the output to a file if you want to keep the report.

Command-line options control how usage reports are generated. For more information about the options, see “Utilization reporting tool command-line options” on page 66.

A default configuration file `ibmurt.cfg` is provided in the `/opt/ibmurt/1.1/config` directory. Entries in this file take the same form as the command-line options and have the same effect. You can also create additional configuration files and use the **-qconfigfile** option to specify their names.

You can specify the options in one or more of the following places:

1. The default configuration file
2. The additional configuration file specified with **-qconfigfile**
3. The command line

The utilization reporting tool uses the options in the default configuration file before it uses the options on the command line. When it encounters a **-qconfigfile** option on the command line, it reads the options in the specified configuration file and puts them on the command line at the place where the **-qconfigfile** option is used.

If an option is specified multiple times, the last specification that the tool encounters takes effect. Exceptions are **-qconfigfile** and **-qsameuser**. For these two options, all specifications take effect.

## Utilization reporting tool command-line options

The utilization reporting tool command-line options control the generation of the compiler utilization report.

Use these command-line options to modify the details of your compiler utilization report.

### **-qreporttype=detail** | **maxconcurrent**

Specifies the type of the usage report to generate.

- **detail** specifies that all invocations of the compiler are listed in the usage report. With this suboption, you can get a detailed report, in which the invocations that have exceeded the allowed maximum number of concurrent users are indicated.
- **maxconcurrent** specifies that only the compiler invocations that have exceeded the allowed maximum number of concurrent users are listed. With this suboption, you can get a compact report, which does not list those invocations within the maximum number of allowed concurrent users.

**Note:** The allowed maximum number of concurrent users is specified with the **-qmaxconcurrentusers** entry in the utilization tracking configuration file.

Default: **-qreporttype=maxconcurrent**.

### **-qrptmaxrecords=num** | **nomax**

Specifies the maximum number of records to list in the report for each product. *num* must be a positive integer.

Default: **-qrptmaxrecords=nomax**, which means all the records are listed.

### **-qusagefileloc=directory\_or\_file\_name**

Specifies the location of the usage files for report generation or pruning. It can be a list of directories or file names, or both, separated by colons.

The following rules apply when you specify this option:

- If one or more directories are specified, all files with the **.cuf** extension in those directories are used. Subdirectories can also be searched by using the **-qmaxsubdirs** option.
- The path of the specified directory can be relative or absolute. Relative paths are relative to the compiler user's current working directory.
- A symlink does not require the **.cuf** extension but the file to which it points must have that extension.

**Note:**

- The **-qusagefileloc** entry in the utilization tracking configuration file tells the compiler which usage files to use for recording compiler utilization. This **-qusagefileloc** option tells the utilization reporting tool where to find those usage files.

Default: **::\$HOME**, which means the utilization reporting tool looks for usage files in your current working directory and your home directory.

### **-qmaxsubdirs=num** | **nomax**

Specifies whether to search subdirectories for usage files, and how many levels of subdirectories to search. *num* must be a non-negative integer.

If **nomax** is specified, all the subdirectories are searched. If 0 is specified, no subdirectories are searched.

Default: 0.

**-qconfigfile=***file\_path*

Specifies the user defined configuration file that you want to use.

For more information about how the utilization reporting tool uses the configuration file, see “Understanding the utilization reporting tool” on page 65.

**Note:** If you specify this option multiple times, all specified instances are used.

**-qsameuser=***user\_account\_info*

Specifies different user accounts that belong to the same compiler user. Use this option when a user accesses the compiler from more than one user ID or machine to avoid having that user reported as multiple users. Invocations of the compiler by these different accounts are counted as a single user instead of multiple different users.

*user\_account\_info* can be any combination of the following items:

- name(*user\_name*)
- uid(*user\_ID*)
- host(*host\_name*)

There are two ways to pass these rules to the utilization reporting tool. You can supply specific lists of the *user\_names*, *user\_IDs* or *host\_names* that are shared by the same user or you can use a more generic (=) syntax.

For example, to indicate that *user1* and *user2* are both user names belonging to the same person who uses the compiler on the *host1* machine, use the syntax in which you specify these user names and the host name explicitly:

```
-qsameuser="name(user1)host(host1) | name(user2)host(host1)"
```

or

```
-qsameuser="name(user1 user2)host(host1)"
```

Both of these examples use specific user names and host names to indicate accounts that belong to the same user, but they do so in slightly different ways. The first example uses a vertical bar to separate the different user accounts that belong to this user, while the second example uses a list of user names within the parentheses instead of repeating the same host information twice. They both convey the same account information, but the second example is more concise.

As an example of the more generic (=) syntax, you can indicate that all user accounts with the same user name and uid belong to the same user as follows:

```
-qsameuser="name(=)uid(=)"
```

With this option, you are not specifying specific user names or uids as you did in the previous example. User accounts that have the same user name and uid are considered as belonging to the same user, regardless of what the specific user names and uids are, and regardless of what the host name is. This establishes a general rule that applies to all accounts in your organization instead of specific ones.

The following rules apply when you specify this option:

- Each instance of the **-qsameuser** option must use either the list or generic (=) syntax. You cannot combine them in the same instance of the option but you can use multiple instances of the **-qsameuser** option to refine the report.
- The utilization reporting tool matches the user information based on the order that the **-qsameuser** option values are specified. Once it finds a match it stops matching the same user information against any subsequent options.

The following examples illustrates the differences:

- If you specify the **-qsameuser** option as follows:

```
-qsameuser="name(user1)" -qsameuser="uid(=)"
```

Specifying the **-qsameuser** option in this order means that user accounts with the user name *user1* matches the first option and is not evaluated against the second option. User accounts *user1* and *user2* are not considered the same user even if they have the same *uid*.

- If you specify the **-qsameuser** option as follows:

```
-qsameuser="uid(=)" -qsameuser="name(user1)"
```

Specifying the **-qsameuser** option in this order means that user accounts with the same *uid* are always considered to be the same user, and in addition, any user accounts with a user name of *user1* should be considered belonging to the same user even if they do not match by *uid*.

**Note:** Specifying this option does not prevent user information from being listed in the usage report. For more information about concurrent users, see “Concurrent user considerations” on page 57.

#### **-qadjusttime=*time\_adjustments***

Adjusts the time that have been recorded in the usage files for the specified machines. *time\_adjustments* is a list of entries with the format of *machine name + | - number of seconds*, separated by colons.

The following rules apply when you use this option:

- An entry of the form *machine name + number of seconds* causes the specified number of seconds to be added to the start and end times of any invocations recorded for the specified machine.
- An entry of the form *machine name - number of seconds* causes the specified number of seconds to be subtracted from the start and end times of any invocations recorded for the specified machine.

For example:

```
-qadjusttime="hostA+5:hostB-3"
```

Five seconds are added to the start and end times of the invocations on *hostA*, and three seconds are subtracted from the start and end times of the invocations on *hostB*.

Only use this option if the usage files contain utilization information from two or more machines, and time is not synchronized across those machines. The adjustments specified by this option compensate for the lack of synchronization

**Notes:**

- The specified adjustments are only used for the current run of the **urt** command. Specifying this option does not change the invocation information recorded in the usage files.
- Do not specify the same machine name more than once with this option.

**-qusagefilemaxage=number\_of\_days | nomax**

Prunes the usage files by removing all invocations older than the specified number of days.

Every usage file specified by the **-qusagefileloc** option is pruned. The usage report contains this information to indicate the number of records that have been pruned.

Default: **-qusagefilemaxage=nomax**, which means no pruning is performed.

**-qusagefilemaxsize=number\_of\_MB | nomax**

Prunes the usage files to keep them under the specified size. It prunes the files by removing the oldest invocations.

Every usage file specified by the **-qusagefileloc** option is pruned. The usage report contains this information to indicate the number of records that have been pruned.

Default: **-qusagefilemaxsize=nomax**, which means no pruning is performed.

**-qtimesort=ascend | descend**

Specifies the chronological order in which the usage report information is sorted.

- Specifying **ascend** means new information is listed after the older information.
- Specifying **descend** means the newest information is at the top of the report.

Default: **-qtimesort=ascend**.

---

## Generating usage reports

You can use the utilization reporting tool to generate compiler usage reports based on the usage information stored in the usage files.

To generate a report, use the command-line options or the **urt** configuration file to specify how you want a report to be generated. For more information about these options, see “Utilization reporting tool command-line options” on page 66.

### Notes:

- You can set up a cron service to run the utilization reporting tool on a regular basis. If the usage files from which the tool generate reports need to be copied to the machine where the tool is running, you can also automate this copying task with cron.
- To reduce contention for read and write access to the usage files, do not run the tool or copy the usage files when the compiler is being used.

The generated report is displayed on the standard output. You can direct the output to a file if you want to keep the report.

After a usage report is generated, the utilization reporting tool uses the following exit codes to indicate the compliance status of your compiler license:

- Exit code = "1".

The utilization reporting tool has detected that the number of Concurrent User license entitlements specified in the **-qmaxconcurrentusers** entry in the utilization tracking configuration file has been exceeded. See the generated report for details and contact your IBM representative to purchase additional Concurrent User licenses.

- Exit code ="0".

The compiler utilization is within the number of Concurrent User license entitlements specified.

For more information about the **urt** command, see "Understanding the utilization reporting tool" on page 65.

## Understanding usage reports

You can use the report that the utilization report tool generates to analyze the compiler usage in your organization.

The report has a REPORT SUMMARY section that lists the following information:

1. The date and time when the report is generated.
2. The .cuf file or a list of all .cuf files used to generate the report.
3. The options that have been passed to the **urt** command, with default values for any unspecified options.
4. Possible messages categorized as ERROR, WARNING, or INFO. For detailed information about possible messages, see "Diagnostic messages from utilization tracking and reporting" on page 73.

After the summary section, there is a REPORT DETAILS section for each compiler version. This section lists all of the compiler invocations recorded in the usage files. The content of these sections varies depending on the report type that you have specified. For detailed information about the report types, see **-qreporttype**.

Here are the sample reports generated with the two different report types:

Sample 1: A sample report generated with **-qreporttype=detail**

```
REPORT SUMMARY
-----

DATE: 12/18/09      TIME: 01:30:24

OPTIONS USED (* indicates that a default value was used):

reporttype=detail
maxsubdirs=0
configfile="/opt/ibmurt/1.1/config/ibmurt.cfg"
rptmaxrecords=nomax
*adjusttime=
usagefileloc="/home/testrun/ibmxlcompiler.cuf"
*sameuser=
timesort=ascend
usagefilemaxsize=nomax
usagefilemaxage=nomax

FILES USED:

/home/testrun/ibmxlcompiler.cuf

REPORT DETAILS
-----
```



USAGE INFORMATION FOR PRODUCT: IBM XL C for AIX 11.1

Max. Concurrent Users Exceeded? : \*\*\* YES \*\*\*

Max. Concurrent Users Allowed: 1                      Max. Concurrent Users Recorded: 5

Exempt Users:

Product invocations:

Start Time	End Time	User	Number of Concurrent Users
12/17/09 16:56:44	12/17/09 16:57:13	user1@host1.ibm.com	1
12/18/09 00:58:29	12/18/09 00:58:32	user2@host2.ibm.com	1
12/18/09 01:16:01	12/18/09 01:16:02	user3@host3.ibm.com	5 ( exceeds max. allowed)
12/18/09 01:16:02	12/18/09 01:16:26	user2@host2.ibm.com	5 ( exceeds max. allowed)
12/18/09 01:16:08	12/18/09 01:16:08	user3@host2.ibm.com	5 ( exceeds max. allowed)
12/18/09 01:16:12	12/18/09 01:16:12	user2@host1.ibm.com	5 ( exceeds max. allowed)
12/18/09 01:16:24	12/18/09 01:16:28	user1@host2.ibm.com	5 ( exceeds max. allowed)
12/18/09 01:26:11	12/18/09 01:27:46	user3@host3.ibm.com	2 ( exceeds max. allowed)
12/18/09 01:26:27	12/18/09 01:27:46	user1@host1.ibm.com	2 ( exceeds max. allowed)
12/18/09 01:29:59	12/18/09 01:30:00	user2@host1.ibm.com	1
12/18/09 01:30:00	12/18/09 01:30:00	user2@host2.ibm.com	3 ( exceeds max. allowed)
12/18/09 01:30:14	12/18/09 01:30:15	user3@host1.ibm.com	3 ( exceeds max. allowed)
12/18/09 01:30:14	12/18/09 01:30:14	user2@host2.ibm.com	3 ( exceeds max. allowed)

Sample 2: A sample report generated with **-qreporttype=maxconcurrent**

REPORT SUMMARY

-----

DATE: 12/18/09                      TIME: 01:32:53

OPTIONS USED (\* indicates that a default value was used):

```
reporttype=maxconcurrent
maxsubdirs=0
configfile="/opt/ibmurt/1.1/config/ibmurt.cfg"
rptmaxrecords=nomax
*adjusttime=
usagefileloc="/home/testrun/ibmxlcompiler.cuf"
*sameuser=
timesort=ascend
usagefilemaxsize=nomax
usagefilemaxage=nomax
```

FILES USED:

/home/testrun/ibmxlcompiler.cuf

REPORT DETAILS

-----

USAGE INFORMATION FOR PRODUCT: IBM XL C for AIX 11.1

Max. Concurrent Users Exceeded? : \*\*\* YES \*\*\*

Max. Concurrent Users Allowed: 1                      Max. Concurrent Users Recorded: 5

Exempt Users:

Dates and times where usage exceeded the maximum allowed:

Date	Time	Number of Concurrent Users	Users
12/18/09	01:16:01	5	user3@host3.ibm.com

12/18/09	01:16:02	5	user2@host2.ibm.com user3@host2.ibm.com user2@host1.ibm.com user1@host2.ibm.com user3@host3.ibm.com user2@host2.ibm.com user3@host2.ibm.com user2@host1.ibm.com user1@host2.ibm.com
12/18/09	01:16:08	5	user3@host3.ibm.com user2@host2.ibm.com user3@host2.ibm.com user2@host1.ibm.com user1@host2.ibm.com
12/18/09	01:16:12	5	user3@host3.ibm.com user2@host2.ibm.com user3@host2.ibm.com user2@host1.ibm.com user1@host2.ibm.com
12/18/09	01:16:24	5	user3@host3.ibm.com user2@host2.ibm.com user3@host2.ibm.com user2@host1.ibm.com user1@host2.ibm.com user3@host3.ibm.com user2@host2.ibm.com user3@host2.ibm.com user2@host1.ibm.com
12/18/09	01:26:11	2	user3@host3.ibm.com user1@host1.ibm.com
12/18/09	01:26:27	2	user3@host3.ibm.com user1@host1.ibm.com
12/18/09	01:30:00	3	user2@host2.ibm.com user2@host1.ibm.com user3@host1.ibm.com
12/18/09	01:30:14	3	user2@host2.ibm.com user2@host1.ibm.com user3@host1.ibm.com
12/18/09	01:30:14	3	user2@host2.ibm.com user2@host1.ibm.com user3@host1.ibm.com

**Note:** There are circumstances under which an end time might not be recorded. These might include:

- There was a major failure of the compiler, for example, power loss during a compilation.
- The invocation had not ended at the time when the report was generated, or at the time when the usage file was being copied.
- The permission to write to the usage file was revoked at some time before the end time of the invocation was recorded.

An invocation with no end time recorded is not included in the count of concurrent users.

---

## Pruning usage files

Usage files grow with each compiler invocation. You can prune the usage files with the utilization report tool.

When you generate a usage report, you can specify the following two options to optionally prune the usage files:

- **-qusagefilemaxage:** Removes the invocations older than the specified number of days. For example, to remove all entries in the usage files older than 30 days, use the following command:

```
urt -qusagefilemaxage=30
```

- **-qusagefilemaxsize:** Removes the oldest invocations to keep the usage files under the specified size. For example, to remove the oldest invocations to keep the usage files under 30 MB, use the following command:

```
urt -qusagefilemaxsize=30
```

When usage files are pruned, the usage report includes an information message that indicates the number of records that have been pruned. If you want to keep the generated report after the files are pruned, you can redirect the output to a file.

To control the size of the usage files, you can prune the usage files on a regular basis. You can create a cron job to do this automatically.

If you do not have the utilization reporting tool installed on each machine where the usage files are located, you have the following options:

- Export the file system from each machine where the usage files exist and mount it on the machine where the utilization reporting tool is installed. Then run the tool to prune the usage files on the mounted network file system.
- After copying the usage files to the machine where the utilization reporting tool is installed, delete the files and use new usage files to capture any subsequent compiler invocations. This approach might also speed up the report generation because the utilization reporting tool is not accessing the usage files remotely and it is not spending time pruning the usage files.

Pruning usage files might slow down the usage report generation process, especially when the number or the size of the usage files is large. If you do not want to prune the files every time you generate reports, you can set the values for the **-qusagefilemaxage** and **-qusagefilemaxsize** options as follows:

- If you generate the report daily, you can specify these two options with very high values so pruning does not occur. The default value **nomax** can be used in this case.
- You can set appropriate values for these two options and use a separate cron job to prune the usage files weekly.

**Note:** To reduce contention for read and write access to the usage files, do not run the utilization report tool or copy the usage files when the compiler is being used.

---

## Diagnostic messages from utilization tracking and reporting

The compiler generates diagnostic messages to indicate utilization tracking issues. These messages can help you to fix the associated problems.

For example:

```
Utilization tracking configuration file could not be read due to  
insufficient permissions.
```

This message indicates that you need read access for utilization tracking configuration file.

When the utilization reporting tool is used to generate usage reports or prune usage files, it also generates diagnostic messages. For example:

```
Unrecognized option -qmaxsubdir.
```

This message indicates that you have specified a wrong option.

**Note:** Possible error, warning, or information messages are also included in the compiler usage report generated by the tool.

---

## Chapter 4. Compiler options reference

The following sections contain a summary of the compiler options available in XL C/C++ by functional category, followed by detailed descriptions of the individual options.

### Related information

- “Specifying compiler options” on page 5
- “Reusing GNU C/C++ compiler options with gxlc and gxlc++” on page 11

---

## Summary of compiler options by functional category

The XL C/C++ options available on the AIX platform are grouped into the following categories. If the option supports an equivalent pragma directive, this is indicated. To get detailed information on any option listed, see the full description for that option.

- “Output control”
- “Input control” on page 76
- “Language element control” on page 77
- “Template control (C++ only)” on page 79
- “Floating-point and integer control” on page 80
- “Error checking and debugging” on page 83
- “Listings, messages, and compiler information” on page 85
- “Optimization and tuning” on page 87
- “Object code control” on page 80
- “Linking” on page 91
- “Portability and migration” on page 91
- “Compiler customization” on page 92
- “Deprecated options” on page 93

### 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 10. Compiler output options

Option name	Equivalent pragma name	Description
“-c” on page 119	None.	Prevents the completed object from being sent to the linker. With this option, the output is a .o file for each source file.
“-C, -C!” on page 120	None.	When used in conjunction with the -E or -P options, preserves or removes comments in preprocessed output.

Table 10. Compiler output options (continued)

Option name	Equivalent pragma name	Description
"-E" on page 142	None.	Preprocesses the source files named in the compiler invocation, without compiling, and writes the output to the standard output.
"-G" on page 172	None.	Generates a shared object enabled for runtime linking.
"-qmakedep, -M" on page 252	None.	Creates an output file containing targets suitable for inclusion in a description file for the <b>make</b> command.
"-MF" on page 258	None.	Specifies the target for the output generated by the <b>-qmakedep</b> or <b>-M</b> options.
"-qmkshrobj" on page 260	None.	Creates a shared object from generated object files.
"-o" on page 268	None.	Specifies a name for the output object, assembler, or executable file.
"-P" on page 280	None.	Preprocesses the source files named in the compiler invocation, without compiling, and creates an output preprocessed file for each input file.
"-S" on page 310	None.	Generates an assembler language file for each source file.
"-qshowmacros" on page 315	None.	Emits macro definitions to preprocessed output.
"-qtimestamps" on page 351	None.	Controls whether or not implicit time stamps are inserted into an object file.

## Input control

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

Table 11. Compiler input options

Option name	Equivalent pragma name	Description
"-+ (plus sign) (C++ only)" on page 95	None.	Compiles any file as a C++ language file.
"-qcinc (C++ only)" on page 126	None.	Places an extern "C" { } wrapper around the contents of include files located in a specified directory.
"-I" on page 181	None.	Adds a directory to the search path for include files.

Table 11. Compiler input options (continued)

Option name	Equivalent pragma name	Description
"-qidirfirst" on page 182	#pragma options idirfirst	Specifies whether the compiler searches for user include files in directories specified by the <code>-I</code> option <i>before</i> or <i>after</i> searching any other directories.
"-qinclude" on page 185	None.	Specifies additional header files to be included in a compilation unit, as though the files were named in an <code>#include</code> statement in the source file.
"-qsourcectype" on page 325	None.	Instructs the compiler to treat all recognized source files as a specified source type, regardless of the actual file name suffix.
"-qstdinc" on page 332	#pragma options stdinc	Specifies whether the standard include directories are included in the search paths for system and user header files.


## 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 12. Language element control options

Option name	Equivalent pragma name	Description
"-qaltivec" on page 105	None	Enables compiler support for vector data types and operators.
"-qasm" on page 110	None	Controls the interpretation and subsequent generation of code for assembler language extensions.
"-qcpluscmt (C only)" on page 129	None.	Enables recognition of C++-style comments in C source files.
"-D" on page 132	None.	Defines a macro as in a <code>#define</code> preprocessor directive.
"-qdfp" on page 137	None.	Enables compiler support for decimal floating-point types and literals.
"-qdigraph" on page 138	#pragma options digraph	Enables recognition of digraph key combinations or keywords to represent characters not found on some keyboards.
"-qdollar" on page 140	#pragma options dollar	Allows the dollar-sign (\$) symbol to be used in the names of identifiers.
"-qignprag" on page 184	#pragma options ignprag	Instructs the compiler to ignore certain pragma statements.

Table 12. Language element control options (continued)

Option name	Equivalent pragma name	Description
"-qkeyword" on page 211	None.	Controls whether the specified name is treated as a keyword or as an identifier whenever it appears in your program source.
"-qlanglvl" on page 215	 #pragma options langlvl, #pragma langlvl	Determines whether source code and compiler options should be checked for conformance to a specific language standard, or subset or superset of a standard.
"-qlonglong" on page 249	#pragma options long long	Allows IBM long long integer types in your program.
"-qmacpstr" on page 251	#pragma options macpstr	Converts Pascal string literals (prefixed by the \p escape sequence) into null-terminated strings in which the first byte contains the length of the string.
"-qmbcs, -qdbcs" on page 257	#pragma options mbcs, #pragma options dbcs	Enables support for multibyte character sets (MBCS) and Unicode characters in your source code.
"-qstaticinline (C++ only)" on page 330	None.	Controls whether inline functions are treated as having static or extern linkage.
"-qtabsize" on page 343	None.	Sets the default tab length, for the purposes of reporting the column number in error messages.
"-qtrigraph" on page 356	None.	Enables the recognition of trigraph key combinations to represent characters not found on some keyboards.
"-U" on page 361	None.	Undefines a macro defined by the compiler or by the <b>-D</b> compiler option.
"-qutf" on page 367	None.	Enables recognition of UTF literal syntax.



## Template control (C++ only)

You can use these options to control how the C++ compiler handles templates.

Table 13. C++ template options

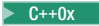
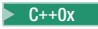
Option name	Equivalent pragma name	Description
"-qtempinc (C++ only)" on page 345	None.	Generates separate template include files for template functions and class declarations, and places these files in a directory which can be optionally specified.
"-qtemplatedepth (C++ only)" on page 346	None.	Specifies the maximum number of recursively instantiated template specializations that will be processed by the compiler.
"-qtemplaterecompile (C++ only)" on page 347	None.	Helps manage dependencies between compilation units that have been compiled using the <b>-qtemplateregistry</b> compiler option.
"-qtemplateregistry (C++ only)" on page 348	None.	Maintains records of all templates as they are encountered in the source and ensures that only one instantiation of each template is made.
"-qtempmax (C++ only)" on page 349	None.	Specifies the maximum number of template include files to be generated by the <b>-qtempinc</b> option for each header file.
"-qtmplinst (C++ only)" on page 353	None.	Manages the implicit instantiation of templates.
"-qtmplparse (C++ only)" on page 354	None.	Controls whether parsing and semantic checking are applied to template definitions.
 <b>-qlanglvl=[no]externtemplate</b>	None.	Suppresses the implicit instantiation of a template specialization or its members.
<b>-qlanglvl=[no]gnu_externtemplate</b>	None.	Suppresses the implicit instantiation of a template specialization or its members. This option is deprecated in XL C/C++ V12.1; you can use the option <b>-qlanglvl=[no]externtemplate</b> instead.

Table 13. C++ template options (continued)

Option name	Equivalent pragma name	Description
 <code>-q[lang]vl=[no]variadic[templates]</code>	None.	Defines class or function templates that can have any number (including zero) of parameters.

## 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.

Table 14. Floating-point and integer control options

Option name	Equivalent pragma name	Description
"-qbitfields" on page 117	None.	Specifies whether bit fields are signed or unsigned.
"-qchars" on page 123	#pragma options chars, #pragma chars	Determines whether all variables of type char are treated as either signed or unsigned.
"-qenum" on page 144	#pragma options enum, #pragma enum	Specifies the amount of storage occupied by enumerations.
"-qfloat" on page 154	#pragma options float	Selects different strategies for speeding up or improving the accuracy of floating-point calculations.
"-qldbl128, -qlongdouble" on page 239	#pragma options ldbl128	Increases the size of long double types from 64 bits to 128 bits.
"-qlonglit" on page 249	None.	In 64-bit mode, when determining the implicit types for integer literals, the compiler behaves as if an l or L suffix were added to integral literals with no suffix or with a suffix consisting only of u or U.
"-y" on page 381	None.	Specifies the rounding mode for the compiler to use when evaluating constant floating-point expressions at compile time.

## Object code control

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

Table 15. Object code control options

Option name	Equivalent pragma name	Description
"-q32, -q64" on page 96	None.	Selects either 32-bit or 64-bit compiler mode.

Table 15. Object code control options (continued)

Option name	Equivalent pragma name	Description
"-qalloca, -ma (C only)" on page 104	#pragma alloca	Provides an inline definition of system function <code>alloca</code> when it is called from source code that does not include the <code>alloca.h</code> header.
"-qconcurrentupdate (C only)" on page 128	None.	Supports updating the operating system while the kernel is running.
"-qeh (C++ only)" on page 144	None.	Controls whether exception handling is enabled in the module being compiled.
"-qexpfile" on page 148	None.	When used together with the <b>-qmkshrobj</b> or <b>-G</b> option, saves all exported symbols in a designated file.
"-qfuncsect" on page 165	#pragma options funcsect	Places instructions for each function in a separate object file control section or CSECT which may reduce the size of your program. Placing each function in its own CSECT enables the linker to perform garbage collection on a per function basis rather than per object file.
"-qinlglue" on page 196	#pragma options inlglue	When used with <b>-O2</b> or higher optimization, inlines glue code that optimizes external function calls in your application.
"-qkeepinlines (C++ only)" on page 209	None.	Keeps or discards definitions for unreferenced extern inline functions.
"-qpic" on page 291	None.	Generates position-independent code suitable for use in shared libraries.
"-qppline" on page 293	None.	When used in conjunction with the <b>-E</b> or <b>-P</b> options, enables or disables the generation of <code>#line</code> directives.
"-qpriority (C++ only)" on page 297	#pragma options priority, #pragma priority	Specifies the priority level for the initialization of static objects.
"-qproto (C only)" on page 300	#pragma options proto	Specifies the linkage conventions for passing floating-point arguments to functions that have not been prototyped.

Table 15. Object code control options (continued)

Option name	Equivalent pragma name	Description
"-qreserved_reg" on page 304	None.	Indicates that the given list of registers cannot be used during the compilation except as a stack pointer, frame pointer or in some other fixed role.
"-qro" on page 305	#pragma options ro, #pragma strings	Specifies the storage type for string literals.
"-qroconst" on page 306	#pragma options roconst	Specifies the storage location for constant values.
"-qroptr" on page 307	None.	Specifies the storage location for constant pointers.
"-qrtti (C++ only)" on page 308	None.	Generates runtime type identification (RTTI) information for exception handling and for use by the typeid and dynamic_cast operators.
"-s" on page 309	None.	Strips the symbol table, line number information, and relocation information from the output file.
"-qsaveopt" on page 311	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 329	None.	Provides protection against malicious code or programming errors that overwrite or corrupt the stack.
"-qstatsym" on page 331	None.	Adds user-defined, nonexternal names that have a persistent storage class, such as initialized and uninitialized static variables, to the symbol table of the object file.
"-qtbtable" on page 344	#pragma options tbtable	Controls the amount of debugging traceback information that is included in the object files.
"-qthreaded" on page 350	None.	Indicates to the compiler whether it must generate threadsafe code.

Table 15. Object code control options (continued)

Option name	Equivalent pragma name	Description
"-qtls" on page 352	None.	Enables recognition of the <code>__thread</code> storage class specifier, which designates variables that are to be allocated threadlocal storage; and specifies the threadlocal storage model to be used.
"-qunique" on page 362	None.	Generates unique names for static constructor/destructor file compilation units.
"-qweakexp" on page 377	None.	When used with the <code>-qmkshrobj</code> or <code>-G</code> option, includes or excludes global symbols marked as weak from the export list generated when you create a shared object.
"-qweaksymbol" on page 378	None.	Enables the generation of weak symbols.
"-qxcall" on page 378	None.	Generates code to treat static functions within a compilation unit as if they were external functions.

## Error checking and debugging

The options in this category 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 the options in "Listings, messages, and compiler information" on page 85.

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

Table 16. Error checking and debugging options

Option name	Equivalent pragma name	Description
"-# (pound sign)" on page 96	None.	Previews the compilation steps specified on the command line, without actually invoking any compiler components.
"-qcheck" on page 125	#pragma options check	Generates code that performs certain types of runtime checking.
"-qdbgfmt" on page 135	None	Specifies the stabstring or DWARF format for the debugging information produced in object files.

Table 16. Error checking and debugging options (continued)

Option name	Equivalent pragma name	Description
"-qdbxextra (C only)" on page 136	#pragma options dbxextra	When used with the <b>-g</b> option, specifies that debugging information is generated for unreferenced typedef declarations, struct, union, and enum type definitions.
"-qdpc1" on page 140	None.	Generates symbols that tools based on the IBM Dynamic Probe Class Library (DPCL) can use to see the structure of an executable file.
"-qextchk" on page 149	#pragma options extchk	Generates link-time type checking information and checks for compile-time consistency.
"-qflttrap" on page 159	#pragma options flttrap	Determines what types of floating-point exceptions to detect at run time.
"-qformat" on page 163	None.	Warns of possible problems with string input and output format specifications.
"-qfullpath" on page 164	#pragma options fullpath	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 166	None.	Calls the tracing routines to trace the entry and exit points of the specified functions in a compilation unit.
"-g" on page 169	None.	Generates debugging information for use by a symbolic debugger, and makes the program state available to the debugging session at selected source locations.
"-qhalt" on page 174	#pragma options 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 175	None.	Stops compilation before producing any object files, executable files, or assembler source files if a specified error message is generated.
"-qheapdebug" on page 177	None.	Enables debug versions of memory management functions.

Table 16. Error checking and debugging options (continued)

Option name	Equivalent pragma name	Description
"-qinfo" on page 187	#pragma options info, #pragma info	Produces or suppresses groups of informational messages.
"-qinitauto" on page 193	#pragma options initauto	Initializes uninitialized automatic variables to a specific value, for debugging purposes.
"-qkeepparm" on page 210	None.	When used with <b>-O2</b> or higher optimization, specifies whether procedure parameters are stored on the stack.
"-qlinedebug" on page 242	None.	Generates only line number and source file name information for a debugger.
"-qmaxerr" on page 254	None.	Stops compilation when the number of error messages of a specified severity level or higher reaches a specified number.
"-qoptdebug" on page 275	None.	When used with high levels of optimization, produces files containing optimized pseudocode that can be read by a debugger.
"-qsymtab (C only)" on page 340	None.	Determines the information that appears in the symbol table.
"-qsyntaxonly (C only)" on page 341	None.	Performs syntax checking without generating an object file.
"-qwarn0x (C++0x)" on page 374	None.	Controls whether to inform users with messages about differences in their programs caused by migration from the C++98 standard to the C++0x standard.
"-qwarn64" on page 376	None.	Enables checking for possible data conversion problems between 32-bit and 64-bit compiler modes.

## Listings, messages, and compiler information

The options in this category 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 83 to provide a more robust overview of your application when checking for errors and unexpected behavior.

Table 17. Listings and messages options

Option name	Equivalent pragma name	Description
"-qattr" on page 114	#pragma options attr	Produces a compiler listing that includes the attribute component of the attribute and cross-reference section of the listing.
"-qflag" on page 153	#pragma options flag, C++ "#pragma report (C++ only)" on page 430	Limits the diagnostic messages to those of a specified severity level or higher.
"-qlist" on page 243	#pragma options list	Produces a compiler listing file that includes an object listing.
"-qlistfmt" on page 245	None.	Creates an XML or HTML report to assist with finding optimization opportunities.
"-qlistopt" on page 248	None.	Produces a compiler listing file that includes all options in effect at the time of compiler invocation.
"-qphsinfo" on page 290	None.	Reports the time taken in each compilation phase to standard output.
"-qprint" on page 296	None.	Enables or suppresses listings.
"-qreport" on page 302	None.	Produces listing files that show how sections of code have been optimized.
"-qshowinc" on page 313	#pragma options showinc	When used with <b>-qsource</b> option to generate a listing file, selectively shows user or system header files in the source section of the listing file.
"-qskipsrc" on page 318	None.	When a listing file is generated using the <b>-qsource</b> option, <b>-qskipsrc</b> can be used to determine whether the source statements skipped by the compiler are shown in the source section of the listing file. Alternatively, the <b>-qskipsrc=hide</b> option is used to hide the source statements skipped by the compiler.
"-qsource" on page 324	#pragma options source	Produces a compiler listing file that includes the source section of the listing and provides additional source information when printing error messages.



Table 17. Listings and messages options (continued)

Option name	Equivalent pragma name	Description
"-qsrcmsg (C only)" on page 328	#pragma options srcmsg	Adds the corresponding source code lines to diagnostic messages generated by the compiler.
"-qsuppress" on page 338	None.	Prevents specific informational or warning messages from being displayed or added to the listing file, if one is generated.
"-v, -V" on page 368	None.	Reports the progress of compilation, by naming the programs being invoked and the options being specified to each program.
"-qversion" on page 370	None.	Displays the version and release of the compiler being invoked.
"-w" on page 371	None.	Suppresses informational, language-level and warning messages.
"-qxref" on page 379	#pragma options xref	Produces a compiler listing that includes the cross-reference component of the attribute and cross-reference section of the listing.

## Optimization and tuning

The options in this category allow you to control the optimization and tuning process, which can improve the performance of your application at run time.

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.

You can also control some of these options, such as **Optimize**, **-qcompact**, or **-qstrict**, with an **option\_override** pragma.

In addition to the option descriptions in this section, consult the *XL C/C++ Optimization and Programming Guide* for a details on the optimization and tuning process as well as writing optimization-friendly source code.

Table 18. Optimization and tuning options

Option name	Equivalent pragma name	Description
"-qaggrcopy" on page 97	None.	Enables destructive copy operations for structures and unions.

Table 18. Optimization and tuning options (continued)

Option name	Equivalent pragma name	Description
"-qalias" on page 98	None.	Indicates whether a program contains certain categories of aliasing or does not conform to C/C++ 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 106	None.	Specifies the processor architecture for which the code (instructions) should be generated.
"-qcache" on page 121	None.	When specified with <b>-O4</b> , <b>-O5</b> , or <b>-qipa</b> , specifies the cache configuration for a specific execution machine.
"-qcompact" on page 127	#pragma options compact	Avoids optimizations that increase code size.
"-qdataimported, -qdatalocal, -qtocdata" on page 134	None.	Marks data as local or imported in 64-bit compilations.
"-qdirectstorage" on page 139	None.	Informs the compiler that a given compilation unit may reference write-through-enabled or cache-inhibited storage.
"-qfdpr" on page 152	None.	Provides object files with information that the IBM Feedback Directed Program Restructuring (FDPR <sup>®</sup> ) performance-tuning utility needs to optimize the resulting executable file.
"-qhot" on page 178	#pragma nosimd, #pragma novector	Performs high-order loop analysis and transformations (HOT) during optimization.
"-qignerrno" on page 183	#pragma options ignerrno	Allows the compiler to perform optimizations that assume errno is not modified by system calls.
"-qipa" on page 199	None.	Enables or customizes a class of optimizations known as interprocedural analysis (IPA).
"-qisolated_call" on page 206	#pragma options isolated_call, #pragma isolated_call	Specifies functions in the source file that have no side effects other than those implied by their parameters.

Table 18. Optimization and tuning options (continued)

Option name	Equivalent pragma name	Description
"-qlargepage" on page 238	None.	Takes advantage of large pages provided on POWER4 and higher systems, for applications designed to execute in a large page memory environment.
"-qlibansi" on page 241	#pragma options libansi	Assumes that all functions with the name of an ANSI C library function are in fact the system functions.
"-qlibmpi" on page 241	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 256	#pragma options maxmem	Limits the amount of memory that the compiler allocates while performing specific, memory-intensive optimizations to the specified number of kilobytes.
"-qminimaltoc" on page 259	None.	Controls the generation of the table of contents (TOC), which the compiler creates for an executable file.
"-O, -qoptimize" on page 269	#pragma options optimize	Specifies whether to optimize code during compilation and, if so, at which level.
"-p, -pg, -qprofile" on page 279	None.	Prepares the object files produced by the compiler for profiling.
"-qpdf1, -qpdf2" on page 283	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 293	None.	Inserts prefetch instructions automatically where there are opportunities to improve code performance.
"-qprocimported, -qproclocal, -qprocunknown" on page 298	#pragma options procimported, #pragma options proclocal, #pragma options procunknown	Marks functions as local, imported, or unknown in 64-bit compilations.
"-qinline" on page 197	None.	Attempts to inline functions instead of generating calls to those functions, for improved performance.

Table 18. Optimization and tuning options (continued)

Option name	Equivalent pragma name	Description
"-qshowpdf" on page 316	None.	When used with <b>-qpdf1</b> and a minimum optimization level of <b>-O2</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 316	#pragma nosimd	Controls whether the compiler can automatically take advantage of vector instructions for processors that support them.
"-qsmallstack" on page 319	None.	Reduces the size of the stack frame.
"-qsmp" on page 320	None.	Enables parallelization of program code.
"-qspeculateabsolutes" on page 326	None.	Works with the <b>-qtocmerge -bl:file</b> for non-IPA links and with the <b>-bl:file</b> for IPA links to disable speculation at absolute addresses.
"-qstrict" on page 333	#pragma options strict	Ensures that optimizations done by default at optimization levels <b>-O3</b> and higher, and, optionally at <b>-O2</b> , do not alter the semantics of a program.
"-qstrict_induction" on page 337	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.
"-qtocmerge" on page 355	None.	Enables TOC merging to reduce TOC pointer loads and improves the scheduling of external loads.
"-qtune" on page 357	#pragma options tune	Tunes instruction selection, scheduling, and other architecture-dependent performance enhancements to run best on a specific hardware architecture.
"-qunroll" on page 363	#pragma options unroll, #pragma unroll	Controls loop unrolling, for improved performance.
"-qunwind" on page 365	None.	Specifies whether the call stack can be unwound by code looking through the saved registers on the stack.

## Linking

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

Table 19. Linking options

Option name	Equivalent pragma name	Description
"-b" on page 115	None.	Controls how shared objects are processed by the linker.
"-bmaxdata" on page 117	None.	Sets the maximum size of the area shared by the static data (both initialized and uninitialized) and the heap.
"-brtl" on page 118	None.	Controls runtime linking for the output file.
"-qcr" on page 129	None.	Specifies whether system startup files are to be linked.
"-e" on page 141	None.	When used together with the <b>-qmkshrobj</b> or <b>-G</b> option, specifies an entry point for a shared object.
"-f" on page 150	None.	Names a file that stores a list of object files for the compiler to pass to the linker.
"-L" on page 214	None.	Searches the directory path for library files specified by the <b>-I</b> option.
"-l" on page 213	None.	Searches for the specified library file, <i>libkey.so</i> , and then <i>libkey.a</i> for dynamic linking, or just for <i>libkey.a</i> for static linking.
"-qlib" on page 240	None.	Specifies whether standard system libraries and XL C/C++ libraries are to be linked.
"-qtwolink (C++ only)" on page 359	None.	Minimizes the number of static constructors included from libraries and object files.
"-Z" on page 382	None.	Specifies a prefix for the library search path to be used by the linker.

## 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 20. Portability and migration options

Option name	Equivalent pragma name	Description
"-qalign" on page 100	#pragma options align, #pragma align	Specifies the alignment of data objects in storage, which avoids performance problems with misaligned data.
"-qalignrulefor (C++ only)" on page 103	None.	When used with <b>-qalign=power</b> , determines whether a 4-byte alignment ceiling is applied to non-first members of structures that are of type typedef to array of element type that exceeds the alignment ceiling.
"-qgenproto (C only)" on page 172	None.	Produces prototype declarations from K&R function definitions or function definitions with empty parentheses, and displays them to standard output.
"-qnamemangling (C++ only)" on page 262	#pragma namemangling	Chooses the name mangling scheme for external symbol names generated from C++ source code.
"-qobjmodel (C++ only)" on page 272	#pragma object_model	Sets the object model to be used for structures, unions, and classes.
"-qoldpassbyvalue (C++ only)" on page 274	#pragma pass_by_value	Specifies how classes containing const or reference members are passed in function arguments.
"-qupconv (C only)" on page 366	#pragma options upconv	Specifies whether the unsigned specification is preserved when integral promotions are performed.
"-qvecnv" on page 369	None.	Specifies whether to use volatile or non-volatile vector registers.

## Compiler customization

The options in this category 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 21. Compiler customization options

Option name	Equivalent pragma name	Description
"-qasm_as" on page 112	None.	Specifies the path and flags used to invoke the assembler in order to handle assembler code in an asm assembly statement.
"-B" on page 116	None.	Determines substitute path names for XL C/C++ executables such as the compiler, assembler, linker, and preprocessor.
"-qc_stdinc (C only)" on page 130	None.	Changes the standard search location for the XL C and system header files.

Table 21. Compiler customization options (continued)

Option name	Equivalent pragma name	Description
“-qcpp_stdinc (C++ only)” on page 131	None.	Changes the standard search location for the XL C++ and system header files.
“-F” on page 150	None.	Names an alternative configuration file or stanza for the compiler.
“-qpath” on page 281	None.	Determines substitute path names for XL C/C++ executables such as the compiler, assembler, linker, and preprocessor.
“-qoptfile” on page 277	None.	Specifies a file containing a list of additional command line options to be used for the compilation.
“-qspill” on page 327	#pragma options spill	Specifies the size (in bytes) of the register spill space, the internal program storage areas used by the optimizer for register spills to storage.
“-t” on page 342	None.	Applies the prefix specified by the <b>-B</b> option to the designated components.
“-W” on page 372	None.	Passes the listed options to a component that is executed during compilation.

## Deprecated options

The compiler still accepts options listed in the following table. Options without an asterisk have been replaced by other options that provide the same functionality. Options with an asterisk are obsolete, or can produce unexpected results and are not guaranteed to perform as previously documented. Use with discretion.

Table 22. Deprecated options

Option name	Replacement option
-Q	-qinline
-qansialias	-qalias=ansi
-qarch = 601   602   603   pwr   pwr2   p2sc   pwr2s   com	-qfloat=nosingle:norndsngl
-qassert	-qalias
-qenablevmx	-qsimd
-qfloat=emulate*	
-qfold	-qfloat=fold
-qhsflt	-qfloat=hsflt
-qhssngl	-qfloat=hssngl
-qhot=simd   nosimd	-qsimd
-qinfo=private	-qreport

Table 22. Deprecated options (continued)

Option name	Replacement option
-qinfo=reduction	-qreport
-qipa=clonearch   noclonearch	-qtune
-qipa=cloneproc   nocloneproc	-qtune
-qipa=inline   noinline	-qinline
-qipa=pdfname	-qpdf1=pdfname, -qpdf2=pdfname
-qlanglvl=[no]gnu_externtemplate	-qlanglvl=[no]externtemplate
-qmaf	-qfloat=maf
-qrrm	-qfloat=rrm
-qspnans	-qfloat=spnans

## Individual option descriptions

This section contains descriptions of the individual compiler options available in XL C/C++.

For each option, the following information is provided:

### Category

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

### Pragma equivalent

Many compiler options allow you to use an equivalent pragma 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 even selected sections of code. Where an option supports the **#pragma options** *option\_name* and/or **#pragma name** form of the directive, this is indicated.

### Purpose

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

### Syntax

This section provides the syntax for the option, and where an equivalent **#pragma name** is supported, the specific syntax for the pragma. Syntax for **#pragma options** *option\_name* forms of the pragma is not provided, as this is normally identical to that of the option. Note that you can also use the C99-style `_Pragma` operator form of any pragma; although this syntax is not provided in the option descriptions. For complete details on pragma syntax, see "Pragma directive syntax" on page 383

### 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 and pragma equivalents, where applicable. For suboptions that are specific to the command-line option or to the pragma directive, this is indicated in the descriptions.

**Usage** This section describes any rules or usage considerations you should be



aware of when using the option. These can include restrictions on the option's applicability, valid placement of pragma directives, precedence rules for multiple option specifications, and so on.

#### **Predefined macros**

Many compiler options set macros that are protected (that is, cannot be undefined or redefined by the user). Where applicable, any macros that are predefined by the option, and the values to which they are defined, are listed in this section. A reference list of these macros (as well as others that are defined independently of option setting) is provided in Chapter 6, "Compiler predefined macros," on page 463

#### **Examples**

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

## **-+ (plus sign) (C++ only)**

### **Category**

Input control

### **Pragma equivalent**

None.

### **Purpose**

Compiles any file as a C++ language file.

This option is equivalent to the **-qsource=c++** option.

### **Syntax**

►► -+ ◄◄

### **Usage**

You can use **-+** to compile a file with any suffix other than **.a**, **.o**, **.so**, **.S** or **.s**. If you do not use the **-+** option, files must have a suffix of **.C** (uppercase C), **.cc**, **.cp**, **.cpp**, **.cxx**, or **.c++** to be compiled as a C++ file. If you compile files with suffix **.c** (lowercase c) without specifying **-+**, the files are compiled as a C language file.

The **-+** option should not be used together with the **-qsource** option.

### **Predefined macros**

None.

### **Examples**

To compile the file `myprogram.cplsp1s` as a C++ source file, enter:

```
xlc++ -+ myprogram.cplsp1s
```

### **Related information**

- "**-qsource**" on page 325

## **-# (pound sign)**

### **Category**

Error checking and debugging

### **Pragma equivalent**

None.

### **Purpose**

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

When this option is enabled, information is written to standard output, showing the names of the programs within the preprocessor, compiler, and linker that would be invoked, and the default options that would be specified for each program. The preprocessor, compiler, and linker are not invoked.

### **Syntax**

▶▶ -# —————▶▶

### **Usage**

You can use this command to determine the commands and files that will be involved in a particular compilation. It avoids the overhead of compiling the source code and overwriting any existing files, such as .lst files.

This option displays the same information as `-v`, but does not invoke the compiler. The `-#` option overrides the `-v` option.

### **Predefined macros**

None.

### **Examples**

To preview the steps for the compilation of the source file `myprogram.c`, enter:

```
xlc myprogram.c -#
```

### **Related information**

- “`-v`, `-V`” on page 368

## **-q32, -q64**

### **Category**

Object code control

### **Pragma equivalent**

None.

## Purpose

Selects either 32-bit or 64-bit compiler mode.

Use the `-q32` and `-q64` options, along with the `-qarch` and `-qtune` compiler options, to optimize the output of the compiler to the architecture on which that output will be used.

## Syntax

►► -q 32  
64 \_\_\_\_\_ ►►

## Defaults

`-q32`

## Usage

The `-q32` and `-q64` options override the compiler mode set by the value of the `OBJECT_MODE` environment variable, if it exists.

## Predefined macros

When `-q64` is in effect, `__64BIT__` is defined to 1; otherwise, it is undefined.

## Examples

To specify that the executable program `testing` compiled from `myprogram.c` is to run on a computer with a 64-bit Power architecture, enter:

```
xlc -o testing myprogram.c -q64 -qarch=ppc
```

## Related information

- Specifying compiler options for architecture-specific compilation
- “`-qarch`” on page 106
- “`-qtune`” on page 357

## **-qaggrcopy**

### Category

Optimization and tuning

### Pragma equivalent

None.

### Purpose

Enables destructive copy operations for structures and unions.

## Syntax

►► -q-aggrcopy=nooverlap | overlap ◀◀

## Defaults

-qaggrcopy=nooverlap

## Parameters

**overlap** | **nooverlap**

**nooverlap** assumes that the source and destination for structure and union assignments do not overlap, allowing the compiler to generate faster code. **overlap** inhibits these optimizations.

## Predefined macros

None.

## -qalias

### Category

Optimization and tuning

### Pragma equivalent

None

### Purpose


Indicates whether a program contains certain categories of aliasing or does not conform to C/C++ 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.

## Syntax

►► -qalias=: | notypeptr | restrict | global | noallptrs | ansi | noaddrtaken | addrtaken | noansi | allptrs | noglobal | norestrict | typeptr ◀◀

## Defaults

- C++ -qalias=noaddrtaken:noallptrs:ansi:global:restrict:notypeptr

-  **-qalias=noaddrtaken:noallptrs:ansi:global:restrict:notypeptr** for all invocation commands except **cc**.  
**-qalias=noaddrtaken:noallptrs:noansi:global:restrict:notypeptr** for the **cc** invocation command.

## Parameters

### **addrtaken** | **noaddrtaken**

When **addrtaken** is in effect, variables are disjoint from pointers unless their address is taken. Any class of variable for which an address has *not* been recorded in the compilation unit will be considered disjoint from indirect access through pointers.

When **noaddrtaken** is specified, the compiler generates aliasing based on the aliasing rules that are in effect.

### **allptrs** | **noallptrs**

When **allptrs** is in effect, pointers are never aliased (this also implies **-qalias=typeptr**). Specifying **allptrs** is an assertion to the compiler that no two pointers point to the same storage location. These suboptions are only valid if **ansi** is also specified.

### **ansi** | **noansi**

When **ansi** is in effect, type-based aliasing is used during optimization, which restricts the lvalues that can be safely used to access a data object. The optimizer assumes that pointers can *only* point to an object of the same type. This suboption has no effect unless you also specify an optimization option.

When **noansi** is in effect, the optimizer makes worst case aliasing assumptions. It assumes that a pointer of a given type can point to an external object or any object whose address is already taken, regardless of type.

### **global** | **noglobal**

When **global** is in effect, type-based aliasing rules are enabled during IPA link-time optimization across compilation units. Both **-qipa** and **-qalias=ansi** must be enabled for **-qalias=global** to have an effect. Specifying **noglobal** disables type-based aliasing rules.

**-qalias=global** produces better performance at higher optimization levels and also better link-time performance. If you use **-qalias=global**, it is recommended that you compile as much as possible of the application with the same version of the compiler to maximize the effect of the suboption on performance.

### **restrict** | **norestrict**

When **restrict** is in effect, optimizations for pointers qualified with the **restrict** keyword are enabled. Specifying **norestrict** disables optimizations for **restrict**-qualified pointers.

**-qalias=restrict** is independent from other **-qalias** suboptions. Using the **-qalias=restrict** option will usually result in performance improvements for code that uses **restrict**-qualified pointers. Note, however, that using **-qalias=restrict** requires that restricted pointers be used correctly; if they are not, compile-time and runtime failures may result. You can use **norestrict** to preserve compatibility with code compiled with versions of the compiler previous to V9.0.

### **typeptr** | **notypeptr**

When **typeptr** is in effect, pointers to different types are never aliased. Specifying **typeptr** is an assertion to the compiler that no two pointers of different types point to the same storage location. These suboptions are only valid if **ansi** is also specified.

## Usage

`-qalias` makes assertions to the compiler about the code that is being compiled. If the assertions about the code are false, then the code generated by the compiler may result in unpredictable behaviour when the application is run.

The following are not subject to type-based aliasing:

- Signed and unsigned types. For example, a pointer to a signed `int` can point to an unsigned `int`.
- Character pointer types can point to any type.
- Types qualified as `volatile` or `const`. For example, a pointer to a `const int` can point to an `int`.

The `-qalias=[no]ansi` option replaces the deprecated `-q[no]ansialias` option. Use `-qalias=[no]ansi` in your new applications.

## Predefined macros

None.

## Examples

To specify worst-case aliasing assumptions when compiling `myprogram.c`, enter:

```
xlc myprogram.c -O -qalias=noansi
```

## Related information

- “`-qipa`” on page 199
- `-qinfo=als`
- “`#pragma disjoint`” on page 395
- *Type-based aliasing* in the *XL C/C++ Language Reference*
- *The restrict type qualifier* in the *XL C/C++ Language Reference*

## `-qalign`

### Category

Portability and migration

### Pragma equivalent


`#pragma options align`, `#pragma align`

### Purpose


Specifies the alignment of data objects in storage, which avoids performance problems with misaligned data.

### Syntax

#### Option syntax



## Pragma syntax



## Defaults

`-qalign=power`

## Parameters

### `bit_packed` | `packed`

Bit field data is packed on a bitwise basis without respect to byte boundaries.

### `power`

Uses the RISC System/6000 alignment rules. This is the default.

### `full`

Uses the RISC System/6000 alignment rules.

**Note:** `-qalign=full` is equivalent to `-qalign=power`.

### `mac68k` | `twobyte`

Uses the Macintosh alignment rules. Valid only for 32-bit compilations.

### `natural`

Structure members are mapped to their natural boundaries. This has the same effect as the `power` suboption, except that it also applies alignment rules to double and long double members that are not the first member of a structure or union.

### `reset (pragma only)`

Discards the current pragma setting and reverts to the setting specified by the previous pragma directive. If no previous pragma was specified, reverts to the command-line or default option setting.

## Usage

If you use the `-qalign` option more than once on the command line, the last alignment rule specified applies to the file.

The `full` suboption is the default to ensure compatibility with existing objects. If compatibility with earlier versions is not necessary, you should consider using `natural` alignment to improve potential application performance.

The pragma directives override the **-qalign** compiler option setting for a specified section of program source code. The pragmas affect all aggregate definitions that appear after a given pragma directive; if a pragma is placed inside a nested aggregate, it applies only to the definitions that follow it, not to any containing definitions. Any aggregate variables that are declared use the alignment rule that applied at the point at which the aggregate was *defined*, regardless of pragmas that precede the declaration of the variables. See below for examples.

**Note:** When using **-qalign**, all system headers are also compiled with **-qalign**. For a complete explanation of the option and pragma parameters, as well as usage considerations, see "Aligning data" in the *XL C/C++ Optimization and Programming Guide*.

## Predefined macros

None.

## Examples

The following examples show the interaction of the option and pragmas. Assuming compilation with the command `xlc file2.c`, the following example shows how the pragma affects only an aggregate *definition*, not subsequent declarations of variables of that aggregate type.

```
/* file2.c The default alignment rule is in effect */

typedef struct A A2;

#pragma options align=bit_packed /* bit_packed alignment rules are now in effect */
struct A {
  int a;
  char c;
}; #pragma options align=reset /* Default alignment rules are in effect again */

struct A A1; /* A1 and A3 are aligned using bit_packed alignment rules since */
A2 A3;      /* this rule applied when struct A was defined */
```

Assuming compilation with the command `xlc file.c -qalign=bit_packed`, the following example shows how a pragma embedded in a nested aggregate definition affects only the definitions that follow it.

```
/* file2.c The default alignment rule in effect is bit_packed */

struct A {
  int a;
#pragma options align=power /* Applies to B; A is unaffected */
  struct B {
    char c;
    double d;
  } BB; /* BB uses power alignment rules */
} AA; /* AA uses bit_packed alignment rules */
```

## Related information

- “#pragma pack” on page 424
- "Aligning data" in the *XL C/C++ Optimization and Programming Guide*
- "The `__align` specifier" in the *XL C/C++ Language Reference*
- "The aligned variable attribute" in the *XL C/C++ Language Reference*
- "The packed variable attribute" in the *XL C/C++ Language Reference*



## -qalignrulefor (C++ only)

### Category

Portability and migration

### Pragma equivalent

None

### Purpose

When used with **-qalign=power**, determines whether a 4-byte alignment ceiling is applied to non-first members of structures that are of type typedef to array of element type that exceeds the alignment ceiling.

### Syntax

```
►► -q-alignrulefor=power==typedefrespectsrule | nottypedefrespectsrule ►►
```

### Defaults

-qalignrulefor=power=typedefrespectsrule

### Parameters

#### typedefrespectsrule | nottypedefrespectsrule

When **typedefrespectsrule** is in effect, the member follows the normal alignment rules for **-qalign=power**. This suboption provides compatibility with code compiled with **-qalign=power** with XL C++ V6.0 and earlier.

For XL C++ V9.0, the default is **typedefrespectsrule**.

When **nottypedefrespectsrule** is in effect, a member that exceeds the alignment ceiling of 4 bytes is aligned on 4-byte boundaries. This suboption provides compatibility with code compiled with **-qalign=power** with XL C++ V7.0 and V8.0.

### Predefined macros

None.

### Examples

The following example uses a typedef declaration for an array of structures containing a member of long long type, which is not normally subject to a 4-byte alignment ceiling, and then uses the typedef as the non-first member of a structure variable declaration. The table shows the differing alignment results depending on the setting of the **-qalignrulefor=power** option.

Sample code	Alignment results	
	typedefrespectsrule	nottypedefrespectsrule
<pre> struct A {     long long a1; } a;  typedef struct A ten_A[10];  struct B {     char dummy[116];     struct A ten_a[10]; } b;  struct C {     char dummy[116];     ten_A ten_a; } c; </pre>	<pre> alignment of b.ten_a = 8 alignment of b = 8 alignment of c.ten_a = 8 alignment of c = 8 </pre>	<pre> alignment of b.ten_a = 8 alignment of b = 8 alignment of c.ten_a = 4 alignment of c = 4 </pre>

### Related information

- “-qalign” on page 100

## -qalloca, -ma (C only)

### Category

Object code control

### Pragma equivalent

#pragma alloca

### Purpose

Provides an inline definition of system function `alloca` when it is called from source code that does not include the `alloca.h` header.

The function `void* alloca(size_t size)` dynamically allocates memory, similarly to the standard library function `malloc`. The compiler automatically substitutes calls to the system `alloca` function with an inline built-in function `__alloca` in any of the following cases:

- You include the header file `alloca.h`
- You compile with **-Dalloca=\_\_alloca**
- You directly call the built-in function using the form `__alloca`

The **-qalloca** and **-ma** options and **#pragma alloca** directive provide the same functionality in C only, if any of the above methods are not used.

### Syntax

#### Option syntax



## Pragma syntax

▶ `#pragma alloca` ◀

## Defaults

Not applicable.

## Usage

If you do not use any of the above-mentioned methods to ensure that calls to `alloca` are replaced with `__alloca`, `alloca` is treated as a user-defined identifier rather than as a built-in function.

Once specified, `#pragma alloca` applies to the rest of the file and cannot be disabled. If a source file contains any functions that you want compiled without `#pragma alloca`, place these functions in a different file.

You may want to consider using a C99 variable length array in place of `alloca`.

## Predefined macros

None.

## Examples

To compile `myprogram.c` so that calls to the function `alloca` are treated as inline, enter:

```
xlc myprogram.c -qalloca
```

## Related information

- “-D” on page 132
- “\_\_alignx” on page 603

## -qaltivec

### Category

Language element control

### Pragma equivalent

None.

### Purpose

Enables compiler support for vector data types and operators.

See the *XL C/C++ Language Reference* for complete documentation of vector data types.

## Syntax

►► -q noaltivec  
altivec ◀◀

## Defaults

-qnoaltivec

## Usage

This option has effect only when you set or imply **-qarch** to be an architecture that supports vector instructions. Otherwise, the compiler ignores **-qaltivec** and issues a warning message.

## Predefined macros

`__ALTIVEC__` is defined to 1 and `__VEC__` is defined to 10205 when **-qaltivec** is in effect; otherwise, they are undefined.

## Examples

To enable compiler support for vector programming, enter the following command:

```
xlc myprogram.c -qarch=ppc64v -qaltivec
```

## Related information

- “-qarch”
- “-qsimd” on page 316
- “-qvecnvol” on page 369
- *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


### Pragma equivalent

None.

### Purpose

Specifies the processor architecture for which the code (instructions) should be generated.

### Syntax



## Defaults

- `-qarch=ppc` when `-q32` is in effect.
- `-qarch=ppc64` when `-q64` is in effect.
- `-qarch=auto` when `-O4` or `-O5` 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 `-O4` or `-O5` option is set or implied.

### 403

Produces object code containing instructions that will run on the PowerPC 403 hardware platform.

### 604

Produces object code containing instructions that will run on the PowerPC 604 hardware platform. This suboption is not valid if `-q64` is in effect.

### pwr3

Produces object code containing instructions that will run on the POWER3, POWER4, POWER5, POWER5+, POWER6®, POWER7®, or PowerPC® 970 hardware platforms.

### pwr4

Produces object code containing instructions that will run on the POWER4, POWER5, POWER5+, POWER6, POWER7, or PowerPC 970 hardware platforms.

### pwr5

Produces object code containing instructions that will run on the POWER5, POWER5+, POWER6, or POWER7 hardware platforms.

### pwr5x

Produces object code containing instructions that will run on the POWER5+, POWER6, or POWER7 hardware platforms.

**pwr6**

Produces object code containing instructions that will run on the POWER6 or POWER7 hardware platforms running in POWER6 or POWER7 architected mode. If you would like support for decimal floating-point instructions, be sure to specify this suboption during compilation.

**pwr6e**

Produces object code containing instructions that will run on the POWER6 hardware platforms running in POWER6 enhanced mode.

**pwr7**

Produces object code containing instructions that will run on the POWER7 hardware platforms.

**ppc**

In 32-bit mode, produces object code containing instructions that will run on any of the 32-bit PowerPC hardware platforms. This suboption causes the compiler to produce single-precision instructions to be used with single-precision data. Specifying **-qarch=ppc** together with **-q64** silently upgrades the architecture setting to **-qarch=ppc64**.

**ppc64**

Produces object code that will run on any of the 64-bit PowerPC hardware platforms. This suboption can be selected when compiling in 32-bit mode, but the resulting object code may include instructions that are not recognized or behave differently when run on 32-bit PowerPC platforms.

**ppcgr**

In 32-bit mode, produces object code for PowerPC processors that support optional graphics instructions. Specifying **-qarch=ppcgr** together with **-q64** silently upgrades the architecture setting to **-qarch=ppc64gr**.

**ppc64gr**

Produces code for any 64-bit PowerPC hardware platform that supports optional graphics instructions.

**ppc64grsq**

Produces code for any 64-bit PowerPC hardware platform that supports optional graphics and square root instructions.

**ppc64v**

Generates instructions for generic PowerPC chips with vector processors, such as the PowerPC 970. Valid in 32-bit or 64-bit mode.

**ppc970**

Generates instructions specific to the PowerPC 970 architecture.

**rs64a**

Produces object code that will run on RS64I platforms.

**rs64b**

Produces object code that will run on RS64II platforms.

**rs64c**

Produces object code that will run on RS64III platforms.

**Note:** The **com** suboption, and suboptions representing POWER<sup>®</sup> and POWER2 architectures, are no longer supported. If you would like similar floating-point behavior to that provided by this suboption, use the **-qfloat=nosingle:norndsnegl** option. See “-qfloat” on page 154 for details.

## Usage

All PowerPC machines share a common set of instructions, but may also include additional instructions unique to a given processor or processor family. Using the **-qarch** option to target a specific architecture for the compilation results in code that may not run on other architectures, but provides the best performance for the selected architecture. If you want maximum performance on a specific architecture and will not be using the program on other architectures, use the appropriate architecture option. If you want to generate code that can run on more than one architecture, specify a **-qarch** suboption that supports a group of architectures. Table 23 shows the features supported by the different processor architectures and their representative **-qarch** suboptions:

Table 23. Feature support in processor architectures

Architecture	Graphics support	Square root support	64-bit support	Vector processing support	Large page support
604	yes	no	no	no	no
rs64a	no	no	yes	no	no
rs64b	yes	yes	yes	no	no
rs64c	yes	yes	yes	no	no
pwr3	yes	yes	yes	no	no
pwr4	yes	yes	yes	no	yes
pwr5	yes	yes	yes	no	yes
pwr5x	yes	yes	yes	no	yes
ppc	no	no	no	no	yes
ppc64	no	no	yes	no	yes
ppc64gr	yes	no	yes	no	yes
ppc64grsq	yes	yes	yes	no	yes
ppc64v	yes	yes	yes	VMX	yes
ppc970	yes	yes	yes	VMX	yes
pwr6	yes	yes	yes	VMX	yes
pwr6e	yes	yes	yes	VMX	yes
pwr7	yes	yes	yes	VMX, VSX	yes

**Note:** Vector Multimedia Extension (VMX) and Vector Scalar Extension (VSX) are processor instructions for vector processing.

For any given **-qarch** setting, the compiler defaults to a specific, matching **-qtune** setting, which can provide additional performance improvements. Alternatively, if you specify **-qarch** with a group argument, you can specify **-qtune** as either **auto** or provide a specific architecture in the group. For detailed information on using **-qarch** and **-qtune** together, see “-qtune” on page 357.

Specifying **-q64** changes the effective **-qarch** setting as follows:

Original <b>-qarch</b> setting	Effective setting when <b>-q64</b> is specified
ppc	ppc64
ppcgr	ppc64gr

For a given application program, make sure that you specify the same **-qarch** setting when you compile each of its source files. Although the linker and loader may detect object files that are compiled with incompatible **-qarch** settings, you should not rely on it.

## Predefined macros

See “Macros related to architecture settings” on page 469 for a list of macros that are predefined by `-qarch` suboptions.

## Examples

To specify that the executable program `testing` compiled from `myprogram.c` is to run on a computer with a 32-bit PowerPC architecture, enter:

```
xlc -o testing myprogram.c -q32 -qarch=ppc
```

## Related information

- “Specifying compiler options for architecture-specific compilation” on page 9
- `-qfloat`
- `-qprefetch`
- “`-qtune`” on page 357
- “`-q32`, `-q64`” on page 96
- “Macros related to architecture settings” on page 469
- “Optimizing your applications” in the *XL C/C++ Optimization and Programming Guide*

## `-qasm`

### Category

Language element control

### Pragma equivalent

None.

### Purpose

Controls the interpretation and subsequent generation of code for assembler language extensions.

When `-qasm` is in effect, the compiler generates code for assembly statements in the source code. Suboptions specify the syntax used to interpret the content of the assembly statement.

**Note:** The system assembler program must be available for this command to have effect.

### Syntax

`-qasm syntax — C`



`-qasm syntax — C++`





## Defaults

- ▶ **C** `-qasm=gcc`
- ▶ **C++** `-qasm=gcc` at all language levels except **compat366** or **strict98**.  
`-qnoasm=stdcpp` when `-qlanglvl=compat366` or `-qlanglvl=strict98` is in effect.

`-qasm=gcc`

## Parameters

### gcc

Instructs the compiler to recognize the extended GCC syntax and semantics for assembly statements.

### ▶ **C++** `stdcpp`

Reserved for possible future use.

Specifying `-qasm` without a suboption is equivalent to specifying the default.

## Usage

▶ **C** The token `asm` is not a C language keyword. Therefore, at language levels **stdc89** and **stdc99**, which enforce strict compliance to the C89 and C99 standards, respectively, the option `-qkeyword=asm` must also be specified to compile source that generates assembly code. At all other language levels, the token `asm` is treated as a keyword unless the option `-qnokeyword=asm` is in effect. In C, the compiler-specific variants `__asm` and `__asm__` are keywords at all language levels and cannot be disabled.

▶ **C++** The tokens `asm`, `__asm`, and `__asm__` are keywords at all language levels. Suboptions of `-qnokeyword=token` can be used to disable each of these reserved words individually.

For detailed information on the syntax and semantics of inline `asm` statements, see "Inline assembly statements" in the *XL C/C++ Language Reference*.

## Predefined macros

- ▶ **C** `__IBM_GCC_ASM` is predefined to 1 when `asm` is recognized as a keyword and assembler code is generated; that is, at all language levels except **stdc89** | **stdc99**, or when `-qkeyword=asm` is in effect, and when `-qasm[=gcc]` is in effect. It is predefined to 0 when `asm` is recognized as a keyword but assembler code is not generated; that is, at all language levels except **stdc89** | **stdc99**, or when `-qkeyword=asm` is in effect, and when `-qnoasm` is in effect. It is undefined when the **stdc89** | **stdc99** language level or `-qnokeyword=asm` is in effect.
- ▶ **C++** `__IBM_GCC_ASM` is predefined to 1 when `asm` is recognized as a keyword and assembler code is generated; that is, at all language levels except **compat366** | **strict98**, and when `-qasm[=gcc]` is in effect. It is predefined to 0 when `asm` is recognized as a keyword but assembler code is not generated; that

is, at all language levels except `compat366` | `strict98`, and when `-qnoasm[=gcc]` is in effect. It is undefined when `-qlanglvl=compat366` | `strict98` or `-qnoasm=stdcpp` is in effect. `__IBM_STDCPP_ASM` is predefined to 0 when `-qlanglvl=compat366` | `strict98` or `-qnoasm=stdcpp` is in effect; otherwise it is undefined.

## Examples

The following code snippet shows an example of the GCC conventions for `asm` syntax in inline statements:

```
int a, b, c;
int main() {
    asm("add %0, %1, %2" : "=r"(a) : "r"(b), "r"(c) );
}
```

## Related information

- “`-qasm_as`”
- “`-qlanglvl`” on page 215
- “`-qkeyword`” on page 211
- "Inline assembly statements" in the *XL C/C++ Language Reference*
- "Keywords for language extensions"

## **-qasm\_as**

### Category

Compiler customization

### Pragma equivalent

None.

### Purpose

Specifies the path and flags used to invoke the assembler in order to handle assembler code in an `asm` assembly statement.

Normally the compiler reads the location of the assembler from the configuration file; you can use this option to specify an alternate assembler program and flags to pass to that assembler.

### Syntax

▶▶ `-qasm_as=` path `"` path `"` flags

### Defaults

By default, the compiler invokes the assembler program defined for the `as` command in the compiler configuration file.

### Parameters

#### *path*

The full path name of the assembler to be used.

### *flags*

A space-separated list of options to be passed to the assembler for assembly statements. Quotation marks must be used if spaces are present.

### Predefined macros

None.

### Examples

To instruct the compiler to use the assembler program at `/bin/as` when it encounters inline assembler code in `myprogram.c`, enter:

```
xlc myprogram.c -qasm_as=/bin/as
```

To instruct the compiler to pass some additional options to the assembler at `/bin/as` for processing inline assembler code in `myprogram.c`, enter:

```
xlc myprogram.c -qasm_as="/bin/as -a64 -l a.lst"
```

### Related information

- “-qasm” on page 110

## -qassert

### Category

Optimization and tuning

### Pragma equivalent


None

### Purpose

Provides information about the characteristics of the files that can help to fine-tune optimizations.

### Syntax

#### Option:



### Defaults

-qassert=norefalign

### Parameters

#### **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`.

## **-qattr**

### **Category**

Listings, messages, and compiler information

### **Pragma equivalent**

`#pragma options [no]attr`

### **Purpose**

Produces a compiler listing that includes the attribute component of the attribute and cross-reference section of the listing.

### **Syntax**



### **Defaults**

`-qnoattr`

### **Parameters**

#### **full**

Reports all identifiers in the program. If you specify `attr` without this suboption, only those identifiers that are used are reported.

### **Usage**

If `-qattr` is specified after `-qattr=full`, it has no effect; the full listing is produced.

This option does not produce a cross-reference listing unless you also specify `-qxref`.

The `-qnoprint` option overrides this option.

**Note:** Specifying `-qattr` does not list the `#define` directives. You can use `"-qshowmacros"` on page 315 instead.

### **Predefined macros**

None.

### **Examples**

To compile the program `myprogram.c` and produce a compiler listing of all identifiers, enter:

```
xlc myprogram.c -qxref -qattr=full
```

## Related information

- “-qshowmacros” on page 315
- “-qprint” on page 296
- “-qxref” on page 379

## -b

### Category

Linking

### Pragma equivalent

None.

### Purpose

Controls how shared objects are processed by the linker.

### Syntax



### Defaults

-bdynamic

### Parameters

#### dynamic | shared

Causes the linker to process subsequent shared objects in dynamic mode. In dynamic mode, shared objects are not statically included in the output file. Instead, the shared objects are listed in the loader section of the output file.

#### static

Causes the linker to process subsequent shared objects in static mode. In static mode, shared objects are statically linked in the output file.

### Usage

The default option, **-bdynamic**, ensures that the C library (libc) links dynamically. To avoid possible problems with unresolved linker errors when linking the C library, you must add the **-bdynamic** option to the end of any compilation sections that use the **-bstatic** option.

### Predefined macros

Not applicable.

### Related information

- “-brtl” on page 118

## -B

### Category

Compiler customization

### Pragma equivalent

None.

### Purpose

Determines substitute path names for XL C/C++ executables such as the compiler, assembler, linker, and preprocessor.

You can use this option if you want to keep multiple levels of some or all of the XL C/C++ executables and have the option of specifying which one you want to use. However, it is recommended that you use the **-qpath** option to accomplish this instead.

### Syntax

►► -B prefix ◄◄

### Defaults

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

### Parameters

*prefix*

Defines part of a path name for programs you can name with the **-t** option. You must add a slash (/). If you specify the **-B** option without the *prefix*, the default prefix is `/lib/o`.

### Usage

The **-t** option specifies the programs to which the **-B** prefix name is to be appended; see “-t” on page 342 for a list of these. If you use the **-B** option without **-tprograms**, the prefix you specify applies to all of the compiler executables.

The **-B** and **-t** options override the **-F** option.

### Predefined macros

None.

### Examples

In this example, an earlier level of the compiler components is installed in the default installation directory. To test the upgraded product before making it available to everyone, the system administrator restores the latest installation image under the directory `/home/jim` and then tries it out with commands similar to:

```
xlc -tcbI -B/home/jim/usr/vacpp/bin/ test_suite.c
```

Once the upgrade meets the acceptance criteria, the system administrator installs it in the default installation directory.

### Related information

- “-qpath” on page 281
- “-t” on page 342
- “Invoking the compiler” on page 1

## -qbitfields

### Category

Floating-point and integer control

### Pragma equivalent

None.

### Purpose

Specifies whether bit fields are signed or unsigned.

### Syntax

►► -qbitfields=signed | unsigned ◀◀

### Defaults

-qbitfields=unsigned

### Parameters

#### **signed**

Bit fields are signed.

#### **unsigned**

Bit fields are unsigned.

### Predefined macros

None.

## -bmaxdata

### Category

Linking

### Pragma equivalent

None

## Purpose

Sets the maximum size of the area shared by the static data (both initialized and uninitialized) and the heap.

## Syntax

►► — `-bmaxdata—:—number—` —————►►

## Defaults

`-bmaxdata:0`

## Parameters

*number*

The number of bytes used representing the soft **ulimit** set by the system loader. Valid values are 0 and multiples of 0x10000000 (0x10000000, 0x20000000, 0x30000000, ...). The maximum value allowed by the system is 0x80000000. If the value is 0, a single 256MB (0x10000000 byte) data segment (segment 2) will be shared by the static data, the heap, and the stack. If the value is non-zero, a data area of the specified size (starting in segment 3) will be shared by the static data and the heap, while a separate 256 MB data segment (segment 2) will be used by the stack. So, the total data size when 0 is specified is 256MB, and the total size when 0x10000000 is specified is 512MB, with 256MB for the stack and 256MB for static data and the heap.

## Predefined macros

None.

## **-brtl**

### Category

Linking

### Pragma equivalent

None.

### Purpose

Controls runtime linking for the output file.

Runtime linking is the ability to resolve undefined and non-deferred symbols in shared modules after the program execution has already begun. It is a mechanism for providing runtime definitions (these function definitions are not available at link-time) and symbol rebinding capabilities. Compiling with `-brtl` adds a reference to the runtime linker to your program, which will be called by your program's startup code (`/lib/crt0.o`) when program execution begins. Shared object input files are listed as dependents in the program loader section in the same order as they are specified on the command line. When the program execution begins, the system loader loads these shared objects so their definitions are available to the runtime linker.



## Syntax

►► -brtl ◀◀

## Usage

The main application must be built to enable runtime linking. The system loader must be able to load and resolve all symbols referenced in the main program and called modules, or the program will not execute.

DCE thread libraries and heap debug libraries are not compatible with runtime linking. Do not specify the **-brtl** compiler option if you are invoking the compiler with **xlC\_r4** or **xlC++\_r4**, or if the **-qheapdebug** compiler option is specified.

## Predefined macros

None.

## Related information

- “-b” on page 115
- “-G” on page 172

## -C

### Category

Output control

### Pragma equivalent

None.

### Purpose

Prevents the completed object from being sent to the linker. With this option, the output is a .o file for each source file.

## Syntax

►► -c ◀◀

## Defaults

By default, the compiler invokes the linker to link object files into a final executable.

## Usage

When this option is in effect, the compiler creates an output object file, *file\_name.o*, for each valid source file, such as *file\_name.c*, *file\_name.i*, *file\_name.C*, *file\_name.cpp*. You can use the **-o** option to provide an explicit name for the object file.

The **-c** option is overridden if the **-E**, **-P**, or **-qsyntaxonly** options are specified.

## Predefined macros

None.

## Examples

To compile `myprogram.c` to produce an object file `myprogram.o`, but no executable file, enter the command:

```
xlc myprogram.c -c
```

To compile `myprogram.c` to produce the object file `new.o` and no executable file, enter:

```
xlc myprogram.c -c -o new.o
```

## Related information

- “-E” on page 142
- “-o” on page 268
- “-P” on page 280
- “-qsyntaxonly (C only)” on page 341

## -C, -C!

### Category

Output control

### Pragma equivalent

None.

### Purpose

When used in conjunction with the `-E` or `-P` options, preserves or removes comments in preprocessed output.

When `-C` is in effect, comments are preserved. When `-C!` is in effect, comments are removed.

### Syntax



### Defaults

`-C`

### Usage

The `-C` option has no effect without either the `-E` or the `-P` option. If `-E` is specified, continuation sequences are preserved in the output. If `-P` is specified, continuation sequences are stripped from the output, forming concatenated output lines.

You can use the `-C!` option to override the `-C` option specified in a default makefile or configuration file.

## Predefined macros

None.

## Examples

To compile `myprogram.c` to produce a file `myprogram.i` that contains the preprocessed program text including comments, enter:

```
xlc myprogram.c -P -C
```

## Related information

- “-E” on page 142
- “-P” on page 280

## -qcache

### Category

Optimization and tuning

### Pragma equivalent

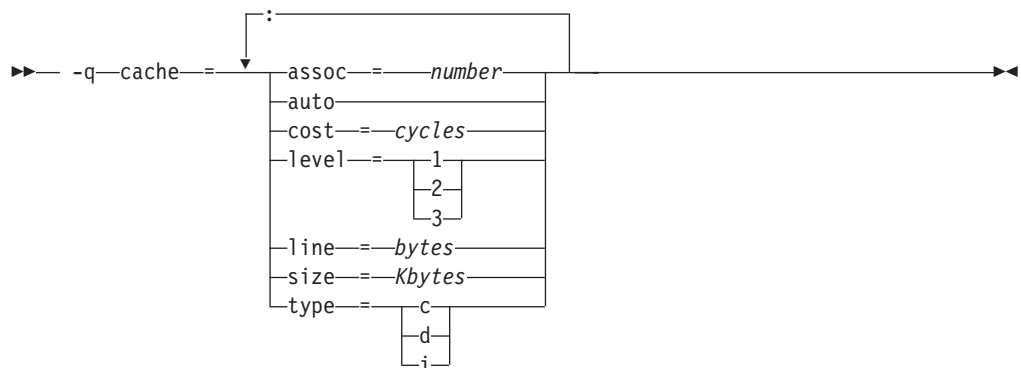
None.

### Purpose

When specified with `-O4`, `-O5`, or `-qipa`, specifies the cache configuration for a specific execution machine.

If you know the type of execution system for a program, and that system has its instruction or data cache configured differently from the default case, use this option to specify the exact cache characteristics. The compiler uses this information to calculate the benefits of cache-related optimizations.

### Syntax



### Defaults

Automatically determined by the setting of the `-qtune` option.

## Parameters

### **assoc**

Specifies the set associativity of the cache.

### *number*

Is one of:

- 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. This assumes that the execution environment will be the same as the compilation environment.

### **cost**

Specifies the performance penalty resulting from a cache miss.

### *cycles*

### **level**

Specifies the level of cache affected. If a machine has more than one level of cache, use a separate **-qcache** option.

### *level*

Is one of:

- 1 Basic cache
- 2 Level-2 cache or, if there is no level-2 cache, the table lookaside buffer (TLB)
- 3 TLB

### **line**

Specifies the line size of the cache.

### *bytes*

An integer representing the number of bytes of the cache line.

### **size**

Specifies the total size of the cache.

### *Kbytes*

An integer representing the number of kilobytes of the total cache.

### **type**

Specifies that the settings apply to the specified *cache\_type*.

### *cache\_type*

Is one of:

- c Combined data and instruction cache
- d Data cache
- i Instruction cache

## Usage

The **-qtune** setting determines the optimal default **-qcache** settings for most typical compilations. You can use the **-qcache** to override these default settings. However,

if you specify the wrong values for the cache configuration, or run the program on a machine with a different configuration, the program will work correctly but may be slightly slower.

You must specify **-O4**, **-O5**, or **-qipa** with the **-qcache** option.

Use the following guidelines when specifying **-qcache** suboptions:

- Specify information for as many configuration parameters as possible.
- If the target execution system has more than one level of cache, use a separate **-qcache** option to describe each cache level.
- If you are unsure of the exact size of the cache(s) on the target execution machine, specify an estimated cache size on the small side. It is better to leave some cache memory unused than it is to experience cache misses or page faults from specifying a cache size larger than actually present.
- The data cache has a greater effect on program performance than the instruction cache. If you have limited time available to experiment with different cache configurations, determine the optimal configuration specifications for the data cache first.
- If you specify the wrong values for the cache configuration, or run the program on a machine with a different configuration, program performance may degrade but program output will still be as expected.
- The **-O4** and **-O5** optimization options automatically select the cache characteristics of the compiling machine. If you specify the **-qcache** option together with the **-O4** or **-O5** options, the option specified last takes precedence.

## Predefined macros

None.

## Examples

To tune performance for a system with a combined instruction and data level-1 cache, where cache is 2-way associative, 8 KB in size and has 64-byte cache lines, enter:

```
xlc -O4 -qcache=type=c:level=1:size=8:line=64:assoc=2 file.c
```

## Related information

- “-qcache” on page 121
- “-O, -qoptimize” on page 269
- “-qtune” on page 357
- “-qipa” on page 199
- “Optimizing your applications” in the *XL C/C++ Optimization and Programming Guide*

## -qchars

### Category

Floating-point and integer control

### Pragma equivalent

#pragma options chars, #pragma chars

## Purpose

Determines whether all variables of type `char` are treated as either signed or unsigned.

## Syntax

### Option syntax

►► `-qchars=`unsigned  
signed\_\_\_\_\_►►

### Pragma syntax

►► `#pragma chars` (unsigned  
signed)\_\_\_\_\_►►

## Defaults

`-qchars=unsigned`

## Parameters

### unsigned

Variables of type `char` are treated as unsigned char.

### **signed**

Variables of type `char` are treated as signed char.

## Usage

Regardless of the setting of this option or pragma, the type of `char` is still considered to be distinct from the types `unsigned char` and `signed char` for purposes of type-compatibility checking or C++ overloading.

The pragma must appear before any source statements. If the pragma is specified more than once in the source file, the first one will take precedence. Once specified, the pragma applies to the entire file and cannot be disabled; if a source file contains any functions that you want to compile without **#pragma chars**, place these functions in a different file.

## Predefined macros

- `__CHAR_SIGNED` and `__CHAR_SIGNED__` are defined to 1 when **signed** is in effect; otherwise, it is undefined.
- `__CHAR_UNSIGNED` and `__CHAR_UNSIGNED__` are defined to 1 when **unsigned** is in effect; otherwise, they are undefined.

## Examples

To treat all `char` types as signed when compiling `myprogram.c`, enter:

```
xlc myprogram.c -qchars=signed
```

## -qcheck

### Category

Error checking and debugging

### Pragma equivalent


#pragma options [no]check

### Purpose

Generates code that performs certain types of runtime checking.

If a violation is encountered, a runtime error is raised by sending a SIGTRAP signal to the process. Note that the runtime checks may result in slower application execution.

### Syntax



### Defaults

-qnocheck

### Parameters

**all**  
Enables all suboptions.

#### **bounds | nobounds**

Performs runtime checking of addresses for subscripting within an object of known size. The index is checked to ensure that it will result in an address that lies within the bounds of the object's storage. A trap will occur if the address does not lie within the bounds of the object.

This suboption has no effect on accesses to a variable length array.

#### **divzero | nodivzero**

Performs runtime checking of integer division. A trap will occur if an attempt is made to divide by zero.

#### **nullptr | nonullptr**

Performs runtime checking of addresses contained in pointer variables used to reference storage. The address is checked at the point of use; a trap will occur if the value is less than 512.

Specifying the **-qcheck** option with no suboptions is equivalent to **-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:

```
xlc myprogram.c -qcheck=all:nonnullptr
```

provides checking for everything except for addresses contained in pointer variables used to reference storage. If you use **all** with the **no...** form of the suboptions, **all** should be the first suboption.

## Predefined macros

None.

## Examples

The following code example shows the effect of **-qcheck=nullptr:bounds**:

```
void func1(int* p) {
    *p = 42;          /* Traps if p is a null pointer */
}

void func2(int i) {
    int array[10];
    array[i] = 42;   /* Traps if i is outside range 0 - 9 */
}
```

The following code example shows the effect of **-qcheck=divzero**:

```
void func3(int a, int b) {
    a / b;          /* Traps if b=0 */
}
```

## **-qcinc (C++ only)**

### Category

Input control

### Pragma equivalent

None.

### Purpose

Places an extern "C" { } wrapper around the contents of include files located in a specified directory.

### Syntax

```
➤ -q nocinc cinc = directory_path ➤
```



## Defaults

-qnocinc

## Parameters

*directory\_path*

The directory where the include files to be wrapped with an extern "C" linkage specifier are located.

## Predefined macros

None.

## Examples

Assume your application `myprogram.C` includes header file `foo.h`, which is located in directory `/usr/tmp` and contains the following code:

```
int foo();
```

Compiling your application with:

```
xlc++ myprogram.C -qcinc=/usr/tmp
```

will include header file `foo.h` into your application as:

```
extern "C" {  
int foo();  
}
```

## -qcompact

### Category

Optimization and tuning

### Pragma equivalent

```
#pragma options [no]compact
```

### Purpose

Avoids optimizations that increase code size.

Code size is reduced by inhibiting optimizations that replicate or expand code inline, such as inlining or loop unrolling. Execution time may increase.

### Syntax

►► — -q —  compact —►►

### Defaults

-qnocompact

## Usage

This option only has an effect when specified with an optimization option.

## Predefined macros

`__OPTIMIZE_SIZE__` is predefined to 1 when `-qcompact` and an optimization level are in effect. Otherwise, it is undefined.

## Examples

To compile `myprogram.c`, instructing the compiler to reduce code size whenever possible, enter:

```
xlc myprogram.c -O -qcompact
```

## **-qconcurrentupdate (C only)**

### Category

Object code control

### Pragma equivalent

None.

### Purpose

Supports updating the operating system while the kernel is running.

### Syntax

►► — `-q` — `noconcurrentupdate` — `concurrentupdate` — ◀◀

### Defaults

`-qnoconcurrentupdate`

### Usage

If you want to use AIX Concurrent Update (hot-patch), you must use `-qconcurrentupdate` to compile your code. For details about Concurrent Update, see the AIX Concurrent Update documentation.

**Note:** This is a C-only option. If you compile your code with `-qconcurrentupdate` using a C++ compiler, the compiler issues a message: The option "`-qconcurrentupdate`" is not supported.

### Predefined macros

None.

### Examples

```
xlc myprogram.c -qconcurrentupdate
```

## **-qcpluscmt (C only)**

### **Category**

Language element control

### **Pragma equivalent**

None.

### **Purpose**

Enables recognition of C++-style comments in C source files.

### **Syntax**

►► -q cpluscmt  
nocpluscmt ◀◀

### **Defaults**

- **-qcpluscmt** when the **xlc** or **c99** and related invocations are used, or when the **stdc99** | **extc99** language level is in effect.
- **-qnocpluscmt** for all other invocation commands and language levels.

### **Predefined macros**

`__C99_CPLUSCMT` is predefined to 1 when **-qcpluscmt** is in effect; otherwise, it is undefined.

### **Examples**

To compile `myprogram.c` so that C++ comments are recognized as comments, enter:

```
xlc myprogram.c -qcpluscmt
```

Note that `//` comments are *not* part of C89. The result of the following valid C89 program will be incorrect:

```
main() {  
    int i = 2;  
    printf("%i\n", i /* 2 */  
           + 1);  
}
```

The correct answer is 2 (2 divided by 1). When **-qcpluscmt** is in effect (as it is by default), the result is 3 (2 plus 1).

### **Related information**

- “-C, -C!” on page 120
- “-qlanglvl” on page 215
- “Comments” in the *XL C/C++ Language Reference*

## **-qcrt**

### **Category**

Linking

## Pragma equivalent

None.

## Purpose

Specifies whether system startup files are to be linked.

When `-qcrt` is in effect, the system startup routines are automatically linked. When `-qnocrt` is in effect, the system startup files are not used at link time; only the files specified on the command line with the `-l` flag will be linked.

This option can be used in system programming to disable the automatic linking of the startup routines provided by the operating system.

## Syntax

The diagram shows the syntax for the `-qcrt` and `-qnocrt` options. It starts with a right-pointing arrow followed by a hyphen and a lowercase 'q'. A bracket above the 'q' is labeled 'crt', and a bracket below the 'q' is labeled 'nocrt'. A long horizontal line with arrowheads at both ends follows, indicating that the option can be repeated.

## Defaults

`-qcrt`

## Predefined macros

None.

## Related information

- “`-qlib`” on page 240

## `-qc_stdinc` (C only)

### Category

Compiler customization

### Pragma equivalent

None.

### Purpose

Changes the standard search location for the XL C and system header files.

### Syntax

The diagram shows the syntax for the `-qc_stdinc` option. It starts with a right-pointing arrow followed by a hyphen, a lowercase 'c', an underscore, and 'stdinc'. This is followed by a double quote character, a bracketed placeholder for a directory path, another double quote character, and a colon. A long horizontal line with arrowheads at both ends follows, indicating that the option can be repeated. A vertical arrow points from the colon to the text 'directory\_path'.

## Defaults

By default, the compiler searches the directories specified in the configuration file for the XL C header files (this is normally `/usr/vac/include/`) and for the system header files (this is normally `/usr/include/`).

## Parameters

### *directory\_path*

The path for the directory where the compiler should search for the XL C and system header files. The *directory\_path* can be a relative or absolute path. You can surround the path with quotation marks to ensure it is not split up by the command line.

## Usage

This option allows you to change the search paths for specific compilations. To permanently change the default search paths for the XL C and system headers, you use a configuration file to do so; see “Directory search sequence for include files” on page 13 for more information.

If this option is specified more than once, only the last instance of the option is used by the compiler.

This option is ignored if the `-qnostdinc` option is in effect.

## Predefined macros

None.

## Examples

To override the default search path for the XL C headers with `mypath/headers1` and `mypath/headers2`, enter:

```
xlc myprogram.c -qc_stdinc=mypath/headers1:mypath/headers2
```

## Related information

- “`-qstdinc`” on page 332
- “`-qinclude`” on page 185
- “Directory search sequence for include files” on page 13
- “Specifying compiler options in a configuration file” on page 7

## **-qcpp\_stdinc (C++ only)**

### Category

Compiler customization


### Pragma equivalent

None.

### Purpose

Changes the standard search location for the XL C++ and system header files.

## Syntax



## Defaults

By default, the compiler searches the directories specified in the configuration file for the XL C++ header files (this is normally `/usr/vacpp/include/`) and for the system header files (this is normally `/usr/include/`).

## Parameters

### *directory\_path*

The path for the directory where the compiler should search for the XL C++ and system header files. The *directory\_path* can be a relative or absolute path. You can surround the path with quotation marks to ensure it is not split up by the command line.

## Usage

This option allows you to change the search paths for specific compilations. To permanently change the default search paths for the XL C++ and system headers, you use a configuration file to do so; see “Directory search sequence for include files” on page 13 for more information.

If this option is specified more than once, only the last instance of the option is used by the compiler.

This option is ignored if the `-qnostdinc` option is in effect.

## Predefined macros

None.

## Examples

To override the default search path for the XL C++ headers with `mypath/headers1` and `mypath/headers2`, enter:

```
xlc++ myprogram.C -qcpp_stdinc=mypath/headers1:mypath/headers2
```

## Related information

- “-qstdinc” on page 332
- “-qinclude” on page 185
- “Directory search sequence for include files” on page 13
- “Specifying compiler options in a configuration file” on page 7

## -D

### Category

Language element control

## Pragma equivalent

None.

## Purpose

Defines a macro as in a `#define` preprocessor directive.

## Syntax

```
▶▶ -D name [=definition]
```

## Defaults

Not applicable.

## Parameters

*name*

The macro you want to define. `-Dname` is equivalent to `#define name`. For example, `-DCOUNT` is equivalent to `#define COUNT`.

*definition*

The value to be assigned to *name*. `-Dname=definition` is equivalent to `#define name definition`. For example, `-DCOUNT=100` is equivalent to `#define COUNT 100`.

## Usage

Using the `#define` directive to define a macro name already defined by the `-D` option will result in an error condition.

To aid in program portability and standards compliance, the operating system provides several header files that refer to macro names you can set with the `-D` option. You can find most of these header files either in the `/usr/include` directory or in the `/usr/include/sys` directory. To ensure that the correct macros for your source file are defined, use the `-D` option with the appropriate macro name. For example, if your source file includes the `/usr/include/sys/stat.h` header file, you must compile with the option `-D_POSIX_SOURCE` to pick up the correct definitions for that file.

The `-Uname` option, which is used to undefine macros defined by the `-D` option, has a higher precedence than the `-Dname` option.

## Predefined macros

The compiler configuration file uses the `-D` option to predefine several macro names for specific invocation commands. For details, see the configuration file for your system.

## Examples

AIX 4.2 and later provides support for files greater than 2 gigabytes in size so you can store large quantities of data in a single file. To allow large file manipulation in your application, compile with the `-D_LARGE_FILES` and `-qlonglong` compiler options. For example:

```
xlc myprogram.c -D_LARGE_FILES -qlonglong
```

To specify that all instances of the name `COUNT` be replaced by `100` in `myprogram.c`, enter:

```
xlc myprogram.c -DCOUNT=100
```

## Related information

- “-U” on page 361
- Chapter 6, “Compiler predefined macros,” on page 463
- “Header files” in the *AIX Files Reference*

## **-qdataimported, -qdatalocal, -qtocdata**

### Category

Optimization and tuning

### Pragma equivalent

None.

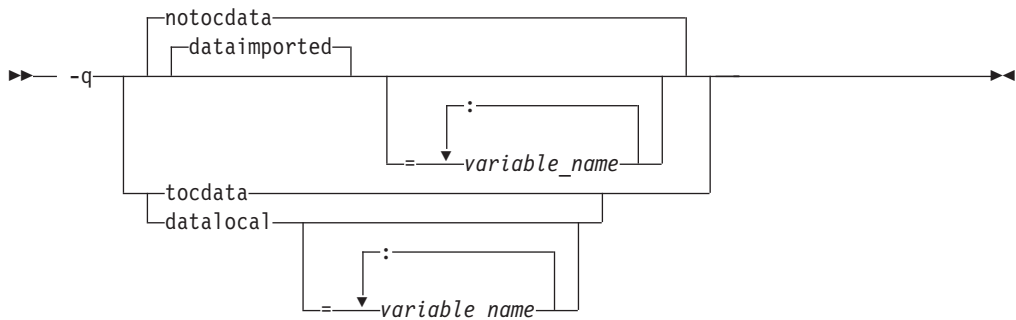
### Purpose

Marks data as local or imported in 64-bit compilations.

Local variables are statically bound with the functions that use them. You can use the `-qdatalocal` option to name variables that the compiler can assume are local. Alternatively, you can use the `-qtocdata` option to instruct the compiler to assume all variables are local.

Imported variables are dynamically bound with a shared portion of a library. You can use the `-qdataimported` option to name variables that the compiler can assume are imported. Alternatively, you can use the `-qnotocdata` option to instruct the compiler to assume all variables are imported.

### Syntax






## Defaults

**-qdataimported** or **-qnotocdata**: The compiler assumes all variables are imported.

## Parameters

*variable\_name*

The name of a variable that the compiler should assume is local or imported (depending on the option specified).

 Names must be specified using their mangled names. To obtain C++ mangled names, compile your source to object files only, using the **-c** compiler option, and use the **nm** operating system command on the resulting object file. You can also use the **c++filt** utility provided by the compiler for a side-by-side listing of source names and mangled names; see "Demangling compiled C++ names" in the *XL C/C++ Optimization and Programming Guide* for details. (See also "Name mangling" in the *XL C/C++ Language Reference* for details on using the extern "C" linkage specifier on declarations to prevent name mangling.)

Specifying **-qdataimported** without any *variable\_name* is equivalent to **-qnotocdata**: all variables are assumed to be imported. Specifying **-qdatalocal** without any *variable\_name* is equivalent to **-qtocdata**: all variables are assumed to be local.

## Usage

These options apply to 64-bit compilations only.

If any variables that are marked as local are actually imported, performance may decrease.

If you specify any of these options with no variables, the last option specified is used. If you specify the same variable name on more than one option specification, the last one is used.

## Predefined macros

None.

## Related information

- “-qprocimported, -qproclocal, -qprocunknown” on page 298

## **-qdbgfmt**

### Category

Error checking and debugging

### Pragma equivalent

None.

### Purpose

Specifies the stabstring or DWARF format for the debugging information produced in object files.

DWARF is a standard that defines the format of debugging information in programs. It is used on a wide variety of operating systems and is extensible and compact.

## Syntax

```
►► -q-dbgfmt=stabstring  
                  └─ dwarf ─┘
```

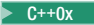


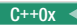
## Defaults

- `-qdbgfmt=stabstring`

## Parameters

### stabstring

Generates debugging information in stabstring format.

  **Note:** This suboption does not generate debugging information for C++0x or C1X features. Use the **dwarf** suboption instead for these features.  

### **dwarf**

Generates debugging information in DWARF format.

#### Notes:

- To use `-qdbgfmt=dwarf`, the program must be compiled and linked on AIX V7.1 or above.
- To debug programs built with this suboption, a DWARF-enabled debugger such as **dbx** is required.

## Usage

`-qdbgfmt` does not imply any of the debugging options, such as “**-g**” on page 169. For example:

- To generate debugging information in stabstring format, use **-g -qdbgfmt=stabstring**.
- To generate debugging information in DWARF format, use **-g -qdbgfmt=dwarf**.

`-qdbgfmt` also applies to “**-qlinedebug**” on page 242, which generates a subset of “**-g**” on page 169 information. You can use **-qlinedebug -qdbgfmt=dwarf** to generate line number information in DWARF format.

## Related information

- “**-g**” on page 169
- “**-qlinedebug**” on page 242

## **-qdbxextra (C only)**

### Category

Error checking and debugging

### Pragma equivalent

#pragma options dbxextra

## Purpose

When used with the `-g` option, specifies that debugging information is generated for unreferenced typedef declarations, struct, union, and enum type definitions.

To minimize the size of object and executable files, the compiler only includes information for typedef declarations, struct, union, and enum type definitions that are referenced by the program. When you specify the `-qdbxextra` option, debugging information is included in the symbol table of the object file. This option is equivalent to the `-qsyntab=unref` option.

## Syntax

►► `-q` nodbxextra  
dbxextra \_\_\_\_\_ ►►

## Defaults

`-qnodbxextra`: Unreferenced typedef declarations, struct, union, and enum type definitions are not included in the symbol table of the object file.

## Usage

Using `-qdbxextra` may make your object and executable files larger.

## Predefined macros

None.

## Examples

To compile `myprogram.c` so that unreferenced typedef, structure, union, and enumeration declarations are included in the symbol table for use with a debugger, enter:

```
xlc myprogram.c -g -qdbxextra
```

## Related information

- “`-qfullpath`” on page 164
- “`-qlinedebug`” on page 242
- “`-g`” on page 169
- “`#pragma options`” on page 420
- “`-qsyntab (C only)`” on page 340

## **-qdfp**

### Category

Language element control

### Pragma equivalent

None.

## Purpose

Enables compiler support for decimal floating-point types and literals.

## Syntax

►► `-q` nodfp  
dfp ►►

## Defaults

`-qnodfp`

## Usage

If you enable `-qdfp` for a `-qarch` value that does not support decimal floating-point instructions, `-qfloat=dfpemulate` is automatically enabled, and the decimal floating-point operations are performed by software. This may cause a slowdown in the application's runtime performance.

Note that runtime support for decimal floating-point operations is available only on AIX for POWER version 5.3 with the 5300-06 Technology Level or later. If you enable `-qdfp` on a version of the operating system that does not provide runtime support, your application will compile, but it may not link or run.

Programs that use decimal floating-point functions or macros defined in the `<math.h>` include file should not be compiled on AIX 5.2 or on older levels of AIX 5.3 or 5.4, because those functions and macros won't be retrofitted into 5.2.

## Predefined macros

When `-qdfp` is in effect, `__IBM_DFP__` is predefined to a value of 1; otherwise it is undefined.

## Related information

- “`-qarch`” on page 106
- “`-qfloat`” on page 154

## **-qdigraph**

### Category

Language element control

### Pragma equivalent

`#pragma options [no]digraph`

### Purpose

Enables recognition of digraph key combinations or keywords to represent characters not found on some keyboards.

## Syntax

►► -q digraph  
nodigraph ◀◀

## Defaults

- C **-qdigraph** when the `extc89` | `extended` | `extc99` | `stdc99` language level is in effect. **-qnodigraph** for all other language levels.
- C++ **-qdigraph**

## Usage

A digraph is a keyword or combination of keys that lets you produce a character that is not available on all keyboards. For details on digraphs, see "Digraph characters" in the *XL C/C++ Language Reference*.

## Predefined macros

`__DIGRAPHS__` is predefined to 1 when **-qdigraph** is in effect; otherwise it is not defined.

## Examples

To disable digraph character sequences when compiling your program, enter:

```
xlc myprogram.c -qnodigraph
```

## Related information

- “-qlanglvl” on page 215
- “-qtrigraph” on page 356

## -qdirectstorage

### Category

Optimization and tuning

### Pragma equivalent

None.

### Purpose

Informs the compiler that a given compilation unit may reference write-through-enabled or cache-inhibited storage.

## Syntax

►► -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. To ensure that your application will execute correctly on all implementations, you should assume that separate instruction and data caches exist and program your application accordingly.

## -qdollar

### Category

Language element control

### Pragma equivalent

#pragma options [no]dollar

### Purpose

Allows the dollar-sign (\$) symbol to be used in the names of identifiers.

When **dollar** is in effect, the dollar symbol \$ in an identifier is treated as a base character.

### Syntax

►► -q ———— ◀◀

### Defaults

-qnodollar

### Usage

If **nodollar** and the **ucs** language level are both in effect, the dollar symbol is treated as an extended character and translated into `\u0024`.

### Predefined macros

None.

### Examples

To compile `myprogram.c` so that \$ is allowed in identifiers in the program, enter:

```
xlc myprogram.c -qdollar
```

### Related information

- “-qlanglvl” on page 215

## -qdpcl

### Category

Error checking and debugging

## Pragma equivalent

None.

## Purpose

Generates symbols that tools based on the IBM Dynamic Probe Class Library (DPCL) can use to see the structure of an executable file.

DPCL is an open-source set of libraries used by application performance analysis tools (for more information, visit <http://dpcl.sourceforge.net>). When **-qdpcl** is in effect, the compiler emits symbols to define blocks of code in a program; you can then use tools that use the DPCL interface to examine performance information such as memory usage for object files compiled with this option.

## Syntax

► `-q` nodpcl  
dpcl ►

## Defaults

`-qnodpcl`

## Usage

You must specify **-qdpcl** together with the **-g** option to ensure that the compiler generates debugging information required by debugging and program analysis tools.

**-qdpcl** is not supported for any optimization level except zero. If a non-zero optimization level is specified or implied by other options, **-qdpcl** will be disabled.

You cannot specify the **-qipa** or **-qsmp** options together with **-qdpcl**.

## Predefined macros

None.

## Related information

- “-g” on page 169
- “-qipa” on page 199
- “-qsmp” on page 320

**-e**

## Category

Linking

## Pragma equivalent

None.

## Purpose

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

## Syntax

▶▶ **-e** *entry\_name* ▶▶

## Defaults

Not applicable.

## Parameters

*name*

The name of the entry point for the shared executable.

## Usage

Specify the **-e** option only with the **-qmkshrobj** or **-G** option. For more information, see the description for **-qmkshrobj**.

**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.

## Predefined macros

None.

## Related information

- “**-qmkshrobj**” on page 260
- “**-G**” on page 172

## **-E**

### Category

Output control

### Pragma equivalent

None.

### Purpose

Preprocesses the source files named in the compiler invocation, without compiling, and writes the output to the standard output.

### Syntax

▶▶ **-E** ▶▶



## Defaults

By the default, source files are preprocessed, compiled, and linked to produce an executable file.

## Usage

The **-E** option accepts any file name. Source files with unrecognized file name suffixes are treated and preprocessed as C files, and no error message is generated.

Unless **-qnopline** is specified, `#line` directives are generated to preserve the source coordinates of the tokens. Continuation sequences are preserved.

Unless **-C** is specified, comments are replaced in the preprocessed output by a single space character. New lines and `#line` directives are issued for comments that span multiple source lines.

The **-E** option overrides the **-P**, **-o**, and **-qsyntaxonly** options.

## Predefined macros

None.

## Examples

To compile `myprogram.c` and send the preprocessed source to standard output, enter:

```
xlc myprogram.c -E
```

If `myprogram.c` has a code fragment such as:

```
#define SUM(x,y) (x + y)
int a ;
#define mm 1 /* This is a comment in a
              preprocessor directive */
int b ;      /* This is another comment across
              two lines */
int c ;
              /* Another comment */
c = SUM(a, /* Comment in a macro function argument*/
        b) ;
```

the output will be:

```
#line 2 "myprogram.c"
int a ;
#line 5
int b ;

int c ;

c = a + b ;
```

## Related information

- “`-qppline`” on page 293
- “`-C`, `-C!`” on page 120
- “`-P`” on page 280
- “`-qsyntaxonly` (C only)” on page 341

## **-qeh (C++ only)**

### **Category**

Object code control

### **Pragma equivalent**

None.

### **Purpose**

Controls whether exception handling is enabled in the module being compiled.

When **-qeh** is in effect, exception handling is enabled. If your program does not use C++ structured exception handling, you can compile with **-qnoeh** to prevent generation of code that is not needed by your application.

### **Syntax**



### **Defaults**

**-qeh=v6**

### **Parameters**

**v6** Generates exception handling code, compatible with VisualAge® C++ V6.0, that correctly handles try-catch blocks nested within other catch blocks.

**v5** Generate exception handling code that is compatible with VisualAge C++ V5.0.

Specifying **-qeh** with no suboption is equivalent to **-qeh=v6**.

### **Predefined macros**

`__EXCEPTIONS` is predefined to 1 when **-qeh** is in effect; otherwise, it is undefined.

### **Related information**

- “-qrtti (C++ only)” on page 308

## **-qenum**

### **Category**

Floating-point and integer control

### **Pragma equivalent**

`#pragma options enum`, `#pragma enum`

## Purpose

Specifies the amount of storage occupied by enumerations.


► C++0x The **-qenum** option affects only unscoped enumerations that have no fixed underlying type. For enumerations with a fixed underlying type, the **-qenum** option is ignored. ◀ C++0x

## Syntax

### Option syntax



### Pragma syntax



## Defaults


`-qenum=intlong`

## Parameters

- 1 Specifies that enumerations occupy 1 byte of storage, are of type signed char if the range of enumeration values falls within the limits of signed char, and unsigned char otherwise.
- 2 Specifies that enumerations occupy 2 bytes of storage, are of type short if the range of enumeration values falls within the limits of signed short, and unsigned short otherwise. ► C Values cannot exceed the range of signed int.
- 4 Specifies that enumerations occupy 4 bytes of storage, are of type int if the range of enumeration values falls within the limits of signed int, and unsigned int otherwise.
- 8 Specifies that enumerations occupy 8 bytes of storage. In 32-bit compilation mode, the enumeration is of type long long if the range of enumeration values falls within the limits of signed long long, and unsigned long long otherwise. In 64-bit compilation mode, the enumeration is of type long if the range of enumeration values falls within the limits of signed long, and unsigned long otherwise.

## int

 Specifies that enumerations occupy 4 bytes of storage and are of type int.

 Specifies that enumerations occupy 4 bytes of storage, are of type int if the range of enumeration values falls within the limits of signed int, and unsigned int otherwise.

## intlong

Specifies that enumerations occupy 8 bytes of storage, as with the 8 suboption, if the range of values in the enumeration cannot be represented by one of int or unsigned int. Otherwise, the enumerations occupy 4 bytes of storage as with the 4 suboption.

## small

Specifies that enumerations occupy the smallest amount of space (1, 2, 4, or 8 bytes of storage) that can accurately represent the range of values in the enumeration. Signedness is unsigned, unless the range of values includes negative values. If an 8-byte enum results, the actual enumeration type used is dependent on compilation mode.

## pop | reset (pragma only)

Discards the current pragma setting and reverts to the setting specified by the previous pragma directive. If no previous pragma was specified, reverts to the command-line or default option setting.

## Usage

The tables that follow show the priority for selecting a predefined type. The table also shows the predefined type, the maximum range of enum constants for the corresponding predefined type, and the amount of storage that is required for that predefined type, that is, the value that the sizeof operator would yield when applied to the minimum-sized enum. All types are signed unless otherwise noted.

Table 24. Enumeration sizes and types




	enum=1		enum=2		enum=4		enum=8			
	var	const	var	const	var	const	32-bit compilation mode		64-bit compilation mode	
Range	var	const	var	const	var	const	var	const	var	const
0..127	signed char	int	short	int	int	int	long long	long long	long	long
-128..127	signed char	int	short	int	int	int	long long	long long	long	long
0..255	unsigned char	int	short	int	int	int	long long	long long	long	long
0..32767	ERROR <sup>1</sup>	int	short	int	int	int	long long	long long	long	long
-32768..32767	ERROR <sup>1</sup>	int	short	int	int	int	long long	long long	long	long
0..65535	ERROR <sup>1</sup>	int	unsigned short	int	int	int	long long	long long	long	long
0..2147483647	ERROR <sup>1</sup>	int	ERROR <sup>1</sup>	int	int	int	long long	long long	long	long
-(2147483647+1).. 2147483647	ERROR <sup>1</sup>	int	ERROR <sup>1</sup>	int	int	int	long long	long long	long	long
0..4294967295	ERROR <sup>1</sup>	unsigned int <sup>2</sup>	ERROR <sup>1</sup>	unsigned int <sup>2</sup>	unsigned int <sup>2</sup>	unsigned int <sup>2</sup>	long long	long long	long	long
0..(2 <sup>63</sup> -1)	ERROR <sup>1</sup>	long <sup>2</sup>	ERROR <sup>1</sup>	long <sup>2</sup>	ERROR <sup>1</sup>	long <sup>2</sup>	long long <sup>2</sup>	long long <sup>2</sup>	long <sup>2</sup>	long <sup>2</sup>

Table 24. Enumeration sizes and types (continued)

$-2^{63}..(2^{63}-1)$	ERROR <sup>1</sup>	long <sup>2</sup>	ERROR <sup>1</sup>	long <sup>2</sup>	ERROR <sup>1</sup>	long <sup>2</sup>	long long <sup>2</sup>	long long <sup>2</sup>	long <sup>2</sup>	long <sup>2</sup>
$0..2^{64}$	ERROR <sup>1</sup>	unsigned long <sup>2</sup>	ERROR <sup>1</sup>	unsigned long <sup>2</sup>	ERROR <sup>1</sup>	unsigned long <sup>2</sup>	unsigned long long <sup>2</sup>	unsigned long long <sup>2</sup>	unsigned long <sup>2</sup>	unsigned long <sup>2</sup>

	enum=int		enum=intlong				enum=small			
			32-bit compilation mode		64-bit compilation mode		32-bit compilation mode		64-bit compilation mode	
Range	var	const	var	const	var	const	var	const	var	const
0..127	int	int	int	int	int	int	unsigned char	int	unsigned char	int
-128..127	int	int	int	int	int	int	signed char	int	signed char	int
0..255	int	int	int	int	int	int	unsigned char	int	unsigned char	int
0..32767	int	int	int	int	int	int	unsigned short	int	unsigned short	int
-32768..32767	int	int	int	int	int	int	short	int	short	int
0..65535	int	int	int	int	int	int	unsigned short	int	unsigned short	int
0..2147483647	int	int	int	int	int	int	unsigned int	unsigned int	unsigned int	unsigned int
$-(2147483647+1) .. 2147483647$	int	int	int	int	int	int	int	int	int	int
0..4294967295	unsigned int <sup>1</sup>	unsigned int <sup>2</sup>	unsigned int <sup>2</sup>	unsigned int <sup>2</sup>	unsigned int <sup>2</sup>	unsigned int <sup>2</sup>	unsigned int <sup>2</sup>	unsigned int <sup>2</sup>	unsigned int <sup>2</sup>	unsigned int <sup>2</sup>
$0..(2^{63}-1)$	ERR <sup>2</sup>	ERR <sup>2</sup>	long long <sup>2</sup>	long long <sup>2</sup>	long <sup>2</sup>	long <sup>2</sup>	unsigned long long <sup>2</sup>	unsigned long long <sup>2</sup>	unsigned long <sup>2</sup>	unsigned long <sup>2</sup>
$-2^{63}..(2^{63}-1)$	ERR <sup>2</sup>	ERR <sup>2</sup>	long long <sup>2</sup>	long long <sup>2</sup>	long <sup>2</sup>	long <sup>2</sup>	long long <sup>2</sup>	long long <sup>2</sup>	long <sup>2</sup>	long <sup>2</sup>
$0..2^{64}$	ERR <sup>2</sup>	ERR <sup>2</sup>	unsigned long long <sup>2</sup>	unsigned long long <sup>2</sup>	unsigned long <sup>2</sup>	unsigned long <sup>2</sup>	unsigned long long <sup>2</sup>	unsigned long long <sup>2</sup>	unsigned long <sup>2</sup>	unsigned long <sup>2</sup>

**Notes:**

- These enumerations are too large for the `-qenum=1|2|4| int ` setting. A Severe error is issued and compilation stops. To correct this condition, you should reduce the range of the enumerations, choose a larger `-qenum` setting, or choose a dynamic `-qenum` setting, such as `small` or `intlong`.
-  Enumeration types must not exceed the range of `int` when compiling C applications to ISO C 1989 and ISO C 1999 Standards. With the `stdc89` | `stdc99` language level in effect, the compiler will behave as follows if the value of an enumeration exceeds the range of `int` and the `-qenum` option in effect supports this value:
  - If `-qenum=int` is in effect, a severe error message is issued and compilation stops.
  - For all other settings of `-qenum`, an informational message is issued and compilation continues.
- When `-qenum=8`, for all ranges of enumerator values up to  $2^{32}-1$ , the table identifies the underlying type to be `long long` for 32-bit compilation mode and

long for 64-bit compilation mode. This violates the rule in the standard: The underlying type should not be larger than int if enumerator values fit in int or unsigned int.

- When `-qenum=small`, for enumerators in the range of 0-2147483647 with at least one enumerator having a value bigger than 65535, the table identifies the underlying type to be unsigned int, which cannot be promoted to int. This violates the rule in the standard: Promotion could be to int, because it can hold the enumeration values in the range of 0-2147483647.

The **#pragma enum** directive must precede the declaration of enum variables that follow; any directives that occur within a declaration are ignored and diagnosed with a warning.

For each **#pragma enum** directive that you put in a source file, it is good practice to have a corresponding **#pragma enum=reset** before the end of that file. This should prevent one file from potentially changing the setting of another file that includes it.

## Examples

If the following fragment is compiled with the **enum=small** option:

```
enum e_tag {a, b, c} e_var;
```

the range of enumeration constants is 0 through 2. This range falls within all of the ranges described in the table above. Based on priority, the compiler uses predefined type unsigned char.

If the following fragment is compiled with the **enum=small** option:

```
enum e_tag {a=-129, b, c} e_var;
```

the range of enumeration constants is -129 through -127. This range only falls within the ranges of short (signed short) and int (signed int). Because short (signed short) is smaller, it will be used to represent the enum.

The following code segment generates a warning and the second occurrence of the **enum** pragma is ignored:

```
#pragma enum=small
enum e_tag {
    a,
    b,
    #pragma enum=int /* error: cannot be within a declaration */
    c
} e_var;
#pragma enum=reset
#pragma enum=reset /* second reset isn't required */
```

## Predefined macros

None.

## -qexpfile

### Category

Object code control

## Pragma equivalent

None.

## Purpose

When used together with the `-qmkshrobj` or `-G` option, saves all exported symbols in a designated file.

## Syntax

►► `-q-expfile=filename` ◀◀

## Parameters

*filename*

The name of the file to which exported symbols are written.

## Usage

This option is valid only when used with the `-qmkshrobj` or `-G` option.

## Predefined macros

None.

## Related information

- “`-qmkshrobj`” on page 260
- “`-G`” on page 172

## **-qextchk**

### Category

Error checking and debugging

### Pragma equivalent

`#pragma options [no]extchk`

### Purpose

Generates link-time type checking information and checks for compile-time consistency.

### Syntax

►► `-q`

noextchk
extchk

 ◀◀

### Defaults

`-qnoextchk`

## Usage

This option does not perform type checking on functions or objects that contain references to incomplete types.

## Predefined macros

None.

## Examples

To compile `myprogram.c` so that link-time checking information is produced, enter:

```
xlc myprogram.c -qextchk
```

## -f

### Category

Linking

### Pragma equivalent

None.

### Purpose

Names a file that stores a list of object files for the compiler to pass to the linker.

### Syntax

▶▶ `-f—filelistname`—————▶▶

## Usage

The *filelistname* file should contain only the names of object files. There should be one object file per line.

This option is the same as the `-f` option for the `ld` command.

## Predefined macros

None.

## Examples

To pass the list of files contained in `myobjlistfile` to the linker, enter:

```
xlc -f/usr/tmp/myobjlistfile
```

## -F

### Category

Compiler customization



## Pragma equivalent

None.

## Purpose

Names an alternative configuration file or stanza for the compiler.

## Syntax



## Defaults

By default, the compiler uses the configuration file that is supplied at installation time, and uses the stanza defined in that file for the invocation command currently being used.

## Parameters

*file\_path*

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 `xlc`, but you specify the `c99` stanza, the compiler will use all the settings specified in the `c99` stanza.

## Usage

Note that any file names or stanzas that you specify with the `-F` option override the defaults specified in the system configuration file. If you have specified a custom configuration file with the `XLC_USR_CONFIG` environment variable, that file is processed before the one specified by the `-F` option.

The `-B`, `-t`, and `-W` options override the `-F` option.

## Predefined macros

None.

## Examples

To compile `myprogram.c` using a stanza called `debug` that you have added to the default configuration file, enter:

```
xlc myprogram.c -F:debug
```

To compile `myprogram.c` using a configuration file called `/usr/tmp/myconfig.cfg`, enter:

```
xlc myprogram.c -F/usr/tmp/myconfig.cfg
```

To compile `myprogram.c` using the stanza `c99` you have created in a configuration file called `/usr/tmp/myconfig.cfg`, enter:

```
xlc myprogram.c -F/usr/tmp/myconfig.cfg:xlf95c99
```

### Related information

- “Using custom compiler configuration files” on page 40
- “-B” on page 116
- “-t” on page 342
- “-W” on page 372
- “Specifying compiler options in a configuration file” on page 7
- “Compile-time and link-time environment variables” on page 26

## -qfdpr

### Category

Optimization and tuning

### Pragma equivalent

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.

### Predefined macros

None.

## Examples

To compile `myprogram.c` so it includes data required by the FDPR utility, enter:

```
xlc myprogram.c -qfdpr
```

## Related information

- “-qpdf1, -qpdf2” on page 283

## -qflag

### Category

Listings, messages, and compiler information

### Pragma equivalent

#pragma options flag, “#pragma report (C++ only)” on page 430


### Purpose

Limits the diagnostic messages to those of a specified severity level or higher.

The messages are written to standard output and, optionally, to the listing file if one is generated.

### Syntax


#### -qflag syntax – C



#### Notes:

- 1 Minimum severity level of messages reported in listing
- 2 Minimum severity level of messages reported on terminal

#### -qflag syntax – C++



#### Notes:

- 1 Minimum severity level of messages reported in listing
- 2 Minimum severity level of messages reported on terminal

## Defaults

`-qflag=i` : `i`, which shows all compiler messages

## Parameters

- `i` Specifies that all diagnostic messages are to display: warning, error and informational messages. Informational messages (I) are of the lowest severity.
- `w` Specifies that warning (W) and all types of error messages are to display.
- `e` Specifies that only error (E), severe error (S), and unrecoverable error (U) messages are to display.
- `s` Specifies that only severe error (S) and unrecoverable error (U) messages are to display.

## Usage

You must specify a minimum message severity level for both listing and terminal reporting.

You must specify a minimum message severity level for the listing. If you do not specify a suboption for the terminal, the compiler assumes the same severity as for the listing.

Note that using `-qflag` does not enable the classes of informational message controlled by the `-qinfo` option; see `-qinfo` for more information.

The `-qhaltonmsg` option has precedence over the `-qflag` option. If both `-qhaltonmsg` and `-qflag` are specified, messages that are not selected by `-qflag` are also printed.

## Predefined macros

None.

## Examples

To compile `myprogram.c` so that the listing shows all messages that were generated and your workstation displays only error and higher messages (with their associated information messages to aid in fixing the errors), enter:

```
xlc myprogram.c -qflag=i:e
```

## Related information

- “`-qinfo`” on page 187
- “`-qhaltonmsg`” on page 175
- “`-w`” on page 371
- “Compiler messages” on page 17

## `-qfloat`

### Category

Floating-point and integer control

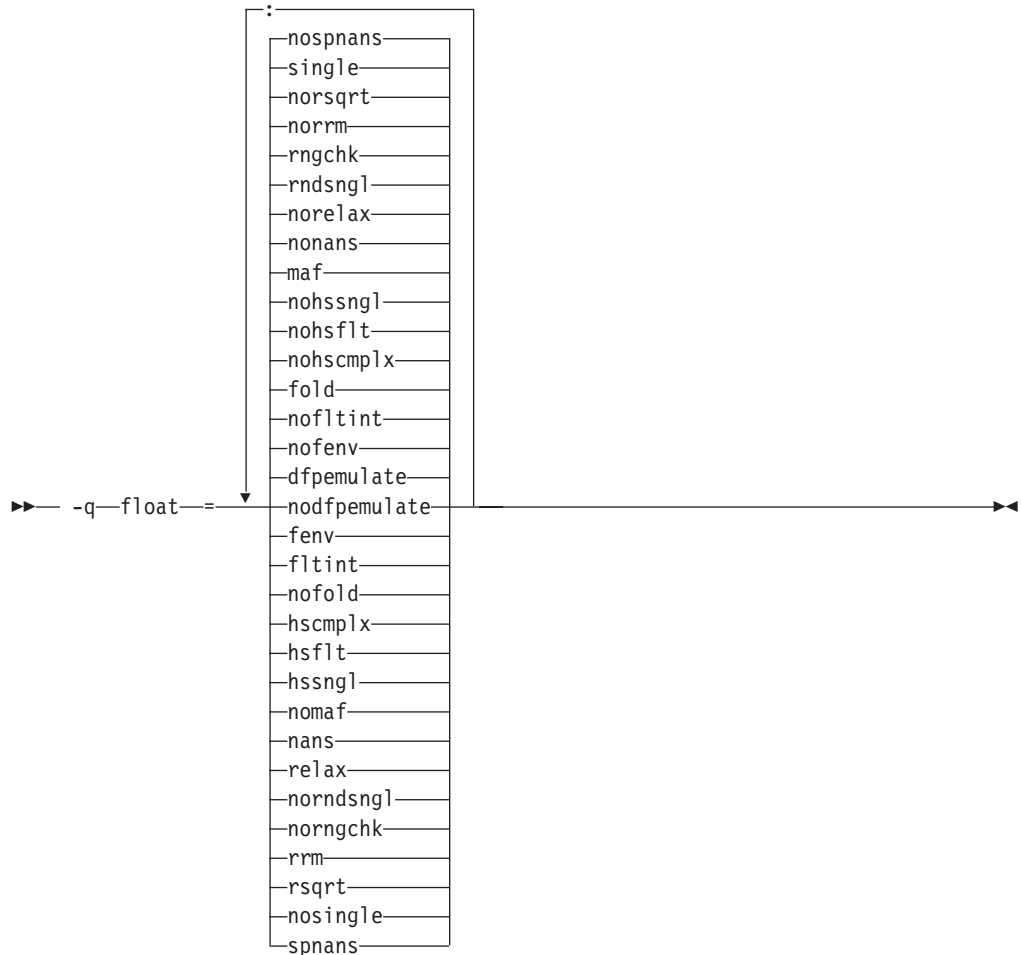
## Pragma equivalent

#pragma options float

## Purpose

Selects different strategies for speeding up or improving the accuracy of floating-point calculations.

## Syntax



## Defaults

- `-qfloat=dfpemulate:nofenv:nofltint:fold:nohscmplx:nohsflt:nohssngl:maf:nonans:norelax:rndsngl:rngchk:norm:norsqrt:single:nospnans`
- `-qfloat=fltint:rsqrt:norngchk` when `-qnostrict`, `-qstrict=nooperationprecision:noexceptions`, or `-O3` or higher optimization level is in effect.
- `-qfloat=nodfpemulate` when `-qarch=pwr6` is in effect.

## Parameters

### dfpemulate | nodfpemulate

Specifies whether decimal floating-point computations are implemented in hardware instructions or emulated in software by calls to library functions.

**nodfpemulate** is only valid on a system that supports decimal floating-point instructions; that is, on AIX 5.3 and higher, and with **-qarch=pwr6** in effect. **nodfpemulate** is the recommended setting for those systems, and results in improved performance of decimal floating-point operations and overall program runtime performance. **dfpemulate** is required for all other **-qarch** values.

Note that **-qdfp** must also be enabled for either suboption to have any effect. Otherwise, **nodfpemulate** is set.

#### **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.

Any directives specified in the source code (such as the standard C `FENV_ACCESS` pragma) take precedence over the option setting.

#### **fltint** | **nofltint**

Speeds up floating-point-to-integer conversions by using an inline sequence of code instead of a call to a library function. The library function, which is called when **nofltint** is in effect, checks for floating-point values outside the representable range of integers and returns the minimum or maximum representable integer if passed an out-of-range floating-point value.

If **-qarch** is set to a processor that has an instruction to convert from floating point to integer, that instruction will be used regardless of the **[no]fltint** setting. This conversion also applies to all Power processors in 64-bit mode.

If you compile with **-O3** or higher optimization level, **fltint** is enabled automatically. To disable it, also specify **-qstrict**, **-qstrict=operationprecision**, or **-qstrict=exceptions**.

#### **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**.

The **-qfloat=[no]fold** option replaces the deprecated **-q[no]fold** option. Use **-qfloat=[no]fold** in your new applications.

#### **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. It also uses the same technique as the **fltint** suboption for floating-point-to-integer conversions. **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**.

### **hssngl** | **nohssngl**

Specifies that single-precision expressions are rounded only when the results are stored into memory locations, but not after expression evaluation. Using **hssngl** can improve runtime performance and is safer than using **hsflt**.

This option only affects double-precision (double) expressions cast to single-precision (float) and used in an assignment operator for which a store instruction is generated, when **-qfloat=nosingle** is in effect. Do not use this option if you are compiling with the default **-qfloat=single**.

### **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. 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.

The **-qfloat=[no]maf** option replaces the deprecated **-q[no]maf** option. Use **-qfloat=[no]maf** in your new applications.

### **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.

The **hsflt** option overrides the **nans** option.

The **-qfloat=[no]nans** option replaces the deprecated **-qfloat=[no]spnans** option and the **-q[no]spnans** option. Use **-qfloat=[no]nans** in your new applications.

### **relax** | **norelax**

Relaxes strict IEEE conformance slightly for greater speed, typically by removing some trivial floating-point arithmetic operations, such as adds and subtracts involving a zero on the right. These changes are allowed if either **-qstrict=noieee** or **-qfloat=relax** is specified.

### **norndsngl** | **rndsngl**

Rounds the result of each single-precision operation to single-precision, rather

than waiting until the full expression is evaluated. It sacrifices speed for consistency with results from similar calculations on other types of computers.

This option only affects double-precision expressions cast to single-precision. You can only specify **norndsngl** when **-qfloat=nosingle** is in effect.

The **hsflt** suboption overrides the **rndsngl** option.

#### **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.

The **-qfloat=[no]rrm** option replaces the deprecated **-q[no]rrm** option. Use **-qfloat=[no]rrm** in your new applications.

#### **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.

**rsqrt** has no effect unless **-qignerrno** is also specified; **errno** will *not* be set for any sqrt function calls.

If you compile with **-O3** or higher optimization level, **rsqrt** is enabled automatically. To disable it, also specify **-qstrict**, **-qstrict=nans**, **-qstrict=infinities**, **-qstrict=zerosigns**, or **-qstrict=exceptions**.

#### **single** | **nosingle**

Allows single-precision arithmetic instructions to be generated for single-precision floating-point values. All Power processors support single-precision instructions; however, if you wish to preserve the behavior of applications compiled for earlier architectures, in which all floating-point arithmetic was performed in double-precision and then truncated to single-precision, you can use **-qfloat=nosingle:norndsngl**. This suboption



provides computation precision results compatible with those provided by the deprecated options `-qarch=com | pwr | pwr2 | p2sc | 601 | 602 | 603`. `-qfloat=nosingle` can be specified in 32-bit mode only.

### **spnans** | nospnans

Generates extra instructions to detect signalling NaN on conversion from single-precision to double-precision.

The `hsflt` suboption overrides the `spnans` suboption.

### **Note:**

- As of the V9.0 release of the compiler, the `emulate` | `noemulate` suboptions are deprecated.
- For details about the relationship between `-qfloat` suboptions and their `-qstrict` counterparts, see “`-qstrict`” on page 333.

## **Usage**

Using `-qfloat` suboptions other than the default settings may produce incorrect results in floating-point computations if not all required conditions for a given suboption are met. For this reason, you should only use this option if you are experienced with floating-point calculations involving IEEE floating-point values and can properly assess the possibility of introducing errors in your program. See also “Handling floating-point operations” in the *XL C/C++ Optimization and Programming Guide* for more information.

If the `-qstrict` | `-qnostrict` and `float` suboptions conflict, the last setting specified is used.

## **Predefined macros**

`__IBM_DFP_SW_EMULATION__` is predefined to a value of 1 when `-qfloat=dfpemulate` is in effect; otherwise it is undefined.

## **Examples**

To compile `myprogram.c` so that constant floating-point expressions are evaluated at compile time and multiply-add instructions are not generated, enter:

```
xlc myprogram.c -qfloat=fold:nomaf
```

## **Related information**

- “`-qarch`” on page 106
- “`-qflttrap`”
- “`-qldbl128`, `-qlongdouble`” on page 239
- “`-qstrict`” on page 333

## **-qflttrap**

### **Category**

Error checking and debugging

### **Pragma equivalent**

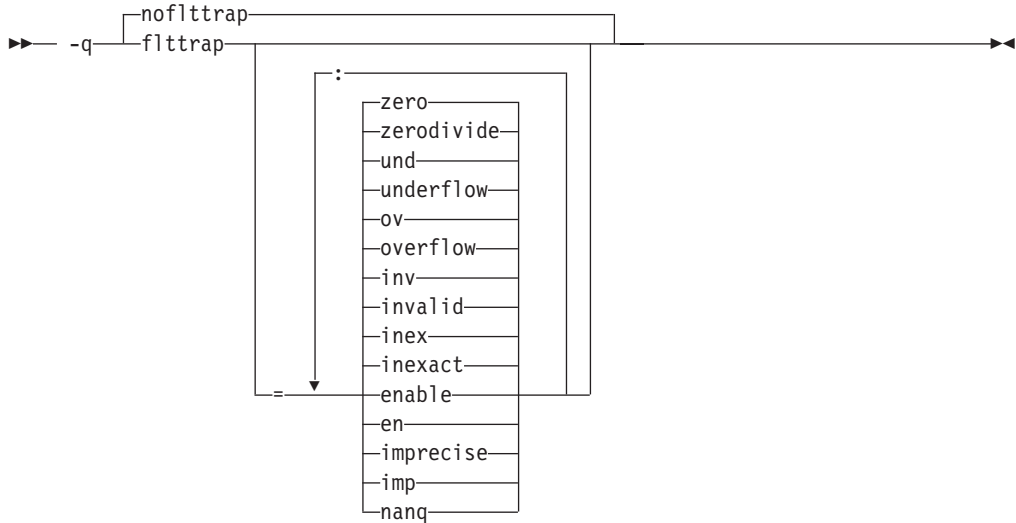
```
#pragma options [no]flttrap
```

## Purpose

Determines what types of floating-point exceptions to detect at run time.

The program receives a **SIGTRAP** signal when the corresponding exception occurs.

## Syntax



## Defaults

`-qnofltrap`

## Parameters

### **enable, en**

Inserts a trap when the specified exceptions (**overflow**, **underflow**, **zerodivide**, **invalid**, or **inexact**) occur. You must specify this suboption if you want to turn on exception trapping without modifying your source code. If any of the specified exceptions occur, a **SIGTRAP** or **SIGFPE** signal is sent to the process with the precise location of the exception. If **imprecise** is in effect, traps will not report exactly where the exception occurred.

### **imprecise, imp**

Enables imprecise detection of the specified exceptions. The compiler generates instructions after a block of code and just before the `main` program returns, to check if any of the specified exceptions (**overflow**, **underflow**, **zerodivide**, **invalid**, or **inexact**) have occurred. If an exception has occurred, an exception status flag is set in the Floating-Point Status and Control Register, but the exact location of the exception is not determined. Because instructions are not generated after each floating-point operation and function call to check for exceptions, this suboption can result in a slight performance improvement.

### **inexact, inex**

Enables the detection of floating-point inexact operations. If **imprecise** is not also specified, the compiler generates instructions after each floating-point operation and function call to check if an inexact operation exception has

occurred. If a floating-point inexact operation occurs, an inexact operation exception status flag is set in the Floating-Point Status and Control Register (FPSCR).

#### **invalid, inv**

Enables the detection of floating-point invalid operations. If **imprecise** is not also specified, the compiler generates instructions after each floating-point operation and function call to check if an invalid operation exception has occurred. If a floating-point invalid operation occurs, an invalid operation exception status flag is set in the FPSCR.

#### **nanq**

Generates code to detect Not a Number Quiet (NaNQ) and Not a Number Signalling (NaNS) exceptions before and after each floating-point operation, including assignment, and after each call to a function returning a floating-point result to trap if the value is a NaN. Trapping code is generated regardless of whether the **enable** suboption is specified.

#### **overflow, ov**

Enables the detection of floating-point overflow. If **imprecise** is not also specified, the compiler generates instructions after each floating-point operation and function call to check if an overflow exception has occurred. If a floating-point overflow occurs, an overflow exception status flag is set in the FPSCR.

#### **underflow, und**

Enables the detection of floating-point underflow. If **imprecise** is not also specified, the compiler generates instructions after each floating-point operation and function call to check if an underflow exception has occurred. If a floating-point underflow occurs, an underflow exception status flag is set in the FPSCR.

#### **zerodivide, zero**

Enables the detection of floating-point division by zero. If **imprecise** is not also specified, the compiler generates instructions after each floating-point operation and function call to check if a zero-divide exception has occurred. If a floating-point zero-divide occurs, a zero-divide exception status flag is set in the FPSCR.

## **Usage**

Specifying **-qfltrap** option with no suboptions is equivalent to **-qfltrap=overflow:underflow:zerodivide:invalid:inexact**

Exceptions will be detected by the hardware, but trapping is not enabled.

It is recommended that you use the **enable** suboption whenever compiling the main program with **-qfltrap**. This ensures that the compiler will generate the code to automatically enable floating-point exception trapping, without requiring that you include calls to the appropriate floating-point exception library functions in your code.

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.

If your program contains signalling NaNs, you should use the `-qfloat=nans` option along with `-qflttrap` to trap any exceptions.

The compiler exhibits behavior as illustrated in the following examples when the `-qflttrap` option is specified together with an optimization option:

- with `-O2`:
  - `1/0` generates a `div0` exception and has a result of infinity
  - `0/0` generates an invalid operation
- with `-O3` or greater:
  - `1/0` generates a `div0` exception and has a result of infinity
  - `0/0` returns zero multiplied by the result of the previous division.

If you use `-qflttrap=inv:en` to compile a program containing an IEEE invalid SQRT operation and you specify a `-qarch` target that does not implement the `sqrt` instruction set, the expected `SIGTRAP` signal will not occur when you run the program. You can fix this problem by specifying the following command before running the program:

```
export SQRT_EXCEPTION=3.1
```

**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.

## Predefined macros

None.

## Example

```
#include <stdio.h>

int main()
{
    float x, y, z;
    x = 5.0;
    y = 0.0;
    z = x / y;
    printf("%f", z);
}
```

When you compile this program with the following command, the program stops when the division is performed.

```
xlc -qflttrap=zerodivide:enable divide_by_zero.c
```

The `zerodivide` suboption identifies the type of exception to guard against. The `enable` suboption causes a `SIGTRAP` signal to be generated when the exception occurs.

## Related information

- “`-qfloat`” on page 154
- “`-qarch`” on page 106

## -qformat

### Category

Error checking and debugging

### Pragma equivalent

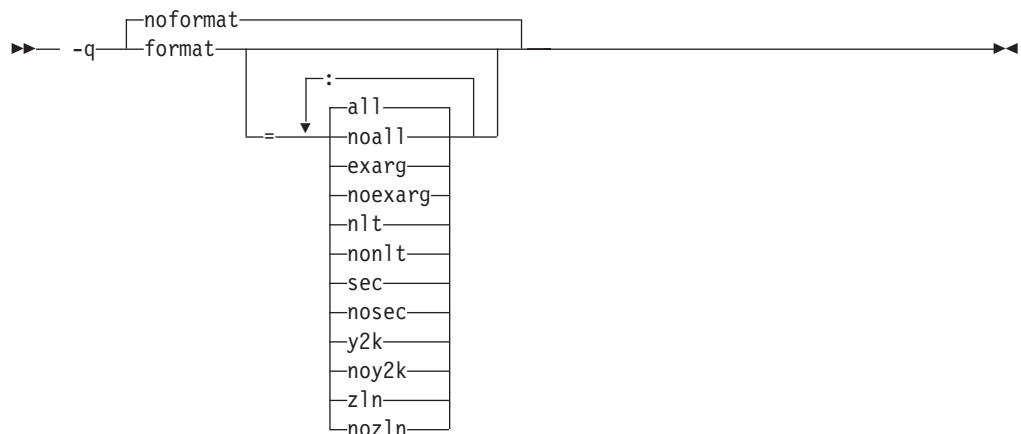
None.

### Purpose

Warns of possible problems with string input and output format specifications.

Functions diagnosed are `printf`, `scanf`, `strftime`, `strfmon` family functions and functions marked with format attributes.

### Syntax



### Defaults

`-qnoformat`

### Parameters

**all** | **noall**

Enables or disables all format diagnostic messages.

**exarg** | **noexarg**

Warns if excess arguments appear in `printf` and `scanf` style function calls.

**n1t** | **non1t**

Warns if a format string is not a string literal, unless the format function takes its format arguments as a `va_list`.

**sec** | **nosec**

Warns of possible security problems in use of format functions.

**y2k** | **noy2k**

Warns of `strftime` formats that produce a 2-digit year.

**z1n** | **noz1n**

Warns of zero-length formats.

Specifying **-qformat** with no suboptions is equivalent to **-qformat=all**.

**-qnoformat** is equivalent to **-qformat=noall**.

## Predefined macros

None.

## Examples

To enable all format string diagnostics, enter either of the following:

```
xlc myprogram.c -qformat=all
```

```
xlc myprogram.c -qformat
```

To enable all format diagnostic checking except that for y2k date diagnostics, enter:

```
xlc myprogram.c -qformat=all:noy2k
```

## **-qfullpath**

### Category

Error checking and debugging

### Pragma equivalent

```
#pragma options [no]fullpath
```

### Purpose

When used with the **-g** or **-qlinedebug** 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.

When **fullpath** is in effect, the absolute (full) path names of source files are preserved. When **nofullpath** is in effect, the relative path names of source files are preserved.

### Syntax

►► -q nofullpath  
fullpath ◀◀

### Defaults

**-qnofullpath**

### Usage

If your executable file was moved to another directory, the debugger would be unable to find the file unless you provide a search path in the debugger. You can use **fullpath** to ensure that the debugger locates the file successfully.

## Predefined macros

None.

## Related information

- “-qlinedebug” on page 242
- “-g” on page 169

## -qfuncsect

### Category

Object code control

### Pragma equivalent

#pragma options [no]funcsect

### Purpose


Places instructions for each function in a separate object file control section or CSECT which may reduce the size of your program. Placing each function in its own CSECT enables the linker to perform garbage collection on a per function basis rather than per object file.

When **-qfuncsect** is specified the compiler generates references from each function to the static data area, if one exists, in order to ensure that if any function from that object file is included in the final executable, the static data area also is included. This is done to ensure that any static strings or strings from a pragma comment, possible containing copyright information, are also included in the executable. This can, in some cases, cause code bloat or unresolved symbols at link time.

When **-qnofuncsect** is in effect, each object file consists of a single control section combining all functions defined in the corresponding source file. You can use **-qfuncsect** to place each function in a separate control section.

In prior releases, **-qfuncsect** had minimal size reductions for C++ programs. You should see an improvement in the current release.

### Syntax



### Defaults

-qnofuncsect

## Parameters

### **`implicitstaticref` | `noimplicitstaticref`**

Specifies whether references to the static data section of the object file by functions contained in static variables, virtual function tables, or exception handling tables, are maintained.

In releases before XL C/C++ for AIX V11.1, all exception handling tables were placed in one static data section. Including one exception handling table meant all the other tables were also included. Therefore, references to functions in the unused exception handling tables prevented linker garbage collection of those functions, which would otherwise have been cleaned up. As of XL C/C++ for AIX, V11.1, this problem is solved by allocating each exception handling table its own TOC entry. As a result, the size of the final executable might be reduced.

**Note:** The XL C/C++ for AIX, V11.1 enhancement enables large TOC access, which sets no limit on the number of TOC entries.

When your code contains a **#pragma comment** directive or a static string for copyright information purposes, the compiler automatically places these strings in the static data area, and generates references to these static data areas in the object code.

When **implicitstaticref** is in effect, a reference to the static area is generated even if not otherwise referenced.

When **noimplicitstaticref** is in effect, a reference to the static area is only generated if referenced by the program.

Specifying **-qfuncsect** with no suboption implies **implicitstaticref**.

## Usage

Using multiple control sections increases the size of the object file, but can reduce the size of the final executable by allowing the linker to remove functions that are not called or that have been inlined by the optimizer at all places they are called.

The pragma directive must be specified before the first statement in the compilation unit.

## Predefined macros

None.

## Related information

- “#pragma comment” on page 393
- “-qkeepinlines (C++ only)” on page 209
- “-qtwolink (C++ only)” on page 359

## **-qfunctrace**

### **Category**

Error checking and debugging

### **Pragma equivalent**


None.




## Purpose

Calls the tracing routines to trace the entry and exit points of the specified functions in a compilation unit.

## Syntax



## Pragma syntax



## Defaults

-qnofunctrace

**Note:** When **-qfunctrace** is specified for a C++ program, the functions in the *std* namespace are not traced by default.

## Parameters

- + Instructs the compiler to trace *function\_name*, *classes*, or *namespace*, and all its internal functions.
- Instructs the compiler not to trace *function\_name*, *classes*, or *namespace*, or any of its internal functions.

*function\_name*

Indicates the named functions to be traced.

*classname*

Indicates the named class to be traced.

*namespace*

Indicates the namespace to be traced.

## Usage

**-qfunctrace** enables tracing for all functions in your program. **-qnofunctrace** disables tracing that was enabled by **-qfunctrace**.

The **-qfunctrace+** and **-qfunctrace-** suboptions enable tracing for a specific list of functions and are not affected by **-qnofunctrace**. The list of functions is cumulative.

This option inserts calls to the tracing functions that you have defined. These functions must be provided at the link step. For details about the interface of tracing functions, as well as when they are called, see the Tracing functions in your code section in the *XL C/C++ Optimization and Programming Guide*.

Use `+` or `-` to indicate the function, classname, or namespace to be traced by the compiler. For example, if you want to trace function `x`, use `-qfunctrace+x`. To trace a list of functions, you must use a colon `:` to separate them.

Two colons in a row `::` is a scope qualifier, you can use it to indicate C++ qualified names. For example, use `-qfunctrace+A::B:C` traces functions that begin with qualifiers `A::B` or `C`.

If you want to trace functions in your code, you can write tracing functions in your code by using the following C function prototypes:

- Use `void __func_trace_enter(const char *const function_name, const char *const file_name, int line_number, void **const user_data);` to define the entry point tracing routine.
- Use `void __func_trace_exit(const char *const function_name, const char *const file_name, int line_number, void **const user_data);` to define the exit point tracing routine.
- Use `void __func_trace_catch(const char *const function_name, const char *const file_name, int line_number, void **const user_data);` to define the catch tracing routine.

You must define your functions when you write the preceding function prototypes in your code.

For details about these function prototypes as well as when they are called, see the **Tracing functions in your code** section in the *XL C/C++ Optimization and Programming Guide*.

#### Note:

- You can only use `+` and `-` one at a time. Do not use both of them together in the same `-qfunctrace` invocation.
- Definition of an inline function is traced. It is only the calls that have been inlined are not traced.

## Predefined macros

None.

## Examples

To trace functions `x`, `y`, and `z`, use `-qfunctrace+x:y:z`.

To trace all functions except for `x`, use `-qfunctrace -qfunctrace-x`.

The `-qfunctrace+` and `-qfunctrace-` suboptions only enable or disable tracing on the given list of cumulative functions. When functions, classes, and namespaces are used, the most completely specified option is in effect. The following is a list of examples:

- `-qfunctrace+x -qfunctrace+y` or `-qfunctrace+x -qnofunctrace -qfunctrace+y` enables tracing for only `x` and `y`.
- `-qfunctrace-x -qfunctrace` or `-qfunctrace -qfunctrace-x` traces all functions in the compilation unit except for `x`.
- `-qfunctrace -qfunctrace+x` traces all functions.
- `-qfunctrace+y -qnofunctrace` traces `y` only.
- If `functionX` is a member function of `classX`, then `-qfunctrace-classX::functionX` `-qfunctrace+classX` or `-qfunctrace+classX`

`-qfunctrace-classX::functionX` traces all member functions of `classX` but not `functionX`. This is because `classX::functionX` is more completely specified than `classX`. The more completely specified option has precedence over the less completely specified option.

- `-qfunctrace+MyClass` traces all member functions in `MyClass`.
- `-qfunctrace+std::vector` traces all instantiations of `std::vector`.
- `-qfunctrace+ABC -qfunctrace-ABC::foo` traces all functions defined in namespace `ABC` except for `foo`.

### Related information

- For details about `#pragma nofunctrace`, see “`#pragma nofunctrace`” on page 418.
- For detailed information about how to implement function tracing routines in your code, as well as detailed examples and a list of rules for using them, see **Tracing functions in your code** in the *XL C/C++ Optimization and Programming Guide*.

## -g

### Category

Error checking and debugging

### Pragma equivalent

None.

### Purpose

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

Program state refers to the values of user variables at certain points during the execution of a program.

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

## Syntax



## Defaults

`-g0`

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

## Parameters

`-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.

`-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.
- Function parameter values are available to the debugger at the beginning of each function.

`-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, function definitions, and function calls. For details, see “Usage” on page 171.
- Function parameter values are available to the debugger at the beginning of each function.

`-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.

- Function parameter values are available to the debugger at the beginning of each function.
- 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.
  - Function parameter values are available to the debugger at the beginning of each function.

## 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, for, or while statement, namely at the line where the do, for, or while keyword is specified. It is also available at the beginning of the next executable statement right after the do, for, or while construct.
- Function definitions
 

The debugging information is available at the first executable statement in the body of the function.
- Function calls
 

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

## Examples

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

```
xlc myprogram.c -o testing -g
```

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

```
xlc myprogram.c -O2 -g8
```

## Related information

- “-qdbxextra (C only)” on page 136
- “-qsymtab (C only)” on page 340
- “#pragma ibm snapshot” on page 406
- “-qlinedebug” on page 242
- “-qfullpath” on page 164
- “-O, -qoptimize” on page 269
- “-qkeepparm” on page 210

## -G

### Category

Output control

### Pragma equivalent

None.

### Purpose

Generates a shared object enabled for runtime linking.

### Syntax

►► — -G ————— ◀◀

### Usage

The compiler automatically exports all global symbols from the shared object unless you specify which symbols to export by using **-bE:**, **-bexport:**, or **-bnoexpall**. You can also prevent weak symbols from being exported by using the **-qnoweakexp** option. To save the export list to a file, use the **-qexpfile** option.

### Predefined macros

None.

### Related information

- “-b” on page 115
- “-brtl” on page 118
- “-qexpfile” on page 148
- “-qmkshrobj” on page 260
- “-qweakexp” on page 377
- Summary of compiler options by functional category: Linking
- “Shared Objects and Runtime Linking” in *AIX General Programming Concepts: Writing and Debugging Programs*
- **ld** in *AIX Commands Reference, Volume 3: i through m*

## -qgenproto (C only)

### Category

Portability and migration

## Pragma equivalent


None.

## Purpose

Produces prototype declarations from K&R function definitions or function definitions with empty parentheses, and displays them to standard output.

The compiler accepts and compiles K&R function definitions or definitions with a function declarator with empty parentheses; however, these function definitions are considered by the C standard to be obsolete (the compiler will diagnose them if you enable the **-qinfo=obs** option). When **-qgenproto** is in effect, the compiler generates the corresponding prototype declarations and displays them to standard output. You can use this option to help you identify obsolete function definitions and automatically obtain equivalent prototypes.

## Syntax



## Defaults

-qnogenproto

## Parameters

### parmnames

Parameter names are included in the prototype. If you do not specify this suboption, parameter names will not be included in the prototype.

## Predefined macros

None.

## Examples

Compiling with **-qgenproto** for the following function definitions:

```
int foo(a, b) // K&R function
  int a, b;
  {
}

int faa(int i) { } // prototyped function

main() { // missing void parameter
}
```

produces the following output on the display:

```
int foo(int, int);
int main(void);
```

Specifying **-qgenproto=parmnames** produces:

```
int foo(int a, int b);
int main(void);
```

## -qhalt

### Category

Error checking and debugging

### Pragma equivalent

#pragma options halt

### 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

#### -qhalt syntax — C



#### -qhalt syntax — C++



### Defaults

-qhalt=s

### Parameters

**i** Specifies that compilation is to stop for all types of errors: warning, error and informational. Informational diagnostics (I) are of the lowest severity.

**w** Specifies that compilation is to stop for warnings (W) and all types of errors.

**C** **e**

Specifies that compilation is to stop for errors (E), severe errors (S), and unrecoverable errors (U).

**S** **C** Specifies that compilation is to stop for severe errors (S) and unrecoverable errors (U). **C++** Specifies that compilation is to stop for severe errors (S).

### Usage

When the compiler stops as a result of the **halt** option, the compiler return code is nonzero. For a list of return codes, see “Compiler return codes” on page 19.

When **-qhalt** is specified more than once, the lowest severity level is used.



Diagnostic messages may be controlled by the **-qflag** option.

You can also instruct the compiler to stop compilation based on the number of errors of a type of severity by using the **-qmaxerr** option, which overrides **-qhalt**.

You can also use the **-qhaltmsg** option to stop compilation according to error message number.

## Predefined macros

None.

## Examples

To compile `myprogram.c` so that compilation stops if a warning or higher level message occurs, enter:

```
xlc myprogram.c -qhalt=w
```

## Related information

- “-qhaltmsg”
- “-qflag” on page 153
- “-qmaxerr” on page 254

## -qhaltmsg

### Category

Error checking and debugging

### Pragma equivalent


None.

### Purpose

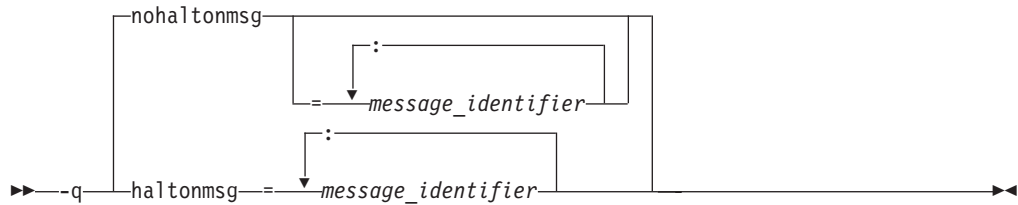
Stops compilation before producing any object files, executable files, or assembler source files if a specified error message is generated.

### Syntax

#### -qhaltmsg syntax - C



## -qhaltonmsg syntax - C++



## Defaults

-qnohaltonmsg

## Parameters

*message\_identifier*

Represents a message identifier. The message identifier must be in the following format:

*15dd-number*

where:

**15** Is the compiler product identifier.

*dd* Is the two-digit code representing the compiler component that produces the message. See “Compiler message format” on page 18 for descriptions of these codes.

*number*

Is the message number.

## 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.

You cannot specify **-qnohaltonmsg** to resume compilation if a message whose severity level is S has been issued.

**C** The **-qnohaltonmsg** compiler option cancels previous settings of **-qhaltonmsg**. **C**

**C++** When you specify **-qnohaltonmsg** with message identifiers, the previous **-qhaltonmsg** instances with the same message identifiers lose effect. When you specify **-qnohaltonmsg** without specific message identifiers, all previous **-qhaltonmsg** instances lose effect.

If you specify two or three of the following options, the last option specified has precedence:

**-qhaltonmsg=***message\_identifier*

**-qnohaltonmsg=***message\_identifier*

**-qnohaltonmsg**

**C++**

`-qhaltonmsg` has precedence over `-qsuppress` and `-qflag`.

## Predefined macros

None.

## Related information

- “`-qflag`” on page 153
- “`-qhalt`” on page 174
- “Compiler messages” on page 17
- “`-qsuppress`” on page 338

## `-qheapdebug`

### Category

Error checking and debugging

### Pragma equivalent

None.

### Purpose

Enables debug versions of memory management functions.

The compiler ships a set of "debug" versions of the standard memory management functions defined in `stdlib.h` (such as `_debug_malloc` and `_debug_malloc`); the header files for these functions are found in the product include directory (`usr/vacpp/include`). By default, the compiler uses the regular memory management functions (such as `malloc` and `malloc`) and does not preinitialize their local storage. When `-qheapdebug` is in effect, the compiler searches for header files first in the product include directory, where the debug versions of memory management functions are stored, and then in the system include directory.

### Syntax

→ `-q` noheapdebug  
heapdebug →

### Defaults

`-qnoheapdebug`

### Usage

For complete information on the debug memory management functions, see "Memory debug library functions" in the *XL C/C++ Optimization and Programming Guide*.

### Predefined macros

`__DEBUG_ALLOC__` is defined to 1 when `-qheapdebug` is in effect; otherwise, it is undefined.

## Examples

To compile `myprogram.c` with the debug versions of memory management functions, enter:

```
xlc -qheapdebug myprogram.c -o testing
```

## Related information

- "Debugging memory heaps" in the *XL C/C++ Optimization and Programming Guide*

## -qhot

### Category

Optimization and tuning

### Pragma equivalent

```
#pragma novector
```

### Purpose

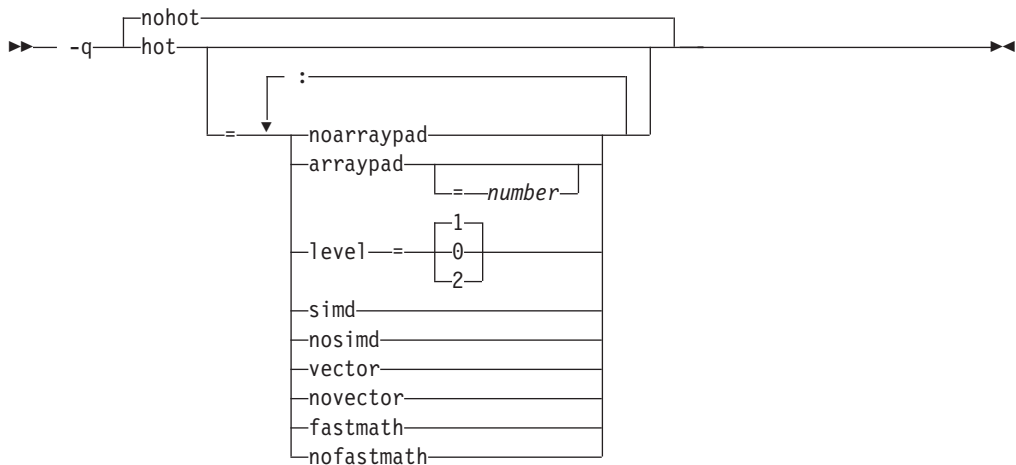
Performs high-order loop analysis and transformations (HOT) during optimization.

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.

You can use the pragma directives to disable these transformations for selected sections of code.

## Syntax

### Option syntax



### Pragma syntax



## 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` (option only)

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* (option only)

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` (option only)

Performs a subset of the high-order transformations and sets the default to `novector:noarraypad:fastmath`.

### `level=1` (option only)

Performs the default set of high-order transformations.

### `level=2` (option only)

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.

### `simd` (option only) | `nosimd`

This suboption has been deprecated. Consider using the `-qsimd` compiler option.

### `vector` (option only) | `novector`

When specified with `-qnostrict` and `-qignerrno`, or an optimization level of `-O3` 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 **-O4** 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.

For C/C++, you must use this suboption together with **-qignerrno**, unless **-qignerrno** is already enabled by other options.

**-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 **-O2**.

If you want to override the default **level** setting of 1 when using **-qsmp**, **-O4** or **-O5**, be sure to specify **-qhot=level=0** or **-qhot=level=2** *after* the other options.

The pragma directives apply only to **while**, **do while**, and **for** loops that immediately follow the placement of the directives. They have no effect on other loops that may be nested within the specified loop.

You can use the **-qreport** option in conjunction with **-qhot** or any optimization option that implies **-qhot** to produce a pseudo-C report showing how the loops were transformed. The loop transformations are included in the listing report if either the option **-qreport** or **-qlistfmt** 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 **-O3 -qhot** or higher. For more information, see “-qreport” on page 302.

## **Predefined macros**

None.

## **Related information**

- “-qarch” on page 106
- “-qsimd” on page 316
- “-qprefetch” on page 293
- “-qreport” on page 302
- “-O, -qoptimize” on page 269
- “-qstrict” on page 333
- “-qsmp” on page 320
- *Using the Mathematical Acceleration Subsystem (MASS) in the XL C/C++ Optimization and Programming Guide*

**-I**

## Category

Input control

## Pragma equivalent

None.

## Purpose

Adds a directory to the search path for include files.

## Syntax

►— `-I`—*directory\_path*—◄

## Defaults

See “Directory search sequence for include files” on page 13 for a description of the default search paths.

## Parameters

*directory\_path*

The path for the directory where the compiler should search for the header files.

## Usage

If `-qnostdinc` is in effect, the compiler searches *only* the paths specified by the `-I` option for header files, and not the standard search paths as well. If `-qidirfirst` is in effect, the directories specified by the `-I` option are searched before any other directories.

If the `-I` directory option is specified both in the configuration file and on the command line, the paths specified in the configuration file are searched first. The `-I` directory option can be specified more than once on the command line. If you specify more than one `-I` option, directories are searched in the order that they appear on the command line.

The `-I` option has no effect on files that are included using an absolute path name.

## Predefined macros

None.

## Examples

To compile `myprogram.c` and search `/usr/tmp` and then `/oldstuff/history` for included files, enter:

```
xlc myprogram.c -I/usr/tmp -I/oldstuff/history
```

## Related information

- “-qidirfirst”
- “-qstdinc” on page 332
- “-qinclude” on page 185
- “Directory search sequence for include files” on page 13
- “Specifying compiler options in a configuration file” on page 7

## -qidirfirst

### Category

Input control

### Pragma equivalent

#pragma options [no]idirfirst

### Purpose

Specifies whether the compiler searches for user include files in directories specified by the **-I** option *before* or *after* searching any other directories.

When **-qidirfirst** is in effect, the compiler first searches the directories specified by the **-I** option before searching any other directories. When **-qnoidirfirst** is in effect, before searching directories named on the **-I** option, the compiler first searches a) the directories in which source files named on the **-qinclude** option are located; and b) the directories in which the including files are located.

### Syntax

→ -q noidirfirst  
idirfirst →

### Defaults

-qnoidirfirst

### Usage

This option only affects files included with the `#include "file_name"` directive or the **-qinclude** option; **-qidirfirst** is independent of the **-qnostdinc** option and has no effect on the search order for XL C/C++ or system header files. (For the search order of header files, see “Directory search sequence for include files” on page 13.) This option also has no effect on files that are included using an absolute path name.

The last valid pragma directive remains in effect until replaced by a subsequent pragma.

### Predefined macros

None.



## Examples

To compile `myprogram.c` and search `/usr/tmp/myinclude` for included files before searching the current directory (where the source file resides), enter:

```
xlc myprogram.c -I/usr/tmp/myinclude -qidirfirst
```

## Related information

- “-I” on page 181
- “-qinclude” on page 185
- “-qstdinc” on page 332
- “-qc\_stdinc (C only)” on page 130
- “-qcpp\_stdinc (C++ only)” on page 131
- “Directory search sequence for include files” on page 13

## -qignerrno

### Category

Optimization and tuning

### Pragma equivalent

```
#pragma options [no]ignerrno
```

### Purpose

Allows the compiler to perform optimizations that assume `errno` is not modified by system calls.

Some system library functions set `errno` when an exception occurs. When **ignerrno** is in effect, the setting and subsequent side effects of `errno` are ignored. This allows the compiler to perform optimizations that assume `errno` is not modified by system calls.

### Syntax

```
► -q [noignerrno | ignerrno] ◄
```

### Defaults

- `-qnoignerrno`
- `-qignerrno` when `-O3` or higher optimization is in effect.

### Usage

If you require both `-O3` or higher and the ability to set `errno`, you should specify `-qnoignerrno` *after* the optimization option on the command line.

### Predefined macros

► C++ `__IGNERRNO__` is defined to 1 when **ignerrno** is in effect; otherwise, it is undefined.

### Related information

- “-O, -qoptimize” on page 269

## -qignprag

### Category

Language element control

### Pragma equivalent


#pragma options [no]ignprag

### Purpose

Instructs the compiler to ignore certain pragma statements.

This option is useful for detecting aliasing pragma errors. Incorrect aliasing gives runtime errors that are hard to diagnose. When a runtime error occurs, but the error disappears when you use **ignprag** with the **-O** option, the information specified in the aliasing pragmas is likely incorrect.

### Syntax



### Defaults

Not applicable.

### Parameters

#### all

Ignores all **#pragma isolated\_call** and **#pragma disjoint** directives in the source file.

#### disjoint

Ignores all **#pragma disjoint** directives in the source file.

#### ibm

 Ignores all **#pragma ibm snapshot** directives and all IBM SMP directives (such as **#pragma ibm parallel\_loop** and **#pragma ibm schedule**) in the source file.

#### isolated\_call

Ignores all **#pragma isolated\_call** directives in the source file.

#### omp

Ignores all OpenMP parallel processing directives in the source file, such as **#pragma omp parallel**, **#pragma omp critical**.

### Predefined macros

None.

## Examples

To compile `myprogram.c` and ignore any `#pragma isolated_call` directives, enter:  
`xlc myprogram.c -qignprag=isolated_call`

## Related information

- “`#pragma disjoint`” on page 395
- “`-qisolated_call`” on page 206
- “`#pragma ibm snapshot`” on page 406
- “Pragma directives for parallel processing” on page 439

## -qinclude

### Category

Input control

### Pragma equivalent

None.

### Purpose

Specifies additional header files to be included in a compilation unit, as though the files were named in an `#include` statement in the source file.

The headers are inserted before all code statements and any headers specified by an `#include` preprocessor directive in the source file. This option is provided for portability among supported platforms.

### Syntax

```
►► -q ┌ noinclude ───────────────────►  
     │ include--file_path ─────────►
```

### Defaults

`-qnoinclude`

### Parameters

*file\_path*

The absolute or relative path and name of the header file to be included in the compilation units being compiled. If *file\_path* is specified with a relative path, the search for it follows the sequence described in “Directory search sequence for include files” on page 13.

### Usage

The usage of the `-qinclude` option is similar to that of the `#include` directive. This section describes the differences between using `-qinclude` and `#include`.


The `-qinclude` option applies only to the files specified in the same compilation in which the option is specified. It is not passed to any compilations that occur during the link step, nor to any implicit compilations, such as those invoked by the option `-qtemplateregistry`, or to the files generated by `-qtempinc`.

When the option is specified multiple times in an invocation, the header files are included in order of appearance on the command line. If the same header file is specified multiple times with this option, the header is treated as if included multiple times by `#include` directives in the source file, in order of appearance on the command line.

Specifying `-qnoinclude` ignores any previous specification of `-qinclude`. Only the specifications of `-qinclude` after `-qnoinclude` are effective.

Any pragma directives that must appear before noncommentary statements in a source file will be affected; you cannot use `-qinclude` to include files if you need to preserve the placement of these pragmas.

The following rules apply when you use `-qinclude` with other options:

-  When used with `-qtemplateregistry`, `-qinclude` is recorded in the template registry file, along with the source files affected by it. When these file dependencies initiate recompilation of the template registry, the `-qinclude` option is passed to the dependent files only if it had been specified for them when they were added to the template registry.
- If you generate a listing file with `-qsource`, the header files included by `-qinclude` do not appear in the source section of the listing. Use `-qshowinc=usr` or `-qshowinc=all` in conjunction with `-qsource` if you want these header files to appear in the listing.
- After searching the directory from which the compiler was invoked, `-qinclude` searches additional search paths added to the search chain by the `-I` option. You can specify the `-I` option before or after the `-qinclude` option.
- Files specified with `-qinclude` are included as dependencies in the `-M` output. However, the paths are different in the dependency file for the `-qinclude` option and the `#include` directive, because the files specified with the `-qinclude` option are searched in the invocation path first, whereas files included by the `#include` directive are not.

When a dependency file is created as a result of a first build with the `-qinclude` option, a subsequent build without the `-qinclude` option will trigger recompile if the header file on the `-qinclude` option was touched between the two builds.

- In the compiler listing file generated by the `-qlistopt` option, each use of the `-qinclude` option has a separate entry in the OPTION SECTION.
- If both the `-qsource` option and the `-qinclude` option are used, header files specified with `-qinclude` are not included in the program source listing as `#include` directives. However, the files specified on `#include` directives in source programs are included.

## Predefined macros

None.

## Examples

To include the files `test1.h` and `test2.h` in the source file `test.c`, enter the following command:

```
xlc -qinclude=test1.h test.c -qinclude=test2.h
```

## Related information

- “Directory search sequence for include files” on page 13

# -qinfo

## Category

Error checking and debugging

## Pragma equivalent

#pragma options [no]info, #pragma info

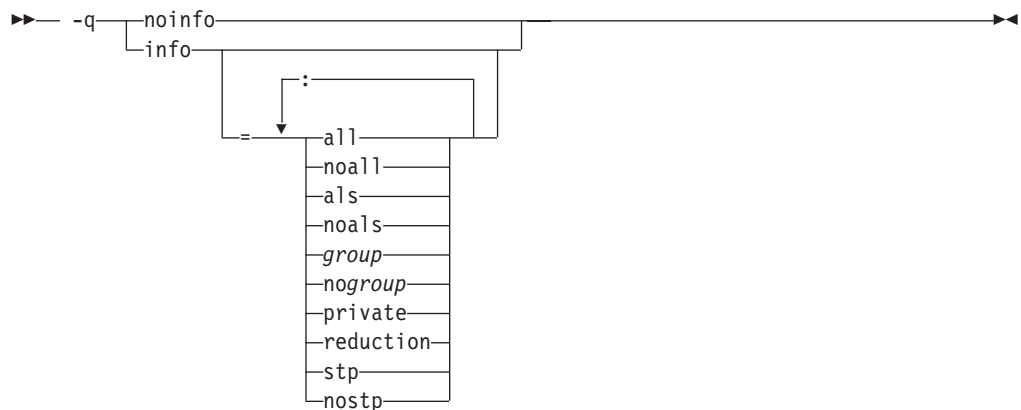
## Purpose

Produces or suppresses groups of informational messages.

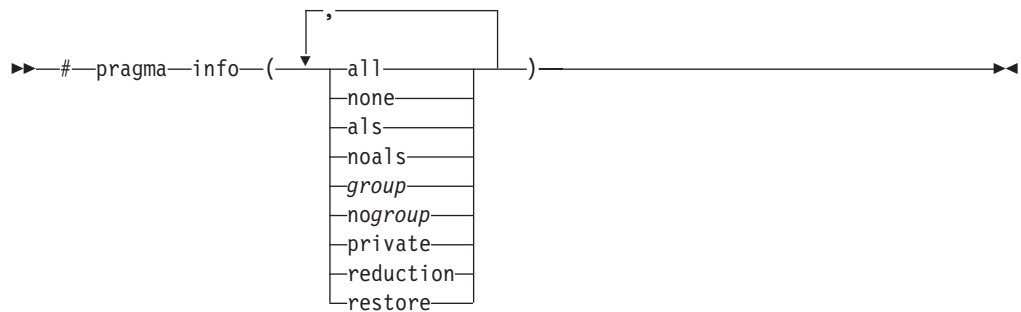
The messages are written to standard output and, optionally, to the listing file if one is generated.

## Syntax

### Option syntax



### Pragma syntax



## Defaults

-qnoinfo

- `C` -qnoinfo
- `C++` -qinfo=lan:trx

## Parameters

### **all**

Enables all diagnostic messages for all groups except **als** and **ppt**.

### **noall (option only)**

Disables all diagnostic messages for all groups.

### **none (pragma only)**

Disables all diagnostic messages for all groups.

### **als**

Enables reporting possible violations of the ANSI aliasing rule in effect.

### **noals**

Disables reporting possible aliasing-rule violations.

### **group | nogroup**

Enables or disables specific groups of messages, where *group* can be one or more of:

#### **group**

Type of informational messages returned or suppressed.

#### **c99 | noc99**

C code that may behave differently between C89 and C99 language levels.

#### **cls | nocls**

C++ classes.

#### **cmp | nocmp**

Possible redundancies in unsigned comparisons.

#### **cnd | nocnd**

Possible redundancies or problems in conditional expressions.

#### **cns | nocns**

Operations involving constants.

#### **cnv | nocnv**

Conversions.

#### **dcl | nodcl**

Consistency of declarations.

#### **eff | noeff**

Statements and pragmas with no effect.

#### **enu | noenu**

Consistency of enum variables.

#### **ext | noext**

Unused external definitions.

#### **gen | nogen**

General diagnostic messages.

#### **gnr | nognr**

Generation of temporary variables.

#### **got | nogot**

Use of goto statements.

#### **ini | noini**

Possible problems with initialization.

**lan** | **no1an**  
Language level effects.

**obs** | **noobs**  
Obsolete features.

**ord** | **noord**  
Unspecified order of evaluation.

**par** | **nopar**  
Unused parameters.

**por** | **nopor**  
Nonportable language constructs.

**ppc** | **noppc**  
Possible problems with using the preprocessor.

**ppt** | **noppt**  
Trace of preprocessor actions.

**pro** | **nopro**  
Missing function prototypes.

**rea** | **norea**  
Code that cannot be reached.

**ret** | **noret**  
Consistency of return statements.

**trd** | **notrd**  
Possible truncation or loss of data or precision.

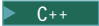
**tru** | **notru**  
Variable names truncated by the compiler.

**trx** | **notrx**  
Hexadecimal-floating point constants rounding.

**uni** | **nouni**  
Uninitialized variables.

**upg** | **noupg**  
Generates messages describing new behaviors of the current compiler release as compared to the previous release.

**use** | **nouse**  
Unused auto and static variables.

 **vft** | **novft**  
Generation of virtual function tables.

**zea** | **nozea**  
Zero-extent arrays.

#### **private**

This suboption is deprecated. **-qreport** replaces it. For details, see “-qreport” on page 302 and the “Deprecated options” on page 93 section in the *XL C/C++ Compiler Reference*.

#### **reduction**

This suboption is deprecated. **-qreport** replaces it. For details, see “-qreport” on page 302 and the “Deprecated options” on page 93 section in the *XL C/C++ Compiler Reference*.

## **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.

## **restore (pragma only)**

Discards the current pragma setting and reverts to the setting specified by the previous pragma directive. If no previous pragma was specified, reverts to the command-line or default option setting.

## **Usage**

Specifying **-qinfo** with no suboptions is equivalent to **-qinfo=all**.

Specifying **-qnoinfo** is equivalent to **-qinfo=noall**.

Consider the following when enabling the reporting of aliasing-rule violations:

- **-qalias=ansi** must be set before reporting of aliasing-rule violations (**-qinfo=als**) can occur.
- Any level of optimization or inlining implies **-qinfo=noals** and a warning will be issued.
- Diagnostics are heuristic and may emit false positives. Points-to analysis cannot be evaluated deterministically in static compilation. The points-to analysis used for diagnostics is evaluated in a context-and-flow, insensitive manner. The sequence of traceback messages in diagnostics is such that if executed in the order specified, the indirect expression will point to the offending object. If that execution sequence cannot occur in the application, the diagnostic is a false positive. (See the **Examples** section for the types of diagnostics that can occur.)


## **Predefined macros**

None.

## **Examples**

To compile `myprogram.c` to produce informational message about all items except conversions and unreachable statements, enter:

```
xlc myprogram.c -qinfo=all -qinfo=nocnv:norea
```

 The following example shows code constructs that the compiler detects when the code is compiled with **-qinfo=cnd:eff:got:obs:par:pro:rea:ret:uni** in effect:

```
#define COND 0

void faa() // Obsolete prototype (-qinfo=obs)
{
    printf("In faa\n"); // Unprototyped function call (-qinfo=pro)
}

int foo(int i, int k)
{
    int j; // Uninitialized variable (-qinfo=uni)

    switch(i) {
    case 0:
        i++;
        if (COND) // Condition is always false (-qinfo=cnd)
            i--; // Unreachable statement (-qinfo=rea)
    }
```



```

break;

case 1:
    break;
    i++; // Unreachable statement (-qinfo=rea)
default:
    k = (i) ? (j) ? j : i : 0;
}

goto L; // Use of goto statement (-qinfo=got)
return 3; // Unreachable statement (-qinfo=rea)
L:
    faa(); // faa() does not have a prototype (-qinfo=pro)

// End of the function may be reached without returning a value
// because of there may be a jump to label L (-qinfo=ret)

} //Parameter k is never referenced (-qinfo=ref)

int main(void) {
({ int i = 0; i = i + 1; i; }); // Statement does not have side effects (-qinfo=eff)

    return foo(1,2);
}

```

► **C++** The following example shows code constructs that the compiler detects, with this code is compiled with **-qinfo=cls:cnf:eff:use** in effect:

```

#pragma abc // pragma not supported (-qinfo=eff or -qinfo=gen)

int bar() __attribute__((xyz)); // attribute not supported (-qinfo=eff)
int j();

class A {
public:
    A(): x(0), y(0), z(0) { }; // this constructor is in the correct order
                                // hence, no info message.
    A(int m): y(0), z(0)
    { x=m; }; // suggest using member initialization list
                // for x (-qinfo=cls)

    A(int m, int n):
    x(0), z(0) { }; // not all data members are initialized
                    // namely, y is not initialized (-qinfo=cls)

    A(int m, int n, int* l):
    x(m), z(l), y(n) { }; // order of class initialization (-qinfo=cls)

private:
    int x;
    int y;
    int *z; // suggest having user-defined copy constructor/
            // assignment operator to handle the pointer data member
            // (-qinfo=cls)
};

int foo() {
    int j=5;
    j; // null statement (-qinfo=eff)
        // The user may mean to call j().

return j;
}

void boo() {
    int x;

```

```

int *i = &x;
float *f;           // f is not used (-qinfo=use)
f = (float *) i;   // incompatible type (-qinfo=eff)
                   // With ansi aliasing mode, a float pointer
                   // is not supposed to point to an int
}

void cond(int y) {
    const int i=0;
    int j;
    int k=0;

    if (i) {        // condition is always false (-qinfo=cnd)
        j=3;
    }

    if (1) {        // condition is always true (-qinfo=cnd)
        j=4;
    }

    j=0;
    if (j==0) {    // cond. is always true (-qinfo=cnd)
        j=5;
    }

    if (y) {
        k+=5
    }

    if (k==5) {    // This case cannot be determined, because k+=5
                   // is in a conditional block.
        j=6;
    }
}

```

In the following example, the **#pragma info(eff, nouni)** directive preceding MyFunction1 instructs the compiler to generate messages identifying statements or pragmas with no effect, and to suppress messages identifying uninitialized variables. The **#pragma info(restore)** directive preceding MyFunction2 instructs the compiler to restore the message options that were in effect before the **#pragma info(eff, nouni)** directive was specified.

```

#pragma info(eff, nouni)
int MyFunction1()
{
    .
    .
    .
}

#pragma info(restore)
int MyFunction2()
{
    .
    .
    .
}

```

The following example shows a valid diagnostic for an aliasing violation:

```

t1.c:
int main() {
    short s = 42;
    int *pi = (int*) &s;
    *pi = 63;
}

```

```

    return 0;
}
x1C -qinfo=als t1.c
"t1.c", line 4.3: 1540-0590 (I) Dereference may not conform to the current
aliasing rules.
"t1.c", line 4.3: 1540-0591 (I) The dereferenced expression has type "int".
"pi" may point to "s" which has incompatible
type "short".
"t1.c", line 4.3: 1540-0592 (I) Check assignment at line 3 column 11 of t1.c.

```

In the following example, the analysis is context insensitive in that the two calls to `floatToInt` are not distinguished. There is no aliasing violation in this example, but a diagnostic is still issued.

```

t2.c:
int* floatToInt(float *pf) { return (int*)pf; }

```

```

int main() {
    int i;
    float f;
    int* pi = floatToInt((float*)&i);
    floatToInt(&f);
    return *pi;
}

```

```

x1C -qinfo=als t2.c
"t2.c", line 8.10: 1540-0590 (I) Dereference may not conform to the current
aliasing rules.
"t2.c", line 8.10: 1540-0591 (I) The dereferenced expression has type "int".
"pi" may point to "f" which has incompatible
type "float".
"t2.c", line 8.10: 1540-0592 (I) Check assignment at line 7 column 14 of t2.c.
"t2.c", line 8.10: 1540-0592 (I) Check assignment at line 1 column 37 of t2.c.
"t2.c", line 8.10: 1540-0592 (I) Check assignment at line 6 column 11 of t2.c.

```

```

t3.c:
int main() {
    float f;
    int i = 42;
    int *p = (int*) &f;
    p = &i;
    return *p;
}

```

```

x1C -qinfo=als t3.c
"t3.c", line 6.10: 1540-0590 (I) Dereference may not conform to the current
aliasing rules.
"t3.c", line 6.10: 1540-0591 (I) The dereferenced expression has type "int".
"p" may point to "f", which has incompatible type "float".
"t3.c", line 6.10: 1540-0592 (I) Check assignment at line 4 column 10 of t3.c.

```

## Related information

- “-qflag” on page 153
- “-qreport” on page 302
- “-qstackprotect” on page 329
- For a list of deprecated options, see the “Deprecated options” on page 93 section in the *XL C/C++ Compiler Reference*.

## -qinitauto

### Category

Error checking and debugging

## Pragma equivalent

```
#pragma options [no]initauto
```

### Purpose

Initializes uninitialized automatic variables to a specific value, for debugging purposes.

### Syntax

```
►► -q noinitauto  
initauto == hex_value ►►
```

### Defaults

```
-qnoinitauto
```

### Parameters

#### **hex\_value**

A one- to eight-digit hexadecimal number.

- To initialize each byte of storage to a specific value, specify one or two digits for the *hex\_value*.
- To initialize each word of storage to a specific value, specify three to eight digits for the *hex\_value*.
- In the case where less than the maximum number of digits are specified for the size of the initializer requested, leading zeros are assumed.
- In the case of word initialization, if an automatic variable is smaller than a multiple of 4 bytes in length, the *hex\_value* is truncated on the left to fit. For example, if an automatic 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. See Example 1.
- If an automatic variable is larger than the *hex\_value* in length, the compiler repeats the *hex\_value* and assigns it to the variable. See Example 1.
- If the automatic variable is an array, the *hex\_value* is copied into the memory location of the array in a repeating pattern, beginning at the first memory location of the array. See Example 2.
- You can specify alphabetic digits as either uppercase or lowercase.
- The *hex\_value* can be optionally prefixed with 0x, in which x is case-insensitive.

### Usage

The `-qinitauto` option provides the following benefits:

- Setting *hex\_value* to zero ensures that all automatic 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.

This option generates extra code to initialize the value of automatic variables. It reduces the runtime performance of the program and is to be used for debugging purposes only.

**Restriction:** Objects that are equivalenced, structure components, and array elements are not initialized individually. Instead, the entire storage sequence is initialized collectively.

## Predefined macros

- `__INITAUTO__` is defined to the *hex\_value* that is specified on the `-qinitauto` option or pragma; otherwise, it is undefined.
- `__INITAUTO_W__` is defined to the *hex\_value*, repeated 4 times, specified on the `-qinitauto` option or pragma; otherwise, it is undefined.

## Examples

**Example 1:** Use the `-qinitauto` option to initialize automatic variables of scalar types.

```
#include <stdio.h>

int main()
{
    char a;
    short b;
    int c;
    long long int d;

    printf("char a = 0x%X\n", (char)a);
    printf("short b = 0x%X\n", (short)b);
    printf("int c = 0x%X\n", c);
    printf("long long int d = 0x%X\n", d);
}
```

If you compile the program with `-qinitauto=AABBCCDD`, for example, the result is as follows:

```
char a = 0xDD
short b = 0xFFFFCCDD
int c = 0xAABBCCDD
long long int d = 0xAABBCCDDAABBCCDD
```

**Example 2:** Use the `-qinitauto` option to initialize automatic array variables.

```
#include <stdio.h>
#define ARRAY_SIZE 5

int main()
{
    char a[5];
    short b[5];
    int c[5];
    long long int d[5];

    printf("array of char: ");
    for (int i = 0; i < ARRAY_SIZE; i++)
        printf("0x%1X ", (unsigned)a[i]);
    printf("\n");

    printf("array of short: ");
    for (int i = 0; i < ARRAY_SIZE; i++)
        printf("0x%1X ", (unsigned)b[i]);
    printf("\n");

    printf("array of int: ");
    for (int i = 0; i < ARRAY_SIZE; i++)
        printf("0x%1X ", (unsigned)c[i]);
    printf("\n");
}
```

```

printf("array of long long int: ");
for (int i = 0; i<ARRAY_SIZE; i++)
    printf("0x%1X ",(unsigned)d[i]);
printf("\n");
}

```

If you compile the program with `-qinitauto=AABBCCDD`, for example, the result is as follows:

```

array of char: 0xAA 0xBB 0xCC 0xDD 0xAA
array of short: 0xAABB 0xCCDD 0xAABB 0xCCDD 0xAABB
array of int: 0xAABBCCDD 0xAABBCCDD 0xAABBCCDD 0xAABBCCDD 0xAABBCCDD
array of long long int: 0xAABBCCDDAABBCCDD 0xAABBCCDDAABBCCDD 0xAABBCCDDAABBCCDD
0xAABBCCDDAABBCCDD 0xAABBCCDDAABBCCDD

```

## **-qinlglue**

### **Category**

Object code control

### **Pragma equivalent**

`#pragma options [no]inlglue`

### **Purpose**

When used with `-O2` or higher optimization, inlines glue code that optimizes external function calls in your application.

*Glue 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**



### **Defaults**

- `-qnoinlglue`
- `-qinlglue` when `-qtune=pwr4`, `-qtune=pwr5`, `-qtune=pwr6`, `-qtune=ppc970`, `-qtune=auto`, or `-qtune=balanced` is in effect.

### **Usage**

If you use the `-qtune` option with any of the suboptions that imply `-qinlglue` and you want to disable inlining of glue code, make sure to specify `-qnoinlglue` as well.

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.

The **-qinlglue** option only affects function calls through pointers or calls to an external compilation unit. For calls to an external function, you should specify that the function is imported by using, for example, the **-qprocimported** option.

## Predefined macros

None.

## Related information

- “-qcompact” on page 127
- “-qprocimported, -qproclocal, -qprocunknown” on page 298
- “-qtune” on page 357

## -qinline

### Category

Optimization and tuning

### Pragma equivalent


None.


### Purpose

Attempts to inline functions instead of generating calls to those functions, for improved performance.

#### Note:

- **-qinline** replaces **-Q** and its suboptions.
- **-Q**, **-Q!**, **-Q=threshold**, **-Q+name**, and **-Q-name** are all deprecated options and suboptions.
- **-qipa=inline** and all of its associated suboptions are deprecated. **-qinline** replaces them all.


 Specifying **-qinline** enables automatic inlining by the compiler front end. Specifying **-qinline** with **-O** provides additional inlining by enabling inlining by the low-level optimizer. In both cases, the compiler attempts to inline all functions, in addition to those defined inside a class declaration or explicitly marked with the `inline` specifier.

 You must specify a minimum optimization level of **-O2** along with **-qinline** to enable inlining of functions, including those declared with the `inline` specifier. You can also use the **-qinline** option to specify restrictions on the functions that should or should not be inlined.

In all cases where **-qinline** is in effect, the compiler uses heuristics to determine whether inlining a specific function will result in a performance benefit. That is, whether a function is appropriate for inlining is subject to limits on the number of inlined calls and the amount of code size increase as a result. Therefore, simply enabling inlining does not guarantee that a given function will be inlined.

Specifying **-qnoinline** disables all inlining, including that performed by the high-level optimizer with the **-qipa** option, and functions declared explicitly as `inline`.

## Syntax



## Defaults

- **-qnoinline**
- At an optimization level of **-O0**, the default is **-qinline=noauto**
- At optimization levels of **-O2** and higher, the default is **-qinline=auto**
- **-qinline=auto:level=5** is the default suboption of **-qinline**

## Parameters

### **noauto** | **auto**

Enables or disables automatic inlining. If you do not specify any **-qinline** suboptions, **-qinline=auto** is the default.

**Note:** At optimization levels of **-O2** and higher, the default is **-qinline=auto**

### **level=number**

Provides guidance to the compiler about the relative value of inlining. The values you specify for *number* must be positive integers between 0 and 10 inclusive. The default value for *number* is 5. If you specify a value less than 5, it implies less inlining. A value greater than 5 implies more inlining than the default.

### **C** *autothreshold*

Represents the number of executable statements in a function. The number of executable statements in a function must be fewer than or equal to *autothreshold* for it to be considered for inlining. The value you specify for *autothreshold* must be a positive integer. The default value for *autothreshold* is 20. If you specify a value of 0, no functions are inlined. As you can see in the following example:

```
increment()
{
    int a, b, i;
    for (i=0; i<10; i++) /* statement 1 */
    {
        a=i;           /* statement 2 */
        b=i;           /* statement 3 */
    }
}
```

## Usage

To maximize inlining, specify optimization (**-O**) and also specify the appropriate **-qinline** options.

Because inlining does not always improve runtime performance, you should test the effects of this option on your code. Do not attempt to inline recursive or mutually recursive functions.



If you specify the `-g` option to generate debugging information, inlining may be suppressed.

## Predefined macros

None.

## Examples

To compile `myprogram.c` so that no functions are inlined, enter:

```
xlc myprogram.c -O2 -qnoinline
```

Assuming you have functions `salary`, `taxes`, `expenses`, and `benefits`, to compile `myprogram.c` so that the compiler tries to inline these functions, you enter:

```
xlc myprogram.c -O2 -qinline+salary:taxes:expenses:benefits
```

If you do not want the functions `salary`, `taxes`, `expenses`, and `benefits` to be inlined when you compile `myprogram.c`, you enter:

```
xlc myprogram.c -O2 -qinline-salary:taxes:expenses:benefits
```

**C** In general, you can turn off automatic inlining and request that specific functions be inlined by naming them with the `+` form:

```
-O2 -qinline=noauto -qinline+salary:taxes:benefits
```

This causes only the functions named `salary`, `taxes`, or `benefits` to be inlined, if possible, and no others.

If you want to use the automatic inlining function, you use the `auto` suboption:

```
-O2 -qinline=auto
```

You can specify an inlining level between 6 and 10 to perform 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 (For example: 7), you enter:

```
-O2 -qinline=level=7
```

## Related information

- “`-g`” on page 169
- “`-qipa`”
- “`-O, -qoptimize`” on page 269
- “The inline function specifier” in the *XL C/C++ Language Reference*
- For a list of deprecated compiler options, see [Deprecated options](#)

## **-qipa**

### Category

Optimization and tuning

### Pragma equivalent

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.

You can generate relinkable objects while preserving IPA information by specifying **-r -qipa=relink**. This creates a nonexecutable package that contains all object files. By using this suboption, you can postpone linking until the very last stage.

If you want to use your own archive files, you can use the **ar** tool and set the **XL\_AR** environment variable to point to its location. If you do not specify a location, the compiler sets the environment variable according to the information contained in the configuration file.

### Note:

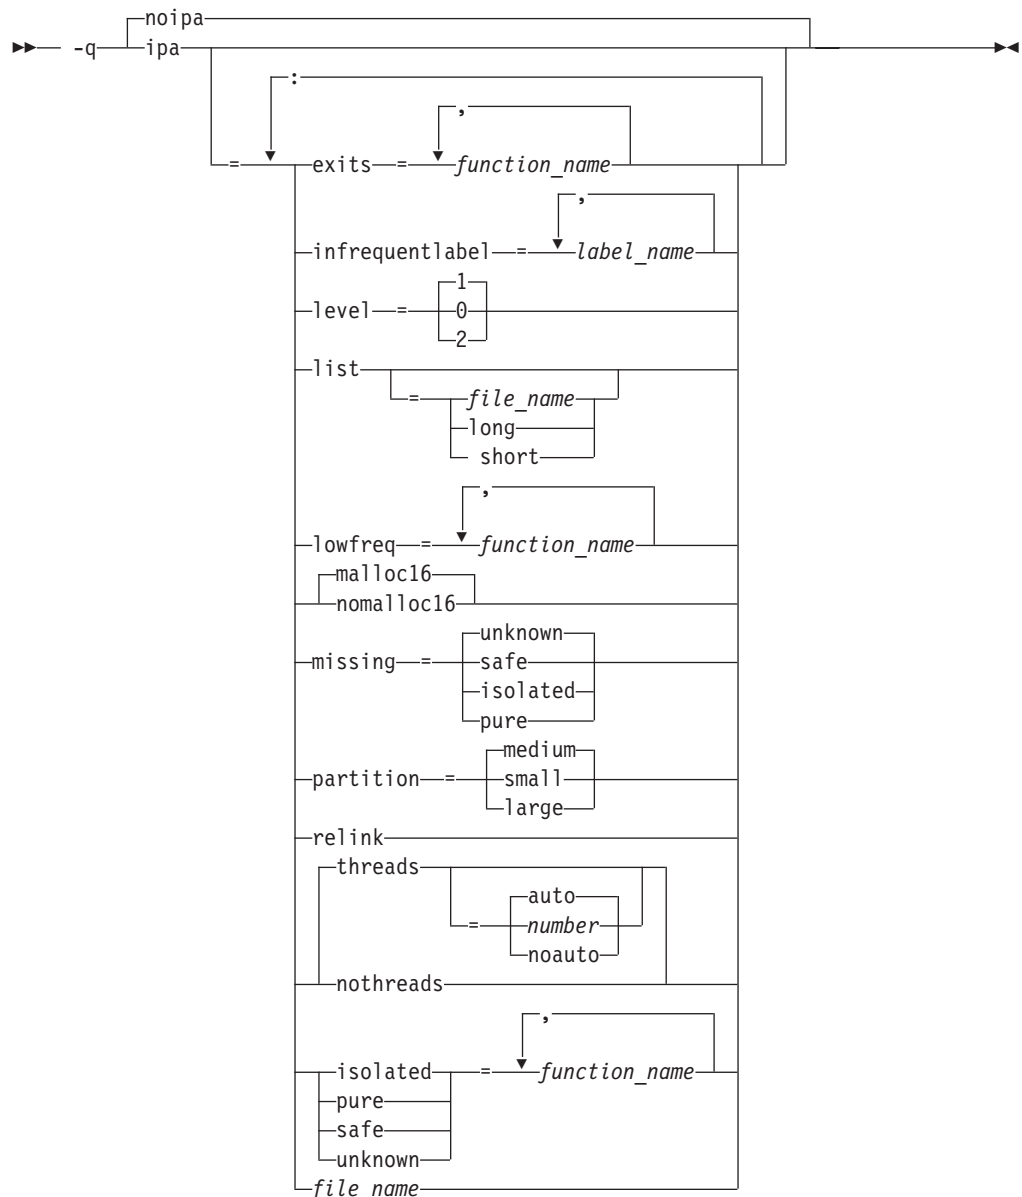
- This suboption does not link the objects; instead, it only aggregates them. As a result, the compiler does not report any error or warning messages; furthermore, the compiler ignores linker or binder options when you use this suboption.
- You must use the **-r** suboption with **-qipa=relink**. Without **-r**, **-qipa=relink** is ignored.
- The **-Wm** option cannot be used with **-qipa**.

## 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 compilation and link step in the same compiler invocation, or during a separate link step only:

**clonearch** | **noclonearch**

This suboption is no longer supported. Consider using **-qtune=balanced**.

**cloneproc** | **nocloneproc**

This suboption is no longer supported. Consider using **-qtune=balanced**.

**exits**

Specifies names of functions which represent program exits. Program exits are calls which can never return and can never call any function which has been compiled with IPA pass 1. The compiler can optimize calls to these functions (for example, by eliminating save/restore sequences), because the calls never return to the program. These functions must not call any other parts of the program that are compiled with **-qipa**.

**infrequentlabel**

Specifies user-defined labels that are likely to be called infrequently during a program run.

*label\_name*

The name of a label, or a comma-separated list of labels.

**isolated**

Specifies a comma-separated list of functions that are not compiled with **-qipa**. Functions that you specify as *isolated* or functions within their call chains cannot refer directly to any global variable.

**level**

Specifies the optimization level for interprocedural analysis. Valid suboptions are as follows:

- 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 **-O5** together with **-qreport**. During the IPA link phase, the data reorganization messages for program variable data are produced in the data reorganization section of the listing file. Reorganizations include 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.c, 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 functions that are likely to be called infrequently. These are typically error handling, trace, or initialization functions. The compiler may be able to make other parts of the program run faster by doing less optimization for calls to these functions.

### **malloc16 | nomalloc16**

Informs the compiler that the dynamic memory allocation routines will return 16-byte aligned memory addresses. The compiler can then optimize the code based on that assertion.

In 64-bit mode, AIX always returns 16-byte aligned addresses and therefore by default **-qipa=malloc16** is in effect. You can use **-qipa=nomalloc16** to override the default setting.

**Note:** You must make sure that the executables generated with **-qipa=malloc16** run in an environment in which dynamic memory allocations return 16-byte aligned addresses, otherwise, wrong results can be generated. For example, in 32-bit mode, addresses are not 16-byte aligned. In this case, you must set the `MALLOCALIGN=16` runtime environment variable.

### **missing**

Specifies the interprocedural behavior of functions 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 functions do not indirectly call a visible (not missing) function either through direct call or through a function pointer.

#### **isolated**

Specifies that the missing functions do not directly reference global variables accessible to visible function. Functions bound from shared libraries are assumed to be *isolated*.

**pure** Specifies that the missing functions are *safe* and *isolated* and do not indirectly alter storage accessible to visible functions. *pure* functions also have no observable internal state.

#### **unknown**

Specifies that the missing functions are not known to be *safe*, *isolated*, or *pure*. This suboption greatly restricts the amount of interprocedural optimization for calls to missing functions.

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 functions, 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* functions that are not compiled with **-qipa**. Any function specified as *pure* must be *isolated* and *safe*, and must not alter the internal state nor have side-effects, defined as potentially altering any data visible to the caller.

#### **relink**

Creates relinkable objects by packaging them into a nonexecutable file. When using this suboption, you must also use the **-r** option along with it. Otherwise, the compiler ignores **-qipa=relink**.

#### **Note:**

- If you use **-qipa=noobject** (either directly or indirectly) and use the **relink** suboption, you must link the resulting object files with **-qipa**. Otherwise, unresolved references to your object files can occur.
- You might indirectly use **-qipa=noobject** if you link and compile your object files in one step. In addition, you cannot use the shared objects with **-qipa=relink**, they must be used at the last link step together with the prelink output.

#### **safe**

Specifies *safe* functions that are not compiled with **-qipa** and do not call any other part of the program. Safe functions can modify global variables, but may not call functions compiled with **-qipa**.

#### **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 one of the following suboptions:

#### **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 spawns 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* functions that are not compiled with **-qipa**. Any function specified as *unknown* can make calls to other parts of the program compiled with **-qipa**, and modify global variables.

#### *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

**Note:**

- **-qipa=inline** and all of its associated suboptions are deprecated. **-qinline** replaces them all. For details, see “-qinline” on page 197 and “Deprecated options” on page 93.
- As of the V9.0 release of the compiler, the **pdfname** suboption is deprecated; you should use **-qpdf1=pdfname** or **-qpdf2=pdfname** in your new applications. See “-qpdf1, -qpdf2” on page 283 for details.

## 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 function to multiple functions (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.

You can use **-r -qipa=relink** to create a relinkable package that contains all object files without generating an executable program. If you want to use your archive files, set the path to your **ar** tool using the **XL\_AR** environment variable.

Some symbols which are clearly referenced or set in the source code may be optimized away by IPA, and may be lost to **debug**, **dump**, 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 C/C++ Optimization and Programming Guide*.

## Predefined macros

None.

## Examples

The following example shows how you might compile a set of files with interprocedural analysis:

```
xlc -c *.c -qipaxlc
-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:

```
xlc -c *.c -qipa=noobject
xlc -c *.o -qipa=lowfreq=user_trace[123]:exit=user_abort
```

The following example demonstrates how you can create a relinkable package that includes your object files:

```
xlc -O5 -o -r -qipa=relink result obj1.o obj2.o obj3.o

ls -l result

-rw-r--r-- result

xlc -O5 -o res result obj4.o obj5.o
```

Here is how you can generate a relinkable package using your own archive files:

```
ar -X64 -r arch1.a object11.o object12.o

ar -X64 -r arch2.a object21.o object22.o

xlc -O5 -o -r -qipa=relink -q64 result obj1.o obj2.o obj3.o arch1.a arch2.a
xlc -O5 -o res result obj4.o obj5.o
```

## Related information

- “-qinline” on page 197
- “-qisolated\_call”
- “-qlibmpi” on page 241
- “#pragma execution\_frequency” on page 398
- -qpdf1, -qpdf2
- -r
- “-S” on page 310
- Deprecated options
- “Optimizing your applications” in the *XL C/C++ Optimization and Programming Guide*
- Runtime environment variables

## -qisolated\_call

### Category

Optimization and tuning

### Pragma equivalent

#pragma options `isolated_call`, #pragma `isolated_call`



## Purpose

Specifies functions in the source file that have no side effects other than those implied by their parameters.

Essentially, any change in the state of the runtime environment is considered a side effect, including:

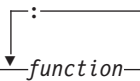
- Accessing a volatile object
- Modifying an external object
- Modifying a static object
- Modifying a file
- Accessing a file that is modified by another process or thread
- Allocating a dynamic object, unless it is released before returning
- Releasing a dynamic object, unless it was allocated during the same invocation
- Changing system state, such as rounding mode or exception handling
- Calling a function that does any of the above

Marking a function as isolated indicates to the optimizer that external and static variables cannot be changed by the called function and that pessimistic references to storage can be deleted from the calling function where appropriate. Instructions can be reordered with more freedom, resulting in fewer pipeline delays and faster execution in the processor. Multiple calls to the same function with identical parameters can be combined, calls can be deleted if their results are not needed, and the order of calls can be changed.

## Syntax

### Option syntax

```
►► -q-isolated_call=function ◀◀
```



### Pragma syntax


```
►► #pragma isolated_call(function) ◀◀
```

## Defaults

Not applicable.

## Parameters

*function*

The name of a function that does not have side effects or does not rely on functions or processes that have side effects. *function* is a primary expression that can be an identifier, operator function, conversion function, or qualified name. An identifier must be of type function or a typedef of function. 

If the name refers to an overloaded function, all variants of that function are marked as isolated calls.

## Usage

The only side effect that is allowed for a function named in the option or pragma is modifying the storage pointed to by any pointer arguments passed to the function, that is, calls by reference. The function is also permitted to examine non-volatile external objects and return a result that depends on the non-volatile state of the runtime environment. Do not specify a function that causes any other side effects; that calls itself; or that relies on local static storage. If a function is incorrectly identified as having no side effects, the program behavior might be unexpected or produce incorrect results.

The **#pragma options isolated\_call** directive must be placed at the top of a source file, before any statements. The **#pragma isolated\_call** directive can be placed at any point in the source file, before or after calls to the function named in the pragma.

The **-qignprag** compiler option causes aliasing pragmas to be ignored; you can use **-qignprag** to debug applications containing the **#pragma isolated\_call** directive.

## Predefined macros

None.

## Examples

To compile `myprogram.c`, specifying that the functions `myfunction(int)` and `classfunction(double)` do not have side effects, enter:

```
xlc myprogram.c -qisolated_call=myfunction:classfunction
```

The following example shows you when to use the **#pragma isolated\_call** directive (on the `addmult` function). It also shows you when not to use it (on the `same` and `check` functions):

```
#include <stdio.h>
#include <math.h>

int addmult(int op1, int op2);
#pragma isolated_call(addmult)

/* This function is a good candidate to be flagged as isolated as its */
/* result is constant with constant input and it has no side effects. */
int addmult(int op1, int op2) {
    int rslt;

    rslt = op1*op2 + op2;
    return rslt;
}

/* The function 'same' should not be flagged as isolated as its state */
/* (the static variable delta) can change when it is called. */
int same(double op1, double op2) {
    static double delta = 1.0;
    double temp;

    temp = (op1-op2)/op1;
    if (fabs(temp) < delta)
        return 1;
    else {
        delta = delta / 2;
        return 0;
    }
}
```

```

}

/* The function 'check' should not be flagged as isolated as it has a */
/* side effect of possibly emitting output. */
int check(int op1, int op2) {
    if (op1 < op2)
        return -1;
    if (op1 > op2)
        return 1;
    printf("Operands are the same.\n");
    return 0;
}

```

### Related information

- “-qignprag” on page 184

## -qkeepinlines (C++ only)

### Category

Object code control

### Pragma equivalent

None.

### Purpose

Keeps or discards definitions for unreferenced extern inline functions.

When **-qnokeepinlines** is in effect, the compiler discards the definitions of unreferenced external inline functions. When **-qkeepinlines** is in effect, the compiler keeps the definitions of unreferenced external inline functions.

### Syntax



### Defaults

-qnokeepinlines

### Parameters

#### exports

Ensures that the compiler does not discard the inline functions that are in the export lists.

### Usage

**-qnokeepinlines** reduces the size of the object files. **-qkeepinlines** provides the same behavior as VisualAge C++ compilers previous to the v5.0.2.1 update level, allowing compatibility with shared libraries and object files built with the earlier releases of the compiler.

If you want the compiler to keep the list of symbols and their definitions that were built by using an earlier version of the compiler, you can use **-qkeepinlines=exports** to make sure that the compiler does not discard these symbols and their definitions while inlining program functions. However, if you do not specify an export file, or the export file does not contain any symbols, the compiler generates the same object file as **-qnokeepinlines**.

When you use **-qkeepinlines=exports** to compile a program, you must use either the **-bE** or the **-bexport** option to specify a file that contains the symbols to export as shown in the following examples:

```
xlc -qmkshrobj -qkeepinlines=exports -bE:file_name source_file
xlc -qmkshrobj -qkeepinlines=exports -bexport:file_name source_file
```

## Predefined macros

None.

## Related information

- “-qmkshrobj” on page 260
- For information about creating a shared library, see the Compiling a shared library section in the *XL C/C++ Optimization and Programming Guide*

## -qkeepparam

### Category

Error checking and debugging

### Pragma equivalent

None.

### Purpose

When used with **-O2** or higher optimization, specifies whether procedure parameters are stored on the stack.

A function usually stores its incoming parameters on the stack at the entry point. However, when you compile code with optimization options enabled, the compiler may remove these parameters from the stack if it sees an optimizing advantage in doing so. When **-qkeepparam** is in effect, parameters are stored on the stack even when optimization is enabled. When **-qnokeepparam** is in effect, parameters are removed from the stack if this provides an optimization advantage.

### Syntax

→ -q ———— nokeepparam ———— keepparam →

### Defaults

-qnokeepparam

## Usage

Specifying `-qkeepparm` that the values of incoming parameters are available to tools, such as debuggers, by preserving those values on the stack. However, this may negatively affect application performance.

## Predefined macros

None.

## Related information

- “-O, -qoptimize” on page 269

## -qkeyword

### Category

Language element control

### Pragma equivalent

None

### Purpose

Controls whether the specified name is treated as a keyword or as an identifier whenever it appears in your program source.

### Syntax

►► -q  = *keyword\_name* ►►

### Defaults

By default, all the built-in keywords defined in the C and C++ language standards are reserved as keywords.

### Usage

You cannot add keywords to the language with this option. However, you can use `-qnokeyword=keyword_name` to disable built-in keywords, and use `-qkeyword=keyword_name` to reinstate those keywords.

#### ► C++

This option can be used with all C++ built-in keywords.

► C++0x This option can be used with the following new keywords introduced by the C++0x standard:

Table 25. New keywords introduced by the C++0x standard

Keyword	Feature	-qlanglvl suboption
constexpr	Generalized constant expressions <sup>1</sup>	-qlanglvl=[no]constexpr

Table 25. New keywords introduced by the C++0x standard (continued)

Keyword	Feature	-q <code>langlvl</code> suboption
<code>decltype</code>	"The <code>decltype(expression)</code> type specifier (C++0x)"	<b>-q<code>langlvl</code>=[no]decltype</b>
<code>static_assert</code>	"static_assert declaration (C++0x)"	<b>-q<code>langlvl</code>=[no]static_assert</b>
<b>Note:</b>		
1. In XL C/C++ V12.1, this feature is a partial implementation of what is defined in the C++0x standard.		

These features introduce new keywords to the C++0x standard in addition to new semantics. To enable each feature semantically, the associated *keyword\_name* (`constexpr`, `decltype`, or `static_assert`) must be recognized as a keyword by the compiler.

You can use the **-q`langlvl`** suboption of each feature or the **-q`langlvl`=extended0x** group option to enable both the feature and the associated keyword. If you enable the **-qkeyword=*keyword\_name*** option, the compiler reserves the *keyword\_name* as a keyword, but the associated feature is not enabled automatically.

**C++0x** ◀

**C++** ◀

**C** ▶ This option can also be used with the following C keywords:

- `asm`
- `inline`
- `restrict`
- `typeof`

**Note:** `asm` is not a keyword when the **-q`langlvl`** option is set to `stdc89` or `stdc99`.

**C** ▶

## Predefined macros

- **C++** ▶ `__BOOL__` is defined to 1 by default; however, it is undefined when **-qnokeyword=bool** is in effect.
- **C** ▶ `__C99_INLINE` is defined to 1 when **-qkeyword=inline** is in effect.
- `__C99_RESTRICT` is defined to 1 when **-qkeyword=restrict** is in effect.
- **C** ▶ `__IBM_GCC_ASM` is defined to 1 when **-qkeyword=asm** is in effect. (In C++ it is defined by default.)
- `__IBM_TYPEDEF__` is defined to 1 when **-qkeyword=typeof** is in effect.

## Examples

**C++** ▶ You can reinstate `bool` with the following invocation:

```
xlc++ -qkeyword=bool
```

**C++** ▶

**C** You can reinstate `typeof` with the following invocation:

```
xlc -qkeyword=typeof
```

**C**

## Related information

- “-qasm” on page 110

**-l**

## Category

Linking

## Pragma equivalent

None.

## Purpose

Searches for the specified library file, *libkey.so*, and then *libkey.a* for dynamic linking, or just for *libkey.a* for static linking.

## Syntax

►► -l—*key*—————►►

## Defaults

The compiler default is to search only 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.

The C and C++ runtime libraries are automatically added.

## Parameters

*key*

The name of the library minus the `lib` characters.

## Usage

You must also provide additional search path information for libraries not located in the default search path. The search path can be modified with the **-L** or **-Z** option. See “-B” on page 116, “-brtl” on page 118, and “-b” on page 115 for information on specifying the types of libraries that are searched (for static or dynamic linking).

The **-l** option is cumulative. Subsequent appearances of the **-l** option on the command line do not replace, but add to, the list of libraries specified by earlier occurrences of **-l**. Libraries are searched in the order in which they appear on the command line, so the order in which you specify libraries can affect symbol resolution in your application.

For more information, refer to the `ld` documentation for your operating system.

## Predefined macros

None.

## Examples

To compile `myprogram.c` and link it with library `mylibrary` (`libmylibrary.a`) found in the `/usr/mylibdir` directory, enter:

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

## Related information

- “-L”
- “-b” on page 115
- “-brtl” on page 118
- “-Z” on page 382
- “Specifying compiler options in a configuration file” on page 7

## -L

### Category

Linking

### Pragma equivalent

None.

### Purpose

Searches the directory path for library files specified by the `-l` option.

### Syntax

►► -L—*directory\_path*—————►►

### Defaults

The default is to search only the standard directories. See the compiler configuration file for the directories that are set by default.

### Parameters

*directory\_path*

The path for the directory which should be searched for library files.

### Usage

When you link shared libraries into an executable, specifying the paths to the libraries with the `-L` option during the link also embeds the path information in the executable, so the shared libraries can be correctly located at run time. If you do not specify any paths with `-L` during this link and you additionally prevent the compiler from automatically passing `-L` arguments to the linker by using the `-bnolibpath` linker option, only paths that are specified by the `LIBPATH` environment variable are embedded in the executable file.



If the `-Ldirectory` option is specified both in the configuration file and on the command line, search paths specified in the configuration file are the first to be searched.

For more information, refer to the `ld` documentation for your operating system.

## Predefined macros

None.

## Examples

To compile `myprogram.c` so that the directory `/usr/tmp/old` is searched for the library `libspfiles.a`, enter:

```
xlc myprogram.c -lspfiles -L/usr/tmp/old
```

## Related information

- “`-l`” on page 213

## `-qlanglvl`

This topic includes the following information:

- “Category”
- “Pragma equivalent”
- “Purpose”
- “Syntax”
- “Defaults” on page 216
- “Parameters for C language programs” on page 219
- “Parameters for C++ language programs” on page 223
- “Usage” on page 237
- “Predefined macros” on page 238

### Category

Language element control

### Pragma equivalent

 `#pragma options langlvl, #pragma langlvl`

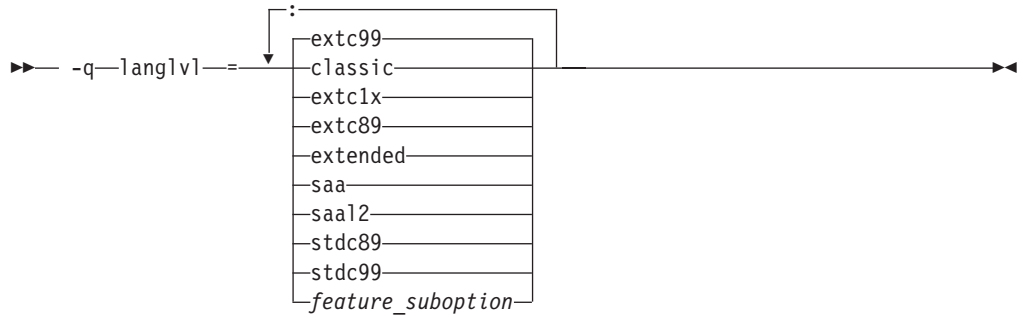
 See also `#pragma operator_new`

### Purpose

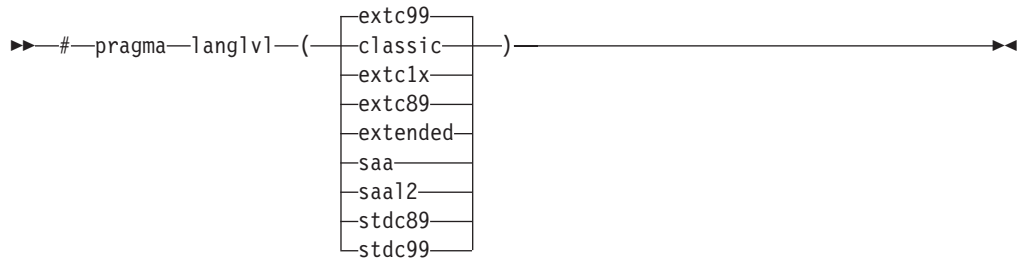
Determines whether source code and compiler options should be checked for conformance to a specific language standard, or subset or superset of a standard.

### Syntax


`-qlanglvl syntax` — C



### #pragma langlvl syntax — C only



### -qlanglvl syntax — C++



## Defaults

- C The default is set according to the command used to invoke the compiler:
  - qlanglvl=extc99:ucs for the `xlc` and related invocation commands
  - qlanglvl=extended:noucs for the `cc` and related invocation commands
  - qlanglvl=stdc89:noucs for the `c89` and related invocation commands
  - qlanglvl=stdc99:ucs for the `c99` and related invocation commands
- C++ The suboptions and their default settings for different language levels (compat366, strict98, extended (C++), and extended0x) are listed in Table 26. The default setting On means that the suboption is enabled; otherwise, the default setting Off means that the suboption is disabled.

Table 26. Default Settings of suboptions for different language levels

Options	Language levels			
	compat366	strict98	extended (C++)	extended0x
-qlanglvl=anonstruct   noanonstruct	Off	Off	On	On
-qlanglvl=anonunion   noanonunion	On	Off	On	On
-qlanglvl=ansifor   noansifor	Off	On	On	On

Table 26. Default Settings of suboptions for different language levels (continued)

Options	Language levels			
	compat366	strict98	extended (C++)	extended0x
-qlanglvl=ansisinit   noansisinit	On	On	On	On
➤ C++0x -qlanglvl=autotypededuction   noautotypededuction	Off	Off	Off	On
-qlanglvl=c1xnoreturn   noc1xnoreturn	Off	Off	On	On
-qlanglvl=c99__func__   noc99__func__	Off	Off	On	On
-qlanglvl=c99complex   noc99complex	Off	Off	Off	Off
-qlanglvl=c99complexheader   noc99complexheader	Off	Off	Off	Off
-qlanglvl=c99compoundliteral   noc99compoundliteral	Off	Off	On	On
-qlanglvl=c99hexfloat   noc99hexfloat	Off	Off	On	On
➤ C++0x -qlanglvl=c99longlong   noc99longlong	Off	Off	Off	On
➤ C++0x -qlanglvl=c99preprocessor   noc99preprocessor	Off	Off	Off	On
-qlanglvl=c99vla   noc99vla	Off	Off	On	On
IBM -qlanglvl=compatrvaluebinding   nocompatrvaluebinding	Off	Off	Off	Off
-qlanglvl=compatzea   nocompatzea	Off	Off	Off	Off
-qlanglvl=complexinit   nocomplexinit	Off	Off	On	On
➤ C++0x -qlanglvl=constexpr   noconstexpr	Off	Off	Off	On
➤ C++0x -qlanglvl=decltype   nodecltype	Off	Off	Off	On
➤ C++0x -qlanglvl=delegatingctors   nodelegatingctors	Off	Off	Off	On
-qlanglvl=dependentbaselookup   nodependentbaselookup	On	On	On	Off
-qlanglvl=emptystruct   noemptystruct	On	On	On	On
➤ C++0x -qlanglvl=explicitconversionoperators   noexplicitconversionoperators	Off	Off	Off	On
➤ C++0x -qlanglvl=extendedfriend   noextendedfriend	Off	Off	Off	On
➤ C++0x IBM -qlanglvl=extendedintegersafe   noextendedintegersafe	Off	Off	Off	Off
➤ C++0x -qlanglvl=externtemplate   noexterntemplate	Off	Off	On	On
-qlanglvl=FileScopeConstExternLinkage   noFileScopeConstExternLinkage	Off	Off	Off	Off

Table 26. Default Settings of suboptions for different language levels (continued)




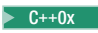
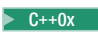
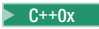



Options	Language levels			
	compat366	strict98	extended (C++)	extended0x
 -qlanglvl=inlinenamespace   noinlinenamespace	Off	Off	Off	On
-qlanglvl=gnu_assert   nognu_assert	Off	Off	On	On
-qlanglvl=gnu_complex   nognu_complex	Off	Off	Off	Off
-qlanglvl=gnu_computedgoto   nognu_computedgoto	Off	Off	On	On
-qlanglvl=gnu_explicitregvar   nognu_explicitregvar	Off	Off	On	On
-qlanglvl=gnu_externtemplate   nognu_externtemplate	Off	Off	On	On
-qlanglvl=gnu_labelvalue   nognu_labelvalue	Off	Off	On	On
-qlanglvl=gnu_locallabel   nognu_locallabel	Off	Off	On	On
-qlanglvl=gnu_include_next   nognu_include_next	On	On	On	On
-qlanglvl=gnu_membernamereuse   nognu_membernamereuse	Off	Off	On	On
-qlanglvl=gnu_suffixij   nognu_suffixij	Off	Off	On	On
-qlanglvl=gnu_varargmacros   nognu_varargmacros	Off	Off	On	On
-qlanglvl=gnu_warning   nognu_warning	Off	Off	On	On
-qlanglvl=illptom   noillptom	On	Off	On	On
-qlanglvl=implicitint   noimplicitint	On	Off	On	On
-qlanglvl=newexcp   nonewexcp	Off	Off	Off	Off
-qlanglvl=offsetnonpod   nooffsetnonpod	On	Off	On	On
-qlanglvl=olddigraph   noolddigraph	Off	On	Off	Off
-qlanglvl=oldfriend   nooldfriend	On	Off	On	Off
-qlanglvl=oldmath   nooldmath	On	Off	Off	Off
-qlanglvl=oldtempacc   nooldtempacc	On	Off	On	On
-qlanglvl=oldtmplalign   nooldtmplalign	On	Off	Off	Off
-qlanglvl=oldtmplspect   nooldtmplspect	On	Off	On	On
-qlanglvl=reddefmac   noreddefmac	Off	Off	Off	Off
 -qlanglvl=referencecollapsing   noreferencecollapsing	Off	Off	Off	On
 -qlanglvl=rightanglebracket   norightanglebracket	Off	Off	Off	On
 -qlanglvl=rvalueresferences   norvalueresferences	Off	Off	Off	On
 -qlanglvl=scopedenum   nosscopedenum	Off	Off	Off	On

Table 26. Default Settings of suboptions for different language levels (continued)

Options	Language levels			
	compat366	strict98	extended (C++)	extended0x
 <code>-qlanglvl=static_assert   nostatic_assert</code>	Off	Off	Off	On
 <code>-qlanglvl=tempsaslocals   notempsaslocals</code>	Off	Off	Off	Off
 <code>-qlanglvl=textafterendif   notextafterendif</code>	Off	Off	Off	Off
<code>-qlanglvl=trailenum   notrailenum</code>	On	Off	On	On
<code>-qlanglvl=typedefclass   notypedefclass</code>	On	Off	On	On
<code>-qlanglvl=noucs   nonoucs</code>	Off	Off	Off	Off
<code>-qlanglvl=varargmacros   novargmacros</code>	Off	Off	On	On
 <code>-qlanglvl=variadic[templates]   novariadic[templates]</code>	Off	Off	Off	On
<code>-qlanglvl=zeroextarray   nozeroextarray</code>	Off	Off	On	On

## Parameters for C language programs

 The following are the `-qlanglvl/#pragma langlvl` parameters for C language programs:

### **classic**

Allows the compilation of nonstandard programs, and conforms closely to the K&R level preprocessor. This language level is not supported by the AIX V5.1 and higher system header files, such as `math.h`. If you use the AIX V5.1 or higher system header files, consider compiling your program to the **stdc89** or **extended** language levels.

The following outlines the differences between the **classic** language level and all other standard-based language levels:

#### **Tokenization**

Tokens introduced by macro expansion may be combined with adjacent tokens in some cases. Historically, this was an artifact of the text-based implementations of older preprocessors, and because, in older implementations, the preprocessor was a separate program whose output was passed on to the compiler.

For similar reasons, tokens separated only by a comment may also be combined to form a single token. Here is a summary of how tokenization of a program compiled in **classic** mode is performed:

1. At a given point in the source file, the next token is the longest sequence of characters that can possibly form a token. For example, `i+++++j` is tokenized as `i ++ ++ + j` even though `i ++ + ++ j` may have resulted in a correct program.
2. If the token formed is an identifier and a macro name, the macro is replaced by the text of the tokens specified on its `#define` directive. Each parameter is replaced by the text of the corresponding argument. Comments are removed from both the arguments and the macro text.

3. Scanning is resumed at the first step from the point at which the macro was replaced, as if it were part of the original program.
4. When the entire program has been preprocessed, the result is scanned again by the compiler as in the first step. The second and third steps do not apply here since there will be no macros to replace. Constructs generated by the first three steps that resemble preprocessing directives are not processed as such.

It is in the third and fourth steps that the text of adjacent but previously separate tokens may be combined to form new tokens.

The `\` character for line continuation is accepted only in string and character literals and on preprocessing directives.

Constructs such as:

```
#if 0
    "unterminated
#endif
#define US "Unterminating string
char *s = US terminated now"
```

will not generate diagnostic messages, since the first is an unterminated literal in a FALSE block, and the second is completed after macro expansion. However:

```
char *s = US;
```

will generate a diagnostic message since the string literal in US is not completed before the end of the line.

Empty character literals are allowed. The value of the literal is zero.

### Preprocessing directives

The `#` token must appear in the first column of the line. The token immediately following `#` is available for macro expansion. The line can be continued with `\` only if the name of the directive and, in the following example, the `(` has been seen:

```
#define f(a,b) a+b
f\
(1,2)      /* accepted */
#define f(a,b) a+b
f(\
1,2)      /* not accepted */
```

The rules concerning `\` apply whether or not the directive is valid. For example,

```
#\
define M 1  /* not allowed */
#def\
ine M 1    /* not allowed */
#define\
M 1        /* allowed */
#dfine\
M 1        /* equivalent to #dfine M 1, even
              though #dfine is not valid */
```

Following are the preprocessor directive differences.

### `#ifndef/#ifndef`

When the first token is not an identifier, no diagnostic message is generated, and the condition is FALSE.

**#else** When there are extra tokens, no diagnostic message is generated.

**#endif**

When there are extra tokens, no diagnostic message is generated.

**#include**

The < and > are separate tokens. The header is formed by combining the spelling of the < and > with the tokens between them. Therefore /\* and // are recognized as comments (and are always stripped), and the " and ' do begin literals within the < and >. (Remember that in C programs, C++-style comments // are recognized when **-qcplusplusmt** is specified.)

**#line** The spelling of all tokens which are not part of the line number form the new file name. These tokens need not be string literals.

**#error**

Not recognized.

**#define**

A valid macro parameter list consists of zero or more identifiers each separated by commas. The commas are ignored and the parameter list is constructed as if they were not specified. The parameter names need not be unique. If there is a conflict, the last name specified is recognized.

For an invalid parameter list, a warning is issued. If a macro name is redefined with a new definition, a warning will be issued and the new definition used.

**#undef**

When there are extra tokens, no diagnostic message is generated.

#### Macro expansion

- When the number of arguments on a macro invocation does not match the number of parameters, a warning is issued.
- If the ( token is present after the macro name of a function-like macro, it is treated as too few arguments (as above) and a warning is issued.
- Parameters are replaced in string literals and character literals.
- Examples:

```
#define M()    1
#define N(a)  (a)
#define O(a,b) ((a) + (b))

M(); /* no error */
N(); /* empty argument */
O(); /* empty first argument
      and too few arguments */
```

#### Text output

No text is generated to replace comments.

#### ▶ C1X

#### extc1x

Compilation is based on the C1X standard, invoking all the currently supported C1X features and other implementation-specific language extensions. For more information about these C1X features, see *Extensions for C1X compatibility* in the *XL C/C++ Language Reference*.

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

C1X <

#### **extc89**

Compilation conforms to the ANSI C89 standard, and accepts implementation-specific language extensions.

#### **extc99**

Compilation conforms to the ISO C99 standard, and accepts implementation-specific language extensions.

#### **extended**

Provides compatibility with the RT compiler and **classic**. This language level is based on C89.

#### **saa**

Compilation conforms to the current SAA C CPI language definition. This is currently SAA C Level 2.

#### **saa12**

Compilation conforms to the SAA C Level 2 CPI language definition, with some exceptions.

#### **stdc89**

Compilation conforms strictly to the ANSI C89 standard, also known as ISO C90.

#### **stdc99**

Compilation conforms strictly to the ISO C99 standard.

**Note:** Not all operating system releases support the header files and runtime library required by C99.

The **-qlanglvl** suboption parameters for individual C features are listed as follows:

*feature\_suboption*

*feature\_suboption* in the syntax diagram represents a colon-separated list of the C options. They can be any of the following options:

**Note:** When multiple **-qlanglvl** group options and suboptions are specified for one individual C feature, the last one takes effect.

IBM

#### **textafterendif | notextafterendif**

Specifies whether to suppress the warning message that is emitted when you are porting code from a compiler that allows extra text after **#endif** or **#else** to the IBM XL C/C++ compiler. The default option is **-qlanglvl=notextafterendif**, indicating that a message is emitted if **#else** or **#endif** is followed by any extraneous text. However, when the language level is **classic**, the default option is **-qlanglvl=textafterendif**, because this language level already allows extra text after **#else** or **#endif** without generating a message.

IBM <



### **ucs | noucs (option only)**

Controls whether Unicode characters are allowed in identifiers, string literals and character literals in program source code. This suboption is enabled by default when **stdc99** or **extc99** is in effect. For details on the Unicode character set, see "The Unicode standard" in the *XL C/C++ Language Reference*.

The following **-qlanglvl** suboptions are accepted but ignored by the C compiler. Use **extended** | **extc99** | **extc89** to enable the functions that these suboptions imply. For other language levels, the functions implied by these suboptions are disabled.

#### **[no]gnu\_assert**

GNU C portability option.

#### **[no]gnu\_explicitregvar**

GNU C portability option.

#### **[no]gnu\_include\_next**

GNU C portability option.

#### **[no]gnu\_locallabel**

GNU C portability option.

#### **[no]gnu\_warning**

GNU C portability option.

## **Parameters for C++ language programs**

**C++** The following are the **-qlanglvl** group option parameters for corresponding C++ language levels:

#### **compat366**

Compilation conforms to some, but not all, IBM C++ Compiler V3.6 features.

#### **strict98**

Compilation conforms strictly to the ISO C++ standard.

#### **extended**

Compilation is based on the ISO C++ standard, with some differences to accommodate extended language features.

#### **C++0x extended0x**

Compilation is based on the C++0x standard, invoking most of the C++ features and all the currently-supported C++0x features. Table 26 on page 216 provides details about the supported features. For more information about C++0x features, see "Extensions for C++0x compatibility" in the *XL C/C++ Language Reference*.

**Note:** C++0x is a new version of the C++ programming language standard. IBM continues to develop and implement the features of the new standard. The implementation of the language level is based on IBM's interpretation of the standard. Until IBM's implementation of all the features of the C++0x standard is complete, including the support of a new C++ standard library, the implementation may change from release to release. IBM makes no attempt to maintain compatibility, in source, binary, or listings and other compiler interfaces, with earlier releases of IBM's implementation of the new features of the C++0x standard and therefore they should not be relied on as a stable programming interface.

**Note:** C++0x has been ratified and published as ISO/IEC 14882:2011. All references to C++0x in this document are equivalent to the ISO/IEC 14882:2011 standard. Corresponding information, including programming interfaces, will be updated in a future release.

The following are the **-qlanglvl** suboption parameters for individual C++ features.

*feature\_suboption*

*feature\_suboption* in the syntax diagram represents a colon-separated list of the remaining C++ options. They can be any of the following:

**Note:** When multiple **-qlanglvl** group options and suboptions are specified for one individual C++ feature, the last one takes effect.

**anonstruct | noanonstruct**

Enables or disables support for anonymous structures and classes. Anonymous structures are typically used in unions, as in the following code fragment:

```
union U {
    struct {
        int i:16;
        int j:16;
    };
    int k;
} u;
// ...
u.j=3;
```

When the default, **-qlanglvl=anonstruct**, is in effect, anonymous structures are supported.

This is an extension to the C++ standard and gives behavior that is designed to be compatible with Microsoft Visual C++. Specify **-qlanglvl=noanonstruct** for compliance with standard C++.

**anonunion | noanonunion**

Controls the members that are allowed in anonymous unions. When the default, **-qlanglvl=anonunion**, is in effect, anonymous unions can have members of all types that standard C++ allows in non-anonymous unions. For example, non-data members, such as structures, typedefs, and enumerations are allowed. Member functions, virtual functions, or objects of classes that have non-trivial default constructors, copy constructors, or destructors cannot be members of a union, regardless of the setting of this option.

This is an extension to standard C++ and gives behavior that is designed to be compatible with previous versions of VisualAge C++ and predecessor products, and Microsoft Visual C++. Specify **-qlanglvl=noanonunion** for compliance with standard C++.

**ansifor | noansifor**

Controls whether scope rules defined in the C++ standard apply to names declared in for loop initialization statements. When the default, **-qlanglvl=ansifor**, is in effect, standard C++ rules are used, and the following code causes a name lookup error:

```
{
    //...
    for (int i=1; i<5; i++) {
        cout << i * 2 << endl;
    }
    i = 10; // error
}
```

The reason for the error is that `i`, or any name declared within a for loop initialization statement, is visible only within the for statement. To correct the error, either declare `i` outside the loop or set **`noansifor`**.

When **`-qlanglvl=noansifor`** is in effect, the old language behavior is used; specify **`-qlanglvl=noansifor`** for compatibility with earlier versions of VisualAge C++ and predecessor products, and Microsoft Visual C++.

#### **`ansisinit`** | **`noansisinit`**

Controls whether standard C++ rules apply for handling static destructors for global and static objects. When the default, **`-qlanglvl=ansisinit`**, is in effect, the standard rules are used.

When **`-qlanglvl=noansisinit`** is in effect, the old language behavior is used; specify **`-qlanglvl=noansisinit`** for compatibility with earlier versions of VisualAge C++ and predecessor products.

#### **`C++0x`** **`autotypededuction`** | **`noautotypededuction`**

Controls whether the auto type deduction feature is enabled. When you specify the **`-qlanglvl=autotypededuction`** option, the auto type deduction feature is enabled, with which you no longer need to specify a type while declaring a variable. Instead, the compiler deduces the type of an auto variable from the type of its initializer expression.

You can also use the **`-qlanglvl=autotypededuction`** option to control the trailing return type feature. This feature is useful when declaring the following types of templates and functions:

- Function templates or member functions of class templates with return types that depend on the types of the function arguments
- Functions or member functions of classes with complicated return types
- Perfect forwarding functions

The **`-qlanglvl=autotypededuction`** option is included in the group option **`-qlanglvl=extended0x`**, so you can also use this group option to enable the auto type deduction feature.

The default option is **`-qlanglvl=noautotypededuction`**.

#### **`c1xnoreturn`** | **`noc1xnoreturn`**

Enables or disables support of the `_Noreturn` function specifier.

The **`-qlanglvl=c1xnoreturn`** option is included in group options **`-qlanglvl=extended`** and **`-qlanglvl=extended0x`**, so you can also use these group options to enable the `_Noreturn` function specifier.

The default option is **`-qlanglvl=noc1xnoreturn`**.

#### **`c99_func__`** | **`noc99_func__`**

Enables or disables support for the C99 `__func__` identifier. For details of this feature, see "func\_predefined identifier" in the *XL C/C++ Language Reference*.

#### **`c99complex`** | **`noc99complex`**

Enables or disables C99 complex data types and related keywords.

#### **`c99complexheader`** | **`noc99complexheader`**

Enables or disables use of the C99 `complex.h` header file.

#### **`c99compoundliteral`** | **`noc99compoundliteral`**

Enables or disables support for C99 compound literals.

#### **`c99hexfloat`** | **`noc99hexfloat`**

Enables or disables support for C99-style hexadecimal floating constants.

**C++0x** **c99longlong** | **noc99longlong**

Controls whether the C99 long long feature is enabled. When you specify the **-qlanglvl=c99longlong** option, the C++ compiler provides the C99 long long feature, which improves source compatibility between the C and C++ languages.

The **-qlanglvl=c99longlong** option conflicts with the **-qlonglong** option. If you specify both these two options, the **-qlonglong** option is ignored. For more information about the **-qlonglong** option, see “-qlonglong” on page 249.

The **-qlanglvl=c99longlong** option is included in the group option **-qlanglvl=extended0x**, so you can also use this group option to enable the C99 long long feature.

The default option is **-qlanglvl=noc99longlong**.

**C++0x** **c99preprocessor** | **noc99preprocessor**

Controls whether the C99 preprocessor features adopted in C++0x are enabled. When **-qlanglvl=c99preprocessor** is in effect, the C99 and C++0x compilers provide a more common preprocessor interface, which can ease porting C source files to the C++ compiler and avoid preprocessor compatibility issues.

The default option is **-qlanglvl=noc99preprocessor**.

**Note:** Specifying **-qlanglvl=c99preprocessor** implicitly sets **-qlanglvl=varargmacros**. Also, specifying **-qlanglvl=noc99preprocessor** implicitly sets **-qlanglvl=novarargmacros**.

The **-qlanglvl=c99preprocessor** option is included in the group option **-qlanglvl=extended0x**, so you can also use this group option to enable the C99 preprocessor feature.

**c99vla** | **noc99vla**

Enables or disables support for C99-type variable length arrays.

**IBM** **compatrvaluebinding** | **nocompatrvaluebinding**

The C++ Standard (2003) indicates that an rvalue can only be bound to a const non-volatile lvalue reference. Non-compliant compilers might allow a non-const or volatile lvalue reference to be bound to an rvalue. When you are porting code to IBM XL C/C++ compiler, you can specify this option to instruct the compiler to allow a non-const or volatile lvalue reference to bind to an rvalue of a user-defined type where an initializer is not required. **IBM**

**C++0x** If both the **-qlanglvl=compatrvaluebinding** and **-qlanglvl=rvalueresferences** options are in effect, the compiler issues an error message. **C++0x**

**compatzea** | **nocompatzea**

Controls whether zero extent arrays have an underlying dimension of 1 or 0. When the default, **-qlanglvl=nocompatzea**, is in effect, zero extent arrays have a dimension of 0. Use **-qlanglvl=compatzea** to specify that zero extent arrays have a dimension of 1, for compatibility with code compiled with VisualAge C++ V6.0 and predecessor products. Specifying **-qlanglvl=compatzea** has effect only if **-qlanglvl=zeroextarray** is also in effect.

**C++** **complexinit** | **nocomplexinit**

Controls whether the C++ compiler uses the C1X style initialization of the C99 complex types.

The **-qlanglvl=complexinit** option is included in the group options **-qlanglvl=extended** and **-qlanglvl=extended0x**, so you can also use these

group options to enable the initialization of complex types. In this case, specify `-qlanglvl=c99complexheader` so that correct header files can be used by the compiler.

The default option is `-qlanglvl=complexinit`. C++

C++0x **constexpr** | **noconstexpr**

Controls whether the generalized constant expressions feature is enabled. When you specify the `-qlanglvl=constexpr` option, the compiler extends the expressions permitted within constant expressions. A constant expression is one that can be evaluated at compile time.

The `-qlanglvl=constexpr` option is included in the group option `-qlanglvl=extended0x`, so you can also use this group option to enable the generalized constant expressions feature.

The default option is `-qlanglvl=noconstexpr`.

**Note:** In XL C/C++ V12.1, this feature is a partial implementation of what is defined in the C++0x standard. Here are the implementation details:

- Support valid syntax using the `constexpr` keyword.
- Support valid definitions of `constexpr` functions whose arguments and return type are of built-in types.
- Diagnostic messages might not be issued for invalid `constexpr` function definitions. A function call to an invalid `constexpr` function definition might prevent the compiler from evaluating the result, and an error is issued depending on the context where a constant expression is required.

C++0x **decltype** | **nodecltype**

Controls whether the `decltype` feature is enabled. With this feature, you can get a type that is based on the resultant type of a possibly type-dependent expression. To enable this feature, you can specify the `-qlanglvl=decltype` option.

The `-qlanglvl=decltype` option is included in the group option `-qlanglvl=extended0x`, so you can also use this group option to enable the `decltype` feature.

The default option is `-qlanglvl=nodecltype`.

C++0x **delegatingctors** | **nodelegatingctors**

Controls whether the delegating constructors feature is enabled. With this feature, you can concentrate on common initializations and post initializations in one constructor, which can make programs more readable and maintainable. To enable this feature, you can specify the `-qlanglvl=delegatingctors` option.

The `-qlanglvl=delegatingctors` option is included in the group option `-qlanglvl=extended0x`, so you can also use this group option to enable the delegating constructors feature.

The default option is `-qlanglvl=nodelegatingctors`.

**DependentBaseLookup** | **noDependentBaseLookup**

Controls whether the name lookup rules for a template base class of dependent type defined in the Technical Corrigendum 1 (TC1) of the C++ Standard apply. Specify `-qlanglvl=noDependentBaseLookup` for compliance with TC1. When `-qlanglvl=noDependentBaseLookup` is in effect, unqualified names in a template class will not be resolved in a base class if that base class is dependent on a template parameter. These names must be qualified with the

base class name in order to be found by name lookup. When the default, **-q`langlvl=DependentBaseLookup`**, is in effect, the behavior of previous XL C++ compilers remains.

**Note:** The default option is **-q`langlvl=noDependentBaseLookup`** at the C++0x language level.

The following example shows code that does not compile with **-q`langlvl=noDependentBaseLookup`**:

```
struct base
{
    int baseName;
};

template <class B> struct derived : public B
{
    void func()
    {
        int i = baseName;    // this name will not be found in the base class
    };
};

int main(void)
{
    derived<base> x;
    x.func();
    return 0;
}
```

The following example shows code that compiles with or without **-q`langlvl=nodependentbaselookup`**:

```
struct base
{
    int baseName;
};

template <class B> struct derived : public B
{
    void func()
    {
        int i = B::baseName; // qualified name will be found in the base class
    };
};

int main(void)
{
    derived<base> x;
    x.func();
    return 0;
}
```

#### **empty\_struct | noempty\_struct**

This option instructs the compiler to tolerate empty member declarations in structs. Empty member declaration in structs is not allowed. For example, when **-q`langlvl=noemptystruct`** is in effect, the following example will be rejected by the compiler:

```
struct S {
    ; // this line is ill-formed
};
```

The default is **-q`langlvl=noemptystruct`**.

#### **C++0x** **explicitconversionoperators | noexplicitconversionoperators**

Controls whether the explicit conversion operators feature is enabled. When

you specify the `-qclanglvl=explicitconversionoperators` option, you can apply the explicit function specifier to the definition of a user-defined conversion function, and thus inhibit unintended implicit conversions through the user-defined conversion function.

The `-qclanglvl=explicitconversionoperators` option is included in the group option `-qclanglvl=extended0x`, so you can also use this group option to enable the explicit conversion operators feature.

The default option is `-qclanglvl=noexplicitconversionoperators`.

**C++0x** `extendedfriend` | `noextendedfriend`

Controls whether the extended friend declarations feature is enabled. When you specify the `-qclanglvl=extendedfriend` option, rules governing friend declarations are relaxed as follows:

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

The `-qclanglvl=extendedfriend` option is included in the group option `-qclanglvl=extended0x`, so you can also use this group option to enable the extended friend declarations feature.

The default option is `-qclanglvl=noextendedfriend`.

**Note:** `-qclanglvl=extendedfriend` is incompatible with the `-qclanglvl=oldfriend` option. When `-qclanglvl=extendedfriend` is in effect, the `-qclanglvl=oldfriend` option is ignored and the setting of `-qclanglvl=[no]oldfriend` is `-qclanglvl=nooldfriend`.

**C++0x** **IBM** `extendedintegersafe` | `noextendedintegersafe`

With this option, if a decimal integer literal that does not have a suffix containing `u` or `U` cannot be represented by the `long long int` type, you can decide whether to use the unsigned `long long int` type to represent the literal or not.

This option takes effect only when the `-qclanglvl=c99longlong` option is specified, otherwise, the compiler issues a warning message to indicate that the option is ignored. When you specify both the `-qclanglvl=c99longlong` and `-qclanglvl=extendedintegersafe` options, if a decimal integer literal that does not have a suffix containing `u` or `U` cannot be represented by the `long long int` type, the compiler issues an error message stating that the value of the literal is out of range.

The default option is `-qclanglvl=noextendedintegersafe` in all the language levels.

**C++0x** `externtemplate` | `noexterntemplate`

Controls whether the explicit instantiation declarations feature is enabled. With this feature, you can suppress the implicit instantiations of a template specialization or its members. To enable this feature, you can specify the `-qclanglvl=externtemplate` option, which is the default option.

The `-qclanglvl=externtemplate` option is included in the group options of `-qclanglvl=extended` and `-qclanglvl=extended0x`, so you can use these two group options to enable this feature.

The following table lists options that interact with the `-qclanglvl=externtemplate` option:

Table 27. Options that interact with `-qlanglvl=externtemplate`

Option	Description
<code>-qtemplateregistry, -qtempinc</code>	Explicit instantiation declarations remain effective. Referenced specializations that are the subjects of explicit instantiation declarations, but not the subjects of explicit instantiation definitions in a translation unit are not instantiated from or because of that translation unit.

The following table lists IBM language extensions that interact with the `-qlanglvl=externtemplate` option:

Table 28. IBM language extensions that interact with `-qlanglvl=externtemplate`

IBM language extension	Description
<code>#pragma instantiate</code>	This pragma is semantically the same as an explicit instantiation definition.
<code>#pragma do_not_instantiate</code>	This pragma provides a subset of the functionality of the explicit instantiation declarations which is introduced in the C++0x standard. It is provided for backwards compatibility purposes only. New applications can use explicit instantiation declarations.
<code>#pragma hashome, #pragma ishome</code>	This pragma causes the generation of the virtual function table (VFT) for a class template specialization irrespective of explicit instantiation declarations of the specialization.

The `-qlanglvl=[no]externtemplate` option replaces the deprecated `-qlanglvl=[no]gnu_externtemplate` option. Use the `-qlanglvl=[no]externtemplate` option in your applications.

#### `FileScopeConstExternLinkage` | `noFileScopeConstExternLinkage`

Controls whether the file scope of constant variables has internal or external linkage when the `static` or `extern` keyword is not specified.

When `-qlanglvl=FileScopeConstExternLinkage` is in effect, all file scope constant variables are marked as externally visible. Otherwise, all file scope constant variables are marked as `static`.

The default is `-qlanglvl=noFileScopeConstExternLinkage`.

#### `gnu_assert` | `nognu_assert`

Enables or disables support for the following GNU C system identification assertions:

- `#assert`
- `#unassert`
- `#cpu`
- `#machine`
- `#system`

#### `gnu_complex` | `nognu_complex`

Enables or disables GNU complex data types and related keywords.

#### `gnu_computedgoto` | `nognu_computedgoto`

Enables or disables support for computed `goto` statements.

#### `gnu_externtemplate` | `nognu_externtemplate`

Enables or disables extern template instantiations. For details of this feature, see "Explicit instantiation" in the *XL C/C++ Language Reference*.

**Note:** The option `-qlanglvl=[no]gnu_externtemplate` is deprecated in XL C/C++ V12.1; you can use the option `-qlanglvl=[no]externtemplate` instead.



**gnu\_include\_next | nognu\_include\_next**

Enables or disables support for the GNU C `#include_next` preprocessor directive.

**gnu\_labelvalue | nognu\_labelvalue**

Enables or disables support for labels as values.

**gnu\_locallabel | nognu\_locallabel**

Enables or disables support for locally-declared labels.

**gnu\_membernamereuse | nognu\_membernamereuse**

Enables or disables reusing a template name in a member list as a typedef.

**gnu\_suffixij | nognu\_suffixij**

Enables or disables support for GNU-style complex numbers. When `-qlanglvl=gnu_suffixij` is in effect, a complex number can be ended with suffix `i/I` or `j/J`.

**gnu\_varargmacros | nognu\_varargmacros**

Enables or disables support for GNU-style macros with variable arguments.

For details of this feature, see "Variadic macro extensions" in the *XL C/C++ Language Reference*.

**gnu\_warning | nognu\_warning**

Enables or disables support for the GNU C `#warning` preprocessor directive.

**illptom | noillptom**

Controls the expressions that can be used to form pointers to members. When the default, `-qlanglvl=illptom`, is in effect, the XL C++ compiler accepts some forms that are in common use but do not conform to the C++ Standard. For example, the following code defines a pointer to a function member, `p`, and initializes it to the address of `C::func`, in the old style:

```
struct C {
void func(int);
};

void (C::*p) (int) = C::func;
```

This is an extension to standard C++ and gives behavior that is designed to be compatible with earlier versions of VisualAge C++ and its predecessor products, and Microsoft Visual C++.

Specify `-qlanglvl=noillptom` for compliance with the C++ standard. The example code above must be modified to use the `&` operator.

```
struct C {
void func(int);
};

void (C::*p) (int) = &C::func;
```

**implicitint | noimplicitint**

Controls whether the compiler accepts missing or partially specified types as implicitly specifying `int`. When the default, `-qlanglvl=implicitint`, is in effect, a function declaration at namespace scope or in a member list will implicitly be declared to return `int`. Also, any declaration specifier sequence that does not completely specify a type will implicitly specify an integer type. The effect is as if the `int` specifier were present.

The following specifiers do not completely specify a type:

- `auto`
- `const`

- extern
- extern "literal"
- inline
- mutable
- friend
- register
- static
- typedef
- virtual
- volatile
- platform-specific types

**C++0x** C++0x has removed the use of auto as a storage class specifier. In C++0x, the keyword auto is used as a type specifier. The compiler deduces the type of an auto variable from the type of its initializer expression. For more information, see "The auto type specifier (C++0x)" in the *XL C/C++ Language Reference*.

For example, the return type of function `MyFunction` is `int` because it was omitted in the following code:

```
MyFunction()
{
    return 0;
}
```

Note that any situation where a type is specified is affected by this suboption. This includes, for example, template and parameter types, exception specifications, types in expressions (eg, casts, `dynamic_cast`, `new`), and types for conversion functions.

This is an extension to the C++ standard and gives behavior that is designed to be compatible with earlier versions of VisualAge C++ and predecessor products, and Microsoft Visual C++.

Specify **-qlanglvl=noimplicitint** for compliance with standard C++. For example, the function declaration above must be modified to:

```
int MyFunction()
{
    return 0;
}
```

**C++0x** **inlinenamespace | noinlinenamespace**

Controls whether inline namespace definitions are enabled, which are namespace definitions preceded by an initial `inline` keyword. A namespace so defined is an inline namespace. When you specify the **-qlanglvl=inlinenamespace** option, members of the inline namespace can be defined and specialized as if they were also members of the enclosing namespace.

The **-qlanglvl=inlinenamespace** option is included in the group option **-qlanglvl=extended0x**, so you can also use this group option to enable the inline namespace definitions feature.

The default option is **-qlanglvl=noinlinenamespace**.

**newexcp | nonewexcp**

Controls whether the `new` operator throws an exception when the requested memory fails. When the default, **-qlanglvl=nonewexcp**, is in effect, the null pointer `0` is returned. When **-qlanglvl=newexcp** is in effect, the standard exception `std::bad_alloc` is thrown. For compatibility with earlier versions of

VisualAge C++ and predecessor products, specify **-qlanglvl=nonewexcp**. For conformance to the C++ standard, which fully supports new exceptions, specify **-qlanglvl=newexcp**.

This suboption does not apply to the nothrow versions of the new operator, new operators with empty throw specifications, class-specific new operators, and new operators with placement arguments.

**Note:** You can also use the equivalent **#pragma operator\_new** directive to specify this suboption for selected portions of code. See “**#pragma operator\_new** (C++ only)” on page 419 for details.

#### **offsetnonpod | nooffsetnonpod**

Controls whether the `offsetof` macro can be applied to classes that are not data-only. C++ programmers often casually call data-only classes “Plain Old Data” (POD) classes. When the default, **-qlanglvl=offsetnonpod**, is in effect, you can apply `offsetof` to a class that contains one of the following:

- user-declared constructors or destructors
- user-declared assignment operators
- private or protected non-static data members
- base classes
- virtual functions
- non-static data members of type pointer to member
- a struct or union that has non-data members
- references

This is an extension to the C++ standard, and gives behavior that is designed to be compatible with VisualAge C++ for OS/2 3.0, VisualAge for C++ for Windows, V3.5, and Microsoft Visual C++. Specify **-qlanglvl=nooffsetnonpod** for compliance with standard C++.

#### **olddigraph | noolddigraph**

Enables or disables support for old-style digraphs. When the default, **-qlanglvl=olddigraph**, is in effect, old-style digraphs are not supported. When **-qlanglvl=olddigraph** is in effect, the following digraphs are supported:

##### **Digraph**

###### **Resulting character**

%% # (pound sign)

%%%%

## (double pound sign, used as the preprocessor macro concatenation operator)

Specify **-qlanglvl=noolddigraph** for compatibility with standard C++ and the extended C++ language level supported by previous versions of VisualAge C++ and predecessor products.

This suboption only has effect when **-qdigraphs** is in effect.

#### **oldfriend | nooldfriend**

Controls whether friend declarations that name classes without elaborated class names are treated as C++ errors. When the default, **-qlanglvl=oldfriend**, is in effect, you can declare a friend class without elaborating the name of the class with the keyword `class`. For example, the statement below declares the class `IFont` to be a friend class:

```
friend IFont;
```

This is an extension to the C++ standard and gives behavior that is designed to be compatible with earlier versions of VisualAge C++ and predecessor products, and Microsoft Visual C++.

Specify the **-qlanglvl=nooldfriend** for compliance with standard C++. The example declaration above must be modified to the following:

```
friend class IFont;
```

**Note:** **-qlanglvl=oldfriend** is incompatible with the **-qlanglvl=extendedfriend** option. When **-qlanglvl=extendedfriend** is in effect, the **-qlanglvl=oldfriend** option is ignored and the setting of **-qlanglvl=[no]oldfriend** is **-qlanglvl=nooldfriend**.

#### **oldmath | nooldmath**

Controls the versions of math function declarations in math.h that are included when you specify math.h as an included or primary source file.

Specify **-qlanglvl=nooldmath** for strict compliance with the C++ standard.

Specify **-qlanglvl=oldmath** for compatibility with earlier versions of VisualAge C++ and predecessor products.

#### **oldtempacc | nooldtempacc**

Controls whether access to a copy constructor to create a temporary object is always checked, even if creation of the temporary object is avoided. When the default, **-qlanglvl=oldtempacc**, is in effect, access checking is suppressed.

This is an extension to the C++ standard and gives behavior that is designed to be compatible with VisualAge C++ for OS/2 3.0, VisualAge for C++ for Windows, V3.5, and Microsoft Visual C++. Specify **-qlanglvl=nooldtempacc** for compliance with standard C++. For example, the throw statement in the following code causes an error because the copy constructor is a protected member of class C:

```
class C {
public:
    C(char *);
protected:
    C(const C&);
};

C func() {return C("test");} // return copy of C object
void f()
{
// catch and throw both make implicit copies of
// the throw object
    throw C("error"); // throw a copy of a C object
    const C& r = func(); // use the copy of a C object
//                          created by func()
}
```

The example code above contains three ill formed uses of the copy constructor C(const C&).

#### **oldtmplalign | nooldtmplalign**

Controls whether alignment rules specified for nested templates are ignored.

When the default, **-qlanglvl=nooldtmplalign**, is in effect, these alignment rules are not ignored. For example, given the following template the size of A<char>::B will be 5 with **-qlanglvl=nooldtmplalign**, and 8 with

**-qlanglvl=oldtmplalign :**

```
template <class T>
struct A {
#pragma options align=packed
```

```

struct B {
    T m;
    int m2;
};
#pragma options align=reset
};

```

Specify **-qlanglvl=oldtmplalign** for compatibility with VisualAge for C++ V4.0 and predecessor products.

#### **oldtmpspec | nooldtmpspec**

Controls whether template specializations that do not conform to the C++ standard are allowed. When the default, **-qlanglvl=oldtmpspec**, is in effect, you can explicitly specialize a template class as in the following example, which specializes the template class `ribbon` for type `char`:

```

template<class T> class ribbon { /*...*/};
class ribbon<char> { /*...*/};

```

This is an extension to standard C++ and gives behavior that is designed to be compatible with VisualAge C++ for OS/2 3.0, VisualAge for C++ for Windows, V3.5, and Microsoft Visual C++.

Specify **-qlanglvl=nooldtmpspec** for compliance with standard C++. In the example above, the template specialization must be modified to:

```

template<class T> class ribbon { /*...*/};
template<> class ribbon<char> { /*...*/};

```

#### **redefmac | noredefmac**

Controls whether a macro can be redefined without a prior `#undef` or `undefine()` statement.

#### **C++0x referencecollapsing | noreferencecollapsing**

Controls whether the reference collapsing feature is enabled. To enable this feature, specify the **-qlanglvl=referencecollapsing** option.

The **-qlanglvl=referencecollapsing** option is included in the group option **-qlanglvl=extended0x**, so you can also use this group option to enable the reference collapsing feature.

When the **-qlanglvl=rvalueresferences** option is in effect, but the **-qlanglvl=referencecollapsing** option is not in effect, the compiler behaves as if the **-qlanglvl=referencecollapsing** option were specified.

The default option is **-qlanglvl=noreferencecollapsing**.

#### **C++0x rightanglebracket | norightanglebracket**

Controls whether the right angle bracket feature is enabled. To enable this feature, you can specify the **-qlanglvl=rightanglebracket** option.

The **-qlanglvl=rightanglebracket** option is included in the group option **-qlanglvl=extended0x**, so you can also use this group option to enable the right angle bracket feature.

The default option is **-qlanglvl=norightanglebracket**.

#### **C++0x rvalueresferences | norvalueresferences**

Controls whether the rvalue references feature is enabled. To enable this feature, specify the **-qlanglvl=rvalueresferences** option.

The **-qlanglvl=rvalueresferences** option is included in the group option **-qlanglvl=extended0x**, so you can also use this group option to enable the rvalue references feature.

If both the `-qlanglvl=compatvaluebinding` and `-qlanglvl=rvaluereferences` options are in effect, the compiler issues an error message.

The default option is `-qlanglvl=norvaluereferences`.

**C++0x** `scopedenum` | `nosscopedenum`

Controls whether the scoped enumeration feature is enabled. To enable this feature, you can specify the `-qlanglvl=scopedenum` option.

The `-qlanglvl=scopedenum` option is included in the group option `-qlanglvl=extended0x`, so you can also use this group option to enable the scoped enumeration feature.

The default option is `-qlanglvl=nosscopedenum`.

**C++0x** `static_assert` | `nostatic_assert`

Controls whether the static assertions feature is enabled. When `-qlanglvl=static_assert` is in effect, this feature can be used to produce compile-time assertions for which a severe error message is issued on failure.

`-qlanglvl=static_assert` is included in the group option `-qlanglvl=extended0x`, so you can also use this group option to enable the static assertions feature.

The default is `-qlanglvl=nostatic_assert`.

**IBM** `tempsaslocals` | `notempsaslocals`

The C++ Language Standard describes the lifetime of temporaries in section Temporary Object [class.temporary]. When you are porting an application from a compiler that implements late temporary destruction, you might need to extend the lifetime of C++ temporaries beyond which is specified in the C++ Language Standard. This option extends the lifetime of temporaries to reduce migration difficulty.

`textafterendif` | `notextafterendif`

Specifies whether to suppress the warning message that is emitted when you are porting code from a compiler that allows extra text after `#endif` or `#else` to IBM XL C/C++ compiler. The default option is `-qlanglvl=notextafterendif`, indicating that a message is emitted if `#else` or `#endif` is followed by any extraneous text. **IBM**

`trailenum` | `notrailenum`

Controls whether trailing commas are allowed in enum declarations. When the default, `-qlanglvl=trailenum`, is in effect, one or more trailing commas are allowed at the end of the enumerator list. For example, the following enum declaration uses this extension:

```
enum grain { wheat, barley, rye, , };
```

This is an extension to the C++ standard, and is intended to provide compatibility with Microsoft Visual C++.

Specify `-qlanglvl=notrailenum` for compliance with standard C++.

`typedefclass` | `notypedefclass`

Controls whether a typedef name can be specified where a class name is expected. When the default, `-qlanglvl=typedefclass`, is in effect, the standard C++ rule applies, and a typedef name cannot be specified where a class name is expected. Specify `-qlanglvl=typedefclass` to allow the use of typedef names in base specifiers and constructor initializer lists, for compatibility with earlier versions of VisualAge for C++ and predecessor products.

`ucs` | `noucs`

Controls whether Unicode characters are allowed in identifiers, string literals

and character literals in program source code. For details on the Unicode character set, see "The Unicode standard" in the *XL C/C++ Language Reference*.

### **varargmacros** | **novargmacros**

Enables or disables support for C99-style variable argument lists in function-like macros.

**Note:** Specifying **-qlanglvl=c99preprocessor** implicitly set **-qlanglvl=varargmacros**. Vice versa, specifying **-qlanglvl=noc99preprocessor** implicitly set **-qlanglvl=novargmacros**.

For details of this feature, see "Function-like macros" in the *XL C/C++ Language Reference*.

### ► **C++0x** **variadic[templates]** | **novariadic[templates]**

Controls whether the variadic templates feature is enabled. With this feature, you can define class and function templates that have any number (including zero) of parameters. To enable this feature, you can specify the **-qlanglvl=variadic[templates]** option. The word *templates* included in the brackets is optional. If you specify only the **-qlanglvl=variadic** option, the compiler assumes that the **-qlanglvl=variadictemplates** option is specified.

The **-qlanglvl=variadic[templates]** option is included in the group option **-qlanglvl=extended0x**, so you can also use this group option to enable the variadic templates feature.

The default option is **-qlanglvl=novariadic[templates]**.

### **zeroextarray** | **nozeroextarray**

Controls whether you can use zero-extent arrays as the last nonstatic data member in a structure definition. When the default, **-qlanglvl=zeroextarray**, is in effect, you can use arrays with zero elements. The following statement declares a zero-extent array `a`.

```
struct S1 { char a[0]; };
```

This is an extension to the C++ standard, and is intended to provide compatibility with Microsoft Visual C++.

Specify **-qlanglvl=nozeroextarray** for compliance with standard C++ or with the ANSI language level supported by previous versions of VisualAge C++ and predecessor products.

## Usage

► **C++** In general, if you specify a suboption with the **no** form of the option, the compiler will diagnose any uses of the feature in your code with a warning, unless you disable the warning with the **-qsuppress** option. Additionally, you can use the **-qinfo=por** option to generate informational messages along with the following suboptions:

- [no]c99complex
- [no]gnu\_complex

**Note:** In the C++0x language level, if you use the **no** form of a suboption to disable the C++0x meaning of `decltype` or `static_assert`, the compiler emits syntax errors but no diagnostic message if the user happens to use the C++0x syntax of `decltype` or `static_assert`.

► **C** Since the `pragma` directive makes your code non-portable, it is recommended that you use the option rather than the `pragma`. If you do use the

pragma, it must appear before any noncommentary lines in the source code. Also, because the directive can dynamically alter preprocessor behavior, compiling with the preprocessing-only options may produce results different from those produced during regular compilation.

## Predefined macros

See "Macros related to language levels" on page 470 for a list of macros that are predefined by **-qlanglvl** suboptions.

## Related information

- "-qsuppress" on page 338
- "The IBM XL C language extensions" in the *XL C/C++ Language Reference* and "The IBM XL C++ language extensions" in the *XL C/C++ Language Reference*

## -qlargepage

### Category

Optimization and tuning

### Pragma equivalent

None.

### Purpose

Takes advantage of large pages provided on POWER4 and higher systems, for applications designed to execute in a large page memory environment.

When **-qlargepage** is in effect to compile a program designed for a large page environment, an increase in performance can occur.

### Syntax

```
►► -q { nolargepage | largepage } _____►►
```

### Defaults

-qnolargepage

### Usage

Note that this option is only useful in the following conditions:

- Large pages must be available and configured on the system.
- You must compile with an option that enables loop optimization, such as **-O3** or **-qhot**.
- You must link with the **-blpdata** option.

See your AIX operating system documentation for more information on using large page support.



## Predefined macros

None.

## Examples

To compile `myprogram.c` to use large page heaps, enter:

```
xlc myprogram.c -qlargepage -blpdata
```

## -qdbl128, -qlongdouble

### Category

Floating-point and integer control

### Pragma equivalent

```
#pragma options [no]dbl128
```

### Purpose

Increases the size of long double types from 64 bits to 128 bits.

### Syntax



### Defaults

```
-qnoldbl128
```

### Usage

Separate libraries are provided that support 128-bit long double types. These libraries will be automatically linked if you use any of the invocation commands with the **128** suffix (`xlc++128`, `xlc128`, `cc128`, `xlc++128_r`, `xlc128_r`, or `cc128_r`). You can also manually link to the 128-bit versions of the libraries using the `-lkey` option, as shown in the following table:

Default (64-bit) long double		128-bit long double	
Library	Form of the <i>-lkey</i> option	Library	Form of the <i>-lkey</i> option
libC.a	-lC	libC128.a	-lC128
libC_r.a	-lC_r	libC128_r.a	-lC128_r

Linking without the 128-bit versions of the libraries when your program uses 128-bit long doubles (for example, if you specify `-qdbl128` alone) may produce unpredictable results.

The `#pragma options` directive must appear before the first C or C++ statement in the source file, and the option applies to the entire file.

## Predefined macros

- `__LONGDOUBLE128` is defined to 1 when `-qldb1128` is in effect; otherwise, it is undefined.
- `__LONGDOUBLE64` is defined to 1 when `-qnoldb1128` is in effect; it is undefined when `-qldb1128` is in effect.

## Examples

To compile `myprogram.c` so that long double types are 128 bits, enter:

```
xlc myprogram.c -qldb1128 -lC128
```

## Related information

- “-l” on page 213

## -qlib

### Category

Linking

### Pragma equivalent

None.

### Purpose

Specifies whether standard system libraries and XL C/C++ libraries are to be linked.

When `-qlib` is in effect, the standard system libraries and compiler libraries are automatically linked. When `-qno lib` is in effect, the standard system libraries and compiler libraries are not used at link time; only the libraries specified on the command line with the `-l` flag will be linked.

This option can be used in system programming to disable the automatic linking of unneeded libraries.

### Syntax

```
► — -q ————— ►
```

### Defaults

`-qlib`

### Usage

Using `-qno lib` specifies that no libraries, including the system libraries as well as the XL C/C++ libraries (these are found in the `lib/aix51/`, `lib/aix52/`, and `lib/aix53/` subdirectories of the compiler installation directory), are to be linked. The system startup files are still linked, unless `-qno crt` is also specified.

Note that if your program references any symbols that are defined in the standard libraries or compiler-specific libraries, link errors will occur. To avoid these

unresolved references when compiling with **-qnolib**, be sure to explicitly link the required libraries by using the command flag **-l** and the library name.

## Predefined macros

None.

## Examples

To compile `myprogram.c` without linking to any libraries except the compiler library `libxlopt.a`, enter:

```
xlc myprogram.c -qnolib -lxlopt
```

## Related information

- “-qcr” on page 129

## -qlibansi

### Category

Optimization and tuning

### Pragma equivalent

#pragma options [no]libansi

### Purpose

Assumes that all functions with the name of an ANSI C library function are in fact the system functions.

When **libansi** is in effect, the optimizer can generate better code because it will know about the behavior of a given function, such as whether or not it has any side effects.

### Syntax

►► -q  

### Defaults

-qnolibansi

### Predefined macros

► C++ `__LIBANSI__` is defined to 1 when **libansi** is in effect; otherwise, it is not defined.

## -qlibmpi

### Category

“Optimization and tuning” on page 87

## Pragma equivalent

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.

**-qlibmpi** 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 **-qlibmpi**, 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.

## Predefined macros

None.

## Examples

To compile `myprogram.c`, enter the following command:

```
xlc -O5 myprogram.c -qlibmpi
```

## Related information

- Message Passing Interface Forum
- “-qipa” on page 199

## **-qlinedebug**

### Category

Error checking and debugging

## Pragma equivalent

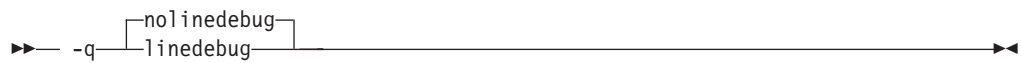
None.

## 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.

## Syntax



► — -q — [no linedebug] — [linedebug] — ►

## Defaults

-qnolinedebug

## Usage

When **-qlinedebug** is in effect, function inlining is disabled.

Avoid using **-qlinedebug** with **-O** (optimization) option. The information produced may be incomplete or misleading.

The **-g** option overrides the **-qlinedebug** option. If you specify **-g** with **-qnolinedebug** on the command line, **-qnolinedebug** is ignored and a warning is issued.

## Predefined macros

None.

## Examples

To compile `myprogram.c` to produce an executable program `testing` so you can step through it with a debugger, enter:

```
xlc myprogram.c -o testing -qlinedebug
```

## Related information

- “-g” on page 169
- “-O, -qoptimize” on page 269

## -q`list`

### Category

Listings, messages, and compiler information

## Pragma equivalent

#pragma options [no]list


## Purpose

Produces a compiler listing file that includes an object listing.

When **list** is in effect, a listing file is generated with a .lst suffix for each source file named on the command line. For details of the contents of the listing file, see “Compiler listings” on page 20.

You can use the object or assembly listing to help understand the performance characteristics of the generated code and to diagnose execution problems.

## Syntax



## Defaults

-qno list

## Parameters

### offset | nooffset

Changes the **offset** of the PDEF header from 00000 to the offset of the start of the text area. Specifying the option 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 only relevant if there are multiple procedures in a compilation unit.

Specifying **list** without the suboption is equivalent to **list=nooffset**.

## Usage

The **-qno print** compiler option overrides this option.

## Predefined macros

None.

## Examples

To compile myprogram.c and to produce a listing (.lst) file that includes an object listing, enter:

```
xlc myprogram.c -qlist
```

## Related information

- “-qlistopt” on page 248
- “-qprint” on page 296
- “-qsource” on page 324

## -qlistfmt

### Category

Listings, messages, and compiler information

### @PROCESS

None.


### Pragma equivalent

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 247.

#### *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 | nottransforms**

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.c` 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,

```
xlc -qlistfmt=xml=all:filename=abc.xml -O3 myfile1.c myfile2.c myfile3.c
```

will result in only one report, `abc.xml` based on the compilation of the last file `myfile3.c`

**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 C/C++ is `xlstyle.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 (`xlstyle.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=xlstyle.xml**, `xlstyle.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 `/usr/vacpp/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 C/C++ for AIX, V12.1 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.

## Predefined macros

None.

## Examples

If you want to compile `myprogram.c` to produce an XML report that shows how loops are optimized, enter:

```
xlc -qhot -O3 -qlistfmt=xml=transforms myprogram.c
```

If you want to compile `myprogram.c` to produce an XML report that shows which functions are inlined, enter:

```
xlc -qinline -qlistfmt=xml=inlines myprogram.c
```

## 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 extension 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 302
- "Using compiler reports to diagnose optimization opportunities" in the *XL C/C++ Optimization and Programming Guide*

## -qlistopt

### Category

Listings, messages, and compiler information

### Pragma equivalent

None.

### 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. For details of the contents of the listing file, see “Compiler listings” on page 20.

### Syntax

```
→ -q [ nolistopt | listopt ] →
```

### Defaults

-qnolistopt

### Usage

Option settings caused by pragma statements in the program source are not shown in the compiler listing.

The **-qnoprint** compiler option overrides this option.

### Predefined macros

None.

### Examples

To compile `myprogram.c` to produce a listing (.lst) file that shows all options in effect, enter:

```
xlc myprogram.c -qlistopt
```

## Related information

- “-q`list`” on page 243
- “-q`print`” on page 296
- “-q`source`” on page 324

## -q`longlit`

### Category

Floating-point and integer control

### Pragma equivalent

None.

### Purpose

In 64-bit mode, when determining the implicit types for integer literals, the compiler behaves as if an `l` or `L` suffix were added to integral literals with no suffix or with a suffix consisting only of `u` or `U`.

### Syntax

→ -q no`longlit` `longlit` →

### Defaults

-q`nolonglit`

### Usage

After you specify the **-q`longlit`** option, if the `int` or unsigned `int` type is contained in the implicit type list of a integer literal, the `int` or unsigned `int` type is replaced with the long `int` or unsigned long `int` type, respectively. For more information about the integer literals, see "Integer literals".

### Predefined macros

None.

### Examples

After you specify the **-q`longlit`** option, the integer literal `0x80000000` has the long `int` type in 64-bit mode. Otherwise, if this option is not specified, the integer literal has the unsigned `int` type in both 32-bit and 64-bit modes.

## -q`longlong`

### Category

Language element control

## Pragma equivalent

#pragma options [no]longlong

## Purpose

Allows IBM long long integer types in your program.

## Syntax

►► -q longlong  
no longlong ►►

## Defaults

- C **-qlonglong** for the xlc, xlc++, xlc, cc and c99 invocation commands; **-qno longlong** for the c89 invocation command.
- C++ **-qlonglong** for the **-qlanglvl= compat366 | extended** option; **-qno longlong** for the **-qlanglvl=strict98 | extended0x** option. When multiple **-qlanglvl** options that imply **-q[no]longlong** or actual **-q[no]longlong** options are specified, the last specified option determines whether **-qlonglong** is in effect.

## Usage

C This option takes effect when the **-qlanglvl=extended | stdc89 | extc89** option is in effect. It is not valid when the **-qlanglvl=stdc99 | extc99** option is in effect, because the long long support provided by this option is incompatible with the semantics of the long long types mandated by the C99 standard.

C++ This option does not take effect when the **-qlanglvl=c99longlong** option is in effect, because the long long support provided by this option is incompatible with the semantics of the long long types mandated by the C99 standard as adopted in C++0x.

## Predefined macros

`_LONG_LONG` is defined to 1 when long long data types are available; otherwise, it is undefined.

## Examples

To compile myprogram.c with support for IBM long long integers, enter:

```
cc myprogram.c -qlonglong
```

AIX v4.2 and later provides support for files greater than 2 gigabytes in size so you can store large quantities of data in a single file. To allow large file manipulation in your application, compile with the **-D\_LARGE\_FILES** and **-qlonglong** compiler options. For example:

```
xlc myprogram.c -D_LARGE_FILES -qlonglong
```

## Related information

- "Integral types" in the *IBM XL C/C++ for AIX, V12.1 Language Reference*

## -ma (C only)

See “-qalloca, -ma (C only)” on page 104.

## -qmacpstr

### Category

Language element control

### Pragma equivalent

#pragma options [no]macpstr

### Purpose

Converts Pascal string literals (prefixed by the `\p` escape sequence) into null-terminated strings in which the first byte contains the length of the string.

For example, when the `-qmacpstr` option is in effect, the compiler converts:

```
"\pABC"
```

to:

```
'\03' , 'A' , 'B' , 'C' , '\0'
```

### Syntax

→ -q nomacpstr  
macpstr →

### Defaults

-qnomacpstr

### Usage

A Pascal string literal always contains the characters `"\p`. The characters `\p` in the middle of a string do not form a Pascal string literal; the characters must be *immediately preceded* by the `"` (double quote) character.

Entering the characters:

```
'\p' , 'A' , 'B' , 'C' , '\0'
```

into a character array does not form a Pascal string literal.

The compiler ignores the `-qmacpstr` option when the `-qmbcs` or `-qdbcs` option is active because Pascal-string-literal processing is only valid for one-byte characters.

The `#pragma options` keyword `macpstr` is only valid at the top of a source file before any C or C++ source statements. If you attempt to use it in the middle of a source file, it is ignored and the compiler issues an error message.

The following describes how Pascal string literals are processed.

- Because there is no Pascal-string-literal processing of wide strings, using the escape sequence `\p` in a wide string literal with the `-qmacpstr` option, generates a warning message and the escape sequence is ignored.
- Concatenating a Pascal string literal to a normal string gives a non-Pascal string. For example, concatenating the strings:  
`"ABC" "\pDEF"`

gives:

`"ABCpDEF"`

- Concatenating two Pascal string literals, for example, `strcat`, does not result in a Pascal string literal. However, as described above, two adjacent Pascal string literals can be concatenated to form one Pascal string literal in which the first byte is the length of the new string literal. For example, concatenating the strings:

`"\p ABC" "\p DEF"`

or

`"\p ABC" "DEF"`

results in:

`"\06ABCDEF"`

- A Pascal string literal cannot be concatenated with a wide string literal.
- The compiler truncates a Pascal string literal that is longer than 255 bytes (excluding the length byte and the terminating `NULL`) to 255 characters.
- The Pascal string literal is *not* a basic type different from other C or C++ string literals. After the processing of the Pascal string literal is complete, the resulting string is treated the same as all other strings. If the program passes a C string to a function that expects a Pascal string, or vice versa, the behavior is undefined.
- Modifying any byte of the Pascal string literal after the processing has been completed does not alter the original length value in the first byte. For example, in the string `"\06ABCDEF"`, substituting a null character for one of the existing characters in the middle of the string does not change the value of the first byte of the string, which contains the length of the string.
- No errors or warnings are issued when the bytes of the processed Pascal string literal are modified.

## Predefined macros

None.

## Examples

To compile `mypascal.c` and convert string literals into Pascal-style strings, enter:

```
xlc mypascal.c -qmacpstr
```

## Related information

- “`-qmbcs`, `-qdbcs`” on page 257

## **-qmakedep, -M** Category

Output control

## Pragma equivalent


None.

## Purpose

Creates an output file containing targets suitable for inclusion in a description file for the **make** command.

The output file is named with a `.u` suffix.

## Syntax



## Defaults

Not applicable.

## Parameters

### gcc (-q-makedep option only)

The format of the generated **make** rule to matches the GCC format: the description file includes a single target listing all of the main source file's dependencies.

If you specify **-q-makedep** with no suboption, or **-M**, the description file specifies a separate rule for each of the main source file's dependencies.

## Usage

For each source file with a `.c`, `.C`, `.cpp`, or `.i` suffix named on the command line, an output file is generated with the same name as the object file and a `.u` suffix. Output files are not created for any other types of input files. If you use the **-o** option to rename the object file, the output file uses the name you specified on the **-o** option. See below for examples.

The output files generated by these options are not **make** files; they must be linked before they can be used with the **make** command. For more information on this command, see your operating system documentation.

The output file contains a line for the input file and an entry for each include file. It has the general form:

```
file_name.o:include_file_name  
file_name.o:file_name.suffix
```

You can also use the following option with **q-makedep** and **-M**:

### **-MF=file\_path**

Sets the name of the output file, where *file\_path* is the full or partial path or file name for the output file. See below for examples.

Include files are listed according to the search order rules for the `#include` preprocessor directive, described in “Directory search sequence for include files” on page 13. If the include file is not found, it is not added to the `.u` file.

Files with no include statements produce output files containing one line that lists only the input file name.

## Predefined macros

None.

## Examples

To compile `mysource.c` and create an output file named `mysource.u`, enter:

```
xlc -c -qmakedep mysource.c
```

To compile `foo_src.c` and create an output file named `mysource.u`, enter:

```
xlc -c -qmakedep foo_src.c -MF mysource.u
```

To compile `foo_src.c` and create an output file named `mysource.u` in the `deps/` directory, enter:

```
xlc -c -qmakedep foo_src.c -MF deps/mysource.u
```

To compile `foo_src.c` and create an object file named `foo_obj.o` and an output file named `foo_obj.u`, enter:

```
xlc -c -qmakedep foo_src.c -o foo_obj.o
```

To compile `foo_src.c` and create an object file named `foo_obj.o` and an output file named `mysource.u`, enter:

```
xlc -c -qmakedep foo_src.c -o foo_obj.o -MF mysource.u
```

To compile `foo_src1.c` and `foo_src2.c` to create two output files, named `foo_src1.u` and `foo_src2.u`, respectively, in the `c:/tmp/` directory, enter:

```
xlc -c -qmakedep foo_src1.c foo_src2.c -MF c:/tmp/
```

## Related information

- “-MF” on page 258
- “-o” on page 268
- “Directory search sequence for include files” on page 13

## **-qmaxerr**

### Category

Error checking and debugging

### Pragma equivalent

None.


### Purpose

Stops compilation when the number of error messages of a specified severity level or higher reaches a specified number.




## Syntax

### -qmaxerr syntax — C



### -qmaxerr syntax — C++



## Defaults

-qnomaxerr

## Parameters

*number*

It specifies the maximum number of messages the compiler generates before it stops. *number* must be an integer with a value of 1 or greater.

**i** Specifies that the severity level is Informational (I) or higher.

**w** Specifies that the severity level is Warning (W) or higher.

**c e**

Specifies that the severity level is Error (E) or higher.

**s** Specifies that the severity level is Severe (S).

## 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.

Diagnostic messages can be controlled by the **-qflag** option.

## Predefined macros

None.

## Examples

To stop compilation of `myprogram.c` when 10 warnings are encountered, enter the command:

```
xlc myprogram.c -qmaxerr=10:w
```

To stop compilation of `myprogram.c` when 5 severe errors are encountered, assuming that the current `-qhalt` option value is `s` (severe), enter the command:  
`xlc myprogram.c -qmaxerr=5`

To stop compilation of `myprogram.c` when 3 informational messages are encountered, enter the command:  
`xlc myprogram.c -qmaxerr=3:i`

or:

```
xlc myprogram.c -qmaxerr=3 -qhalt=i
```

### Related information

- “`-qflag`” on page 153
- “`-qhalt`” on page 174
- “Message severity levels and compiler response” on page 18

## **-qmaxmem**

### Category

Optimization and tuning

### Pragma equivalent

```
#pragma options maxmem
```

### Purpose

Limits the amount of memory that the compiler allocates while performing specific, memory-intensive optimizations to the specified number of kilobytes.

### Syntax

```
►► -qmaxmem=size_limit ◀◀
```

### Defaults

- `-qmaxmem=8192` when `-O2` is in effect.
- `-qmaxmem=-1` when `-O3` or higher optimization is in effect.

### Parameters

*size\_limit*

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

A smaller limit does not necessarily mean that the resulting program will be slower, only that the compiler may finish before finding all opportunities to increase performance. Increasing the limit does not necessarily mean that the

resulting program will be faster, only that the compiler is better able to find opportunities to increase performance if they exist.

Setting a large limit has no negative effect on the compilation of source files when the compiler needs less memory. However, depending on the source file being compiled, the size of subprograms in the source, the machine configuration, and the workload on the system, setting the limit too high, or to -1, might exceed available system resources.

## Predefined macros

None.

## Examples

To compile `myprogram.c` so that the memory specified for local table is 16384 kilobytes, enter:

```
xlc myprogram.c -qmaxmem=16384
```

## -qmbcs, -qdbcs

### Category

Language element control

### Pragma equivalent

`#pragma options [no]mbcs`, `#pragma options [no]dbcs`

### Purpose

Enables support for multibyte character sets (MBCS) and Unicode characters in your source code.

When `mbcs` or `dbcs` is in effect, multibyte character literals and comments are recognized by the compiler. When `nombs` or `nodbcs` is in effect, the compiler treats all literals as single-byte literals.

### Syntax



### Defaults

`-qnombs`, `-qnodbcs`

### Usage

For rules on using multibyte characters in your source code, see "Multibyte characters" in the *XL C/C++ Language Reference*.

In addition, you can use multibyte characters in the following contexts:

- In file names passed as arguments to compiler invocations on the command line; for example:  
`xlc /u/myhome/c_programs/kanji_files/multibyte_char.c -omultibyte_char`
- In file names, as suboptions to compiler options that take file names as arguments
- In the definition of a macro name using the **-D** option; for example:  
`-DMYMACRO="kpsmultibyte_chardcs"`  
`-DMYMACRO='multibyte_char'`

Listing files display the date and time for the appropriate international language, and multibyte characters in the source file name also appear in the name of the corresponding list file. For example, a C source file called:

`multibyte_char.c`

gives a list file called  
`multibyte_char.lst`

## Predefined macros

None.

## Examples

To compile `myprogram.c` if it contains multibyte characters, enter:

`xlc myprogram.c -qmbcs`

## Related information

- “-D” on page 132

## -MF

### Category

Output control

### Pragma equivalent

None.

### Purpose

Specifies the target for the output generated by the **-qmakedep** or **-M** options.

This option is used only together with the **-qmakedep** or **-M** options. See the description for the “-qmakedep, -M” on page 252 for more information.

### Syntax

▶▶ `-MF` *path* ◀◀

### Defaults

Not applicable.

## Parameters

*path*

The target output path. *path* can be a full directory path or file name. If *path* is the name of a directory, the dependency file generated by the compiler is placed into the specified directory. If you do not specify a directory, the dependency file is stored in the current working directory.

## 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 compiling multiple source files, only a single dependency file will be generated containing the **make** rule for the last file specified on the command line.

## Predefined macros

None.

## Related information

- “-qmakedep, -M” on page 252
- “-o” on page 268
- “Directory search sequence for include files” on page 13

## -qminimaltoc

### Category

Optimization and tuning

### Pragma equivalent

None.

### Purpose

Controls the generation of the table of contents (TOC), which the compiler creates for an executable file.

Programs compiled in 64-bit mode have a limit of 8192 TOC entries. As a result, you may encounter "relocation truncation" error messages when linking large programs in 64-bit mode; these error messages are caused by TOC overflow conditions. When **-qminimaltoc** is in effect, the compiler avoids these overflow conditions by placing TOC entries into a separate data section for each object file.

Specifying **-qminimaltoc** ensures that the compiler creates only one TOC entry for each compilation unit. Specifying this option can minimize the use of available TOC entries, but its use impacts performance. Use the **-qminimaltoc** option with discretion, particularly with files that contain frequently executed code.

### Syntax

►► -q nomimaltoc  
minimaltoc \_\_\_\_\_►►

## Defaults

-qnominaltoc

## Usage

This compiler option applies to 64-bit compilations only.

Compiling with **-qminimaltoc** may create slightly slower and larger code for your program. However, these effects may be minimized by specifying optimizing options when compiling your program.

## Predefined macros

None.

## -qmkshrobj

### Category

Output control

### Pragma equivalent

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 (or using the `makeC++SharedLib` utility, in C++) to create a shared object. The advantages of using this option are the automatic handling of link-time C++ template instantiation (using either the template include directory or the template registry), and compatibility with **-qipa** link-time optimizations (such as those performed at **-O5**).

### Syntax

**-qmkshrobj syntax — C**

▶▶ -qmkshrobj —————▶▶

**-qmkshrobj syntax — C++**

▶▶ -qmkshrobj [=-priority] —————▶▶

### Defaults

By default, the output object is linked with the runtime libraries and startup routines to create an executable file.

## Parameters

**C++** *priority*

Specifies the priority level for the initialization order of static C++ objects declared in the shared object (relative to static objects declared in other shared objects). The *priority* may be any number from -214 782 624 (highest priority — initialized first) to 214 783 647 (lowest priority — initialized last). Numbers from -214 783 648 to -214 782 623 are reserved for system use. If no priority is specified a default priority of 0 is used. The priority has no effect if you link shared objects written in C, if you link with the C compiler (**xlc**), or if the shared object has no static initialization.

## Usage

When the **-qmkshrobj** option is specified, the driver program starts the CreateExportList utility to create an export list from the input list of object files.

The compiler automatically exports all global symbols from the shared object unless you explicitly specify which symbols to export with the **-bE:**, **-bexport:** or **-bnoexpall** options, or if you use the **-qnoweakexp** option to prevent weak symbols from being exported.

Specifying **-qmkshrobj** implies **-qpic**.

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 shr.o.

**-qexpfile=***filename*

Saves all exported symbols in *filename*.

**-e** *name*

Sets the entry name for the shared executable to *name*.

**-q[no]weakexp**

Specifies whether symbols marked as weak (with the **#pragma weak** directive) are to be included in the export list. If you do not explicitly set this option, the default is **-qweakexp** (global weak symbols are exported).

For detailed information about using **-qmkshrobj** to create shared libraries, as well as examples of using **-qmkshrobj** with priority values, see "Constructing a library" in the *XL C/C++ Optimization and Programming Guide*.

## Predefined macros

None.

## Examples

To construct the shared library `big_lib.so` from three smaller object files, enter the following command:

```
xlc -qmkshrobj -o big_lib.so lib_a.o lib_b.o lib_c.o
```

## Related information

- “-b” on page 115
- “-e” on page 141
- “-G” on page 172

- “-qexpfile” on page 148
- “-qipa” on page 199
- “-o” on page 268
- “-qplic” on page 291
- “-qpriority (C++ only)” on page 297
- “-qweakexp” on page 377

## **-qnamemangling (C++ only)**

### **Category**

Portability and migration

### **Pragma equivalent**

#pragma namemangling

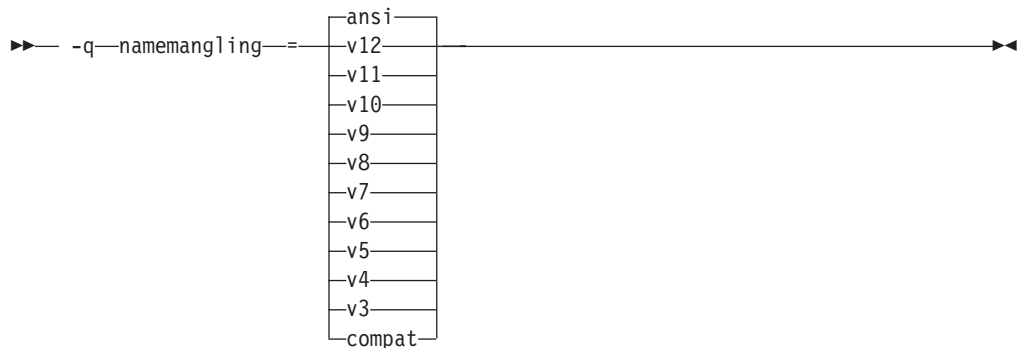
### **Purpose**

Chooses the name mangling scheme for external symbol names generated from C++ source code.

The option and pragma are provided to ensure binary compatibility with link modules created with previous versions of the compiler. If you do not need to ensure compatibility with earlier versions, do not change the default setting of this option.

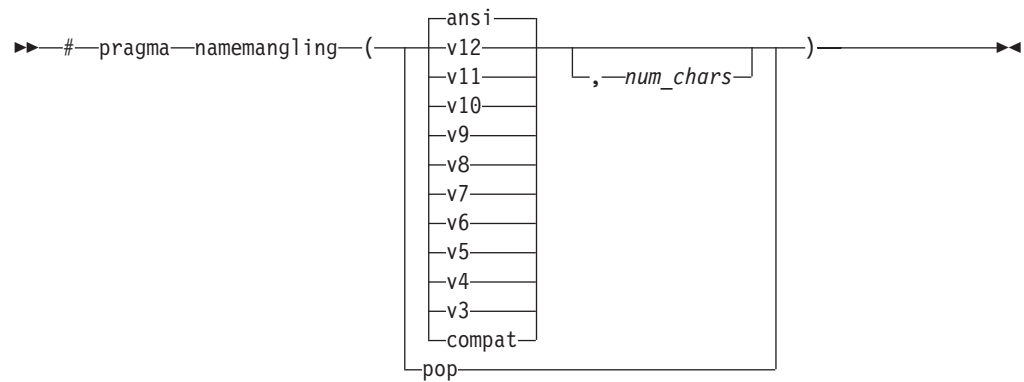
### **Syntax**

#### **Option syntax**



#### **Pragma syntax**





## Defaults

`-qnamemangling(ansi, 64000)`

## Parameters

### ansi

The name mangling scheme supports the most recent Standard C++ language features. This suboption is equivalent to **v12**.

### **v3** | **compat**

The name mangling scheme is compatible with VisualAge C++ V3.0 in 32-bit mode only.

**v4** The name mangling scheme is compatible with VisualAge C++ V4.0. Before this release, a function and a function template specialization with the same name and parameter list were considered to have the same signature, and the following test case would fail to compile:

```

int f(int) {
    return 42;
}

template < class T > int f(T) {
    return 43;
}

int main() {
    f < int > (3); // instantiate int f < int > (int)
    return f(4);
}
  
```

From V4.0 on, the compiler treats a function and a function template specialization with the same name and parameter list as distinct functions. The following examples illustrate this behavior:

Source name	Mangled name before v4	Mangled name in v4 and later
<code>int f (int)</code>	<code>f__Fi</code>	<code>f__Fi</code>
<code>int f &lt;int&gt; (int)</code>	<code>f__Fi</code>	<code>f__Hi_i_i</code>

**v5** The name mangling scheme is compatible with VisualAge C++ V5.0. Same as the **v4** suboption.

**v6** The name mangling scheme is compatible with VisualAge C++ V6.0. Before this release, top-level cv-qualifiers in function arguments were encoded in

mangled names. From V6.0 on, in accordance with the C++ Standard, top-level cv-qualifiers are not considered part of the underlying type of a function argument, and the cv-qualifiers are not encoded in the mangled names. The following examples illustrate this behavior:

Source name	Mangled name before v6	Mangled name in v6 and later
void f (const int)	f_FCi	f_Fi
void f (int* const)	f_FCPi	f_FPi

**Note:** This behavior can also be controlled with the use of the **nameManglingRule(fnparmtyp)** pragma directive. For more information, see “#pragma namemanglingrule (C++ only)” on page 415.

**v7** The name mangling scheme is compatible with IBM XL C/C++ V7.0.

Several changes to the mangling scheme went into effect in IBM XL C/C++ V7.0. First of all, before V7.0, top-level cv-qualifiers were used to distinguish between types in repeated parameters in a function's signature. From V7.0 on, in accordance with the C++ Standard, top-level cv-qualifiers are ignored for determining the equivalence between function parameters. Parameters that are only differentiated by the presence of a top-level cv-qualifier are considered to be equivalent, and are represented in the compressed encoding scheme used for repeated parameters of the same type. The following examples illustrate this behavior:

Source name	Mangled name before v7	Mangled name in v7 and later
void f (int, const int)	f_FiCi (pre-v6) f_Fii (v6)	f_FiT1
void f (int* const, int* const)	f_FCPiCCPi (pre-v6) f_FPiPi (v6)	f_FPiT1

**Note:** This behavior can also be controlled with the use of the **nameManglingRule(fnparmtyp)** pragma directive. For more information, as well as details of the compressed mangling scheme, see “#pragma namemanglingrule (C++ only)” on page 415.

Secondly, before V7.0, non-type integral template arguments were mangled as 32-bit unsigned decimal numbers prefixed by SP. Due to ambiguities introduced by this in mangling 64-bit values, this scheme has been changed to the following:

```

non-type template argument → SM          #single repeat of a previous parameter
                             → SP number #positive internal argument
                             → SN number #negative internal argument

```

When a non-type integral template argument is positive, the number is prefixed with SP. When a non-type integral template argument is negative, the number is prefixed with SN, and the decimal number is written without the minus sign. There is no limit in the range of decimal numbers which can be represented. The following examples illustrate this behavior:

Source name	Mangled template name before v7	Mangled template name in v7 and later
<pre>template &lt;int n&gt; int f() { return N; } int main() { return f&lt;-3&gt;(); }</pre>	f_HxiSP429	f_HxiSN3x_v

**v8** The name mangling scheme is compatible with IBM XL C/C++ V8.0.

Several changes to the mangling scheme went into effect in IBM XL C/C++ V8.0. First of all, before V8.0, intermediate-level cv-qualifiers were not used to distinguish between types in repeated parameters in a function's signature. From V8.0 on, intermediate-level cv-qualifiers are used for determining the equivalence between function parameters. Parameters that are differentiated by the presence of an intermediate-level cv-qualifier are not considered to be equivalent, and are mangled as separate parameters. The following examples illustrate this behavior:

Source name	Mangled name before v8	Mangled name in v8 and later
void f (int**, int* const *)	f_FPPiT1	f_FPPiPCPi

**Note:** This behavior can also be controlled with the use of the **nameManglingRule(fnparmscmp)** pragma directive. For more information, as well as details of the compressed mangling scheme, see “#pragma namemanglingrule (C++ only)” on page 415.

Secondly, before V8.0, only the underlying type in a typedef definition was used to distinguish between types in repeated parameters in a function's signature. From V8.0 on, the name defined in a typedef declaration in a template parameter is encoded as a separate parameter in the mangled name of a template function that uses the typedef as a parameter. The following examples illustrate this behavior:

Source name	Mangled function name before v8	Mangled function name in v8 and later
<pre>template &lt;typename T&gt; struct A { typedef int INT; };  template &lt;typename V&gt; int f (A &lt;V&gt;, int, typename A&lt;V&gt;::INT) {}  A&lt;int&gt; a; int x = f (a, 1, 10);</pre>	f_Hi_1AXTi iT2_i	f_Hi_1AXTi_iQ2_1AXTi_9INT_i

Source name	Mangled function name before v8	Mangled function name in v8 and later
<pre>template &lt;typename T&gt; struct A { typedef A INT; };  template &lt;typename Y&gt; int f (A &lt;int&gt;::INT, const A&lt;Y&gt;) {}  A&lt;int&gt; a; int x = f (10, a);</pre>	f_Hi_1AXTi_T1_i	f_Hi_Q2_1AXTi_INT1AXTi_i

**v9** The name mangling scheme is compatible with IBM XL C/C++ V9.0.

Before this release, the name mangling scheme did not differentiate between different pointer-to-member template arguments in template instantiations, and the following test case would fail to compile:

```
struct pair
{
    int x, y;
    pair(int x_, int y_) : x(x_), y(y_) {}
};

template <int pair::*PtrToPairMember>

struct str
{
    int f(pair& p)
    {
        return p.*PtrToPairMember;
    }
};

template <int pair::*PtrToPairMember> g(pair& p)
{
    return p.*PtrToPairMember;
}

int main()
{
    pair p(0, 1);
    str<&pair::x> fx;
    str<&pair::y> fy;

    if (fx.f(p) != 0 || fy.f(p) != 1) { return 1; }

    if (g<&pair::x>(p) != 0 || g<&pair::y>(p) != 1) { return 2; }

    return 0;
}
```

From V9.0 on, the compiler treats different pointer-to-member template arguments as distinct. The following examples illustrate this behavior:

Source name	Mangled name before v9	Mangled name in v9 and later
int str<&pair::y>::f(pair &)	f_3strXA0_FR4pair	f_3strXAM1y_FR4pair
int str<&pair::x>::f(pair &)	f_3strXA0_FR4pair	f_3strXAM1x_FR4pair
int g<&pair::y>(pair &)	g_HxM4pairiA0x_R4pair_i	g_HxM4pairiA0yx_R4pair_i

Source name	Mangled name before v9	Mangled name in v9 and later
int g<&pair::x>(pair &)	g_HxM4pairiA0x_R4pair_i	g_HxM4pairiA0xx_R4pair_i

### v10

The name mangling scheme is compatible with IBM XL C/C++ V10.1. This suboption has the same effect as the **v9** suboption.

### v11

The name mangling scheme is compatible with IBM XL C/C++ V11.1. This suboption has the same effect as **v10**.

### v12

The name mangling scheme is compatible with IBM XL C/C++ V12.1. This suboption is equivalent to **ansi**.

Before this release, the name mangling scheme did not distinguish function parameters of template types that were cv-qualified and not cv-qualified, because the cv-qualifiers of the template type parameters were ignored. The **v12** fix preserves the cv-qualifiers, so the function parameters that are cv-qualified and not cv-qualified are handled differently. For example:

```
template<typename Element> struct Iterator
{
    Iterator() {}
    Iterator<Element>& operator+=(long d);
    friend Iterator<Element> operator+(Iterator<Element> it, long d)
    {
        it += d;
        return it;
    }
};

int main()
{
    Iterator<int> iter;
    Iterator<const int> c_iter;
    iter = iter+10;
    c_iter=c_iter+10;
}
```

In the preceding example, before the **v12** fix, `Iterator<const int> operator+(Iterator<const int>, long)()` and `Iterator<int> operator+(Iterator<int>, long)()` had an identical name mangling `__pl__F8IteratorXTi_l`. The fix distinguishes the name mangling of these two functions as shown in the following table:

Source name	Mangled name before v12	Mangled name in v12 and later
Iterator<const int> operator+(Iterator<const int>, long)	__pl__F8IteratorXTi_l	__pl__F8IteratorXTCi_l
Iterator<int> operator+(Iterator<int>, long)()	__pl__F8IteratorXTi_l	__pl__F8IteratorXTi_l

### *num\_chars* (pragma only)

Specifies the maximum number of allowable characters in the mangled names. If you do not specify this suboption, the default maximum is 64000 characters for all settings except **v3** and **compat**, for which the default maximum is 255 characters.

## pop

Discards the current pragma setting and reverts to the setting specified by the previous pragma directive. If no previous pragma was specified, reverts to the command-line or default option setting.

## Predefined macros

None.

## Related information

- “#pragma namemanglingrule (C++ only)” on page 415

## -o

### Category

Output control

### Pragma equivalent

None.

### Purpose

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

### Syntax

►► `-o` *path* ◀◀

### Defaults

See “Types of output files” on page 4 for the default file names and suffixes produced by different phases of compilation.

### Parameters

#### *path*

When you are using the option to compile from source files, *path* can be the name of a file or directory. The *path* can be a relative or absolute path name. When you are using the option to link from object files, *path* must be a file name.

If the *path* is the name of an existing directory, files created by the compiler are placed into that directory. If *path* is not an existing directory, the *path* is the name of the file produced by the compiler. See below for examples.

You can not specify a file name with a C or C++ source file suffix (.C, .c, .cpp, or .i), such as `myprog.c` or `myprog.i`; this results in an error and neither the compiler nor the linker is invoked.

### Usage

If you use the `-c` option with `-o` together and the *path* is not an existing directory, you can only compile one source file at a time. In this case, if more than one source file name is listed in the compiler invocation, the compiler issues a warning message and ignores `-o`.

The `-E`, `-P`, and `-qsyntaxonly` options override the `-o` option.

## Predefined macros

None.

## Examples

To compile `myprogram.c` so that the resulting executable is called `myaccount`, assuming that no directory with name `myaccount` exists, enter:

```
xlc myprogram.c -o myaccount
```

To compile `test.c` to an object file only and name the object file `new.o`, enter:

```
xlc test.c -c -o new.o
```

## Related information

- “`-c`” on page 119
- “`-E`” on page 142
- “`-P`” on page 280
- “`-qsyntaxonly (C only)`” on page 341

## `-O`, `-qoptimize`

### Category

Optimization and tuning


### Pragma equivalent

`#pragma options [no]optimize`

### Purpose

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

### Syntax



### Defaults

`-qnooptimize` or `-O0` or `-qoptimize=0`

## Parameters

### **-O0 | nooptimize | noopt | optimize|opt=0**

Performs only quick local optimizations such as constant folding and elimination of local common subexpressions.

This setting implies **-qstrict\_induction** unless **-qnostrict\_induction** is explicitly specified.

### **-O | -O2 | optimize | opt | optimize|opt=2**

Performs optimizations that the compiler developers considered the best combination for compilation speed and runtime performance. The optimizations may change from product release to release. If you need a specific level of optimization, specify the appropriate numeric value.

This setting implies **-qstrict** and **-qnostrict\_induction**, unless explicitly negated by **-qstrict\_induction** or **-qnostrict**.

### **-O3 | optimize|opt=3**

Performs additional optimizations that are memory intensive, compile-time intensive, or both. They are recommended when the desire for runtime improvement outweighs the concern for minimizing compilation resources.

**-O3** applies the **-O2** level of optimization, but with unbounded time and memory limits. **-O3** also performs higher and more aggressive optimizations that have the potential to slightly alter the semantics of your program. The compiler guards against these optimizations at **-O2**. The aggressive optimizations performed when you specify **-O3** are:

1. Aggressive code motion, and scheduling on computations that have the potential to raise an exception, are allowed.

Loads and floating-point computations fall into this category. This optimization is aggressive because it may place such instructions onto execution paths where they *will* be executed when they *may* not have been according to the actual semantics of the program.

For example, a loop-invariant floating-point computation that is found on some, but not all, paths through a loop will not be moved at **-O2** because the computation may cause an exception. At **-O3**, the compiler will move it because it is not certain to cause an exception. The same is true for motion of loads. Although a load through a pointer is never moved, loads off the static or stack base register are considered movable at **-O3**. Loads in general are not considered to be absolutely safe at **-O2** because a program can contain a declaration of a static array `a` of 10 elements and load `a[600000000003]`, which could cause a segmentation violation.

The same concepts apply to scheduling.

#### **Example:**

In the following example, at **-O2**, the computation of `b+c` is not moved out of the loop for two reasons:

- It is considered dangerous because it is a floating-point operation
- `t` does not occur on every path through the loop

At **-O3**, the code is moved.

```
...
int i ;
float a[100], b, c ;
for (i = 0 ; i < 100 ; i++)
{
```



```

    if (a[i] < a[i+1])
        a[i] = b + c ;
    }
    ...

```

2. Conformance to IEEE rules are relaxed.

With **-O2** certain optimizations are not performed because they may produce an incorrect sign in cases with a zero result, and because they remove an arithmetic operation that may cause some type of floating-point exception.

For example,  $X + 0.0$  is not folded to  $X$  because, under IEEE rules,  $-0.0 + 0.0 = 0.0$ , which is  $-X$ . In some other cases, some optimizations may perform optimizations that yield a zero result with the wrong sign. For example,  $X - Y * Z$  may result in a  $-0.0$  where the original computation would produce  $0.0$ .

In most cases the difference in the results is not important to an application and **-O3** allows these optimizations.

3. Floating-point expressions may be rewritten.

Computations such as  $a*b*c$  may be rewritten as  $a*c*b$  if, for example, an opportunity exists to get a common subexpression by such rearrangement. Replacing a divide with a multiply by the reciprocal is another example of reassociating floating-point computations.

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

**-qfloat=fltint:rsqrt** is set by default with **-O3**.

**-qmaxmem=-1** is set by default with **-O3**, allowing the compiler to use as much memory as necessary when performing optimizations.

Built-in functions do not change `errno` at **-O3**.

Aggressive optimizations do *not* include the following floating-point suboptions: **-qfloat=hsflt** | **hssngl**, or anything else that affects the precision mode of a program.

Integer divide instructions are considered too dangerous to optimize even at **-O3**.

Refer to “`-qflttrap`” on page 159 to see the behavior of the compiler when you specify **optimize** options with the **-qflttrap** option.

You can use the **-qstrict** and **-qstrict\_induction** compiler options to turn off effects of **-O3** that might change the semantics of a program. Specifying **-qstrict** together with **-O3** invokes all the optimizations performed at **-O2** as well as further loop optimizations. Reference to the **-qstrict** compiler option can appear before or after the **-O3** option.

The **-O3** compiler option followed by the **-O** option leaves **-qignerrno** on.

When **-O3** and **-qhot=level=1** are in effect, the compiler replaces any calls in the source code to standard math library functions with calls to the equivalent MASS library functions, and if possible, the vector versions.

#### **-O4 | optimize|opt=4**

This option is the same as **-O3**, except that it also:

- Sets the **-qarch** and **-qtune** options to the architecture of the compiling machine
- Sets the **-qcache** option most appropriate to the characteristics of the compiling machine

- Sets the **-qhot** option
- Sets the **-qipa** option

**Note:** Later settings of **-O**, **-qcache**, **-qhot**, **-qipa**, **-qarch**, and **-qtune** options will override the settings implied by the **-O4** option.

### **-O5 | optimize|opt=5**

This option is the same as **-O4**, except that it:

- Sets the **-qipa=level=2** option to perform full interprocedural data flow and alias analysis.

**Note:**

Later settings of **-O**, **-qcache**, **-qipa**, **-qarch**, and **-qtune** options will override the settings implied by the **-O5** option.

## **Usage**

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

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

Optimization can cause statements to be moved or deleted, and generally should not be specified along with the **-g** flag for debugging programs. The debugging information produced may not be accurate.

## **Predefined macros**

- `__OPTIMIZE__` is predefined to 2 when **-O | O2** is in effect; it is predefined to 3 when **-O3 | O4 | O5** is in effect. Otherwise, it is undefined.
- `__OPTIMIZE_SIZE__` is predefined to 1 when **-O | -O2 | -O3 | -O4 | -O5** and **-qcompact** are in effect. Otherwise, it is undefined.

## **Examples**

To compile and optimize `myprogram.c`, enter:

```
xlc myprogram.c -O3
```

## **Related information**

- “-qhot” on page 178
- “-qipa” on page 199
- “-qpdf1, -qpdf2” on page 283
- “-qstrict” on page 333
- “Optimizing your applications” in the *XL C/C++ Optimization and Programming Guide*.

## **-qobjmodel (C++ only)**

### **Category**

Portability and migration

## Pragma equivalent

```
#pragma object_model
```

## Purpose

Sets the object model to be used for structures, unions, and classes.

The object models differ in the following areas:

- Layout for the virtual function table
- Virtual base class support
- Name mangling scheme

## Syntax

### Option syntax

```
►► -qobjmodel=classic  
                  |  
                  i b m
```

### Pragma syntax

```
►► #pragma object_model (classic  
                          |  
                          i b m  
                          |  
                          p o p )
```

## Defaults

```
-qobjmodel=classic
```

## Parameters

### classic

Uses the object model compatible with V3.6 of the IBM C++ Compiler. This suboption can also be specified using the legacy suboption name of **-qobjmodel=compat**, but support for this legacy suboption name may be removed in future releases of the compiler.

### **ibm**

Uses the object model introduced with VisualAge C++ V5.0. Objects compiled with this object model will use less memory and have better performance for deep inheritance with virtual bases.

### **pop (pragma only)**

Discards the current pragma setting and reverts to the setting specified by the previous pragma directive. If no previous pragma was specified, reverts to the command-line or default option setting.

## Usage

All classes in the same inheritance hierarchy must have the same object model.

## Predefined macros

- `__OBJECT_MODEL_CLASSIC__` is predefined to 1 when `-qobjmodel=classic` or `#pragma object_model(classic)` is in effect (the default); otherwise, it is undefined.
- `__OBJECT_MODEL_IBM__` is predefined to 1 when `-qobjmodel=ibm` or `#pragma object_model(ibm)` is in effect; otherwise, it is undefined.

## Examples

To compile `myprogram.C` with the `ibm` object model, enter:

```
xlc++ myprogram.C -qobjmodel=ibm
```

## -qoldpassbyvalue (C++ only)

### Category

Portability and migration

### Pragma equivalent

```
#pragma pass_by_value
```

### Purpose

Specifies how classes containing `const` or reference members are passed in function arguments.

The IBM C++ Compiler V3.6 uses pass by value only if the class has no `const` or reference data members, and the copy constructor is trivial and the destructor is trivial. VisualAge C++ V5.0 and later compilers use pass by value if the copy constructor is trivial and the destructor is trivial, regardless of `const` or reference data members. When `-qoldpassbyvalue` is in effect, the compiler uses the 3.6 behavior so that when a class containing a `const` or reference member is passed as a function argument, it is not passed by value. When `-qnooldpassbyvalue` is in effect, the compiler uses the new behavior so that when a class containing a `const` or reference member is passed as a function argument, it is passed by value.

The `#pragma pass_by_value` directive allows you greater control over this option for specific files or sections of source code.

### Syntax

#### Option syntax

```
» -q nooldpassbyvalue / oldpassbyvalue «
```

```
» | Pragma syntax | «
```

#### Pragma syntax

```
| #pragma pass_by_value ( compat ) |
                             | ansi |
                             | default |
                             | source |
                             | pop |
                             | reset |
```

## Defaults

`-qnooldpassbyvalue`

## Parameters

### `compat` (pragma only)

Sets `-qoldpassbyvalue` for the code that follows it. This instructs the compiler to use the earlier behavior: when a class containing a const or reference member is passed as a function argument, it is not passed by value.

### `ansi` | `default` (pragma only)

Sets `-qnooldpassbyvalue` for the code that follows it. This instructs the compiler to use the new behavior: when a class containing a const or reference member is passed as a function argument, it is passed by value.

### `source` (pragma only)

Reverts to the setting specified by the command-line option; if no option has been specified, reverts to the default setting (`-qnooldpassbyvalue`).

### `pop` | `reset` (pragma only)

Discards the current pragma setting and reverts to the setting specified by the previous pragma directive. If no previous pragma was specified, reverts to the command-line or default option setting.

## Usage

Use this option if you are linking to libraries compiled with IBM C++ Compiler V3.6 or earlier. Otherwise, functions that have const or reference class parameter types will give incorrect behavior when they are called from modules compiled with a later version of the compiler. Library headers compiled with IBM C++ Compiler V3.6 or earlier should be protected with the `#pragma pass_by_value` directive so that users of these libraries will get the correct calling convention for functions in those libraries that use class parameters.

## Predefined macros

None.

## `-qoptdebug`

### Category

Error checking and debugging

### Pragma equivalent

None.

## Purpose

When used with high levels of optimization, produces files containing optimized pseudocode that can be read by a debugger.

An output file with a .optdbg extension is created for each source file compiled with **-qoptdebug**. You can use the information contained in this file to help you understand how your code actually behaves under optimization.

## Syntax



## Defaults

`-qnooptdebug`

## Usage

**-qoptdebug** only has an effect when used with an option that enables the high-level optimizer, namely **-O3** or higher optimization level, or **-qhot**, **-qsmp**, **-qpdf**, or **-qipa**. You can use the option on both compilation and link steps. If you specify it on the compile step, one output file is generated for each source file. If you specify it on the **-qipa** link step, a single output file is generated.

The naming rules of a .optdbg file are as follows:

- If a .optdbg file is generated at the compile step, its name is based on the output file name of the compile step.
- If a .optdbg file is generated at the link step, its name is based on the output file name of the link step.

If you compile and link in the same step using the **-qoptdebug** option with **-qipa**, the .optdbg file is generated only at the link step.

You must still use the **-g** or **-qlinedebug** option to include debugging information that can be used by a debugger.

For more information and examples of using this option, see "Using -qoptdebug to help debug optimized programs" in the *XL C/C++ Optimization and Programming Guide* *XL C/C++ Optimization and Programming Guide*.

## Related information

- "**-O**, **-qoptimize**" on page 269
- "**-qhot**" on page 178
- "**-qipa**" on page 199
- "**-qpdf1**, **-qpdf2**" on page 283
- "**-qsmp**" on page 320
- "**-g**" on page 169
- "**-qlinedebug**" on page 242

## **-qoptfile**

### **Category**

Compiler customization

### **Pragma equivalent**

None.

### **Purpose**

Specifies a file containing a list of additional command line options to be used for the compilation.

### **Syntax**

►► -qoptfile=*filename* ◀◀

### **Defaults**

None.

### **Parameters**

*filename*

Specifies the name of the 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 option file follows these rules:

- Specify the options you want to include in the file with the same syntax as on the command line. The option 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 options 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 an option file. The files that contain another option file are processed in a depth-first manner. The compiler avoids infinite loops by detecting and ignoring cycles in option 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 option file is included representing the contents of each option file. The options contained in the file are saved to the compiled object file.

## Predefined macros

None.

## Examples

**This is an example of specifying an option file.**

```
$ cat options.file
# To perform optimization at -O4 level, and high-order
# loop analysis and transformations during optimization
-O4 -qhot
# To generate position-independent code
-qpic

$ xlc -qlist -qoptfile=options.file -qipa test.c
```

The preceding example is equivalent to the following invocation:

```
$ xlc -qlist -O4 -qhot -qpic -qipa test.c
```

**This is an example of specifying an option file that contains -qoptfile with a cycle.**

```
$ cat options.file2
# To perform optimization at -O4 level, and high-order
# loop analysis and transformations during optimization
-O4 -qhot
# To include the -qoptfile option in the same option file
-qoptfile=options.file2
# To generate position-independent code
-qpic
# To produce a compiler listing file
-qlist

$ xlc -qlist -qoptfile=options.file2 -qipa test.c
```

The preceding example is equivalent to the following invocation:

```
$ xlc -qlist -O4 -qhot -qpic -qlist -qipa test.c
```

**This is an example of specifying an option file that contains -qoptfile without a cycle.**

```
$ cat options.file1
-O4 -qhot
-qoptfile=options.file2
-qfixed

$ cat options.file2
-qfree

$ xlc -qoptfile=options.file1 test.c
```

The preceding example is equivalent to the following invocation:

```
$ xlc -O4 -qhot -qfree -qfixed test.c
```

**This is an example of specifying -qsaveopt and -qoptfile on the same command line.**

```
$ cat options.file3
-O4
-qhot

$ xlc -qsaveopt -qipa -qoptfile=options.file3 test.c -c
```



```
$ what test.o
test.o:
opt f xLC -qsaveopt -qipa -qoptfile=options.file3 test.c -c
optfile options.file3 -O4 -qhot
```

## Related information

- “-qsaveopt” on page 311

## -p, -pg, -qprofile

### Category

Optimization and tuning

### Pragma equivalent

None.

### Purpose

Prepares the object files produced by the compiler for profiling.

When you compile with a profiling option, the compiler produces monitoring code that counts the number of times each routine is called. The compiler replaces the startup routine of each subprogram with one that calls the monitor subroutine at the start. When you execute a program compiled with **-p**, and it ends normally, it writes the recorded information to a mon.out file; a program compiled with **-pg** writes a gmon.out file. You can then use the **prof** or **gprof** command to generate a runtime profile.

### Syntax



### Defaults

Not applicable.

### Usage

When you are compiling and linking in separate steps, you must specify the profiling option in both steps.

If the **-qtbtable** option is not set, the profiling options will generate full traceback tables.

### Predefined macros

None.

## Examples

To compile `myprogram.c` to include profiling data, enter:

```
xlc myprogram.c -p
```

Remember to compile *and* link with one of the profiling options. For example:

```
xlc myprogram.c -p -c  
xlc myprogram.o -p -o program
```

## Related information

- “-qthtable” on page 344
- See your operating system documentation for more information on the **prof** and **gprof** command.

## -P

### Category

Output control

### Pragma equivalent

None.

### Purpose

Preprocesses the source files named in the compiler invocation, without compiling, and creates an output preprocessed file for each input file.

The preprocessed output file has the same name as the input file, with an `.i` suffix.

### Syntax

►► -P ◀◀

### Defaults

By default, source files are preprocessed, compiled, and linked to produce an executable file.

### Usage

The **-P** option accepts any file name, except those with an `.i` suffix. Otherwise, source files with unrecognized file name suffixes are treated and preprocessed as `C` files, and no error message is generated.

Unless **-qppline** is specified, `#line` directives are not generated.

Line continuation sequences are removed and the source lines are concatenated.

The **-P** option retains all white space including line-feed characters, with the following exceptions:

- All comments are reduced to a single space (unless **-C** is specified).
- Line feeds at the end of preprocessing directives are not retained.

- White space surrounding arguments to function-style macros is not retained.

The **-P** option is overridden by the **-E** option. The **-P** option overrides the **-c**, **-o**, and **-qsyntaxonly** option.

## Predefined macros

None.

## Related information

- “**-C**, **-C!**” on page 120
- “**-E**” on page 142
- “**-qpplne**” on page 293
- “**-qsyntaxonly (C only)**” on page 341

## **-qpath**

### Category

Compiler customization

### Pragma equivalent


None.

### Purpose

Determines substitute path names for XL C/C++ executables such as the compiler, assembler, linker, and preprocessor.

You can use this option if you want to keep multiple levels of some or all of the XL C/C++ executables and have the option of specifying which one you want to use. This option is preferred over the **-B** and **-t** options.

### 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 alternate programs are located.

The following table shows the correspondence between **-qpath** parameters and the component executable names:

Parameter	Description	Executable name
a	Assembler	as
b	Low-level optimizer	xlCcode
c	Compiler front end	xlcentry, xlCentry
 C	C++ compiler front end	xlCentry
d	Disassembler	dis
E	CreateExportList utility	CreateExportList
 f	c++filt utility	c++filt
I	High-level optimizer, compile step	ipa
L	High-level optimizer, link step	ipa
l	Linker	ld
 m	Linkage helper	munch
p	Preprocessor	n/a

## Usage

The **-qpath** option overrides the **-F**, **-t**, and **-B** options.

Note that using the **p** suboption causes the source code to be preprocessed separately before compilation, which can change the way a program is compiled.

## Predefined macros

None.

## Examples

To compile `myprogram.c` using a substitute `xlC` compiler in `/lib/tmp/mine/` enter:  
`xlC myprogram.c -qpath=c:/lib/tmp/mine/`

To compile `myprogram.c` using a substitute linker in `/lib/tmp/mine/`, enter:  
`xlC myprogram.c -qpath=l:/lib/tmp/mine/`

## Related information

- “-B” on page 116
- “-F” on page 150
- “-t” on page 342

## -qpdf1, -qpdf2

### Category

Optimization and tuning

### Pragma equivalent


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.c** 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 29. Profiling type supported on each -qpdf1 level

Profiling type	Level		
	0	1	2
Block-counter profiling	+	+	+

Table 29. Profiling type supported on each `-qpdf1` level (continued)

Profiling type	Level		
	0	1	2
Call-counter profiling	+	+	+
Single-pass profiling	+	+	
Value profiling		+	+
Multiple-pass profiling			+
Cache-miss profiling			+

- `-qpdf1=level=0` is the basic level that results in smaller file size and faster compilation than `-qpdf1=level=1`.
- `-qpdf1=level=1` is the default level. It is equivalent to `-qpdf1` in releases earlier than IBM XL C/C++ for AIX, V11.1.
- `-qpdf1=level=2` is more aggressive than `-qpdf1=level=0` and `-qpdf1=level=1`. It is supported at all optimization levels where the PDF feature is enabled.

**Notes:**

- Only one application compiled with the `-qpdf1=level=2` option can be run at a time on a particular computer.
- Cache-miss profiling is enabled on **pSeries**<sup>®</sup> system, and is only available on the POWER5, POWER6, and POWER7 processors.
- 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 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. 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.

**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.

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 C/C++ 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 C/C++ Compiler Reference*.
- 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. If you specify both the **-qpdf[1|2]=exename** and **-qpdf[1|2]=pdfname** options, **-qpdf[1|2]=exename** is ignored.

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 and running 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 run the resulting application again with one or more sets of data, the profiling information is written to a new profile file named `._pdf.1` by default. If you repeat this compiling and running process 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 indicating 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_map.1`, and so on up to `._pdf_map.5`.

- If you use the **-qpdf2=pdfname** option to specify a PDF file, specify a file name that does not end with a numeric extension 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 extension, the compiler looks for (func, func.1, func.2, func.3).

### Other relative 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 316 for more information.

For recommended procedures of using PDF, see “Using profile-directed feedback” in the *XL C/C++ Optimization and Programming Guide*.

The following utility programs, found in `/usr/vacpp/bin/`, are available for managing the directory to which profiling information is written:

#### **cleanpdf**

```

>> cleanpdf [directory_path]

```

Removes all profiling information from the directory specified by *directory\_path*; or if *pathname* is not specified, from the directory set by the PDFDIR environment variable; or if PDFDIR is not set, from the current directory. Removing profiling information reduces runtime overhead if you change the program and then go through the PDF process again.

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 recompile all of the files again with **-qpdf1**.

#### **mergepdf**

```

>> mergepdf [-r-scaling] input -o output [-n] [-v]

```

Merges two or more PDF records into a single PDF output record.

#### **-r scaling**

Specifies the scaling ratio for the PDF record 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 record file, or a directory that contains PDF record files.

**-o** *output*

Specifies the name of the PDF output record file, or a directory to which the merged output is written.

**-n** If specified, PDF record files are not normalized. If not specified, **mergepdf** normalizes records 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

►—resetpdf directory\_path ►

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 C/C++ 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
xlc -qpdf1=level=0 -O3 file1.c file2.c file3.c
```

```
#Run with one set of input data
./a.out < sample.data
```

```
#Recompile all the files with -qpdf2
xlc -qpdf2 -O3 file1.c file2.c file3.c
```

```
#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
xlc -qpdf1 -O3 file1.c file2.c file3.c
```

```
#Run with one set of input data
./a.out < sample.data
```

```
#Recompile all the files with -qpdf2
xlc -qpdf2 -O3 file1.c file2.c file3.c
```

```
#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
xlc -qpdf1=level=2 -O5 file1.c file2.c file3.c

#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
xlc -qpdf2 -O5 file1.c file2.c file3.c

#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
xlc -qpdf1=level=2 -O5 file1.c file2.c file3.c

#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
xlc -qpdf2 -O5 file1.c file2.c file3.c

#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
xlc -qpdf1=level=2 -O5 file1.c file2.c file3.c

#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_map.1
xlc -qpdf1=level=2 -O5 file1.c file2.c file3.c

#Run it again, the profiling information is recorded in
#._pdf.1
./a.out < sample.data

#Recompile all the files with -qpdf2
xlc -qpdf2 -O5 file1.c file2.c file3.c
```

```
#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=2
xlc -qpdf1=level=2 -O5 file1.c file2.c file3.c

#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
xlc -qpdf2 -O5 file1.c file2.c file3.c

#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
xlc -qpdf1=exename -O5 -o final file1.c file2.c file3.c

#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.c
-rw-r--r-- 1 user staff 50 Dec 05 13:18 file2.c
-rw-r--r-- 1 user staff 50 Dec 05 13:18 file3.c
-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
xlc -qpdf2=exename -O5 -o final file1.c file2.c file3.c

#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
xlc -qpdf1=pdfname=final -O5 file1.c file2.c file3.c

#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.c
-rw-r--r-- 1 user staff 50 Dec 05 13:18 file2.c
-rw-r--r-- 1 user staff 50 Dec 05 13:18 file3.c
-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
xlc -qpdf2=pdfname=final -O5 file1.c file2.c file3.c

#The program is now optimized using PDF information
```

## Related information

- “-qshowpdf” on page 316
- “-qipa” on page 199
- -qprefetch
- “-qreport” on page 302
- “Optimizing your applications” in the *XL C/C++ Optimization and Programming Guide*
- “Runtime environment variables” on page 27
- “Profile-directed feedback” in the *XL C/C++ Optimization and Programming Guide*

## -qphsinfo

### Category

Listings, messages, and compiler information

### Pragma equivalent

None.

### Purpose

Reports the time taken in each compilation phase to standard output.

### Syntax

►► -q nophsinfo  
phsinfo ◀◀

### 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.

### Predefined macros

None.

### Examples

 To compile myprogram.c and report the time taken for each phase of the compilation, enter:

```
xlc myprogram.c -qphsinfo
```

The output will look similar to:

```

C Init    - Phase Ends;  0.010/  0.040
IL Gen    - Phase Ends;  0.040/  0.070
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

```

Compiling the same program with **-O4** gives:

```

C Init    - Phase Ends;  0.010/  0.040
IL Gen    - Phase Ends;  0.060/  0.070
IPA       - Phase Ends;  0.060/  0.070
IPA       - Phase Ends;  0.070/  0.110
W-TRANS   - Phase Ends;  0.060/  0.180
OPTIMIZ   - Phase Ends;  0.010/  0.010
REGALLO   - Phase Ends;  0.010/  0.020
AS        - Phase Ends;  0.000/  0.000

```

 To compile myprogram.C and report the time taken for each phase of the compilation, enter:

```
xlc++ myprogram.C -qphsinfo
```

The output will look similar to:

```

Front End - Phase Ends;  0.004/  0.005
W-TRANS   - Phase Ends;  0.010/  0.010
OPTIMIZ   - Phase Ends;  0.000/  0.000
REGALLO   - Phase Ends;  0.000/  0.000
AS        - Phase Ends;  0.000/  0.000

```

Compiling the same program with **-O4** gives:

```

Front End - Phase Ends;  0.004/  0.006
IPA       - Phase Ends;  0.040/  0.040
IPA       - Phase Ends;  0.220/  0.280
W-TRANS   - Phase Ends;  0.030/  0.110
OPTIMIZ   - Phase Ends;  0.030/  0.030
REGALLO   - Phase Ends;  0.010/  0.050
AS        - Phase Ends;  0.000/  0.000

```

## **-qpic**

### **Category**

Object code control

### **Pragma equivalent**

None.

### **Purpose**

Generates position-independent code suitable for use in shared libraries.

### **Syntax**

```

->> -qpic [small] [large]

```

## Defaults

- `-qplic=small`

## Parameters

### small

Instructs the compiler to assume that the size of the Table of Contents (TOC) is no larger than 64 Kb. When `-qplic=small` is in effect, the compiler generates one instruction for each TOC access.

### large

Instructs the compiler to assume that the size of the TOC is larger than 64 Kb. When `-qplic=large` 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.

**Note:** To enable `-qplic=large`, your operating system must be at AIX 5.3 TL 9 or later:

- If your operating system is at AIX 5.3 TL 9 to AIX 5.3 TL 12 or AIX 6.1 TL 1 to AIX 6.1 TL 6, ensure that you have installed the latest fix pack from <https://www.ibm.com/support/docview.wss?uid=isg1fixinfo118013>. Otherwise, an error message might be generated.
- If your system is at AIX 6.1 TL 7 or later, no action is required.

Specifying `-qplic` without any suboptions is equivalent to `-qplic=small`.

## Usage

Specifying `-qplic=large` has the same effect as passing `-bbigtoc` to `ld`.

You can use different TOC access options for different compilation units in an application.

**Note:** For applications whose TOC size is larger than 64K, using `-qplic=large` can improve performance. However, for applications whose TOC is smaller than 64K, using `-qplic=large` slows down the program. To decide whether to use `-qplic=large`, compile the program with `-qplic=small` first. If an overflow error message is generated, use `-qplic=large` instead.

## Predefined macros

None.

## Examples

To compile a shared library `libmylib.so`, use the following commands:

```
xlc mylib.c -qplic=small -c -o mylib.o
xlc -qmksbobj mylib -o libmylib.so.1
```

## Related information

- “`-q32`, `-q64`” on page 96
- “`-G`” on page 172
- “`-qmksbobj`” on page 260

## **-qppline**

### **Category**

Object code control

### **Pragma equivalent**

None.

### **Purpose**

When used in conjunction with the **-E** or **-P** options, enables or disables the generation of `#line` directives.

### **Syntax**

►► -q ppline noppline ◄◄

### **Defaults**

- **-qnoppline** when **-P** is in effect
- **-qppline** when **-E** is in effect

### **Usage**

The **-C** option has no effect without either the **-E** or the **-P** option. With the **-E** option, line directives are written to standard output. With the **-P** option, line directives are written to an output file.

### **Predefined macros**

None.

### **Examples**

To preprocess `myprogram.c` to write the output to `myprogram.i`, and generate `#line` directives:

```
xlc myprogram.c -P -qppline
```

### **Related information**

- “**-E**” on page 142
- “**-P**” on page 280

## **-qprefetch**

### **Category**

Optimization and tuning

### **Pragma equivalent**


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**
- **-qprefetch=noassistthread**
- **-qprefetch=noassistthread:noaggressive**

## 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.

### **aggressive** | **noaggressive**

This suboption guides the compiler to generate aggressive data prefetching at optimization level **-O3 -qhot** or higher. If you do not specify **aggressive**, **-qprefetch=noaggressive** 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.

## 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`

## Examples

Here is how you generate code using assist threads with `__MEM_DELAY`:

Initial code:

```
int y[64], x[1089], w[1024];

void foo(void){
    int i, j;
    for (i = 0; i &1; 64; i++) {
        for (j = 0; j < 1024; j++) {

            /* what to prefetch? y[i]; inserted by the user */
            __mem_delay(&y[i], 10);
            y[i] = y[i] + x[i + j] * w[j];
            x[i + j + 1] = y[i] * 2;
        }
    }
}
```

Assist thread generated code:

```
void foo@clone(unsigned thread_id, unsigned version)

{ if (!1) goto lab_1;

/* version control to synchronize assist and main thread */
if (version == @2version0) goto lab_5;

goto lab_1;

lab_5:
@CIV1 = 0;

do { /* id=1 guarded */ /* ~2 */

if (!1) goto lab_3;

@CIV0 = 0;

do { /* id=2 guarded */ /* ~4 */

/* region = 0 */

/* __dcbt call generated to prefetch y[i] access */
__dcbt(((char *)&y + (4)*(@CIV1)))
@CIV0 = @CIV0 + 1;
} while ((unsigned) @CIV0 < 1024u); /* ~4 */

lab_3:
@CIV1 = @CIV1 + 1;
} while ((unsigned) @CIV1 < 64u); /* ~2 */
```

```
lab_1:  
return;  
}
```

### Related information

- “-qarch” on page 106
- “-qhot” on page 178
- “-qpdf1, -qpdf2” on page 283
- “-qreport” on page 302
- “\_\_mem\_delay” on page 607

## -qprint

### Category

Listings, messages, and compiler information

### Pragma equivalent

None.

### Purpose

Enables or suppresses listings.

When **-qprint** is in effect, listings are enabled if they are requested by other compiler options that produce listings. When **-qnoprint** is in effect, all listings are suppressed, regardless of whether listing-producing options are specified.

### Syntax

►► -q print  
noprint ◀◀

### Defaults

-qprint

### Usage

You can use **-qnoprint** to override all listing-producing options and equivalent pragmas, regardless of where they are specified. These options are:

- -qattr
- -qlist
- -qlistopt
- -qsource
- -qxref

### Predefined macros

None.

## Examples

To compile `myprogram.c` and suppress all listings, even if some files have `#pragma options source` and similar directives, enter:

```
xlc myprogram.c -qnoprint
```

## -qpriority (C++ only)

### Category

Object code control

### Pragma equivalent

`#pragma options priority`, `#pragma priority`

### Purpose

Specifies the priority level for the initialization of static objects.

The C++ standard requires that all global objects within the same translation unit be constructed from top to bottom, but it does not impose an ordering for objects declared in different translation units. The `-qpriority` option and `#pragma priority` directive allow you to impose a construction order for all static objects declared within the same load module. Destructors for these objects are run in reverse order during termination.

### Syntax

#### Option syntax

► `-qpriority=number` ◀

#### Pragma syntax

► `#pragma priority(number)` ◀

### Defaults

The default priority level is 0.

### Parameters

*number*

An integer literal in the range of -2 147 482 624 to 2147483647. A lower value indicates a higher priority; a higher value indicates a lower priority. Numbers from -214 783 648 to -214 782 623 are reserved for system use. If you do not specify a *number*, the compiler assumes 0.

### Usage

More than one `#pragma priority` can be specified within a translation unit. The priority value specified in one pragma applies to the constructions of all global objects declared after this pragma and before the next one. However, in order to be

consistent with the Standard, priority values specified within the same translation unit must be strictly increasing. Objects with the same priority value are constructed in declaration order.

The effect of a **#pragma priority** exists only within one load module. Therefore, **#pragma priority** cannot be used to control the construction order of objects in different load modules. Refer to "Initializing static objects in libraries" in the *XL C/C++ Optimization and Programming Guide* for further discussions on techniques used in handling static object initialization across modules.

## Examples

To compile the file `myprogram.C` to produce an object file `myprogram.o` so that objects within that file have an initialization priority of 2 000, enter:

```
xlc++ myprogram.C -c -qpriority=2000
```

## Related information

- “-qmkshrobj” on page 260
- "Initializing static objects in libraries" in the *XL C/C++ Optimization and Programming Guide*

## -qprocimported, -qproclocal, -qprocunknown

### Category

Optimization and tuning

### Pragma equivalent

#pragma options `proclocal`, #pragma options `procimported`, #pragma options `procunknown`

### Purpose


Marks functions as local, imported, or unknown in 64-bit compilations.

Local functions are statically bound with the functions that call them; smaller, faster code is generated for calls to such functions. You can use the **proclocal** option or pragma to name functions that the compiler can assume are local.

Imported functions are dynamically bound with a shared portion of a library. Code generated for calls to functions marked as imported may be larger, but is faster than the default code sequence generated for functions marked as unknown. You can use the **procimported** option or pragma to name functions that the compiler can assume are imported.

Unknown functions are resolved to either statically or dynamically bound objects during linking. You can use the **procunknown** option or pragma to name functions that the compiler can assume are unknown.

## Syntax



## Defaults

**-qprocunknown:** The compiler assumes that all functions' definitions are unknown.

## Parameters

*function\_name*

The name of a function that the compiler should assume is local, imported, or unknown (depending on the option specified). If you do not specify any *function\_name*, the compiler assumes that *all* functions are local, imported, or unknown.

**C++** Names must be specified using their mangled names. To obtain C++ mangled names, compile your source to object files only, using the **-c** compiler option, and use the **nm** operating system command on the resulting object file. You can also use the **c++filt** utility provided by the compiler for a side-by-side listing of source names and mangled names; see "Demangling compiled C++ names" in the *XL C/C++ Optimization and Programming Guide* for details. (See also "Name mangling" in the *XL C/C++ Language Reference* for details on using the extern "C" linkage specifier on declarations to prevent name mangling.)

## Usage

This option applies to 64-bit compilations only.

If any functions that are marked as local resolve to shared library functions, the linker will detect the error and issue warnings. If any of the functions that are marked as imported resolve to statically bound objects, the generated code may be larger and run more slowly than the default code sequence generated for unknown functions.

If you specify more than one of these options with no function names, the last option specified is used. If you specify the same function name on more than one option specification, the last one is used.

## Predefined macros

None.

## Examples

To compile `myprogram.c` along with the archive library `oldprogs.a` so that:

- Functions `fun` and `sun` are specified as local
- Functions `moon` and `stars` are specified as imported
- Function `venus` is specified as unknown

use the following command:

```
xlc myprogram.c oldprogs.a -qprolocal=fun(int):sun()  
-qprocimported=moon():stars(float) -qprocunknown=venus()
```

If the following example, in which a function marked as local instead resolves to a shared library function, is compiled with **-qprolocal**:

```
int main(void)  
{  
    printf("Just in function foo1()\n");  
    printf("Just in function foo1()\n");  
}
```

a linker error will result. To correct this problem, you should explicitly mark the called routine as being imported from a shared object. In this case, you would recompile the source file and explicitly mark `printf` as imported by compiling with `-qprolocal -qprocimported=printf`.

### Related information

- “-qdataimported, -qdatalocal, -qtocdata” on page 134

## -qproto (C only)

### Category

Object code control

### Pragma equivalent

#pragma options [no]proto

### Purpose

Specifies the linkage conventions for passing floating-point arguments to functions that have not been prototyped.

When **proto** is in effect, the compiler assumes that the arguments in function calls are the same types as the corresponding parameters of the function definition, even if the function has not been prototyped. By asserting that an unprototyped function actually expects a floating-point argument if it is called with one, you allow the compiler to pass floating-point arguments in floating-point registers exclusively. When **noproto** is in effect, the compiler does not make this assumption, and must pass floating-point parameters in floating-point and general purpose registers.

### Syntax

►► — -q — no proto —►►

### Defaults

-qnoproto

### Usage

This option is only valid when the compiler allows unprototyped functions; that is, with the `cc` or `xlc` invocation command, or with the `-qlanglvl` option set to **classic** | **extended** | **extc89** | **extc99**.

## Predefined macros

None.

## Examples

To compile `my_c_program.c` to allow the compiler to use the standard linkage conventions for floating-point parameters, even when functions are not prototyped, enter:

```
xlc my_c_program.c -qproto
```

**-r**

## Category

Object code control

## Pragma equivalent

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.c` and `myprog2.c` into a single object file `mytest.o`, enter:

```
xlc myprogram.c myprog2.c -r -o mytest.o
```

## Related information

- `-qipa`

## **-qreport**

### **Category**

Listings, messages, and compiler information

### **Pragma equivalent**

None.

### **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-C 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. POWER4 and POWER5 support load stream prefetch and POWER6 supports both load and store stream prefetch.

### **Syntax**

```
→ -q [noreport] [report] →
```

### **Defaults**

`-qnoreport`

### **Usage**

For **-qreport** to generate a loop transformation listing, you must also specify one of the following on the command line:

- **-qsimd=auto**
- **-qsmp**
- **-qhot=level=2** and **-qsmp**
- **-O5**
- **-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**
- **-O5**
- **-qipa=level=2**

To generate data reorganization information, specify the optimization level **-qipa=level=2** or **-O5** together with **-qreport**. Reorganizations include 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 parallelizations 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-C code listing is not intended to be compilable. Do not include any of the pseudo-C code in your program, and do not explicitly call any of the internal routines whose names may appear in the pseudo-C code listing.

## Predefined macros

None.

## Examples

To compile `myprogram.c` so the compiler listing includes a report showing how loops are optimized, enter:

```
xlc -qhot -O3 -qreport myprogram.c
```

To compile `myprogram.c` so the compiler listing also includes a report showing how parallelized loops are transformed, enter:

```
xlc_r -qhot -qsmp -qreport myprogram.c
```

## Related information

- “-qhot” on page 178
- “-qsimd” on page 316
- “-qipa” on page 199
- “-qsmp” on page 320
- “-qoptdebug” on page 275
- “-qprefetch” on page 293
- "Using -qoptdebug to help debug optimized programs" in the *XL C/C++ Optimization and Programming Guide*

## **-qreserved\_reg**

### **Category**

Object code control

### **Pragma equivalent**

None.

### **Purpose**

Indicates that the given list of registers cannot be used during the compilation except as a stack pointer, frame pointer or in some other fixed role.

You should use this option in modules that are required to work with other modules that use global register variables or hand-written assembler code.

### **Syntax**

```
→ -qreserved_reg=register_name →
```

### **Defaults**

Not applicable.

### **Parameters**

*register\_name*

A valid register name on the target platform. Valid registers are:

**r0 to r31**

General purpose registers

**f0 to f31**

Floating-point registers

**v0 to v31**

Vector registers (on selected processors only)

### **Usage**

**-qreserved\_reg** is cumulative, for example, specifying **-qreserved\_reg=r14** and **-qreserved\_reg=r15** is equivalent to specifying **-qreserved\_reg=r14:r15**.

Duplicate register names are ignored.

### **Predefined macros**

None.

### **Examples**

To specify that `myprogram.c` reserves the general purpose registers `r3` and `r4`, enter:

```
xlc myprogram.c -qreserved_reg=r3:r4
```

## Related information

- "Variables in specified registers" in the *XL C/C++ Language Reference*

## -qro

### Category

Object code control

### Pragma equivalent

#pragma options ro, #pragma strings

### Purpose

Specifies the storage type for string literals.

When **ro** or **strings=readonly** is in effect, strings are placed in read-only storage. When **norow** or **strings=writeable** is in effect, strings are placed in read/write storage.

### Syntax

#### Option syntax

►► -q ro norow \_\_\_\_\_►►

#### Pragma syntax

►► #pragma strings ( readonly writeable ) \_\_\_\_\_►►

### Defaults

**C** Strings are read-only for all invocation commands except **cc**. If the **cc** invocation command is used, strings are writeable.

**C++** Strings are read-only.

### Parameters

#### **readonly** (pragma only)

String literals are to be placed in read-only memory.

#### **writeable** (pragma only)

String literals are to be placed in read-write memory.

### Usage

Placing string literals in read-only memory can improve runtime performance and save storage. However, code that attempts to modify a read-only string literal may generate a memory error.

The pragmas must appear before any source statements in a file.

## Predefined macros

None.

## Examples

To compile `myprogram.c` so that the storage type is writable, enter:

```
xlc myprogram.c -qnor
```

## Related information

- “-qro” on page 305
- “-qroconst”

## -qroconst

### Category

Object code control

### Pragma equivalent

```
#pragma options [no]roconst
```

### Purpose

Specifies the storage location for constant values.

When **roconst** is in effect, constants are placed in read-only storage. When **norroconst** is in effect, constants are placed in read/write storage.

### Syntax

► -q roconst  
norroconst ►

### Defaults

- C **-qroconst** for all compiler invocations except **cc** and its derivatives.  
**-qnorroconst** for the **cc** invocation and its derivatives.
- C++ **-qroconst**

### Usage

Placing constant values in read-only memory can improve runtime performance, save storage, and provide shared access. However, code that attempts to modify a read-only constant value generates a memory error.

"Constant" in the context of the **-qroconst** option refers to variables that are qualified by `const`, including `const`-qualified characters, integers, floats, enumerations, structures, unions, and arrays. The following constructs are not affected by this option:

- Variables qualified with `volatile` and aggregates (such as a structure or a union) that contain `volatile` variables
- Pointers and complex aggregates containing pointer members

- Automatic and static types with block scope
- Uninitialized types
- Regular structures with all members qualified by `const`
- Initializers that are addresses, or initializers that are cast to non-address values

The `-qroconst` option does not imply the `-qro` option. Both options must be specified if you want to specify storage characteristics of both string literals (`-qro`) and constant values (`-qroconst`).

## Predefined macros

None.

## Related information

- “`-qro`” on page 305
- “`-qroptr`”

## **-qroptr**

### Category

Object code control

### Pragma equivalent

None.

### Purpose

Specifies the storage location for constant pointers.

When `-qroptr` is in effect, constant pointers, virtual function tables, and virtual type tables are placed in read-only storage. When `-qnoroptr` is in effect, pointers, virtual function tables, and virtual type tables are placed in read/write storage.

### Syntax

►► — `-q` noroptr  
roptr —►►

### Defaults

`-qnoroptr`

### Usage

A constant pointer is equivalent to an address constant. For example:

```
int* const p = &n;
```

When `-qnoroptr` is in effect, you can change the values of constant pointers, virtual function tables, and virtual type tables without generating errors.

The **-qroptr** can improve runtime performance, save storage, and provide shared access, but code that attempts to modify a read-only constant value generates a memory error. For example, assume the following code, which attempts to change the address that `c1_ptr` points to:

```
char c1 = 10;  
char c2 = 20;  
char* const c1_ptr = &c1;  
  
int main() {  
    *(char**)&c1_ptr = &c2;  
}
```

Compiling this code with the **-qroptr** option specified will result in a segmentation fault at run time.

You should not use **-qroptr** for compiled code that will become part of a shared library.

## Predefined macros

None.

## Related information

- “-qro” on page 305
- “-groconst” on page 306

## -qrtti (C++ only)

### Category

Object code control


### Pragma equivalent

#pragma options rtti

### Purpose

Generates runtime type identification (RTTI) information for exception handling and for use by the `typeid` and `dynamic_cast` operators.

### Syntax



### Defaults

`-qnortti`

## Parameters

### **all**

The compiler generates the information needed for the RTTI typeid and dynamic\_cast operators.

### **type | typeinfo**

The compiler generates the information needed for the RTTI typeid operator, but the information needed for dynamic\_cast operator is not generated.

### **dyna | dynamiccast**

The compiler generates the information needed for the RTTI dynamic\_cast operator, but the information needed for typeid operator is not generated.

Specifying **-qrtti** with no suboptions is equivalent to **-qrtti=all**.

## Usage

For improved runtime performance, suppress RTTI information generation with the **-qnortti** setting.

You should be aware of the following effects when specifying the **-qrtti** compiler option:

- Contents of the virtual function table will be different when **-qrtti** is specified.
- When linking objects together, all corresponding source files must be compiled with the correct **-qrtti** option specified.
- If you compile a library with mixed objects (**-qrtti** specified for some objects, **-qnortti** specified for others), you may get an undefined symbol error.

## Predefined macros

- `__RTTI_ALL__` is defined to 1 when **-qrtti** or **-qrtti=all** is in effect; otherwise, it is undefined.
- `__RTTI_DYNAMIC_CAST__` is defined to 1 when **-qrtti** or **-qrtti=all | dynamiccast** is in effect; otherwise, it is undefined.
- `__RTTI_TYPE_INFO__` is defined to 1 when **-qrtti** or **-qrtti=all | typeinfo** is in effect; otherwise, it is undefined.
- `__NO_RTTI__` is defined to 1 when **-qnortti** is in effect; otherwise, it is undefined.

## Related information

- “-geh (C++ only)” on page 144

## -S

### Category

Object code control

### Pragma equivalent

None.

### Purpose

Strips the symbol table, line number information, and relocation information from the output file.

This command is equivalent to the operating system **strip** command.

## Syntax

▶▶ -s —————▶▶

## Defaults

The symbol table, line number information, and relocation information are included in the output file.

## Usage

Specifying **-s** saves space, but limits the usefulness of traditional debug programs when you are generating debugging information using options such as **-g**.

## Predefined macros

None.

## Related information

- “-g” on page 169

## -S

### Category

Output control

### Pragma equivalent

None.

### Purpose

Generates an assembler language file for each source file.

The resulting file has an **.s** suffix and can be assembled to produce object **.o** files or an executable file (**a.out**).

## Syntax

▶▶ -S —————▶▶

## Defaults

Not applicable.

## Usage

You can invoke the assembler with any compiler invocation command. For example,

```
xlc myprogram.s
```



will invoke the assembler, and if successful, the linker to create an executable file, a.out.

If you specify **-S** with **-E** or **-P**, **-E** or **-P** takes precedence. Order of precedence holds regardless of the order in which they were specified on the command line.

You can use the **-o** option to specify the name of the file produced only if no more than one source file is supplied. For example, the following is *not* valid:

```
xlc myprogram1.c myprogram2.c -o -S
```

## Predefined macros

None.

## Examples

To compile myprogram.c to produce an assembler language file myprogram.s, enter:

```
xlc myprogram.c -S
```

To assemble this program to produce an object file myprogram.o, enter:

```
xlc myprogram.s -c
```

To compile myprogram.c to produce an assembler language file asmprogram.s, enter:

```
xlc myprogram.c -S -o asmprogram.s
```

## Related information

- “-E” on page 142
- “-P” on page 280

## -qsaveopt

### Category

Object code control

### Pragma equivalent

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 ◀◀

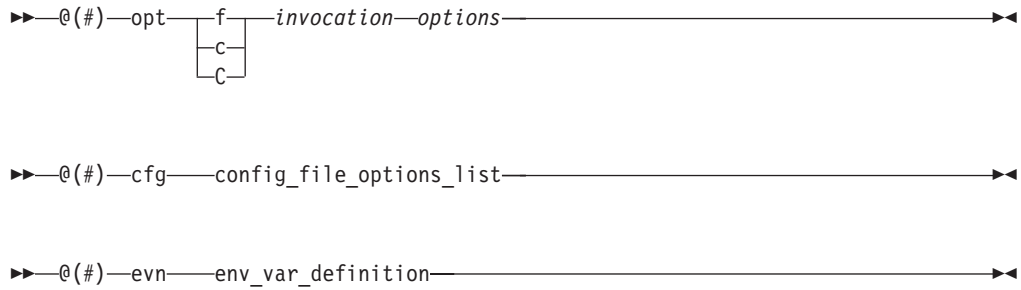
### Defaults

-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 pragma directives are ignored.

Command-line compiler options information is copied as a string into the object file, using the following format:



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, `xlc`.

*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 **XLC\_USR\_CONFIG** is listed.

**Note:** You can always use this option, but the corresponding information is only generated when the environment variable **XLC\_USR\_CONFIG** is set.

For more information about the environment variable **XLC\_USR\_CONFIG**, see Compile-time and link-time environment variables.

**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:



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 (PTF). If the update installed is at the base level, the level is displayed as BASE.

If you want to simply output this information to standard output without writing it to the object file, use the **-qversion** option.

## Predefined macros

None.

## Examples

Compile `t.c` with the following command:

```
xlc t.c -c -qsaveopt -qhot
```

Issuing the **what** command on the resulting `t.o` object file produces information similar to the following:

```
opt c /usr/vac/bin/xlc t.f -c -qsaveopt -qhot
cfg -qlanglvl=extc99 -qcpluscmt -qkeyword=inline -qalias=ansi -D_AIX -D_AIX32
-D_AIX41 -D_AIX43 -D_AIX50 -D_AIX51 -D_AIX52 -D_AIX53 -D_IBMR2 -D_POWER
version IBM XL C/C++ for AIX, V12.1
version Version: 12.01.0000.0000
version Driver Version: 12.01(C/C++) Level: YYMMDD
version Front End Version: 12.01(C/C++) Level: YYMMDDversion C Front End Version : 12.01(C/C++)
version High-Level Optimizer Version: 12.01(C) and 14.01(Fortran) Level: YYMMDD
version Low-Level Optimizer Version: 12.01(C) and 14.01(Fortran) Level: YYMMDD
```

In the first line, `/usr/vac/bin/xlc` identifies the source used as C, `/usr/vacpp/bin/xlc` shows the invocation command used, and `-qhot -qsaveopt` shows the compilation options.

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 (PTFs) you have installed on your system.

## Related information

- “-qversion” on page 370

## -qshowinc

### Category

Listings, messages, and compiler information


### Pragma equivalent

```
#pragma options [no]showinc
```

## Purpose

When used with **-qsource** option to generate a listing file, selectively shows user or system header files in the source section of the listing file.

## Syntax



## Defaults

**-qnoshowinc**: Header files included in source files are not shown in the source listing.

## Parameters

### all

Shows both user and system include files in the program source listing.

### **sys**

Shows system include files (that is, files included with the `#include <filename>` preprocessor directive) in the program source listing.

### **usr**

Shows user include files (that is, files included with the `#include "filename"` preprocessor directive or with **-qinclude**) in the program source listing.

Specifying **showinc** with no suboptions is equivalent to **-qshowinc=sys : usr** and **-qshowinc=all**. Specifying **noshowinc** is equivalent to **-qshowinc=nosys : nousr**.

## Usage

This option has effect only when the **-qlist** or **-qsource** compiler options is in effect.

## Predefined macros

None.

## Examples

To compile `myprogram.c` so that all included files appear in the source listing, enter:  
`xlc myprogram.c -qsource -qshowinc`

## Related information

- “-qsource” on page 324

## -qshowmacros

### Category

“Output control” on page 75

### Pragma equivalent


None

### Purpose

Emits macro definitions to preprocessed output.

Emitting macros to preprocessed output can help determine functionality available in the compiler. The macro listing may prove useful for debugging complex macro expansions, as well.

### Syntax



### Defaults

-qnoshowmacros

### Parameters

#### all

Emits all macro definitions to preprocessed output. This is the same as specifying **-qshowmacros**.

#### pre | nopre

**pre** emits only predefined macro definitions to preprocessed output. **nopre** suppresses appending these definitions.

### Usage

Note the following when using this option:

- This option has no effect unless preprocessed output is generated; for example, by using the **-E** or **-P** options.
- If a macro is defined and subsequently undefined before compilation ends, this macro will not be included in the preprocessed output.
- Only macros defined internally by the preprocessor are considered predefined; all other macros are considered as user-defined.

### Related information

- “-E” on page 142
- “-P” on page 280

## **-qshowpdf**

### **Category**

Optimization and tuning

### **Pragma equivalent**

None.

### **Purpose**

When used with **-qpdf1** and a minimum optimization level of **-O2** at compile and link steps, creates a PDF map file that contains additional profiling information for all procedures in your application.

### **Syntax**

→ -q showpdf noshowpdf →

### **Defaults**

-qshowpdf

### **Usage**

After you run your application with typical data, the profiling information is recorded into a profile-directed feedback (PDF) file (by default, the file is named `._pdf`).

In addition to the PDF file, the compiler also generates a PDF map file that contains static information during the **-qpdf1** phase. With these two files, you can use the **showpdf** utility to view part of the profiling information of your application in text or XML format. For details of the **showpdf** utility, see "Viewing profiling information with showpdf" in the *XL C/C++ 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.

### **Predefined macros**

None.

### **Related information**

- "**-qpdf1**, **-qpdf2**" on page 283
- "Optimizing your applications" in the *XL C/C++ Optimization and Programming Guide*

## **-qsimd**

### **Category**

Optimization and tuning

## Pragma equivalent

```
#pragma nosimd
```

## 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=  
┌ noauto  
└ auto
```

## Defaults

```
-qsimd=noauto
```

## Usage

The **-qsimd=auto** option enables automatic generation of vector instructions for processors that support them. It replaces the **-qenablevmx** option, which has been deprecated.

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 333.

**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:

- Specifying the deprecated **-qenablevmx** option has the same effect as specifying **-qsimd=auto**. The compiler does not issue any warning for this.
- Specifying **-qsimd** without any suboption has the same effect as **-qsimd=auto**.
- This option is available only when you set **-qarch** to a target architecture that supports vector instructions.
- If you specify **-qsimd=auto** to enable IPA at the compile time but specify **-qsimd=noauto** at the link time, the compiler automatically sets **-qsimd=auto** and sets an appropriate value for **-qarch** to match the architecture specified at the compile time.

## Predefined macros

None.

## Examples

### Example 1

The following example shows the usage of **#pragma nosimd** to disable **-qsimd=auto** for a specific for loop:

```
...
#pragma nosimd
for (i=1; i<1000; i++) {
    /* program code */
}
```

## Related information

- “-qarch” on page 106
- “-qstrict” on page 333

## -qskipsrc

### Category

“Listings, messages, and compiler information” on page 85

### Pragma equivalent

None.

### Purpose

When a listing file is generated using the **-qsource** option, **-qskipsrc** can be used to determine whether the source statements skipped by the compiler are shown in the source section of the listing file. Alternatively, the **-qskipsrc=hide** option is used to hide the source statements skipped by the compiler.

### Syntax

►► -q-skipsrc=show  
hide ◀◀

### Defaults

- -qskipsrc=show

### Parameters

**show** | **hide**

When **show** is in effect, the compiler will display all source statements in the listing. This will result in both true and false paths of the preprocessing directives to be shown.

On the contrary, when **hide** is enabled, all source statements that the compiler skipped will be omitted.



## Usage

In general, the **-qskipsrc** option does not control whether the source section is included in the listing file, it only does so when the **-qsource** option is in effect.

To display all source statements in the listing (default option):

```
xlc myprogram.c -qsource -qskipsrc=show
```

To omit source statements skipped by the compiler:

```
xlc myprogram.c -qsource -qskipsrc=hide
```

## Predefined macros

None.

## Related information

- “-qsource” on page 324
- “-qshowinc” on page 313
- “-qsrcmsg (C only)” on page 328

## -qsmallstack

### Category

Optimization and tuning

### Pragma equivalent

None.

### Purpose

Reduces the size of the stack frame.

### Syntax

►► -q nosmallstack smallstack ◀◀

### Defaults

-qnosmallstack

### Usage

AIX limits the stack size to 256 MB. Programs that allocate large amounts of data to the stack, such as threaded programs, may result in stack overflows. This option can reduce the size of the stack frame to help avoid overflows.

This option is only valid when used together with IPA (**-qipa**, **-O4**, **-O5** compiler options).

Specifying this option may adversely affect program performance.

## Predefined macros

None.

## Examples

To compile myprogram.c to use a small stack frame, enter:

```
xlc myprogram.c -qipa -qsmallstack
```

## Related information

- “-g” on page 169
- “-qipa” on page 199
- “-O, -qoptimize” on page 269

## -qsmp

### Category

Optimization and tuning

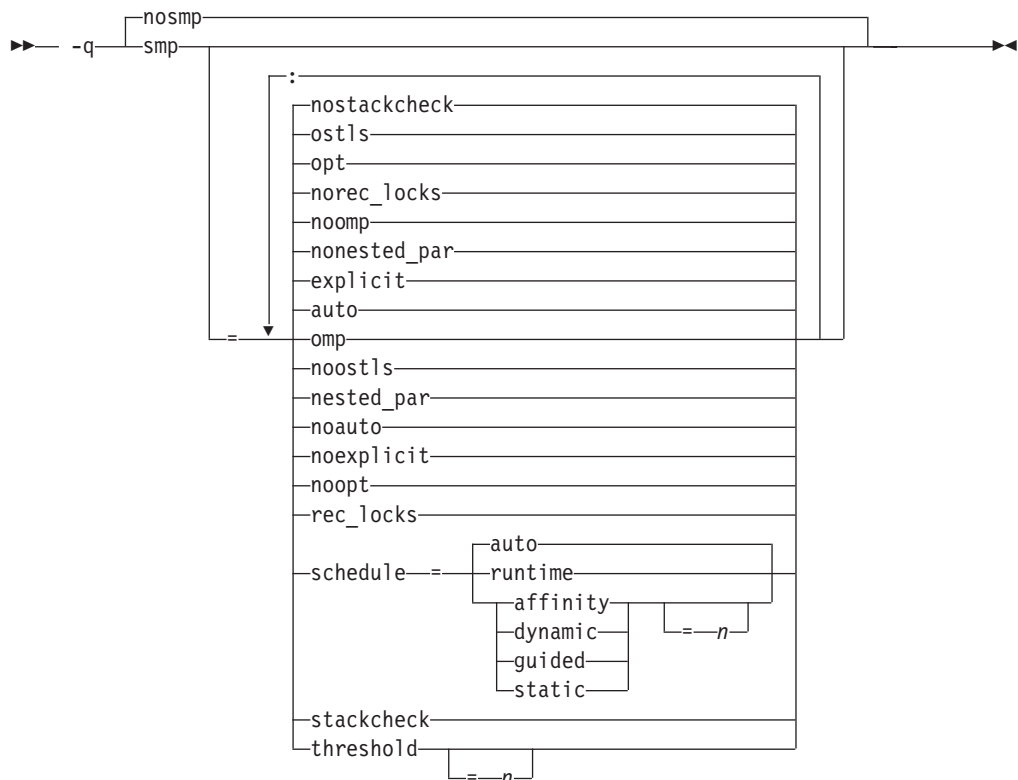
### Pragma equivalent

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. When **noauto** is in effect, only program code explicitly parallelized with SMP or OpenMP directives is optimized. **noauto** is implied if you specify **-qsmp=omp** or **-qsmp=noopt**.

### **explicit** | **noexplicit**

Enables or disables directives controlling explicit parallelization of loops.

### **nested\_par** | **nonested\_par**

By default, the compiler serializes a nested parallel construct. When **nested\_par** is in effect, the compiler parallelizes prescriptive nested parallel constructs. This includes not only the loop constructs that are nested within a scoping unit but also parallel constructs in subprograms that are referenced (directly or indirectly) from within other parallel constructs. Note that this suboption has no effect on loops that are automatically parallelized. In this case, at most one loop in a loop nest (in a scoping unit) will be parallelized. **nested\_par** does not provide true nested parallelism because it does not cause a new team of threads to be created for nested parallel regions. Instead, threads that are currently available are reused.

This suboption should be used with caution. Depending on the number of threads available and the amount of work in an outer loop, inner loops could be executed sequentially even if this option is in effect. Parallelization overhead may not necessarily be offset by program performance gains.

#### Note:

- The implementation of the **nested\_par** suboption does not comply with the OpenMP API.
- If you specify this suboption, the runtime library uses the same threads for the nested constructs that it used for the enclosing constructs.

### **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.

### **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.

### **ostls** | **noostls**

Enables Thread Local Storage (TLS) provided by the operating system to be used for **threadprivate** data. You can use the **noostls** suboption to enable the non-TLS for **threadprivate**. The **noostls** suboption is provided for compatibility with earlier versions.

**Note:** If you want to use this suboption, your operating system must support TLS to implement OpenMP **threadprivate** data. Use **noostls** to disable OS level TLS if your operating system does not support it.

### **rec\_locks** | **norec\_locks**

Determines whether recursive locks are used. When **rec\_locks** is in effect, nested critical sections will not cause a deadlock. 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 does not appear in the OpenMP API standard.

#### **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 containing  $n$  iterations each. If  $n$  is not specified, then the chunks consist of **ceiling**( $number\_of\_iterations / number\_of\_threads$ ). iterations.

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.

If a thread is asleep, its assigned work will be taken over by an active thread once that thread becomes available.

#### **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 **ceiling**(*number\_of\_iterations*/*number\_of\_threads*) iterations. Each thread is assigned one of these chunks. 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**.

**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" on page 28.

**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:explicit:opt:noomp:norec_locks:nested_par:schedule=auto:  
nostackcheck:threshold=100:ostls
```

**Usage**

- Specifying the **omp** suboption always implies **noauto**. Specify **-qsmp=omp:auto** to apply automatic parallelization on OpenMP-compliant applications, as well.
- 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.

- 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.
- The **-qnosmp** default option setting specifies that no code should be generated for parallelization directives, though syntax checking will still be performed. Use **-qignprag=omp:ibm** to completely ignore parallelization directives.
- 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**.

## Predefined macros

 When **-qsmp** is in effect, **\_IBMSMP** is predefined to a value of 1, which indicates that IBM SMP directives are recognized; otherwise, it is not defined.

## Related information

- “-O, -qoptimize” on page 269
- “-qthreaded” on page 350
- “Environment variables for parallel processing” on page 27
- “Pragma directives for parallel processing” on page 439
- “Built-in functions for parallel processing” on page 609

## -qsource

### Category

Listings, messages, and compiler information

### Pragma equivalent

#pragma options [no]source

### Purpose

Produces a compiler listing file that includes the source section of the listing and provides additional source information when printing error messages.

When **source** is in effect, a listing file is generated with a .lst suffix for each source file named on the command line. For details of the contents of the listing file, see “Compiler listings” on page 20.

### Syntax

→ -q nosource  
source →

### Defaults

-qnosource

## Usage

You can selectively print parts of the source by using pairs of **#pragma options source** and **#pragma options nosource** preprocessor directives throughout your source program. The source following **#pragma options source** and preceding **#pragma options nosource** is printed.

The **-qnoprint** option overrides this option.

## Predefined macros

None.

## Examples

To compile `myprogram.c` to produce a compiler listing that includes the source code, enter:

```
xlc myprogram.c -qsource
```

## Related information

- “-qlist” on page 243
- “-qlistopt” on page 248
- “-qprint” on page 296

## -qsourcetype

### Category

Input control

### Pragma equivalent


None.

### Purpose

Instructs the compiler to treat all recognized source files as a specified source type, regardless of the actual file name suffix.

Ordinarily, the compiler uses the file name suffix of source files specified on the command line to determine the type of the source file. For example, a `.c` suffix normally implies C source code, and a `.C` suffix normally implies C++ source code. The **-qsourcetype** option instructs the compiler to not rely on the file name suffix, and to instead assume a source type as specified by the option.

### Syntax



### Defaults

`-qsourcetype=default`

## Parameters

### assembler

All source files following the option are compiled as if they are assembler language source files.

### assembler-with-cpp

All source files following the option are compiled as if they are assembler language source files that need preprocessing.

- c** All source files following the option are compiled as if they are C language source files.

### **c++**

All source files following the option are compiled as if they are C++ language source files. This suboption is equivalent to the `++` option.

### default

The programming language of a source file is implied by its file name suffix.

## Usage

If you do not use this option, files must have a suffix of `.c` to be compiled as C files, and `.C` (uppercase C), `.cc`, `.cp`, `.cpp`, `.cxx`, or `.c++` to be compiled as C++ files.

This option applies whether the file system is case-sensitive or not. That is, even in a case-insensitive file system, where `file.c` and `file.C` refer to the same physical file, the compiler still recognizes the case difference of the file name argument on the command line and determines the source type accordingly.

Note that the option only affects files that are specified on the command line *following* the option, but not those that precede the option. Therefore, in the following example:

```
xlc goodbye.C -qsource=c hello.C
```

`hello.C` is compiled as a C source file, but `goodbye.C` is compiled as a C++ file.

The `-qsource` option should not be used together with the `++` option.

## Predefined macros

None.

## Examples

To treat the source file `hello.C` as being a C language source file, enter:

```
xlc -qsource=c hello.C
```

## Related information

- “`++` (plus sign) (C++ only)” on page 95

## **-qspectateabsolutes**

### Category

Optimization and tuning



## Pragma equivalent

None.

## Purpose

Works with the `-qtocmerge -bl:file` for non-IPA links and with the `-bl:file` for IPA links to disable speculation at absolute addresses.

The `bl:file` is necessary for the compiler to know which addresses are absolutes.

## Syntax

►► `-q` `speculateabsolutes` `nospeculateabsolutes` ►►

## Defaults

`-qspeculateabsolutes`

## Predefined macros

None.

## Related information

- “`-qtocmerge`” on page 355

## **-qspill**

### Category

Compiler customization

### Pragma equivalent

`#pragma options [no]spill`

### 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.

### Syntax

►► `-q` `spill` `=` `size` ►►

### Defaults

`-qspill=512`

### Parameters

*size*

An integer representing the number of bytes for the register allocation spill area.

## Usage

If your program is very complex, or if there are too many computations to hold in registers at one time and your program needs temporary storage, you might need to increase this area. Do not enlarge the spill area unless the compiler issues a message requesting a larger spill area. In case of a conflict, the largest spill area specified is used.

## Predefined macros

None.

## Examples

If you received a warning message when compiling `myprogram.c` and want to compile it specifying a spill area of 900 entries, enter:

```
xlc myprogram.c -qspill=900
```

## -qsrcmsg (C only)

### Category

Listings, messages, and compiler information

### Pragma equivalent

```
#pragma options [no]srcmsg
```

### Purpose

Adds the corresponding source code lines to diagnostic messages generated by the compiler.

When `nosrcmsg` is in effect, the error message simply shows the file, line and column where the error occurred. When `srcmsg` is in effect, the compiler reconstructs the source line or partial source line to which the diagnostic message refers and displays it before the diagnostic message. A pointer to the column position of the error may also be displayed.

### Syntax

```
►► -q nosrcmsg srcmsg ◄◄
```

### Defaults

`-qnosrcmsg`

### Usage

When `srcmsg` is in effect, the reconstructed source line represents the line as it appears after macro expansion. At times, the line may be only partially reconstructed. The characters `"..."` at the start or end of the displayed line indicate that some of the source line has not been displayed.

Use `-qnosrcmsg` to display concise messages that can be parsed.

## Predefined macros

None.

## Examples

To compile `myprogram.c` so that the source line is displayed along with the diagnostic message when an error occurs, enter:

```
xlc myprogram.c -qsrcmsg
```

## **-qstackprotect**

### Category

“Object code control” on page 80


### Pragma equivalent

None.

### Purpose

Provides protection against malicious code or programming errors that overwrite or corrupt the stack.

### Syntax



### 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.

**Note:** When both **all** and **size** are used, the last option wins.

### 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. The default option is `-qnostackprotect`.

To generate code to protect all procedures with vulnerable objects:

```
xlc myprogram.c -qstackprotect=all
```

To generate code to protect procedures with objects of certain bytes:

```
xlc myprogram.c -qstackprotect=size=8
```

**Note:**

- This option cannot be used with `#pragma` options.
- Because of the dependency on `libc.a` in AIX, this option requires the following AIX levels:
  - AIX 5.3/TL11 and up.
  - AIX 6.1/TL4 and up.
- If the link step fails with a message that indicates `__ssp_canary_word` is undefined, you have probably used an unsupported level of AIX.

**Predefined macros**

None.

**Related information**

- “-qinfo” on page 187

**-qstaticinline (C++ only)**

**Category**

Language element control

**Pragma equivalent**

None.

**Purpose**

Controls whether inline functions are treated as having `static` or `extern` linkage.

When `-qnostaticinline` is in effect, the compiler treats inline functions as `extern`: only one function body is generated for a function marked with the `inline` function specifier, regardless of how many definitions of the same function appear in different source files. When `-qstaticinline` is in effect, the compiler treats inline functions as having `static` linkage: a separate function body is generated for each definition in a different source file of the same function marked with the `inline` function specifier.

**Syntax**



**Defaults**

`-qnostaticinline`

## Usage

When `-qnostaticinline` is in effect, any redundant functions definitions for which no bodies are generated are discarded by default; you can use the `-qkeepinlines` option to change this behavior.

## Predefined macros

None.

## Examples

Using the `-qstaticinline` option causes function `f` in the following declaration to be treated as static, even though it is not explicitly declared as such. A separate function body is created for each definition of the function. Note that this can lead to a substantial increase in code size.

```
inline void f() { /*...*/};
```

## Related information

- "Linkage of inline functions" in the *XL C/C++ Language Reference*
- "`-qkeepinlines` (C++ only)" on page 209

## **-qstatsym**

### Category

Object code control

### Pragma equivalent

None.

### Purpose

Adds user-defined, nonexternal names that have a persistent storage class, such as initialized and uninitialized static variables, to the symbol table of the object file.

### Syntax

►► `-q` nostatsym  
statsym ◄◄

### Defaults

`-qnostatsym`: Static variables are not added to the symbol table. However, static functions are added to the symbol table.

### Predefined macros

None.

### Examples

To compile `myprogram.c` so that static symbols are added to the symbol table, enter:  
`xlc myprogram.c -qstatsym`

## -qstdinc

### Category

Input control



### Pragma equivalent

#pragma options [no]stdinc

### Purpose

Specifies whether the standard include directories are included in the search paths for system and user header files.

When **-qstdinc** is in effect, the compiler searches the following directories for header files:

-  The directory specified in the configuration file for the XL C header files (this is normally /usr/vacpp/include/) or by the **-qc\_stdinc** option
-  The directory specified in the configuration file for the XL C and C++ header files (this is normally /usr/vacpp/include/) or by the **-qcpp\_stdinc** option
- The directory specified in the configuration file for the system header files (this is normally /usr/include/), or by the **-qc\_stdinc** and **-qcpp\_stdinc** options or by the **-qgcc\_c\_stdinc** and **-qgcc\_cpp\_stdinc** options

When **-qnostdinc** is in effect, these directories are excluded from the search paths. The only directories to be searched are:

- directories in which source files containing `#include "filename"` directives are located
- directories specified by the **-I** option
- directories specified by the **-qinclude** option

### Syntax



►► -q stdinc  
nostdinc ◀◀

### Defaults

-qstdinc

### Usage

The search order of header files is described in “Directory search sequence for include files” on page 13.

This option only affects search paths for header files included with a relative name; if a full (absolute) path name is specified, this option has no effect on that path name.

The last valid pragma directive remains in effect until replaced by a subsequent pragma.

## Predefined macros

None.

## Examples

To compile `myprogram.c` so that *only* the directory `/tmp/myfiles` (in addition to the directory containing `myprogram.c`) is searched for the file included with the `#include "myinc.h"` directive, enter:

```
xlc myprogram.c -qnostdinc -I/tmp/myfiles
```

## Related information

- “`-qc_stdinc` (C only)” on page 130
- “`-qcpp_stdinc` (C++ only)” on page 131
- “`-I`” on page 181
- “Directory search sequence for include files” on page 13

## -qstrict

### Category

Optimization and tuning

### Pragma equivalent

```
#pragma options [no]strict
```

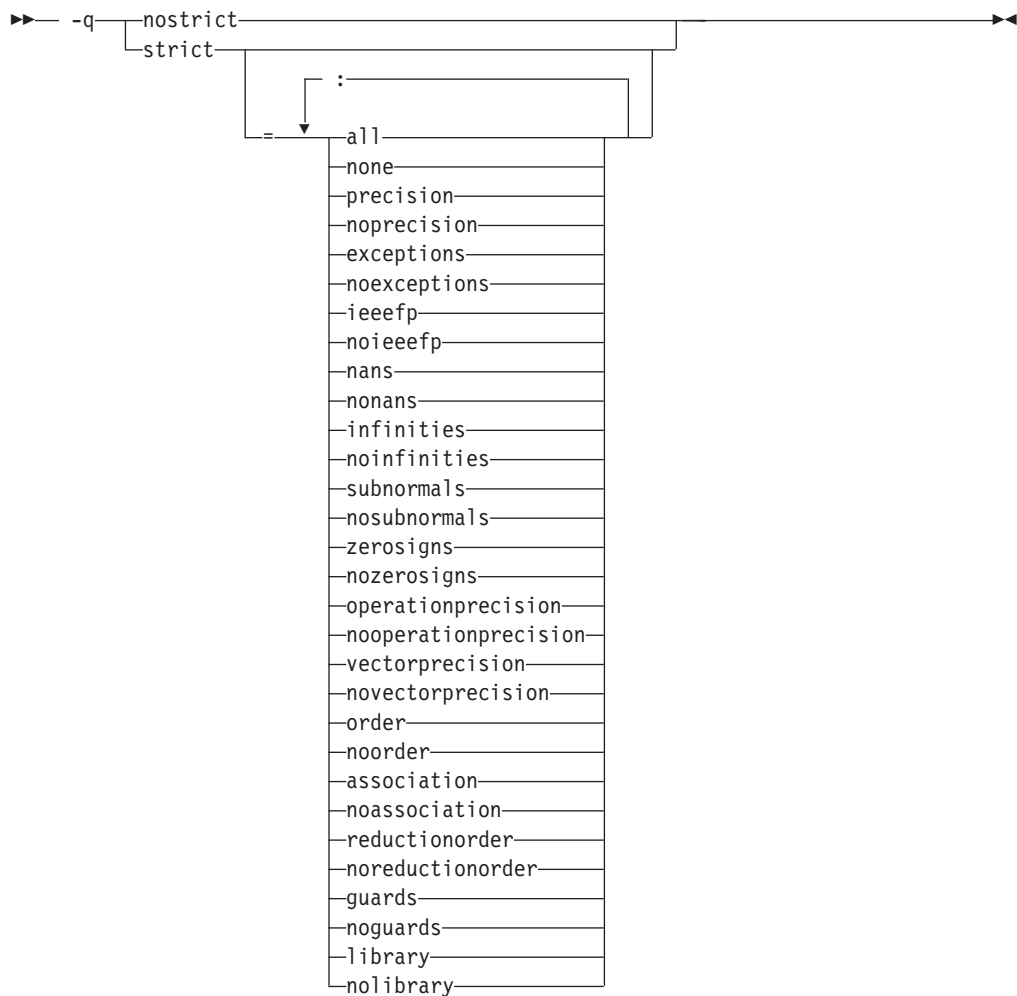
```
#pragma option_override (function_name, "opt (suboption_list)")
```

### Purpose

Ensures that optimizations done by default at optimization levels **-O3** and higher, and, optionally at **-O2**, do not alter the semantics of a program.

This option is intended for situations where the changes in program execution in optimized programs produce different results from unoptimized programs.

### Syntax



## Defaults

- Always **-qstrict** or **-qstrict=all** 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 **-O3** or a 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 **ieeefp**, **order**, **library**, **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**, and **library** suboptions. **noexceptions** enables these transformations.

**ieee754 | noieee754**

**ieee754** disables transformations that affect IEEE floating-point compliance, including those controlled by the **nans**, **infinities**, **subnormals**, **zerosigns**, **vectorprecision**, and **operationprecision** suboptions. **noieee754** 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**, out of loops, or

past function calls that might end the program or throw an exception) 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.

## 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 reassociated. 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 **fltint** and **rsqrt** suboptions of the **-qfloat** option are turned on. You can turn them off again by also using the **-qstrict** option or the **nofltint** and **norsqrt** suboptions of **-qfloat**. With lower-level or no optimization specified, these suboptions are turned off by default.

Specifying various **-qstrict[=suboptions]** or **-qnostrict** combinations sets the following suboptions:

- **-qstrict** or **-qstrict=all** sets **-qfloat=nofltint:norsqrt:rngchk**. **-qnostrict** or **-qstrict=none** sets **-qfloat=fltint:rsqrt:norngchk**.
- **-qstrict=operationprecision** or **-qstrict=exceptions** sets **-qfloat=nofltint**. Specifying both **-qstrict=nooperationprecision** and **-qstrict=noexceptions** sets **-qfloat=fltint**.
- **-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 154.

To override any of these settings, specify the appropriate **-qfloat** suboptions after the **-qstrict** option on the command line.

## Predefined macros

None.

## Examples

To compile `myprogram.c` so that the aggressive optimizations of **-O3** are turned off, range checking is turned off (**-qfloat=fltint**), and division by the result of a square root is replaced by multiplying by the reciprocal (**-qfloat=rsqrt**), enter:

```
xlc myprogram.c -O3 -qstrict -qfloat=fltint:rsqrt
```

To enable all transformations except those affecting precision, specify:

```
xlc myprogram.c -qstrict=none:precision
```

To disable all transformations except those involving NaNs and infinities, specify:

```
xlc myprogram.c -qstrict=all:nonans:noinfinities
```

## Related information

- “-qsimd” on page 316
- “-qfloat” on page 154
- “-qhot” on page 178
- “-O, -qoptimize” on page 269

## **-qstrict\_induction**

### Category

Optimization and tuning

### Pragma equivalent

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

```
➤ -qstrict_inductionnostrict_induction ➤
```

### Defaults

- **-qstrict\_induction**
- **-qnostrict\_induction** when **-O2** or higher optimization level is in effect

## Usage

When using **-O2** or higher optimization, you can specify **-qstrict\_induction** to prevent optimizations that change the result of a program if truncation or sign extension of a loop induction variable should occur as a result of variable overflow or wrap-around. However, use of **-qstrict\_induction** is generally not recommended because it can cause considerable performance degradation.

## Predefined macros

None.

## Related information

- “-O, -qoptimize” on page 269

## -qsuppress

### Category

Listings, messages, and compiler information

### Pragma equivalent


“#pragma report (C++ only)” on page 430

### Purpose


Prevents specific informational or warning messages from being displayed or added to the listing file, if one is generated.

### Syntax

**-qsuppress syntax — C**



**-qsuppress syntax — C++**



### Defaults

**-qnosuppress:** All informational and warning messages are reported, unless set otherwise with the **-qflag** option.

## Parameters

*message\_identifier*

Represents a message identifier. The message identifier must be in the following format:

*15dd-number*

where:

**15** Is the compiler product identifier.

*dd* Is the two-digit code representing the compiler component that produces the message. See “Compiler message format” on page 18 for descriptions of these codes.

*number*

Is the message number.

## Usage

You can only suppress information (I) and warning (W) messages. You cannot suppress other types of messages, such as (S) and (U) level messages. Note that informational and warning messages that supply additional information to a severe error cannot be disabled by this option.

To suppress all informational and warning messages, you can use the **-w** option.

To suppress IPA messages, enter **-qsuppress** before **-qipa** on the command line.

The **-qhaltonmsg** compiler option has precedence over **-qsuppress**. If both **-qhaltonmsg** and **-qsuppress** are specified, messages that are suppressed by **-qsuppress** are also printed.

**C** The **-qnosuppress** compiler option cancels previous settings of **-qsuppress**. **C**

**C++** When you specify **-qnosuppress** with specific message identifiers, the previous **-qsuppress** instances with the same message identifiers lose effect. When you specify **-qnosuppress** without specific message identifiers, all previous **-qsuppress** instances lose effect.

If you specify two or three of the following options, the last option specified has precedence:

**-qsuppress=message\_identifier**  
**-qnosuppress=message\_identifier**  
**-qnosuppress**

**C++**

## Predefined macros

None.

## Examples

If your program normally results in the following output:

```
"myprogram.c", line 1.1:1506-224 (I) Incorrect #pragma ignored
```

you can suppress the message by compiling with:  
`xlc myprogram.c -qsuppress=1506-224`

### Related information

- “-qflag” on page 153
- “-qhaltonmsg” on page 175

## -qsymtab (C only)

### Category

Error checking and debugging

### Pragma equivalent

None.

### Purpose

Determines the information that appears in the symbol table.

### Syntax

►► -q-symtab=unref  
                  └static┘

### Defaults

Static variables and unreferenced typedef, structure, union, and enumeration declarations are not included in the symbol table of the object file.

### Parameters

#### unref

When used with the `-g` option, specifies that debugging information is included for unreferenced typedef declarations, struct, union, and enum type definitions in the symbol table of the object file. This suboption is equivalent to `-qdbxextra`.

Using `-qsymtab=unref` may make your object and executable files larger.

#### static

Adds user-defined, nonexternal names that have a persistent storage class, such as initialized and uninitialized static variables, to the symbol table of the object file. This suboption is equivalent to `-qstatsym`.

### Predefined macros

None.

### Examples

To compile `myprogram.c` so that static symbols are added to the symbol table, enter:  
`xlc myprogram.c -qsymtab=static`

To compile `myprogram.c` so that unreferenced typedef, structure, union, and enumeration declarations are included in the symbol table for use with a debugger, enter:

```
xlc myprogram.c -g -qsymtab=unref
```

### Related information

- “-g” on page 169
- “-qdbxextra (C only)” on page 136
- “-qstatsym” on page 331

## -qsyntaxonly (C only)

### Category

Error checking and debugging

### Pragma equivalent

None.

### Purpose

Performs syntax checking without generating an object file.

### Syntax

►► -q—syntaxonly ◀◀

### Defaults

By default, source files are compiled and linked to generate an executable file.

### Usage

The `-P`, `-E`, and `-C` options override the `-qsyntaxonly` option, which in turn overrides the `-c` and `-o` options.

The `-qsyntaxonly` option suppresses only the generation of an object file. All other files, such as listing files, are still produced if their corresponding options are set.

### Predefined macros

None.

### Examples

To check the syntax of `myprogram.c` without generating an object file, enter:

```
xlc myprogram.c -qsyntaxonly
```

### Related information

- “-C, -C!” on page 120
- “-c” on page 119
- “-E” on page 142
- “-o” on page 268
- “-P” on page 280

## -t

### Category

Compiler customization

### Pragma equivalent

None.

### Purpose

Applies the prefix specified by the **-B** option to the designated components.

### Syntax





### Defaults

The default paths for all of the compiler executables are defined in the compiler configuration file.

### Parameters

The following table shows the correspondence between **-t** parameters and the component executable names:

Parameter	Description	Executable name
a	Assembler	as
b	Low-level optimizer	xlCcode
c	Compiler front end	xlcentry, xlCentry
 C	C++ compiler front end	xlCentry
d	Disassembler	dis
E	CreateExportList utility	CreateExportList
 f	c++filt utility	c++filt
I	High-level optimizer, compile step	ipa
L	High-level optimizer, link step	ipa
l	Linker	ld



Parameter	Description	Executable name
<code>&gt; C++ m</code>	Linkage helper	munch
<code>p</code>	Preprocessor	n/a

## Usage

This option is intended to be used together with the `-Bprefix` option. If `-B` is specified without the *prefix*, the default prefix is `/lib/o`. If `-B` is not specified at all, the prefix of the standard program names is `/lib/n`.

**Note:** If you use the `p` suboption, it can cause the source code to be preprocessed separately before compilation, which can change the way a program is compiled.

## Predefined macros

None.

## Examples

To compile `myprogram.c` so that the name `/u/newones/compilers/` is prefixed to the compiler and assembler program names, enter:

```
xlc myprogram.c -B/u/newones/compilers/ -tca
```

## Related information

- “-B” on page 116

## -qtabsize

### Category

Language element control

### Pragma equivalent

`#pragma options tabsize`

### Purpose

Sets the default tab length, for the purposes of reporting the column number in error messages.

### Syntax

▶▶ `-qtabsize=number` ▶▶

### Defaults

`-qtabsize=8`

### Parameters

*number*

The number of character spaces representing a tab in your source program.

## Usage

This option only affects error messages that specify the column number at which an error occurred.

## Predefined macros

None.

## Examples

To compile `myprogram.c` so the compiler considers tabs as having a width of one character, enter:

```
xlc myprogram.c -qtabsize=1
```

In this case, you can consider one character position (where each character and each tab equals one position, regardless of tab length) as being equivalent to one character column.

## -qthtable

### Category

Object code control

### Pragma equivalent

`#pragma options thtable`

### Purpose

Controls the amount of debugging traceback information that is included in the object files.

Many performance measurement tools require a full traceback table to properly analyze optimized code. If a traceback table is generated, it is placed in the text segment at the end of the object code, and contains information about each function, including the type of function, as well as stack frame and register information.

### Syntax

►► -q—thtable=full  
none  
small ◀◀

### Defaults

- `-qthtable=full`
- `-qthtable=small` when `-O` or higher optimization is in effect

### Parameters

#### full

A full traceback table is generated, complete with name and parameter information.

**none**

No traceback table is generated. The stack frame cannot be unwound so exception handling is disabled.

**small**

The traceback table generated has no name or parameter information, but otherwise has full traceback capability. This suboption reduces the size of the program code.

**Usage**

This option applies only to 64-bit compilations, and is ignored if specified for a 32-bit compilation.

The `#pragma` options directive must be specified before the first statement in the compilation unit.

**Predefined macros**

None.

**Related information**

- “-g” on page 169

**-qtempinc (C++ only)****Category**


Template control

**Pragma equivalent**

None.

**Purpose**

Generates separate template include files for template functions and class declarations, and places these files in a directory which can be optionally specified.

**Syntax****Defaults**

-qnotempinc

**Parameters**

*directory\_path*

The directory in which the generated template include files are to be placed.

## Usage

The `-qtempinc` and `-qtemplateregistry` compiler options are mutually exclusive. Specifying `-qtempinc` implies `-qnotemplateregistry`. Similarly, specifying `-qtemplateregistry` implies `-qnotempinc`. However, specifying `-qnotempinc` does not imply `-qtemplateregistry`.

Specifying either `-qtempinc` or `-qtemplateregistry` implies `-qtmplinst=auto`.

## Predefined macros

`__TEMPINC__` is predefined to 1 when `-qtempinc` is in effect; otherwise, it is not defined.

## Examples

To compile the file `myprogram.C` and place the generated include files for the template functions in the `/tmp/mytemplates` directory, enter:

```
xlc++ myprogram.C -qtempinc=/tmp/mytemplates
```

## Related information

- “`#pragma` implementation (C++ only)” on page 406
- “`-qtmplinst` (C++ only)” on page 353
- “`-qtemplateregistry` (C++ only)” on page 348
- “`-qtemplaterecompile` (C++ only)” on page 347
- “Using C++ templates” in the *XL C/C++ Optimization and Programming Guide*.

## `-qtemplatedepth` (C++ only)

### Category

Template control

### Pragma equivalent

None.

### Purpose

Specifies the maximum number of recursively instantiated template specializations that will be processed by the compiler.

### Syntax

```
►► -qtemplatedepth=number ◀◀
```

### Defaults

```
-qtemplatedepth=300
```

### Parameters

*number*

The maximum number of recursive template instantiations. The number can be a value between 1 and `INT_MAX`. If your code attempts to recursively

instantiate more templates than *number*, compilation halts and an error message is issued. If you specify an invalid value, the default value of 300 is used.

## Usage

Note that setting this option to a high value can potentially cause an out-of-memory error due to the complexity and amount of code generated.

## Predefined macros

None.

## Examples

To allow the following code in `myprogram.cpp` to be compiled successfully:

```
template <int n> void foo() {
    foo<n-1>();
}

template <> void foo<0>() {}

int main() {
    foo<400>();
}
```

Enter:

```
xlc++ myprogram.cpp -qtemplatedepth=400
```

## Related information

- "Using C++ templates" in the *XL C/C++ Optimization and Programming Guide*.

## -qtemplaterecompile (C++ only)

### Category

Template control

### Pragma equivalent

None.

### Purpose

Helps manage dependencies between compilation units that have been compiled using the `-qtemplateregistry` compiler option.

### Syntax

```
→ -q [templaterecompile | notemplaterecompile] →
```

### Defaults

`-qtemplaterecompile`

## Usage

If a source file that has been compiled previously is compiled again, the `-qtemplaterecompile` option consults the template registry to determine whether changes to this source file require the recompile of other compilation units. This can occur when the source file has changed in such a way that it no longer references a given instantiation and the corresponding object file previously contained the instantiation. If so, affected compilation units will be recompiled automatically.

The `-qtemplaterecompile` option requires that object files generated by the compiler remain in the subdirectory to which they were originally written. If your automated build process moves object files from their original subdirectory, use the `-qnotemplaterecompile` option whenever `-qtemplateregistry` is enabled.

## Predefined macros

None.

## Related information

- “-qtmplinst (C++ only)” on page 353
- “-qtempinc (C++ only)” on page 345
- “-qtemplateregistry (C++ only)”
- "Using C++ templates" in the *XL C/C++ Optimization and Programming Guide*.

## `-qtemplateregistry` (C++ only)

### Category

Template control

### Pragma equivalent

None.

### Purpose

Maintains records of all templates as they are encountered in the source and ensures that only one instantiation of each template is made.

The first time that the compiler encounters a reference to a template instantiation, that instantiation is generated and the related object code is placed in the current object file. Any further references to identical instantiations of the same template in different compilation units are recorded but the redundant instantiations are not generated. No special file organization is required to use the `-qtemplateregistry` option.

### Syntax

```
→→ -q [notemplateregistry] [templateregistry] [==file_path] →→
```

### Defaults

`-qnotemplateregistry`

## Parameters

*file\_path*

The path for the file that will contain the template instantiation information. If you do not specify a location the compiler saves all template registry information to the file `templateregistry` stored in the current working directory.

## Usage

Template registry files must not be shared between different programs. If there are two or more programs whose source is in the same directory, relying on the default template registry file stored in the current working directory may lead to incorrect results.

The `-qtempinc` and `-qtemplateregistry` compiler options are mutually exclusive. Specifying `-qtempinc` implies `-qnotemplateregistry`. Similarly, specifying `-qtemplateregistry` implies `-qnotempinc`. However, specifying `-qnotemplateregistry` does not imply `-qtempinc`.

Specifying either `-qtempinc` or `-qtemplateregistry` implies `-qtmplinst=auto`.

## Predefined macros

None.

## Examples

To compile the file `myprogram.C` and place the template registry information into the `/tmp/mytemplateregistry` file, enter:

```
xlc++ myprogram.C -qtemplateregistry=/tmp/mytemplateregistry
```

## Related information

- “`-qtmplinst` (C++ only)” on page 353
- “`-qtempinc` (C++ only)” on page 345
- “`-qtemplaterecompile` (C++ only)” on page 347
- "Using C++ templates" in the *XL C/C++ Optimization and Programming Guide*.

## **-qtempmax (C++ only)**

### Category

Template control

### Pragma equivalent

None.

### Purpose

Specifies the maximum number of template include files to be generated by the `-qtempinc` option for each header file.

### Syntax

►► -qtempmax=*number* ◀◀

## Defaults

-qtempmax=1

## Parameters

*number*

The maximum number of template include files. The number can be a value between 1 and 99 999.

## Usage

This option should be used when the size of files generated by the **-qtempinc** option become very large and take a significant amount of time to recompile when a new instance is created.

Instantiations are spread among the template include files.

## Predefined macros

None.

## Related information

- “-qtempinc (C++ only)” on page 345
- "Using C++ templates" in the *XL C/C++ Optimization and Programming Guide*.

## -qthreaded

### Category

Object code control

### Pragma equivalent

None.

### Purpose

Indicates to the compiler whether it must generate thread-safe code.

Always use this option when compiling or linking multithreaded applications. This option does not make code thread-safe, but it will ensure that code already thread-safe will remain so after compilation and linking. It also ensures that all optimizations are thread-safe.

### Syntax

►► -q 

nothreaded
threaded

 ◀◀



## Defaults

- **-qnothreaded** for all invocation commands except those with the **\_r** suffix
- **-qthreaded** for all **\_r**-suffixed invocation commands

## Usage

This option applies to both compile and linker operations.

To maintain thread safety, a file compiled with the **-qthreaded** option, whether explicitly by option selection or implicitly by choice of **\_r** compiler invocation mode, must also be linked with the **-qthreaded** option.

## Predefined macros

None.

## Related information

- “-qsmp” on page 320

## -qtimestamps

### Category

“Output control” on page 75

### Pragma equivalent

none.

### Purpose

Controls whether or not implicit time stamps are inserted into an object file.

### Syntax

►► -q timestamps  
notimestamps ◀◀

## 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.

## -qtls

### Category

Object code control

### Pragma equivalent

None.

### Purpose


Enables recognition of the `__thread` storage class specifier, which designates variables that are to be allocated threadlocal storage; and specifies the threadlocal storage model to be used.

When this option is in effect, any variables marked with the `__thread` storage class specifier are treated as local to each thread in a multi-threaded application. At run time, a copy of the variable is created for each thread that accesses it, and destroyed when the thread terminates. Like other high-level constructs that you can use to parallelize your applications, thread-local storage prevents race conditions to global data, without the need for low-level synchronization of threads.

Suboptions allow you to specify thread-local storage models, which provide better performance but are more restrictive in their applicability.

**Note:** This option is only supported on AIX for POWER version 5.3 with the 5300-05 Technology Level and higher.

### Syntax



### Defaults

`-qtls=unsupported`

### Parameters

#### **unsupported**

The `__thread` keyword is not recognized and thread-local storage is not enabled. This suboption is equivalent to `-qnotls`.

#### **global-dynamic**

This model is the most general, and can be used for all thread-local variables.

#### **initial-exec**

This model provides better performance than the `global-dynamic` or `local-dynamic` models, and can be used for thread-local variables defined in

dynamically-loaded modules, provided that those modules are loaded at the same time as the executable. That is, it can only be used when all thread-local variables are defined in modules that are not loaded through `dlopen`.

#### **local-dynamic**

This model provides better performance than the global-dynamic model, and can be used for thread-local variables defined in dynamically-loaded modules. However, it can only be used when all references to thread-local variables are contained in the same module in which the variables are defined.

#### **local-exec**

This model provides the best performance of all of the models, but can only be used when all thread-local variables are defined and referenced by the main executable.

#### **default**

Uses the appropriate model depending on the setting of the `-qpik` compiler option, which determines whether position-independent code is generated or not. When `-qpik` is in effect, this suboption results in `-qtls=global-dynamic`. When `-qnopik` is in effect, this suboption results in `-qtls=initial-exec` (`-qpik` is in effect by default)(`-qpik` is in effect by default in 64-bit mode, and cannot be disabled).

Specifying `-qtls` with no suboption is equivalent to `-qtls=default`.

### **Predefined macros**

None.

### **Related information**

- “`-qpik`” on page 291
- “The `__thread` storage class specifier” in the *XL C/C++ Language Reference*

## **-qtmplinst (C++ only)**

### **Category**

Template control


### **Pragma equivalent**

None.

### **Purpose**

Manages the implicit instantiation of templates.

### **Syntax**



### **Defaults**

`-qtmplinst=auto`

## Parameters

### **always**

Instructs the compiler to always perform implicit instantiation. If specified, **-qtempinc** and **-qtemplateregistry** compiler options are ignored.

### **auto**

Manages the implicit instantiations according to the **-qtempinc** and **-qtemplateregistry** options. If both **-qtempinc** and **-qtemplateregistry** are disabled, implicit instantiation will always be performed; otherwise if one of the options is enabled, the compiler manages the implicit instantiation according to that option.

### **noinline**

Instructs the compiler to not perform any implicit instantiations. If specified, the **-qtempinc** and **-qtemplateregistry** compiler options are ignored.

### **none**

Instructs the compiler to instantiate only inline functions. No other implicit instantiation is performed. If specified, **-qtempinc** and **-qtemplateregistry** compiler options are ignored.

## Usage

You can also use **#pragma do\_not\_instantiate** to suppress implicit instantiation of selected template classes. See “**#pragma do\_not\_instantiate (C++ only)**” on page 397.

## Predefined macros

None.

## Related information

- “**-qtemplateregistry (C++ only)**” on page 348
- “**-qtempinc (C++ only)**” on page 345
- “**#pragma do\_not\_instantiate (C++ only)**” on page 397
- “Explicit instantiation” in the *XL C/C++ Language Reference*

## **-qtmplparse (C++ only)**

### **Category**

Template control

### **Pragma equivalent**

None.

### **Purpose**

Controls whether parsing and semantic checking are applied to template definitions.

### **Syntax**

►► -q—tmplparse—= 

no
error
warn

 ◀◀

## Defaults

-qtmplparse=no

## Parameters

### error

Treats problems in template definitions as errors, even if the template is not instantiated.

**no** Do not parse template definitions. This reduces the number of errors issued in code written for previous versions of VisualAge C++ and predecessor products.

### warn

Parses template definitions and issues warning messages for semantic errors.

## Usage

This option applies to template definitions, not their instantiations. Regardless of the setting of this option, error messages are produced for problems that appear outside definitions. For example, messages are always produced for errors found during the parsing or semantic checking of constructs such as the following:

- return type of a function template
- parameter list of a function template

## Predefined macros

None.

## Related information

- "Using C++ templates" in the *XL C/C++ Optimization and Programming Guide*.

## -qtocdata

See "-qdataimported, -qdatalocal, -qtocdata" on page 134.

## -qtocmerge

### Category

Optimization and tuning

### Pragma equivalent

None.

### Purpose

Enables TOC merging to reduce TOC pointer loads and improves the scheduling of external loads.

### Syntax

►► -q notocmerge  
tocmerge \_\_\_\_\_ ►►

## Defaults

-qnotocmerge

## Usage

To use **-qtocmerge**, you must also use the **-bImportfile** linker option to specify the name of the file from which the compiler reads.

## Predefined macros

None.

## -qtrigraph

### Category

Language element control

### Pragma equivalent

None.

### Purpose

Enables the recognition of trigraph key combinations to represent characters not found on some keyboards.

### Syntax

►► -q trigraph  
notrigraph ◀◀

## Defaults

-qtrigraph

## Usage

A trigraph is a combination of three-key character combinations that let you produce a character that is not available on all keyboards. For details, see "Trigraph sequences" in the *XL C/C++ Language Reference*.

► **C++** To override the default **-qtrigraph** setting, you must specify **-qnotrigraph** *after* the **-qlanglvl** option on the command line.

## Predefined macros

None.

## Related information

- "Trigraph sequences" in the *XL C/C++ Language Reference*
- "-qdigraph" on page 138
- "-qlanglvl" on page 215

## **-qtune**

### **Category**

Optimization and tuning


### **Pragma equivalent**

None.

### **Purpose**

Tunes instruction selection, scheduling, and other architecture-dependent performance enhancements to run best on a specific hardware architecture.

### **Syntax**



### **Defaults**

**-qtune=balanced** when the default **-qarch** setting is in effect. Otherwise, the default depends on the effective **-qarch** setting. See Table 30 on page 358 for details.

### **Parameters**

#### **403**

Optimizations are tuned for the PowerPC 403 processor.

#### **604**

Optimizations are tuned for the PowerPC 604 processor.

#### **auto**

Optimizations are tuned for the platform on which the application is compiled.

#### **balanced**

Optimizations are tuned across a selected range of recent hardware.

#### **ppc970**

Optimizations are tuned for the PowerPC 970 processor.

#### **pwr3**

Optimizations are tuned for the POWER3 hardware platforms.

#### **pwr4**

Optimizations are tuned for the POWER4 hardware platforms.

**pwr5**

Optimizations are tuned for the POWER5 hardware platforms.

**pwr6**

Optimizations are tuned for the POWER6 hardware platforms.

**pwr7**

Optimizations are tuned for the POWER7 hardware platforms.

**rs64a**

Optimizations are tuned for the RS64I processor.

**rs64b**

Optimizations are tuned for the RS64II processor.

**rs64c**

Optimizations are tuned for the RS64III processor.

**Note:** As of the V9.0 release of the compiler, suboptions representing 601, 602, 603, POWER and POWER2 architectures are deprecated.

## Usage

If you want your program to run on more than one architecture, but to be tuned to a particular architecture, you can use a combination of the **-qarch** and **-qtune** options. These options are primarily of benefit for floating-point intensive programs.

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.

Acceptable combinations of **-qarch** and **-qtune** are shown in the following table.

*Table 30. Acceptable -qarch/-qtune combinations*

<b>-qarch option</b>	<b>Default -qtune setting</b>	<b>Available -qtune settings</b>
403	403	auto   403
604	604	auto   604
ppc	balanced	auto   604   rs64a   rs64b   rs64c   pwr3   pwr4   pwr5   pwr6   pwr7   ppc970   balanced
ppcgr	balanced	auto   604   rs64b   rs64c   pwr3   pwr4   pwr5   pwr6   pwr7   ppc970   balanced
ppc64	balanced	auto   rs64a   rs64b   rs64c   pwr3   pwr4   pwr5   pwr6   pwr7   ppc970   balanced
ppc64gr	balanced	auto   rs64b   rs64c   pwr3   pwr4   pwr5   pwr6   pwr7   ppc970   balanced
ppc64grsq	balanced	auto   rs64b   rs64c   pwr3   pwr4   pwr5   pwr6   pwr7   ppc970   balanced
ppc64v	ppc970	auto   ppc970   pwr6   balanced
ppc970	ppc970	auto   ppc970   balanced



Table 30. Acceptable **-qarch/-qtune** combinations (continued)

<b>-qarch option</b>	<b>Default -qtune setting</b>	<b>Available -qtune settings</b>
pwr3	pwr3	auto   pwr3   pwr4   pwr5   pwr7   ppc970   balanced
pwr4	pwr4	auto   pwr4   pwr5   pwr7   ppc970   balanced
pwr5	pwr5	auto   pwr5   pwr7   balanced
pwr5x	pwr5	auto   pwr5   pwr7   balanced
pwr6	pwr6	auto   pwr6   pwr7   balanced
pwr6e	pwr6	auto   pwr6   balanced
pwr7	pwr7	auto   pwr7   balanced
rs64a	rs64a	auto   rs64a
rs64b	rs64b	auto   rs64b
rs64c	rs64c	auto   rs64c

## Predefined macros

None.

## Examples

To specify that the executable program `testing` compiled from `myprogram.c` is to be optimized for a POWER7 hardware platform, enter:

```
xlc -o testing myprogram.c -qtune=pwr7
```

## Related information

- “-qarch” on page 106
- “-q32, -q64” on page 96
- “Specifying compiler options for architecture-specific compilation” on page 9
- “Optimizing your applications” in the *XL C/C++ Optimization and Programming Guide*

## **-qtwolink (C++ only)**

### Category

Linking

### Pragma equivalent

None.

### Purpose

Minimizes the number of static constructors included from libraries and object files.

When **-qnotwolink** is in effect, all static constructors in `.o` files and object files are invoked. This generates larger executable files, but ensures that placing a `.o` file in a library does not change the behavior of a program.

Normally, the compiler links in all static constructors defined anywhere in the object (.o) files and library (.a) files. The **-qtwolink** option makes link time longer, but linking is compatible with older versions of C or C++ compilers.

## Syntax

►► -q notwolink  
twolink ►►

## Defaults

-qnotwolink

## Usage

Before using **-qtwolink**, make sure that any .o files placed in an archive do not change the behavior of the program.

## Predefined macros

None.

## Examples

Given the include file foo.h:

```
#include <stdio.h>
struct foo {
    foo() {printf ("in foo\n");}
    ~foo() {printf ("in ~foo\n");}
};
```

and the C++ program t.C:

```
#include "foo.h"
foo bar;
```

and the program t2.C:

```
#include "foo.h"
main() { }
```

Compile t.C and t2.C in two steps, first invoking the compiler to produce object files:

```
xlc++ -c t.C t2.C
```

and then link them to produce the executable file a.out:

```
xlc++ t.o t2.o
```

Invoking a.out produces:

```
in foo
in ~foo
```

If you use the AIX **ar** command with the t.o file to produce an archive file t.a:

```
ar rv t.a t.o
```

and then use the default compiler command:

```
xlc++ t2.o t.a
```

the output from the executable file is the same as above:

```
in foo
in ~foo
```

However, if you use the **-qtwolink** option:

```
xlc++ -qtwolink t2.o t.a
```

there is no output from the executable file `a.out` because the static constructor `foo()` in `t.C` is not found.

## **-U**

### **Category**

Language element control

### **Pragma equivalent**

None.

### **Purpose**

Undefines a macro defined by the compiler or by the **-D** compiler option.

### **Syntax**

►► `-U` *name* ◀◀

### **Defaults**

Many macros are predefined by the compiler; see Chapter 6, “Compiler predefined macros,” on page 463 for those that can be undefined (that is, are not *protected*). The compiler configuration file also uses the **-D** option to predefine several macro names for specific invocation commands; for details, see the configuration file for your system.

### **Parameters**

*name*

The macro you want to undefine.

### **Usage**

The **-U** option is *not* equivalent to the `#undef` preprocessor directive. It *cannot* undefine names defined in the source by the `#define` preprocessor directive. It can only undefine names defined by the compiler or by the **-D** option.

The **-U***name* option has a higher precedence than the **-D***name* option.

### **Predefined macros**

None.

## Examples

Assume that your operating system defines the name `__unix`, but you do not want your compilation to enter code segments conditional on that name being defined, compile `myprogram.c` so that the definition of the name `__unix` is nullified by entering:

```
xlc myprogram.c -U__unix
```

## Related information

- “-D” on page 132

## -qunique

### Category

Object code control

### Pragma equivalent

None.

### Purpose

Generates unique names for static constructor/destructor file compilation units.

### Syntax

```
►► -q nunique unique ◄◄
```

### Defaults

`-qnunique`

### Usage

Unique names are generated with **-qunique** by encoding random numbers into the name of the static constructor and destructor functions. Default behavior is encoding the absolute path name of the source file in the constructor and destructor functions. If the absolute path name will be identical for multiple compilations (for example, if a **make** script is used), the **-qunique** option is necessary.

If you use **-qunique**, you must always link with all `.o` and `.a` files. Do not include an executable file on the link step.

### Predefined macros

None.

### Examples

Suppose you want to compile several files using the same path name, ensuring that static construction works correctly. A makefile may generate the following steps:

```

sqlpreprocess file1.sql > t.C
xlc++ -qunique t.C -o file1.o
rm -f t.C
sqlpreprocess file2.sql > t.C
xlc++ -qunique t.C -o file2.o
rm -f t.C
xlc++ file1.o file2.o

```

Following is a sample makefile for the above example:

```

# rule to get from file.sql to file.o
.SUFFIXES:      .sql
.sql.o:
    sqlpreprocess $< > t.C
    $(CCC) t.C -c $(CCFLAGS) -o $@
    rm -f t.C

```

## Related information

- “#pragma fini (C only)” on page 401
- “#pragma init (C only)” on page 407

## -qunroll

### Category

Optimization and tuning

### Pragma equivalent

#pragma options [no]unroll, #pragma unroll


### Purpose

Controls loop unrolling, for improved performance.


When **unroll** is in effect, the optimizer determines and applies the best unrolling factor for each loop; in some cases, the loop control may be modified to avoid unnecessary branching. The compiler remains the final arbiter of whether the loop is actually unrolled. You can use the **#pragma unroll** directive to gain more control over unrolling.

### Syntax

#### Option syntax



#### Pragma syntax



## Defaults

`-qunroll=auto`

## Parameters

### **auto (option only)**

Instructs the compiler to perform basic loop unrolling.

### **yes (option only)**

Instructs the compiler to search for more opportunities for loop unrolling than that performed with **auto**. In general, this suboption has more chances to increase compile time or program size than **auto** processing, but it may also improve your application's performance.

### **no (option only)**

Instructs the compiler to not unroll loops.

### **number (pragma only)**

Forces *number* - 1 replications of the designated loop body or full unrolling of the loop, whichever occurs first. The value of *number* is unbounded and must be a positive integer. Specifying **#pragma unroll(1)** effectively disables loop unrolling, and is equivalent to specifying **#pragma nounroll**. If *number* is not specified and if **-qhot**, **-qsmp**, or **-O4** or higher is specified, the optimizer determines an appropriate unrolling factor for each nested loop.

Specifying **-qunroll** without any suboptions is equivalent to **-qunroll=yes**.

**-qnounroll** is equivalent to **-qunroll=no**.

## Usage

The pragma overrides the **-q[no]unroll** compiler option setting for a designated loop. However, even if **#pragma unroll** is specified for a given loop, the compiler remains the final arbiter of whether the loop is actually unrolled.

Only one pragma may be specified on a loop. The pragma must appear immediately before the loop or the **#pragma block\_loop** directive to have effect.

The pragma affects only the loop that follows it. An inner nested loop requires a **#pragma unroll** directive to precede it if the desired loop unrolling strategy is different from that of the prevailing **-q[no]unroll** option.

The **#pragma unroll** and **#pragma nounroll** directives can only be used on for loops or **#pragma block\_loop** directives. They cannot be applied to do while and while loops.

The loop structure must meet the following conditions:

- There must be only one loop counter variable, one increment point for that variable, and one termination variable. These cannot be altered at any point in the loop nest.
- Loops cannot have multiple entry and exit points. The loop termination must be the only means to exit the loop.
- Dependencies in the loop must not be "backwards-looking". For example, a statement such as `A[i][j] = A[i - 1][j + 1] + 4` must not appear within the loop.

## Predefined macros

None.

## Examples

In the following example, the `#pragma unroll(3)` directive on the first for loop requires the compiler to replicate the body of the loop three times. The `#pragma unroll` on the second for loop allows the compiler to decide whether to perform unrolling.

```
#pragma unroll(3)
for( i=0; i < n; i++)
{
    a[i] = b[i] * c[i];
}

#pragma unroll
for( j=0; j < n; j++)
{
    a[j] = b[j] * c[j];
}
```

In this example, the first `#pragma unroll(3)` directive results in:

```
i=0;
if (i>n-2) goto remainder;
for (; i<n-2; i+=3) {
    a[i]=b[i] * c[i];
    a[i+1]=b[i+1] * c[i+1];
    a[i+2]=b[i+2] * c[i+2];
}
if (i<n) {
    remainder:
    for (; i<n; i++) {
        a[i]=b[i] * c[i];
    }
}
```

## Related information

- “`#pragma block_loop`” on page 390
- “`#pragma loopid`” on page 410
- “`#pragma stream_unroll`” on page 433
- “`#pragma unrollandfuse`” on page 435

## **-qunwind**

### Category

Optimization and tuning

### Pragma equivalent

None.

### Purpose

Specifies whether the call stack can be unwound by code looking through the saved registers on the stack.

Specifying **-qnounwind** asserts to the compiler that the stack will not be unwound, and can improve optimization of non-volatile register saves and restores.

## Syntax

```
►► — -q —┐┌─ unwind ─┐┌─  
         └─ nounwind ─┘└─►►
```

## Defaults

-qunwind

## Usage

The `setjmp` and `longjmp` families of library functions are safe to use with **-qnounwind**.

 C++ Specifying **-qnounwind** also implies **-qnoeh**.

## Predefined macros

None.

## Related information

- “-qeh (C++ only)” on page 144

## -qupconv (C only)

### Category

Portability and migration

### Pragma equivalent

#pragma options [no]upconv

### Purpose

Specifies whether the unsigned specification is preserved when integral promotions are performed.

When **noupconv** is in effect, any unsigned type smaller than an `int` is converted to `int` during integral promotions. When **upconv** is in effect, these types are converted to unsigned `int` during integral promotions. The promotion rule does not apply to types that are larger than `int`.

## Syntax

```
►► — -q —┐┌─ noupconv ─┐┌─  
         └─ upconv ─┘└─►►
```

## Defaults

- **-qnoupconv** for all language levels except **classic** or **extended**
- **-qupconv** when the **classic** or **extended** language levels are in effect



## Usage

Sign preservation is provided for compatibility with older dialects of C. The ANSI C standard requires value preservation as opposed to sign preservation.

## Predefined macros

None.

## Examples

To compile `myprogram.c` so that all unsigned types smaller than `int` are converted to unsigned `int`, enter:

```
xlc myprogram.c -qupconv
```

The following short listing demonstrates the effect of `-qupconv`:

```
#include <stdio.h>
int main(void) {
    unsigned char zero = 0;
    if (-1 < zero)
        printf("Value-preserving rules in effect\n");
    else
        printf("Unsignedness-preserving rules in effect\n");
    return 0;
}
```

## Related information

- "Usual arithmetic conversions" in the *XL C/C++ Language Reference*
- `"-qlanglvl"` on page 215

## -qutf

### Category

Language element control

### Pragma equivalent

None.

### Purpose

Enables recognition of UTF literal syntax.

### Syntax

►► -q noutf  
          utf ◀◀

### Defaults

- C `-qnoutf`
- C++ `-qutf` for all language levels except `-qlanglvl=strict98`. `-qnoutf` when `-qlanglvl=strict98` is in effect.

## Usage

The compiler uses `iconv` to convert the source file to Unicode. If the source file cannot be converted, the compiler will ignore the `-qutf` option and issue a warning.

## Predefined macros

None.

## Related information

- "UTF literals" in the *XL C/C++ Language Reference*

## **-v, -V**

### Category

Listings, messages, and compiler information

### Pragma equivalent

None.

### Purpose

Reports the progress of compilation, by naming the programs being invoked and the options being specified to each program.

When the `-v` option is in effect, information is displayed in a comma-separated list. When the `-V` option is in effect, information is displayed in a space-separated list.

### Syntax

►► [ -v ] \_\_\_\_\_ ◄◄  
[ -V ]

### Defaults

The compiler does not display the progress of the compilation.

### Usage

The `-v` and `-V` options are overridden by the `-#` option.

### Predefined macros

None.

### Examples

To compile `myprogram.c` so you can watch the progress of the compilation and see messages that describe the progress of the compilation, the programs being invoked, and the options being specified, enter:

```
xlc myprogram.c -v
```

## Related information

- “-# (pound sign)” on page 96

## -qvecnvolf

### Category

Portability and migration

### Pragma equivalent

None.

### Purpose

Specifies whether to use volatile or non-volatile vector registers.

Volatile vector registers are those whose value is not preserved across function calls or across save context, jump or switch context system library functions. When **-qvecnvolf** is in effect, the compiler uses both volatile and non-volatile vector registers. When **-qnovecnvolf** is in effect, the compiler uses only volatile vector registers.

This option is required for programs where there is risk of interaction between modules built with AIX libraries prior to AIX 5.3 with 5300-03 and vector register use. Restricting the compiler to use only volatile registers will make your vector programs safe but it potentially forces the compiler to store vector data to memory more often and therefore results in reducing performance.

### Syntax

►► -q novecnvolf  
          vecnvolf

### Defaults

-qnovecnvolf

### Usage

- To use the **-qvecnvolf** option, you need `bos.adt.include` version 5.3.0.30 or greater installed on your system.
- This option requires platforms that support vector instructions.
- The **-qnovecnvolf** option performs independently from **-qsimd=auto | noauto**, **-qaltivec | -qnoaltivec** and **pragma=nosimd**.
- On AIX 5.3 with 5300-03, by default only 20 volatile registers (vr0-vr19) are used, and 12 non-volatile vector registers (vr20 - vr31) are not used. You can use these registers only when **-qvecnvolf** is in effect.
- **-qvecnvolf** should be enabled only when no legacy code that saves and restores non-volatile registers is involved. Using **-qvecnvolf** and linking with legacy code, may result runtime failure.

### Predefined macros

None.

## Related information

- “-qaltivec” on page 105
- “-qsimd” on page 316

## -qversion

### Category

Listings, messages, and compiler information

### Pragma equivalent

None.

### Purpose

Displays the version and release of the compiler being invoked.

### Syntax



### Defaults

-qnoverversion

### Parameters

#### verbose

Additionally 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_level
```

where:

*component\_name*

Specifies an installed component, such as the low-level optimizer.

*component\_level*

Represents the level of the installed component.

For more details, see Example 2.

## Predefined macros

None.

## Examples

### Example 1:

IBM XL C/C++ for AIX, V12.1  
Version: 12.01.0000.0001

### Example 2:

IBM XL C/C++ for AIX, V12.1  
Version: 12.01.0000.0001  
Driver Version: 12.01(C/C++) Level: 060414  
C Front End Version: 12.01(C/C++) Level: 060419  
C++ Front End Version: 12.01(C/C++) Level: 060420  
High Level Optimizer Version: 12.01(C/C++) and 14.01(Fortran) Level: 060411  
Low Level Optimizer Version: 12.01(C/C++) and 14.01(Fortran) Level: 060418

## Related information

- “-qsaveopt” on page 311

**-W**

## Category

Listings, messages, and compiler information

## Pragma equivalent

None.

## Purpose

Suppresses informational, language-level and warning messages.

 This option is equivalent to specifying **-qflag=e : e**.  This option is equivalent to specifying **-qflag=s : s**.

## Syntax

▶ — -W —▶

## Defaults

All informational and warning messages are reported.

## Usage

Informational and warning messages that supply additional information to a severe error are not disabled by this option.

## Predefined macros

None.

## Examples

To compile `myprogram.c` so that no warning messages are displayed, enter:

```
xlc myprogram.c -w
```

The following example shows how informational messages that result from a severe error, in this case caused by problems with overload resolution in C++, are not disabled :

```
void func(int a){}
void func(int a, int b){}
int main(void)
{
    func(1,2,3);
    return 0;
}
```

The output is as follows:

```
"x.cpp", line 6.4: 1540-0218 (S) The call does not match any parameter list for
"func".
"x.cpp", line 1.6: 1540-1283 (I) "func(int)" is not a viable candidate.
"x.cpp", line 6.4: 1540-0215 (I) The wrong number of arguments have been
specified for "func(int)".
"x.cpp", line 2.6: 1540-1283 (I) "func(int, int)" is not a viable candidate.
"x.cpp", line 6.4: 1540-0215 (I) The wrong number of arguments have been specified
for "func(int, int)".
```

## Related information

- “-qflag” on page 153
- “-qsuppress” on page 338

## -W

### Category

Compiler customization


### Pragma equivalent

None.

### Purpose

Passes the listed options to a component that is executed during compilation.

### Syntax






## Parameters

*option*

Any option that is valid for the component to which it is being passed. Spaces must not appear before the *option*.

The following table shows the correspondence between **-W** parameters and the component executable names:

Parameter	Description	Executable name
a	Assembler	as
b	Low-level optimizer	xlCcode
c	Compiler front end	xlcentry, xlCentry
 C	C++ compiler front end	xlCentry
d	Disassembler	dis
E	CreateExportList utility	CreateExportList
 f	c++filt utility	c++filt
I	High-level optimizer, compile step	ipa
L	High-level optimizer, link step	ipa
l	Linker	ld
 m	Linkage helper	munch
p	Preprocessor	n/a

## 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**.

By default, static objects are initialized in the order of priority specified by **#pragma priority** or the “**-qpriority** (C++ only)” on page 297 option. You can use **-Wm** option to control the initialization order of the objects with the same priorities. Specifying **-Wm -c** instructs the compiler to initialize object files with the same priority in the order in which the files were given on the command line during linking into the library and the static objects within the files are initialized according to their declaration order. **-Wm -r** option, however, specifies that the object files with the same priority are to be initialized in the opposite order in which they were encountered during the linking phase.

**Note:** The **-Wm** option cannot be used with the **-qipa** option.

## Predefined macros

None.

## Examples

To compile the file `file.c` and pass the linker option **-berok** to the linker, enter the following command:

```
xlc -Wl,-berok file.c
```

To compile the file `uses_many_symbols.c` 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:

```
xlc -Wa,-x -Wl,-s produces_warnings.s uses_many_symbols.c
```

## Related information

- “Invoking the compiler” on page 1

## **-qwarn0x (C++0x)**

**Note:** C++0x is a new version of the C++ programming language standard. IBM continues to develop and implement the features of the new standard. The implementation of the language level is based on IBM's interpretation of the standard. Until IBM's implementation of all the features of the C++0x standard is complete, including the support of a new C++ standard library, the implementation may change from release to release. IBM makes no attempt to maintain compatibility, in source, binary, or listings and other compiler interfaces, with earlier releases of IBM's implementation of the new features of the C++0x standard and therefore they should not be relied on as a stable programming interface.

## Category

Error checking and debugging



## Pragma equivalent

None.

## Purpose

Controls whether to inform users with messages about differences in their programs caused by migration from the C++98 standard to the C++0x standard.

For example, when `-qlanglvl=noc99preprocessor` and `-qwarn0x` are specified, the C++0x preprocessor evaluates the controlling expressions in the `#if` and `#elif` conditional inclusion directives, and compare the evaluation results against that of the non-C++0x preprocessor. If they are different, the compiler issues the following warning message:

```
The preprocessor controlling expression evaluates differently between C++0x
and non-C++0x langlvls.
```

For another example, when you specify both the `-qlanglvl=noc99longlong` and `-qwarn0x` options, the compiler might display messages to indicate that the types of an integer literal are different between the non-C++0x and C++0x language levels. In 32-bit mode, when you use the integer literal 2147483648 to initialize a variable, the compiler issues the following message:

```
Integral constant "2147483648" has implied type unsigned long int under
the non-C++0x language levels. It has implied type long long int under C++0x.
```

The compiler issues a similar message for the literal 10000000000000000000 with the same option setting:

```
Integral constant "10000000000000000000" has implied type unsigned long
long or is not allowed with "-qlanglvl=extendedintegersafe" under C++0x.
Its impiled type is not unsigned long long under non-C++0x language levels.
```

The following C++0x keywords are not reserved in non-C++0x mode:

- `constexpr`
- `decltype`
- `static_assert`

For each occurrence of these keywords, the compiler issues a message if the corresponding C++0x features and keywords are disabled and if the `-qwarn0x` option is enabled. For example, when you specify both the `-qlanglvl=nostatic_assert` and `-qwarn0x` options, the compiler emits the following message for each `static_assert` token it encounters:

```
C++0x will reserve "static_assert" as a keyword whose C++0x feature can
be enabled by -qlanglvl=static_assert.
```

## Syntax

►► `-q` nowarn0x  
warn0x ►►

## Defaults

`-qnowarn0x`

## Usage

This option is in effect when `-qwarn0x` is set.

## Predefined macros

None.

## Related information

- “-qlanglvl” on page 215

## **-qwarn64**

### Category

Error checking and debugging

### Pragma equivalent

None.

### Purpose

Enables checking for possible data conversion problems between 32-bit and 64-bit compiler modes.

When `-qwarn64` is in effect, informational messages are displayed where data conversion may cause problems in 64-bit compilation mode, such as:

- Truncation due to explicit or implicit conversion of long types into int types
- Unexpected results due to explicit or implicit conversion of int types into long types
- Invalid memory references due to explicit conversion by cast operations of pointer types into int types
- Invalid memory references due to explicit conversion by cast operations of int types into pointer types
- Problems due to explicit or implicit conversion of constants into long types
- Problems due to explicit or implicit conversion by cast operations of constants into pointer types

### Syntax

```
►► — -q ————— warn64 ————— nowarn64 ————— ◄◄
```

### Defaults

`-qnowarn64`

### Usage

This option functions in either 32-bit or 64-bit compiler modes. In 32-bit mode, it functions as a preview aid to discover possible 32-bit to 64-bit migration problems.

## Predefined macros

None.

## Related information

- `-q32`, `-q64`
- “Compiler messages” on page 17

## **-qweakexp**

### Category

Object code control

### Pragma equivalent

None.

### Purpose

When used with the `-qmkshrobj` or `-G` option, includes or excludes global symbols marked as weak from the export list generated when you create a shared object.

### Syntax

►► `-q` weakexp  
noweakexp ◀◀

### Defaults

`-qweakexp`: weak symbols are exported.

### Usage

See “`-qweaksymbol`” on page 378 for a description of weak symbols.

Use the `-qweakexp` option with the `-qmkshrobj` or `-G` option. See the description of “`-qmkshrobj`” on page 260 or “`-G`” on page 172 for more information.

## Predefined macros

None.

## Examples

To compile `myprogram.c` into a shared object and prevent weak symbols from being exported, enter the following command:

```
xlc myprogram.c -qmkshrobj -qnoweakexp
```

## Related information

- “`-qweaksymbol`” on page 378
- “`#pragma weak`” on page 436
- “`-qmkshrobj`” on page 260
- “`-G`” on page 172

## **-qweaksymbol**

### **Category**

Object code control

### **Pragma equivalent**

None.

### **Purpose**

Enables the generation of weak symbols.

When the **-qweaksymbol** option is in effect, the compiler generates weak symbols for the following:

- Inline functions with external linkage
- Identifiers specified as weak with **#pragma weak** or `__attribute__((weak))`

### **Syntax**

►► -q weaksymbol  
noweaksymbol ◀◀

### **Defaults**

-qweaksymbol

### **Usage**

C++ When compiling C++ programs that contain extern inline functions, you can use **-qweaksymbol** to suppress linker messages warning of duplicate symbols.

### **Predefined macros**

None.

### **Related information**

- “#pragma weak” on page 436
- “-qweakexp” on page 377
- “The weak variable attribute” and “The weak function attribute” in the *XL C/C++ Language Reference*

## **-qxcall**

### **Category**

Object code control

### **Pragma equivalent**

None.

## Purpose

Generates code to treat static functions within a compilation unit as if they were external functions.

## Syntax

►► -q noxcall  
          xcall \_\_\_\_\_►►

## Defaults

-qnoxcall

## Usage

-qxcall generates slower code than -qnoxcall.

## Predefined macros

None.

## Examples

To compile myprogram.c so that all static functions are compiled as external functions, enter:

```
xlc myprogram.c -qxcall
```

## -qxref

### Category

Listings, messages, and compiler information

### Pragma equivalent

#pragma options [no]xref

### Purpose

Produces a compiler listing that includes the cross-reference component of the attribute and cross-reference section of the listing.

When **xref** is in effect, a listing file is generated with a .lst suffix for each source file named on the command line. For details of the contents of the listing file, see “Compiler listings” on page 20.

## Syntax

►► -q noxref  
          xref =full \_\_\_\_\_►►



**-y**

## Category

Floating-point and integer control

## Pragma equivalent

None.

## Purpose

Specifies the rounding mode for the compiler to use when evaluating constant floating-point expressions at compile time.

## Syntax



## Defaults

`-yn`, `-ydn`

## Parameters

The following suboptions are valid for binary floating-point types only:

- m** Round toward minus infinity.
- n** Round to the nearest representable number, ties to even.
- p** Round toward plus infinity.
- z** Round toward zero.

The following suboptions are valid for decimal floating-point types only:

- di** Round toward infinities (away from zero).
- dm** Round toward minus infinity.
- dn** Round to the nearest representable number, ties to even.
- dna** Round to the nearest representable number, ties away from zero.
- dnz** Round to the nearest representable number, ties toward zero.
- dp** Round toward plus infinity.
- dz** Round toward zero.

## Usage

If your program contains operations involving long doubles, the rounding mode must be set to **-yn** (round-to-nearest representable number, ties to even).

## Predefined macros

None.

## Examples

To compile `myprogram.c` so that constant floating-point expressions are rounded toward zero at compile time, enter:

```
xlc myprogram.c -yz -ydz
```

## -Z

### Category

Linking

### Pragma equivalent

None.

### Purpose

Specifies a prefix for the library search path to be used by the linker.

### Syntax

▶▶ — *-Z—string* —————▶▶

### Defaults

By default, the linker searches the `/usr/lib/` directory for library files.

### Parameters

*string*

Represents the prefix to be added to the directory search path for library files.

### Predefined macros

None.



---

## Chapter 5. Compiler pragmas reference

The following sections describe the pragmas available:

- “Pragma directive syntax”
- “Scope of pragma directives” on page 384
- “Summary of compiler pragmas by functional category” on page 384
- “Individual pragma descriptions” on page 389

---

### Pragma directive syntax

XL C/C++ supports three forms of pragma directives:

#### **#pragma options** *option\_name*

These pragmas use exactly the same syntax as their command-line option equivalent. The exact syntax and list of supported pragmas of this type are provided in “#pragma options” on page 420.

#### **#pragma** *name*

This form uses the following syntax:

Diagram illustrating the syntax for the #pragma directive: `#pragma name (—suboptions—)`. The diagram shows a horizontal line with arrows at both ends. Above the line, a box contains the text `name (—suboptions—)`. A vertical line descends from the box to the `name` part of the code, and another vertical line descends from the box to the `(—suboptions—)` part of the code.

The *name* is the pragma directive name, and the *suboptions* are any required or optional suboptions that can be specified for the pragma, where applicable.

#### **\_Pragma ("name")**

This form uses the following syntax:

Diagram illustrating the syntax for the \_Pragma directive: `_Pragma (—"name (—suboptions—)"—)`. The diagram shows a horizontal line with arrows at both ends. Above the line, a box contains the text `name (—suboptions—)`. A vertical line descends from the box to the `name` part of the code, and another vertical line descends from the box to the `(—suboptions—)` part of the code.

For example, the statement:

```
_Pragma ( "pack(1)" )
```

is equivalent to:

```
#pragma pack(1)
```

For all forms of pragma statements, you can specify more than one *name* and *suboptions* in a single **#pragma** statement.

The *name* on a pragma is subject to macro substitutions, unless otherwise stated. The compiler ignores unrecognized pragmas, issuing an informational message indicating this.

If you have any pragmas that are not common to both C and C++ in code that will be compiled by both compilers, you may add conditional compilation directives around the pragmas. (This is not strictly necessary since unrecognized pragmas are

ignored.) For example, **#pragma object\_model** is only recognized by the C++ compiler, so you may decide to add conditional compilation directives around the pragma.

```
#ifdef __cplusplus
#pragma object_model(pop)
#endif
```

---

## Scope of pragma directives

Many pragma directives can be specified at any point within the source code in a compilation unit; others must be specified before any other directives or source code statements. In the individual descriptions for each pragma, the "Usage" section describes any constraints on the pragma's placement.

In general, if you specify a pragma directive before any code in your source program, it applies to the entire compilation unit, including any header files that are included. For a directive that can appear anywhere in your source code, it applies from the point at which it is specified, until the end of the compilation unit.

You can further restrict the scope of a pragma's application by using complementary pairs of pragma directives around a selected section of code. For example, using **#pragma options source** and **#pragma options nosource** directives as follows requests that only the selected parts of your source code be included in your compiler listing:

```
#pragma options source

/* Source code between the source and nosource pragma
   options is included in the compiler listing          */

#pragma options nosource
```

Many pragmas provide "pop" or "reset" suboptions that allow you to enable and disable pragma settings in a stack-based fashion; examples of these are provided in the relevant pragma descriptions.

---

## Summary of compiler pragmas by functional category

The XL C/C++ pragmas available are grouped into the following categories:

- "Language element control" on page 385
- "C++ template pragmas" on page 385
- "Floating-point and integer control" on page 385
- "Error checking and debugging" on page 385
- "Listings, messages and compiler information" on page 386
- "Optimization and tuning" on page 386
- "Object code control" on page 387
- "Portability and migration" on page 388
- "Deprecated directives" on page 388

For descriptions of these categories, see "Summary of compiler options by functional category" on page 75.

## Language element control

Table 32. Language element control pragmas

Pragma	Description
#pragma langlvl (C only)	Determines whether source code and compiler options should be checked for conformance to a specific language standard, or subset or superset of a standard.
"#pragma mc_func" on page 413	Allows you to embed a short sequence of machine instructions "inline" within your program source code.
"#pragma options" on page 420	Specifies compiler options in your source program.

## C++ template pragmas

Table 33. C++ template pragmas

Pragma	Description
"#pragma define, #pragma instantiate (C++ only)" on page 394	Provides an alternative method for explicitly instantiating a template class.
"#pragma do_not_instantiate (C++ only)" on page 397	Prevents the specified template declaration from being instantiated.
"#pragma implementation (C++ only)" on page 406	For use with the <b>-qtempinc</b> compiler option, supplies the name of the file containing the template definitions corresponding to the template declarations contained in a header file.

## Floating-point and integer control

Table 34. Floating-point and integer control pragmas

Pragma	Description
#pragma chars	Determines whether all variables of type char are treated as either signed or unsigned.
#pragma enum	Specifies the amount of storage occupied by enumerations.

## Error checking and debugging

Table 35. Error checking and debugging pragmas

Pragma	Description
"#pragma ibm snapshot" on page 406	Specifies a location at which a breakpoint can be set and defines a list of variables that can be examined when program execution reaches that location.
#pragma info	Produces or suppresses groups of informational messages.
"#pragma operator_new (C++ only)" on page 419	Determines whether the new and new[] operators throw an exception if the requested memory cannot be allocated.

## Listings, messages and compiler information

Table 36. Listings, messages and compiler information pragmas

Pragma	Description
"#pragma report (C++ only)" on page 430	Controls the generation of diagnostic messages.

## Optimization and tuning

Table 37. Optimization and tuning pragmas

Pragma	Description
"#pragma block_loop" on page 390	Marks a block with a scope-unique identifier.
"#pragma STDC cx_limited_range" on page 432	Instructs the compiler that complex division and absolute value are only invoked with values such that intermediate calculation will not overflow or lose significance.
"#pragma disjoint" on page 395	Lists identifiers that are not aliased to each other within the scope of their use.
"#pragma execution_frequency" on page 398	Marks program source code that you expect will be either very frequently or very infrequently executed.
"#pragma expected_value" on page 400	Specifies the value that a parameter passed in a function call is most likely to take at run time. The compiler can use this information to perform certain optimizations, such as function cloning and inlining.
"#pragma ibm iterations" on page 403	Specifies the approximate average number of loop iterations for the chosen loop.
"#pragma ibm max_iterations" on page 404	Specifies the approximate maximum number of loop iterations for the chosen loop.
"#pragma ibm min_iterations" on page 405	Specifies the approximate minimum number of loop iterations for the chosen loop.
#pragma isolated_call	Specifies functions in the source file that have no side effects other than those implied by their parameters.
"#pragma leaves" on page 409	Informs the compiler that a named function never returns to the instruction following a call to that function.
"#pragma loopid" on page 410	Marks a block with a scope-unique identifier.
#pragma nosimd	When used with <b>-qsimd=auto</b> , disables the generation of SIMD instructions for the next loop.
#pragma novector	When used with <b>-qhot=vector</b> , disables auto-vectorization of the next loop.
"#pragma option_override" on page 422	Allows you to specify optimization options at the subprogram level that override optimization options given on the command line.

Table 37. Optimization and tuning pragmas (continued)

Pragma	Description
"#pragma reachable" on page 427	Informs the compiler that the point in the program after a named function can be the target of a branch from some unknown location.
"#pragma reg_killed_by" on page 428	Specifies registers that may be altered by functions specified by <b>#pragma mc_func</b> .
"#pragma simd_level" on page 431	Controls the compiler code generation of vector instructions for individual loops.
"#pragma stream_unroll" on page 433	When optimization is enabled, breaks a stream contained in a for loop into multiple streams.
#pragma unroll	Controls loop unrolling, for improved performance.
"#pragma unrollandfuse" on page 435	Instructs the compiler to attempt an unroll and fuse operation on nested for loops.

## Object code control

Table 38. Object code control pragmas



Pragma	Description
#pragma alloca (C only)	Provides an inline definition of system function <code>alloca</code> when it is called from source code that does not include the <code>alloca.h</code> header.
"#pragma comment" on page 393	Places a comment into the object module.
 "#pragma fini (C only)" on page 401	Specifies the order in which the runtime library calls a list of functions after <code>main()</code> completes or <code>exit()</code> is called.
"#pragma hashome (C++ only)" on page 401	Informs the compiler that the specified class has a home module that will be specified by <b>#pragma ishome</b> .
"#pragma ishome (C++ only)" on page 408	Informs the compiler that the specified class's home module is the current compilation unit.
 "#pragma init (C only)" on page 407	Specifies the order in which the runtime library calls a list of functions before <code>main()</code> is called.
"#pragma map" on page 411	Converts all references to an identifier to another, externally defined identifier.
"#pragma pack" on page 424	Sets the alignment of all aggregate members to a specified byte boundary.
#pragma priority (C++ only)	Specifies the priority level for the initialization of static objects.
"#pragma reg_killed_by" on page 428	Specifies registers that may be altered by functions specified by <b>#pragma mc_func</b> .

Table 38. Object code control pragmas (continued)

Pragma	Description
#pragma strings	Specifies the storage type for string literals.
"#pragma weak" on page 436	Prevents the linker from issuing error messages if it encounters a symbol multiply-defined during linking, or if it does not find a definition for a symbol.

## Portability and migration

Table 39. Portability and migration pragmas

Pragma	Description
#pragma align	Specifies the alignment of data objects in storage, which avoids performance problems with misaligned data.
#pragma namemangling (C++ only)	Chooses the name mangling scheme for external symbol names generated from C++ source code.
"#pragma namemanglingrule (C++ only)" on page 415	Provides fined-grained control over the name mangling scheme in effect for selected portions of source code, specifically with respect to the mangling of cv-qualifiers in function parameters.
#pragma object_model (C++ only)	Sets the object model to be used for structures, unions, and classes.
#pragma pass_by_value (C++ only)	Specifies how classes containing const or reference members are passed in function arguments.

## Deprecated directives

The SMP directives listed in the following table have been deprecated and might be removed in the future release. Use the corresponding OpenMP directives to obtain the same behavior.

Table 40. Deprecated SMP directives

SMP directive name	OpenMP directive/clause name
#pragma ibm critical	"#pragma omp critical" on page 457
#pragma ibm parallel_loop	The "#pragma omp parallel for" on page 454 pragma with the <b>schedule</b> clause.
#pragma ibm schedule	

The following examples show how to replace the deprecated SMP directives with their corresponding OpenMP ones.

For the **critical** pragma:

```
#pragma ibm critical(lck)
{
    ...
}
```

is replaced by

```
#pragma omp critical(lck)
{
...
}
```

For the **schedule** pragma:

```
#pragma ibm parallel_loop
#pragma ibm schedule(static, 5)
for (i=0; i<N; i++)
{
...
}
```

is replaced by

```
#pragma omp parallel for schedule(static, 5)
for (i=0; i<N; i++)
{
...
}
```

---

## Individual pragma descriptions

This section contains descriptions of individual pragmas available in XL C/C++.

For each pragma, the following information is given:

### Category

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

### Purpose

This section provides a brief description of the effect of the pragma, and why you might want to use it.

### Syntax

This section provides the syntax for the pragma. For convenience, the **#pragma name** form of the directive is used in each case. However, it is perfectly valid to use the alternate C99-style `_Pragma` operator syntax; see “Pragma directive syntax” on page 383 for details.

### Parameters

This section describes the suboptions that are available for the pragma, where applicable.

**Usage** This section describes any rules or usage considerations you should be aware of when using the pragma. These can include restrictions on the pragma's applicability, valid placement of the pragma, and so on.

### Examples

Where appropriate, examples of pragma directive use are provided in this section.

## **#pragma align**

See “-qalign” on page 100.

## **#pragma alloca (C only)**

See “-qalloca, -ma (C only)” on page 104.

## #pragma block\_loop

### Category

Optimization and tuning

### Purpose

Marks a block with a scope-unique identifier.

### Syntax

The diagram shows the syntax for the `#pragma block_loop` directive. It consists of the text `#pragma block_loop` followed by an opening parenthesis `(`. Inside the parentheses, there are two elements: `expression` and `name`, separated by a comma `,`. A line with an arrow points from the `name` to a closing parenthesis `)`. The entire construct is enclosed in a larger arrowed box.

### Parameters

*expression*

An integer expression representing the size of the iteration group.

*name*

An identifier that is unique within the scoping unit. If you do not specify a *name*, blocking occurs on the first for loop or loop following the **#pragma block\_loop** directive.

### Usage

For loop blocking to occur, a **#pragma block\_loop** directive must precede a for loop.

If you specify **#pragma unroll**, **#pragma unrollandfuse** or **#pragma stream\_unroll** for a blocking loop, the blocking loop is unrolled, unrolled and fused or stream unrolled respectively, if the blocking loop is actually created. Otherwise, this directive has no effect.

If you specify **#pragma unrollandfuse**, **#pragma unroll** or **#pragma stream\_unroll** directive for a blocked loop, the directive is applied to the blocked loop after the blocking loop is created. If the blocking loop is not created, this directive is applied to the loop intended for blocking, as if the corresponding **#pragma block\_loop** directive was not specified.

You must not specify **#pragma block\_loop** more than once, or combine the directive with **#pragma nounroll**, **#pragma unroll**, **#pragma nounrollandfuse**, **#pragma unrollandfuse**, or **#pragma stream\_unroll** directives for the same for loop. Also, you should not apply more than one **#pragma unroll** directive to a single block loop directive.

Processing of all **#pragma block\_loop** directives is always completed before performing any unrolling indicated by any of the unroll directives

### Examples

The following two examples show the use of **#pragma block\_loop** and **#pragma loop\_id** for loop tiling:



```

#pragma block_loop(50, mymainloop)
#pragma block_loop(20, myfirstloop, mysecondloop)
#pragma loopid(mymainloop)
    for (i=0; i < n; i++)
    {
#pragma loopid(myfirstloop)
        for (j=0; j < m; j++)
        {
#pragma loopid(mysecondloop)
            for (k=0; k < m; k++)
            {
                ...
            }
        }
    }

#pragma block_loop(50, mymainloop)
#pragma block_loop(20, myfirstloop, mysecondloop)
#pragma loopid(mymainloop)
    for (i=0; i < n; i++)
    {
#pragma loopid(myfirstloop)
        for (j=0; j < m; j++)
        {
#pragma loopid(mysecondloop)
            for (k=0; k < m; k++)
            {
                ...
            }
        }
    }
}

```

The following example shows the use **#pragma block\_loop** and **#pragma loop\_id** for loop interchange.

```

    for (i=0; i < n; i++)
    {
        for (j=0; j < n; j++)
        {
#pragma block_loop(1,myloop1)
            for (k=0; k < m; k++)
            {
#pragma loopid(myloop1)
                for (l=0; l < m; l++)
                {
                    ...
                }
            }
        }
    }
}

```

The following example shows the use of **#pragma block\_loop** and **#pragma loop\_id** for loop tiling for multi-level memory hierarchy:

```

#pragma block_loop(l3factor, first_level_blocking)
    for (i=0; i < n; i++)
    {
#pragma loopid(first_level_blocking)
#pragma block_loop(l2factor, inner_space)
        for (j=0; j < n; j++)
        {
#pragma loopid(inner_space)
            for (k=0; k < m; k++)
            {
                for (l=0; l < m; l++)
                {
                    ...
                }
            }
        }
    }
}

```

```

    }
  }
}

```

The following example uses `#pragma unrollandfuse` and `#pragma block_loop` to unroll and fuse a blocking loop.

```

#pragma unrollandfuse
#pragma block_loop(10)
  for (i = 0; i < N; ++i) {
  }

```

In this case, if the block loop directive is ignored, the unroll directives have no effect.

The following example shows the use of `#pragma unroll` and `#pragma block_loop` to unroll a blocked loop.

```

#pragma block_loop(10)
#pragma unroll(2)
  for (i = 0; i < N; ++i) {
  }

```

In this case, if the block loop directive is ignored, the unblocked loop is still subjected to unrolling. If blocking does happen, the unroll directive is applied to the blocked loop.

The following examples show invalid uses of the directive. The first example shows `#pragma block_loop` used on an undefined loop identifier:

```

#pragma block_loop(50, myloop)
  for (i=0; i < n; i++)
  {
  }

```

Referencing `myloop` is not allowed, since it is not in the nest and may not be defined.

In the following example, referencing `myloop` is not allowed, since it is not in the same loop nest:

```

  for (i=0; i < n; i++)
  {
    #pragma loopid(myLoop)
    for (j=0; j < i; j++)
    {
      ...
    }
  }
  #pragma block_loop(myLoop)
  for (i=0; i < n; i++)
  {
    ...
  }

```

The following examples are invalid since the unroll directives conflict with each other:

```

#pragma unrollandfuse(5)
#pragma unroll(2)
  #pragma block_loop(10)
  for (i = 0; i < N; ++i) {
  }

```

```
#pragma block_loop(10)
#pragma unroll(5)
#pragma unroll(10)
  for (i = 0; i < N; ++i) {
  }
```

### Related information

- “#pragma loopid” on page 410
- “-qunroll” on page 363
- “#pragma unrollandfuse” on page 435
- “#pragma stream\_unroll” on page 433

## #pragma chars

See “-qchars” on page 123.

## #pragma comment


### Category

Object code control

### Purpose

Places a comment into the object module.

### Syntax



### Parameters

#### compiler

Appends the name and version of the compiler at the end of the generated object module.

#### date

The date and time of the compilation are appended at the end of the generated object module.

#### timestamp

Appends the date and time of the last modification of the source at the end of the generated object module.

#### copyright

Places the text specified by the *token\_sequence*, if any, into the generated object module. The *token\_sequence* is included in the generated executable and loaded into memory when the program is run.

#### user

Places the text specified by the *token\_sequence*, if any, into the generated object module. The *token\_sequence* is included in the generated executable but is *not* loaded into memory when the program is run.

#### *token\_sequence*

The characters in this field, if specified, must be enclosed in double quotation

marks ("). If the string literal specified in the *token\_sequence* exceeds 32 767 bytes, an information message is emitted and the pragma is ignored.

## Usage

More than one **comment** directive can appear in a translation unit, and each type of **comment** directive can appear more than once, with the exception of **copyright**, which can appear only once.

You can display the object-file comments by using the operating system **strings** command.

## Examples

Assume we have the following program code: `tt.c`:

```
#pragma comment(date)
#pragma comment(compiler)
#pragma comment(timestamp)
#pragma comment(copyright,"My copyright")
int main() { return 0; }
```

Issuing the command:

```
xlc -c tt.c
strings -a tt.o
```

will cause the comment information embedded in `tt.o` to be displayed, along with any other strings that may be found in `tt.o`. For example, assuming the program code shown above:

```
@.text
.data
@.bss
.comment
Thu Sep 24 16:44:25 EDT 2012IBM XL C for AIX ---- Version 12.1.0.0
Thu Sep 24 16:44:09 EDT 2012
main
My copyright
.file
tt.c
.text
.data
.bss
.main
_$STATIC
_$STATIC
main
main
Thu Sep 24 16:44:25 2012
IBM XL C for AIX, Version 12.1.0.0 ---
```

## #pragma define, #pragma instantiate (C++ only)

### Category

Template control

### Purpose

Provides an alternative method for explicitly instantiating a template class.

## Syntax

►► #pragma define  
instantiate (—*template\_class\_name*—) ►►

## Parameters

*template\_class\_name*

The name of the template class to be instantiated.

## Usage

This pragma provides the equivalent functionality to C++ explicit instantiation definitions. It is provided for compatibility with earlier releases only. New applications should use C++ explicit instantiation definitions.

This pragma can appear anywhere an explicit instantiation definition can appear.

## Examples

The following directive:

```
#pragma define(Array<char>)
```

is equivalent to the following explicit instantiation:

```
template class Array<char>;
```

## Related information

- "Explicit instantiation" in the *XL C/C++ Language Reference*
- “#pragma do\_not\_instantiate (C++ only)” on page 397

## #pragma disjoint

### Category

Optimization and tuning


### Purpose

Lists identifiers that are not aliased to each other within the scope of their use.

By informing the compiler that none of the identifiers listed in the pragma shares the same physical storage, the pragma provides more opportunity for optimizations.

### Syntax

►► #pragma disjoint ►►



## Parameters

*variable\_name*

The name of a variable. It must not refer to any of the following:

- A member of a structure, class, or union
- A structure, union, or enumeration tag
- An enumeration constant
- A typedef name
- A label

## Usage

The **#pragma disjoint** directive asserts that none of the identifiers listed in the pragma share physical storage; if any the identifiers *do* actually share physical storage, the pragma may give incorrect results.

The pragma can appear anywhere in the source program that a declaration is allowed. An identifier in the directive must be visible at the point in the program where the pragma appears.

You must declare the identifiers before using them in the pragma. Your program must not dereference a pointer in the identifier list nor use it as a function argument before it appears in the directive.

This pragma can be disabled with the **-qignprag** compiler option.

## Examples

The following example shows the use of **#pragma disjoint**.

```
int a, b, *ptr_a, *ptr_b;

#pragma disjoint(*ptr_a, b) /* *ptr_a never points to b */
#pragma disjoint(*ptr_b, a) /* *ptr_b never points to a */
one_function()
{
    b = 6;
    *ptr_a = 7; /* Assignment will not change the value of b */

    another_function(b); /* Argument "b" has the value 6 */
}
```

External pointer `ptr_a` does not share storage with and never points to the external variable `b`. Consequently, assigning 7 to the object to which `ptr_a` points will not change the value of `b`. Likewise, external pointer `ptr_b` does not share storage with and never points to the external variable `a`. The compiler can assume that the argument to `another_function` has the value 6 and will not reload the variable from memory.

## #pragma do\_not\_instantiate (C++ only)

### Category

Template control

### Purpose

Prevents the specified template declaration from being instantiated.

You can use this pragma to suppress the implicit instantiation of a template for which a definition is supplied.

### Syntax

```
▶▶ #pragma do_not_instantiate template_class_name ▶▶
```

### Parameters

*template\_class\_name*

The name of the template class that should not be instantiated.

### Usage

If you are handling template instantiations manually (that is, `-qnotempinc` and `-qnotemplateregistry` are specified), and the specified template instantiation already exists in another compilation unit, using `#pragma do_not_instantiate` ensures that you do not get multiple symbol definitions during the link step.

▶ C++0x

**Note:** C++0x is a new version of the C++ programming language standard. IBM continues to develop and implement the features of the new standard. The implementation of the language level is based on IBM's interpretation of the standard. Until IBM's implementation of all the features of the C++0x standard is complete, including the support of a new C++ standard library, the implementation may change from release to release. IBM makes no attempt to maintain compatibility, in source, binary, or listings and other compiler interfaces, with earlier releases of IBM's implementation of the new features of the C++0x standard and therefore they should not be relied on as a stable programming interface.

`#pragma do_not_instantiate` on a class template specialization is treated as an explicit instantiation declaration of the template. This pragma provides a subset of the functionality of the explicit instantiation declarations feature, which is introduced by the C++0x standard. It is provided for compatibility purposes only and is not recommended. New applications should use explicit instantiation declarations instead.

C++0x ◀

You can also use the `-qtmplinst` option to suppress implicit instantiation of template declarations for multiple compilation units. See “`-qtmplinst (C++ only)`” on page 353.

## Examples

The following shows the usage of the pragma:

```
#pragma do_not_instantiate Stack < int >
```

## Related information

- “#pragma define, #pragma instantiate (C++ only)” on page 394
- “-qtmplinst (C++ only)” on page 353
- “Explicit instantiation” in the *XL C/C++ Language Reference*
- “-qtempinc (C++ only)” on page 345
- “-qtemplateregistry (C++ only)” on page 348

## #pragma enum

See “-qenum” on page 144.

## #pragma execution\_frequency

### Category

Optimization and tuning

### Purpose

Marks program source code that you expect will be either very frequently or very infrequently executed.

When optimization is enabled, the pragma is used as a hint to the optimizer.

### Syntax

```
▶▶ #pragma execution_frequency ( [very_low] ) ▶▶  
                                  └very_high┘
```

### Parameters

#### very\_low

Marks source code that you expect will be executed very infrequently.

#### very\_high

Marks source code that you expect will be executed very frequently.

### Usage

Use this pragma in conjunction with an optimization option; if optimization is not enabled, the pragma has no effect.

The pragma must be placed within block scope, and acts on the closest point of branching.

## Examples

In the following example, the pragma is used in an if statement block to mark code that is executed infrequently.

```
int *array = (int *) malloc(10000);  
  
if (array == NULL) {
```



```

    /* Block A */
    #pragma execution_frequency(very_low)
    error();
}

```

In the next example, the code block Block B is marked as infrequently executed and Block C is likely to be chosen during branching.

```

if (Foo > 0) {
    #pragma execution_frequency(very_low)
    /* Block B */
    doSomething();
} else {
    /* Block C */
    doAnotherThing();
}

```

In this example, the pragma is used in a switch statement block to mark code that is executed frequently.

```

while (counter > 0) {
    #pragma execution_frequency(very_high)
    doSomething();
} /* This loop is very likely to be executed. */

switch (a) {
    case 1:
        doOneThing();
        break;
    case 2:
        #pragma execution_frequency(very_high)
        doTwoThings();
        break;
    default:
        doNothing();
} /* The second case is frequently chosen. */

```

The following example shows how the pragma must be applied at block scope and affects the closest branching.

```

int a;
#pragma execution_frequency(very_low)
int b;

int foo(boolean boo) {
    #pragma execution_frequency(very_low)
    char c;

    if (boo) {
        /* Block A */
        doSomething();
        {
            /* Block C */
            doSomethingAgain();
            #pragma execution_frequency(very_low)
            doAnotherThing();
        }
    } else {
        /* Block B */
        doNothing();
    }

    return 0;
}

#pragma execution_frequency(very_low)

```

## #pragma expected\_value

### Category

Optimization and tuning

### Purpose

Specifies the value that a parameter passed in a function call is most likely to take at run time. The compiler can use this information to perform certain optimizations, such as function cloning and inlining.

### Syntax

```
▶▶ #pragma expected_value(—argument—, —value—) ▶▶
```

### Parameters

*argument*

The name of the parameter for which you want to provide the expected value. The parameter must be of a simple built-in integral, Boolean, character, or floating-point type.

*value*

A constant literal representing the value that you expect will most likely be taken by the parameter at run time. *value* can be an expression as long as it is a compile time constant expression.

### Usage

The directive must appear inside the body of a function definition, before the first statement (including declaration statements). It is not supported within nested functions.

If you specify an expected value of a type different from that of the declared type of the parameter variable, the value will be implicitly converted only if allowed. Otherwise, a warning is issued.

For each parameter that will be provided the expected value there is a limit of one directive. Parameters that will not be provided the expected value do not require a directive.

### Examples

The following example tells the compiler that the most likely values for parameters *a* and *b* are 1 and 0, respectively:

```
int func(int a,int b)
{
#pragma expected_value(a,1)
#pragma expected_value(b,0)
...
...
}
```

### Related information

- “#pragma execution\_frequency” on page 398

## #pragma fini (C only)

### Category

“Object code control” on page 387

### Purpose

Specifies the order in which the runtime library calls a list of functions after main() completes or exit() is called.

For shared libraries, the fini functions are called when the shared library is loaded from memory. For example, when using dynamic loading, this happens at the point when dlclose() is called.

### Syntax

```
▶▶ #pragma fini ( function_name ) ▶▶
```

### Usage

Any function that is specified in the pragma should have return type void (for example, void fA(); ) and take no parameters. Functions that have a non-void return type are accepted but the return value is discarded.

Functions that take parameters are ignored with a warning since the parameters would contain garbage values.

Within the same compilation unit, the list of functions in pragma fini are called in the order specified. Similarly, within the same compilation unit, functions specified in more than one pragma fini are called in the order in which the pragmas are encountered in the source.

In general, the order of static termination across files and across libraries is nonstandard and therefore, a non-portable behavior. It is not advisable to build any dependency on this behavior. The order of functions across files is undefined, even when using the **-Wm** option.

When mixing C and C++ files, the relative order of init or fini functions in C files with respect to the static constructors/destructors in C++ files is undefined. The **-qunique** option can interact with pragma fini.

**Note:** A C++ invocation, such as **xlC** or the redistributable tools **linkxlC** or **makeC++SharedLib** must be used at link time.

### Related information

- “#pragma init (C only)” on page 407
- “-qunique” on page 362

## #pragma hashome (C++ only)

### Category

Object code control

## Purpose

Informs the compiler that the specified class has a home module that will be specified by **#pragma ishome**.

This class's virtual function table, along with certain inline functions, will not be generated as static. Instead, they will be referenced as externals in the compilation unit of the class in which **#pragma ishome** is specified.

## Syntax

```
▶▶ #pragma ishome ( class_name [ allinlines ] ) ▶▶
```

## Parameters

*class\_name*

The name of a class to be referenced externally. *class\_name* must be a class and it must be defined.

**allinlines**

Specifies that all inline functions from within *class\_name* should be referenced as being external.

## Usage

A warning will be produced if there is a **#pragma ishome** without a matching **#pragma ishome**.

## Examples

In the following example, compiling the code samples will generate virtual function tables and the definition of `S::foo()` only for compilation unit a.o, but not for b.o. This reduces the amount of code generated for the application.

```
// a.h
struct S
{
    virtual void foo() {}

    virtual void bar();
};
```

```
// a.C
#pragma ishome(S)
#pragma ishome (S)
```

```
#include "a.h"
```

```
int main()
{
    S s;
    s.foo();
    s.bar();
}
```

```
// b.C
```

```
#pragma hashome(S)
#include "a.h"

void S::bar() {}
```

## Related information

- “#pragma ishome (C++ only)” on page 408

## #pragma ibm iterations

### Category

Optimization and tuning

### Purpose

The **iterations** pragma specifies the approximate average number of loop iterations for the chosen loop.

### Syntax

```
▶▶ #pragma ibm iterations(iteration_count) ▶▶
```

### Parameters

*iteration\_count*

Specifies the approximate number of loop iterations using a positive integral constant expression.

### Usage

The compiler uses the information in *iteration\_count* for loop optimization. You can specify multiple #pragma ibm iterations(*iteration\_count*).

*iteration\_count* specified in #pragma ibm iterations cannot be smaller than *iteration\_count* specified in #pragma ibm min\_iterations. In addition, it cannot be bigger than *iteration\_count* specified in #pragma ibm max\_iterations. Otherwise, the inconsistent value is ignored with a message.

### Example

```
#pragma ibm iterations(100) // Accepted
#pragma ibm min_iterations(150) // Ignored (150 > 100)
#pragma ibm min_iterations( 30) // Accepted( 30 < 100)
#pragma ibm max_iterations( 60) // Ignored ( 60 < 100)
#pragma ibm iterations( 20) // Ignored ( 20 < 30)
#pragma ibm max_iterations(500) // Accepted(500 > 100 > 30)
#pragma ibm max_iterations(620) // Ignored (Multiple occurrences)
#pragma ibm iterations(200) // Accepted( 30 < 200 < 500)
#pragma ibm min_iterations( 15) // Ignored (Multiple occurrences)

for (int i=0; i < n; ++i)
{
    #pragma ibm max_iterations(130) // Accepted
    #pragma ibm min_iterations( 90) // Accepted( 90 < 130)
    #pragma ibm iterations( 60) // Ignored ( 60 < 90)
    #pragma ibm iterations(100) // Accepted( 90 < 100 < 130)

    for (int j=0; j < m; ++j) b[j] += a[i];
}
```

Related reference:

"#pragma ibm max\_iterations"

"#pragma ibm min\_iterations" on page 405

## #pragma ibm max\_iterations

### Category

Optimization and tuning

### Purpose

The **max\_iterations** pragma specifies the approximate maximum number of loop iterations for the chosen loop.

### Syntax

```
▶▶ #pragma ibm max_iterations(iteration_count) ▶▶
```

### Parameters

*iteration\_count*

Specifies the approximate number of maximum loop iterations using a positive integral constant expression.

### Usage

The compiler uses the information in *iteration\_count* for loop optimization. You can specify #pragma ibm max\_iterations(*iteration\_count*) only once. If you specify #pragma ibm max\_iterations(*iteration\_count*) more than once, the first specified pragma is accepted, and the subsequent pragmas are ignored with a message.

*iteration\_count* specified in #pragma ibm max\_iterations cannot be smaller than *iteration\_count* specified in #pragma ibm iterations or #pragma ibm min\_iterations. Otherwise, the inconsistent value is ignored with a message.

### Example

```
#pragma ibm iterations(100) // Accepted
#pragma ibm min_iterations(150) // Ignored (150 > 100)
#pragma ibm min_iterations( 30) // Accepted( 30 < 100)
#pragma ibm max_iterations( 60) // Ignored ( 60 < 100)
#pragma ibm iterations( 20) // Ignored ( 20 < 30)
#pragma ibm max_iterations(500) // Accepted(500 > 100 > 30)
#pragma ibm max_iterations(620) // Ignored (Multiple occurrences)
#pragma ibm iterations(200) // Accepted( 30 < 200 < 500)
#pragma ibm min_iterations( 15) // Ignored (Multiple occurrences)

for (int i=0; i < n; ++i)
{
    #pragma ibm max_iterations(130) // Accepted
    #pragma ibm min_iterations( 90) // Accepted( 90 < 130)
    #pragma ibm iterations( 60) // Ignored ( 60 < 90)
    #pragma ibm iterations(100) // Accepted( 90 < 100 < 130)

    for (int j=0; j < m; ++j) b[j] += a[i];
}
```

Related reference:

"#pragma ibm iterations" on page 403

"#pragma ibm min\_iterations"

## #pragma ibm min\_iterations

### Category

Optimization and tuning

### Purpose

The **min\_iterations** pragma specifies the approximate minimum number of loop iterations for the chosen loop.

### Syntax

```
▶▶ #pragma ibm min_iterations(iteration_count) ▶▶
```

### Parameters

*iteration\_count*

Specifies the approximate minimum number of loop iterations using a positive integral constant expression.

### Usage

The compiler uses the information in *iteration\_count* for loop optimization. You can specify #pragma ibm min\_iterations(*iteration\_count*) only once. If you specify #pragma ibm min\_iterations(*iteration\_count*) more than once, the first specified pragma is accepted, and the subsequent pragmas are ignored with a message.

*iteration\_count* specified in #pragma ibm min\_iterations cannot be bigger than *iteration\_count* specified in #pragma ibm iterations or #pragma ibm max\_iterations. Otherwise, the inconsistent value is ignored with a message.

### Example

```
#pragma ibm iterations(100) // Accepted
#pragma ibm min_iterations(150) // Ignored (150 > 100)
#pragma ibm min_iterations( 30) // Accepted( 30 < 100)
#pragma ibm max_iterations( 60) // Ignored ( 60 < 100)
#pragma ibm iterations( 20) // Ignored ( 20 < 30)
#pragma ibm max_iterations(500) // Accepted(500 > 100 > 30)
#pragma ibm max_iterations(620) // Ignored (Multiple occurrences)
#pragma ibm iterations(200) // Accepted( 30 < 200 < 500)
#pragma ibm min_iterations( 15) // Ignored (Multiple occurrences)

for (int i=0; i < n; ++i)
{
    #pragma ibm max_iterations(130) // Accepted
    #pragma ibm min_iterations( 90) // Accepted( 90 < 130)
    #pragma ibm iterations( 60) // Ignored ( 60 < 90)
    #pragma ibm iterations(100) // Accepted( 90 < 100 < 130)

    for (int j=0; j < m; ++j) b[j] += a[i];
}
```

**Related reference:**

“#pragma ibm iterations” on page 403

“#pragma ibm max\_iterations” on page 404

## #pragma ibm snapshot

### Category

Error checking and debugging

### Purpose

Specifies a location at which a breakpoint can be set and defines a list of variables that can be examined when program execution reaches that location.

You can use this pragma to facilitate debugging optimized code produced by the compiler.

### Syntax

```
▶▶ #pragma ibm snapshot ( (variable_name) ) ▶▶
```

### Parameters

*variable\_name*

A variable name. It must not refer to structure, class, or union members.

### Usage

During a debugging session, you can place a breakpoint on the line at which the directive appears, to view the values of the named variables. When you compile with optimization and the `-g` option, the named variables are guaranteed to be visible to the debugger.

This pragma does not consistently preserve the contents of variables with a static storage class at high optimization levels. Variables specified in the directive should be considered read-only while being observed in the debugger, and should not be modified. Modifying these variables in the debugger may result in unpredictable behavior.

### Examples

```
#pragma ibm snapshot(a, b, c)
```

### Related information

- “-g” on page 169
- “-O, -qoptimize” on page 269

## #pragma implementation (C++ only)

### Category

Template control



## Purpose

For use with the **-qtempinc** compiler option, supplies the name of the file containing the template definitions corresponding to the template declarations contained in a header file.

## Syntax

```
▶▶ #pragma implementation ("file_name") ▶▶
```

## Parameters

*file\_name*

The name of the file containing the definitions for members of template classes declared in the header file.

## Usage

This pragma is not normally required if your template implementation file has the same name as the header file containing the template declarations, and a .c extension. You only need to use the pragma if the template implementation file does not conform to this file-naming convention. For more information about using template implementation files, see "Using C++ templates"

**#pragma implementation** is only effective if the **-qtempinc** option is in effect. Otherwise, the pragma has no meaning and is ignored.

The pragma can appear in the header file containing the template declarations, or in a source file that includes the header file. It can appear anywhere that a declaration is allowed.

## Related information

- "-qtempinc (C++ only)" on page 345
- "Using C++ templates"

## #pragma info

See "-qinfo" on page 187.

## #pragma init (C only)

### Category

"Object code control" on page 387

### Purpose

Specifies the order in which the runtime library calls a list of functions before `main()` is called.

For shared libraries, the init functions are called when the shared library is loaded to memory. For example, when using dynamic loading, this happens at the point when `dlopen()` is called.

## Syntax

```
▶▶ #pragma init ( function_name ) ▶▶
```

## Usage

Any function that is specified in the pragma should have return type void (for example, `void fA();`) and take no parameters. Functions that have a non-void return type are accepted but the return value is discarded.

Functions that take parameters are ignored with a warning since the parameters would contain garbage values.

Within the same compilation unit, the list of functions in `pragma init` are called in the order specified. Similarly, within the same compilation unit, functions specified in more than one `pragma init` are called in the order in which the pragmas are encountered in the source.

In general, the order of static initialization across files and across libraries is nonstandard and therefore, a non-portable behavior. It is not advisable to build any dependency on this behavior. The order of functions across files is undefined, even when using the `-Wm` option).

When mixing C and C++ files, the relative order of `init` functions in C files with respect to the static constructors/destructors in C++ files is undefined. The `-qunique` option can interact with `pragma init`.

**Note:** A C++ invocation, such as `x1C` or the redistributable tools `linkx1C` or `makeC++SharedLib` must be used at link time.

## Related information

- “`#pragma fini` (C only)” on page 401
- “`-qunique`” on page 362

## #pragma ishome (C++ only)

### Category

Object code control

### Purpose

Informs the compiler that the specified class's home module is the current compilation unit.

The home module is where items, such as the virtual function table, are stored. If an item is referenced from outside of the compilation unit, it will not be generated outside its home. This can reduce the amount of code generated for the application.

## Syntax

▶▶ #pragma ishome (—*class\_name*—) ▶▶

## Parameters

*class\_name*

The name of the class whose home will be the current compilation unit.

## Usage

A warning will be produced if there is a **#pragma ishome** without a matching **#pragma hashome**.

## Examples

See “#pragma hashome (C++ only)” on page 401

## Related information

- “#pragma hashome (C++ only)” on page 401

## #pragma isolated\_call

See “-qisolated\_call” on page 206.

## #pragma langlvl (C only)

See “-qlanglvl” on page 215.

## #pragma leaves

### Category

Optimization and tuning

### Purpose

Informs the compiler that a named function never returns to the instruction following a call to that function.

By informing the compiler that it can ignore any code after the function, the directive allows for additional opportunities for optimization.

This pragma is commonly used for custom error-handling functions, in which programs can be terminated if a certain error is encountered.

**Note:** The compiler automatically inserts **#pragma leaves** directives for calls to the `longjmp` family of functions (`longjmp`, `_longjmp`, `siglongjmp`, and `_siglongjmp`) when you include the `setjmp.h` header.

## Syntax

▶▶ #pragma leaves (—*function\_name*—) ▶▶

## Parameters

*function\_name*

The name of the function that does not return to the instruction following the call to it.

## Defaults

Not applicable.

## Examples

```
#pragma leaves(handle_error_and_quit)
void test_value(int value)
{
    if (value == ERROR_VALUE)
    {
        handle_error_and_quit(value);
        TryAgain(); // optimizer ignores this because
                   // never returns to execute it
    }
}
```

## Related information

- “#pragma reachable” on page 427.

## #pragma loopid

### Category

Optimization and tuning

### Purpose

Marks a block with a scope-unique identifier.

### Syntax

```
▶▶ #pragma loopid (—name—) ◀◀
```

### Parameters

*name*

An identifier that is unique within the scoping unit.

### Usage

The **#pragma loopid** directive must immediately precede a **#pragma block\_loop** directive or for loop. The specified name can be used by **#pragma block\_loop** to control transformations on that loop. It can also be used to provide information on loop transformations through the use of the **-qreport** compiler option.

You must not specify **#pragma loopid** more than once for a given loop.

### Examples

For examples of **#pragma loopid** usage, see “#pragma block\_loop” on page 390.

## Related information

- “-qunroll” on page 363
- “#pragma block\_loop” on page 390
- “#pragma unrollandfuse” on page 435

## #pragma map

### Category

Object code control

### Purpose

Converts all references to an identifier to another, externally defined identifier.

### Syntax

#### #pragma map syntax – C



```
▶▶ #pragma map (—name1—, —"name2"—) —————▶▶
```

#### #pragma map syntax – C++

```
▶▶ #pragma map (—name1—(—argument_list—), —"name2"—) —————▶▶
```

### Parameters

*name1*

The name used in the source code.  *name1* can represent a data object or function with external linkage.  *name1* can represent a data object, a non-overloaded or overloaded function, or overloaded operator, with external linkage. If the name to be mapped is not in the global namespace, it must be fully qualified.


*name1* should be declared in the same compilation unit in which it is referenced, but should not be defined in any other compilation unit. *name1* must not be used in another **#pragma map** directive or any assembly label declaration anywhere in the program.

 *argument\_list*

The list of arguments for the overloaded function or operator function designated by *name1*. If *name1* designates an overloaded function, the function must be parenthesized and must include its argument list if it exists. If *name1* designates a non-overloaded function, only *name1* is required, and the parentheses and argument list are optional.

*name2*

The name that will appear in the object code.  *name2* can represent a data object or function with external linkage.

 *name2* can represent a data object, a non-overloaded or overloaded function, or overloaded operator, with external linkage. *name2* must be specified using its mangled name. To obtain C++ mangled names, compile your source to object files only, using the **-c** compiler option, and use the **nm** operating system command on the resulting object file. You can also use the **c++filt** utility provided by the compiler for a side-by-side listing of source

names and mangled names; see "Demangling compiled C++ names" in the *XL C/C++ Optimization and Programming Guide* for details. (See also "Name mangling" in the *XL C/C++ Language Reference* for details on using the extern "C" linkage specifier on declarations to prevent name mangling.)

If the name exceeds 65535 bytes, an informational message is emitted and the pragma is ignored.

*name2* may or may not be declared in the same compilation unit in which *name1* is referenced, but must not be defined in the same compilation unit. Also, *name2* should not be referenced anywhere in the compilation unit where *name1* is referenced. *name2* must not be the same as that used in another **#pragma map** directive or any assembly label declaration in the same compilation unit.

## Usage

The **#pragma map** directive can appear anywhere in the program. Note that in order for a function to be actually mapped, the map target function (*name2*) must have a definition available at link time (from another compilation unit), and the map source function (*name1*) must be called in your program.

You cannot use **#pragma map** with compiler built-in functions.

## Examples

The following is an example of **#pragma map** used to map a function name (using the mangled name for the map name in C++):

```
/* Compilation unit 1: */
#include <stdio.h>

void foo();
extern void bar(); /* optional */

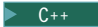
#ifdef __cplusplus
#pragma map (foo, "bar__Fv")
#else
#pragma map (foo, "bar")
#endif
int main()
{
    foo();
}

/* Compilation unit 2: */
#include <stdio.h>

void bar()
{
    printf("Hello from foo bar!\n");
}
```

The call to `foo` in compilation unit 1 resolves to a call to `bar`:

Hello from foo bar!

 The following is an example of **#pragma map** used to map an overloaded function name (using C linkage, to avoid using the mangled name for the map name):

```

// Compilation unit 1:

#include <iostream>
#include <string>

using namespace std;

void foo();
void foo(const string&);
extern "C" void bar(const string&); // optional

#pragma map (foo(const string&), "bar")

int main()
{
foo("Have a nice day!");
}

// Compilation unit 2:

#include <iostream>
#include <string>

using namespace std;

extern "C" void bar(const string& s)
{
cout << "Hello from foo bar!" << endl;
cout << s << endl;
}

```

The call to `foo(const string&)` in compilation unit 1 resolves to a call to `bar(const string&)`:

```

Hello from foo bar!
Have a nice day!

```

### Related information

- "Assembly labels" in the *XL C/C++ Language Reference*

## #pragma mc\_func

### Category

Language element control

### Purpose

Allows you to embed a short sequence of machine instructions "inline" within your program source code.

The pragma instructs the compiler to generate specified instructions in place rather than the usual linkage code. Using this pragma avoids performance penalties associated with making a call to an assembler-coded external function. This pragma is similar in function to inline `asm` statements supported in this and other compilers; see "Inline assembly statements" in the *XL C/C++ Language Reference* for more information.

## Syntax

```
▶▶ #pragma mc_func function_name { instruction_sequence } ▶▶
```

## Parameters

### *function\_name*

The name of a previously-defined function containing machine instructions. If the function is not previously-defined, the compiler will treat the pragma as a function definition.

### *instruction\_sequence*

A string containing a sequence of zero or more hexadecimal digits. The number of digits must comprise an integral multiple of 32 bits. If the string exceeds 16384 bytes, a warning message is emitted and the pragma is ignored.

## Usage

This pragma defines a function and should appear in your program source only where functions are ordinarily defined.

The compiler passes parameters to the function in the same way as to any other function. For example, in functions taking integer-type arguments, the first parameter is passed to GPR3, the second to GPR4, and so on. Values returned by the function will be in GPR3 for integer values, and FPR1 for float or double values.

Code generated from *instruction\_sequence* may use any and all volatile registers available on your system unless you use **#pragma reg\_killed\_by** to list a specific register set for use by the function. See “#pragma reg\_killed\_by” on page 428 for a list of volatile registers available on your system.

Inlining options do not affect functions defined by **#pragma mc\_func**. However, you may be able to improve runtime performance of such functions with **#pragma isolated\_call**.

## Examples

In the following example, **#pragma mc\_func** is used to define a function called `add_logical`. The function consists of machine instructions to add 2 integers with so-called *end-around carry*; that is, if a carry out results from the add then add the carry to the sum. This formula is frequently used in checksum computations.

```
int add_logical(int, int);
#pragma mc_func add_logical {"7c632014" "7c630194"}
    /* addc    r3 <- r3, r4      */
    /* addze   r3 <- r3, carry bit */

main() {
    int i,j,k;

    i = 4;
```



```

    k = -4;
    j = add_logical(i,k);
    printf("\n\nresult = %d\n\n",j);
}

```

The result of running the program is:

```
result = 1
```

### Related information

- “-qisolated\_call” on page 206
- “#pragma reg\_killed\_by” on page 428
- “Inline assembly statements” in the *XL C/C++ Language Reference*

## #pragma namemangling (C++ only)

See “-qnamemangling (C++ only)” on page 262.

## #pragma namemanglingrule (C++ only)

### Category

Portability and migration

### Purpose

Provides fined-grained control over the name mangling scheme in effect for selected portions of source code, specifically with respect to the mangling of cv-qualifiers in function parameters.

When a function name is mangled, repeated function arguments of the same type are encoded according to the following compression scheme:

```

parameter → T param number [ ] #single repeat of a previous parameter
           → N repetition digit param number [ ] #2 to 9 repetitions

```

where:

*param number*

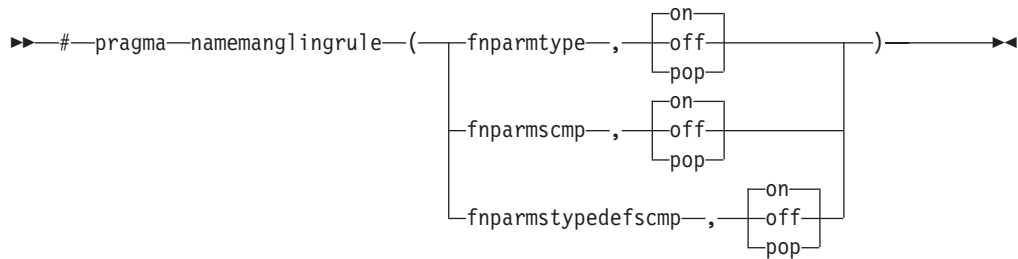
Indicates the number of the previous parameter which is repeated. It is followed by an underscore ( `_` ) if *param number* contains multiple digits.

*repetition digit*

Must be greater than 1 and less than 10. If an argument is repeated more than 9 times, this rule is applied multiple times. For example, a sequence of 38 parameters that are the same as parameter 1 mangles to N91N91N91N91N21.

The `#pragma namemanglingrule` directive allows you to control whether top-level cv-qualifiers are mangled in function parameters or whether intermediate-level cv-qualifiers are to be considered when the compiler compares repeated function parameters for equivalence.

### Syntax



## Defaults

- `fnparmtyp`, **on** when `-qnamemangling=ansi | v6` or higher or `#pragma namemangling (ansi | v6)` or higher is in effect. Otherwise, the default is `fnparmtyp`, **off**.
- `fnparmscmp`, **on** when `-qnamemangling=ansi | v8` or higher or `#pragma namemangling (ansi | v8)` or higher is in effect. Otherwise, the default is `fnparmscmp`, **off**.
- `fnparmstypedefscmp`, **on** when `-qnamemangling=compat` or `#pragma namemangling (compat)` or higher is in effect. Otherwise, the default is `fnparmstypedefscmp`, **off**.

## Parameters

### `fnparmtyp`, **on**

Top-level cv-qualifiers are not encoded in the mangled name of a function parameter. Also, top-level cv-qualifiers are ignored when repeated function parameters are compared for equivalence; function parameters that differ only by the use of a top-level cv-qualifier are considered equivalent and are mangled according to the compressed encoding scheme. This setting is compatible with VisualAge C++ V6.0 and higher.

### `fnparmtyp`, **off**

Top-level cv-qualifiers are encoded in the mangled name of a function parameter. Also, repeated function parameters that differ by the use of cv-qualifiers are not considered equivalent and are mangled as separate parameters. This setting is compatible with VisualAge C++ V5.0 and earlier.

### `fnparmtyp`, **pop**

Reverts to the previous `fnparmtyp` setting in effect. If no previous settings are in effect, the default `fnparmtyp` setting is used.

**Note:** This pragma fixes function signature ambiguities in 32-bit mode, but it is not needed in 64-bit mode since those ambiguities do not exist in 64-bit mode.

### `fnparmscmp`, **on**

Intermediate-level cv-qualifiers are considered when repeated function parameters are compared for equivalence; repeated function parameters that differ by the use of intermediate-level cv-qualifiers are mangled as separate parameters. This setting is compatible with XL C++ V8.0 and higher.

### `fnparmscmp`, **off**

Intermediate-level cv-qualifiers are ignored when repeated function parameters are compared for equivalence; function parameters that differ only by the use of an intermediate-level cv-qualifier are considered equivalent and are mangled according to the compressed encoding scheme. This setting is compatible with XL C++ V7.0 and earlier.

### **fnparmscmp, pop**

Reverts to the previous **fnparmscmp** setting in effect. If no previous settings are in effect, the default **fnparmscmp** setting is used.

## **Usage**

**#pragma namemanglingrule** is allowed in global, class, and function scopes. It has no effect on a block scope function declaration with external linkage.

Different pragma settings can be specified in front of function declarations and definitions. If **#pragma namemanglingrule** settings in subsequent declarations and definitions conflict, the compiler ignores those settings and issues a warning message.

## **Examples**

The following tables show the effects of this pragma applied to different function signatures.

*Table 41. Mangling of function parameters with top-level cv-qualifiers*

Source name	Mangled name	
	fnparmscmp, off	fnparmscmp, on
void foo (const int)	foo_FCi	foo_Fi
void foo (int* const)	foo_FCPi	foo_FPi
void foo (int** const)	foo_FCPPi	foo_FPPi
void foo (int, const int)	foo_FiCi	foo_FiT1

*Table 42. Mangling of function parameters with intermediate level cv-qualifiers*

Source name	Mangled name	
	fnparmscmp, on	fnparmscmp, off
void foo (int** a, int* const * b)	foo_FPPiPCPi	foo_FPPiT1
void bar (int* const* a, int** b)	bar_FPCPiPPi	bar_FPCPiT1

*Table 43. Mangling of function parameters with top-level and intermediate-level cv-qualifiers*

Source name	Mangled name			
	fnparmscmp, on fnparmscmp, on fnparmscmp, on	fnparmscmp, off fnparmscmp, off fnparmscmp, off	fnparmscmp, on fnparmscmp, on fnparmscmp, on	fnparmscmp, off fnparmscmp, off fnparmscmp, off
void foo (int** const, int* const *)	foo_FPPiPCPi	foo_FPPiT1	foo_FCPPiPCPi	foo_FPPiT1

## **Related information**

- “-qnamemangling (C++ only)” on page 262

## #pragma nofunctrace

### Category

Error checking and debugging

### Purpose

Disables tracing for a given function or a list of specified functions.

### Syntax

```
▶ #pragma nofunctrace ( function_name ) ▶
```

### Parameters

*function\_name*

The name of the function for which you want to disable tracing.

### Usage

When you use **#pragma nofunctrace** to specify a list of functions for which you want to disable tracing, use parenthesis **()** and encapsulate the functions in it. For a list of functions, use a comma **,** to separate them. For example, to disable tracing for function *a*, use `#pragma nofunctrace(a)`. To disable tracing for functions *a*, *b*, and *c*, use `#pragma nofunctrace(a,b,c)`.

If you have two functions: `foo(int)` and `foo(double)`, use `#pragma nofunctrace(foo(int))` disables tracing for `foo(int)` but not `foo(double)`.

Two colons in a row `::` are considered scope qualifiers. For example, when you call `-qfunctrace+A::B:C`, the compiler traces functions that begin with the qualifiers `A::B` or `C`.

**Note:** If you want to use the compiler option `-qfunctrace` to disable tracing for a given function or a list of functions, you must use its suboption `-` followed by the names of the functions. For details about how to use `-qfunctrace` and its related suboptions, see “`-qfunctrace`” on page 166.

### Examples

```
#pragma nofunctrace(a,b,c)
```

### Related information

- “`-qfunctrace`” on page 166

## #pragma nosimd

See “`-qhot`” on page 178.

## #pragma novector

See “`-qhot`” on page 178.

## #pragma object\_model (C++ only)

See “-qobjmodel (C++ only)” on page 272.

## #pragma operator\_new (C++ only)

### Category

Error checking and debugging

### Purpose

Determines whether the `new` and `new[]` operators throw an exception if the requested memory cannot be allocated.

This pragma is equivalent to the `-qlanglvl=newexcp` option.

### Syntax

```
▶▶ #pragma operator_new ( [returnsnull] [throwexception] ) ▶▶
```

### Defaults

`returnsnull`

### Parameters

#### `returnsnull`

If the memory requested by the `new` operator cannot be allocated, the compiler returns 0, the null pointer. Use this option for compatibility with versions of the XL C++ compiler previous to Visual C++ V6.0.

#### `throwexception`

If the memory requested by the `new` operator cannot be allocated, the compiler throws a standard exception of type `std::bad_alloc`. Use this option in new applications, for conformance with the C++ standard.

### Usage

The pragma can be specified only once in a source file. It must appear before any statements in the source file. This pragma takes precedence over the `-qlanglvl=newexcp` compiler option.

### Restrictions

This pragma applies only to versions of the `new` operator that throw exceptions; it does not apply to the `nothrow` or empty throw versions of the `new` operator (for the prototypes of all the `new` operator versions, see the description of the `<new>` header in the *Standard C++ Library Reference*). It also does not apply to class-specific `new` operators, user-defined `new` operators, and `new` operators with placement arguments.

### Related information

- "Allocation and deallocation functions" in the *XL C/C++ Language Reference*
- "The `new` operator" in the *XL C/C++ Language Reference*

- “-qlanglvl” on page 215
- The <new> header in the *Standard C++ Library Reference*

## #pragma options


### Category

Language element control

### Purpose



Specifies compiler options in your source program.

### Syntax



### Parameters

The settings in the table below are valid *options* for **#pragma options**. For more information, see the pages of the equivalent compiler option.

Valid settings for #pragma options <i>option_keyword</i>	Compiler option equivalent
align= <i>option</i>	“-qalign” on page 100
[no]attr attr=full	“-qattr” on page 114
chars= <i>option</i>	“-qchars” on page 123
[no]check	“-qcheck” on page 125
[no]compact	“-qcompact” on page 127
[no]dbcs	“-qmbcs, -qdbc” on page 257
 [no]dbxextra	“-qdbxextra (C only)” on page 136
[no]digraph	“-qdigraph” on page 138
[no]dollar	“-qdollar” on page 140
enum= <i>option</i>	“-qenum” on page 144
[no]extchk	“-qextchk” on page 149
flag= <i>option</i>	“-qflag” on page 153
float=[no] <i>option</i>	“-qfloat” on page 154
[no]flttrap	“-qflttrap” on page 159
[no]fullpath	“-qfullpath” on page 164
 [no]funcsect	“-qfuncsect” on page 165
halt	“-qhalt” on page 174

Valid settings for #pragma options <i>option_keyword</i>	Compiler option equivalent
[no]idirfirst	"-qidirfirst" on page 182
[no]ignerrno	"-qignerrno" on page 183
ignprag= <i>option</i>	"-qignprag" on page 184
[no]info= <i>option</i>	"-qinfo" on page 187
initauto= <i>value</i>	"-qinitauto" on page 193
[no]inlglue	"-qinlglue" on page 196
isolated_call= <i>names</i>	"-qisolated_call" on page 206
 langlvl	"-qlanglvl" on page 215
[no]dbl128	"-qdbl128, -qlongdouble" on page 239
[no]libansi	"-qlibansi" on page 241
[no]list	"-qlist" on page 243
[no]longlong	"-qlonglong" on page 249
 [no]macpstr	"-qmacpstr" on page 251
[no]maxmem= <i>number</i>	"-qmaxmem" on page 256
[no]mbcs	"-qmbcs, -qdbc" on page 257
[no]optimize= <i>number</i>	"-O, -qoptimize" on page 269
 priority= <i>number</i>	"-qpriority (C++ only)" on page 297
proclcal, procimported, procunknown	"-qprocimported, -qproclcal, -qprocunknown" on page 298
 [no]proto	"-qproto (C only)" on page 300
[no]ro	"-qro" on page 305
[no]roconst	"-qroconst" on page 306
[no]showinc	"-qshowinc" on page 313
[no]source	"-qsource" on page 324
spill= <i>number</i>	"-qspill" on page 327
 [no]srcmsg	"-qsrcmsg (C only)" on page 328
[no]stdinc	"-qstdinc" on page 332
[no]strict	"-qstrict" on page 333
thtable= <i>option</i>	"-qthtable" on page 344
tune= <i>option</i>	"-qtune" on page 357
[no]unrollunroll= <i>number</i>	"-qunroll" on page 363
 [no]upconv	"-qupconv (C only)" on page 366
[no]xref	"-qxref" on page 379

## Usage

Most **#pragma options** directives must come before any statements in your source program; only comments, blank lines, and other pragma specifications can precede them. For example, the first few lines of your program can be a comment followed by the **#pragma options** directive:

```

/* The following is an example of a #pragma options directive: */
#pragma options langlvl=stdc89 halt=s spill=1024 source
/* The rest of the source follows ... */

```

To specify more than one compiler option with the **#pragma options** directive, separate the options using a blank space. For example:

```

#pragma options langlvl=stdc89 halt=s spill=1024 source

```

## #pragma option\_override

### Category

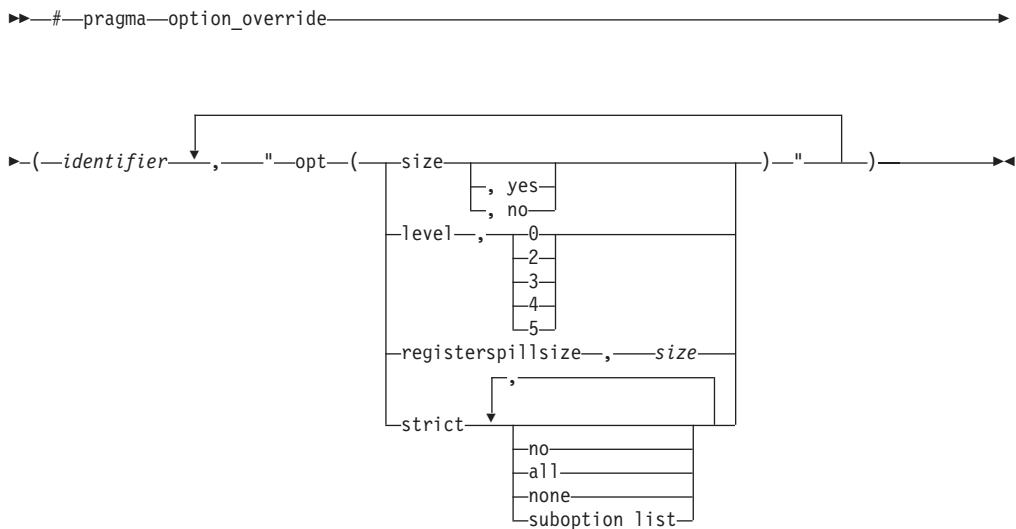
Optimization and tuning

### Purpose

Allows you to specify optimization options at the subprogram level that override optimization options given on the command line.

This enables finer control of program optimization, and can help debug errors that occur only under optimization.

### Syntax



### Parameters

*identifier*

The name of a function for which optimization options are to be overridden.

The following table shows the equivalent command line option for each pragma suboption.

#pragma option_override value	Equivalent compiler option
level, 0	-O
level, 2	-O2
level, 3	-O3



#pragma option_override value	Equivalent compiler option
level, 4	-O4
level, 5	-O5
registerspillsize, <i>size</i>	-qspill= <i>size</i>
size	-qcompact
size, yes	
size, no	-qnocompact
strict, all	-qstrict, -qstrict=all
strict, no, none	-qnostrict
strict, <i>suboption_list</i>	-qstrict= <i>suboption_list</i>

## Defaults

See the descriptions for the options listed in the table above for default settings.

## Usage

The pragma takes effect only if optimization is already enabled by a command-line option. You can only specify an optimization level in the pragma *lower* than the level applied to the rest of the program being compiled.

The **#pragma option\_override** directive only affects functions that are defined in the same compilation unit. The pragma directive can appear anywhere in the translation unit. That is, it can appear before or after the function definition, before or after the function declaration, before or after the function has been referenced, and inside or outside the function definition.

 This pragma cannot be used with overloaded member functions.

## Examples

Suppose you compile the following code fragment containing the functions `foo` and `faa` using `-O2`. Since it contains the `#pragma option_override(faa, "opt(level, 0)")`, function `faa` will not be optimized.

```
foo(){
    .
    .
    .
}

#pragma option_override(faa, "opt(level, 0)")

faa(){
    .
    .
    .
}
```

## Related information

- “`-O`, `-qoptimize`” on page 269
- “`-qcompact`” on page 127
- “`-qspill`” on page 327
- “`-qstrict`” on page 333

# #pragma pack

## Category

Object code control


## Purpose

Sets the alignment of all aggregate members to a specified byte boundary.

If the byte boundary number is smaller than the natural alignment of a member, padding bytes are removed, thereby reducing the overall structure or union size.

## Syntax

### Default #pragma pack syntax



## Defaults

Members of aggregates (structures, unions, and classes) are aligned on their natural boundaries and a structure ends on its natural boundary. The alignment of an aggregate is that of its strictest member (the member with the largest alignment requirement).

## Parameters

### nopack

Disables packing. A warning message is issued and the pragma is ignored.

### number

is one of the following:

- 1 Aligns structure members on 1-byte boundaries, or on their natural alignment boundary, whichever is less.
- 2 Aligns structure members on 2-byte boundaries, or on their natural alignment boundary, whichever is less.
- 4 Aligns structure members on 4-byte boundaries, or on their natural alignment boundary, whichever is less.
- 8 Aligns structure members on 8-byte boundaries, or on their natural alignment boundary, whichever is less.
- 16 Aligns structure members on 16-byte boundaries, or on their natural alignment boundary, whichever is less.

### pop

Removes the previous value added with **#pragma pack**. Specifying **#pragma pack()** with no parameters is equivalent to **pop**.

## Usage

The **#pragma pack** directive applies to the definition of an aggregate type, rather than to the declaration of an instance of that type; it therefore automatically applies to all variables declared of the specified type.

The **#pragma pack** directive modifies the current alignment rule for only the members of structures whose declarations follow the directive. It does not affect the alignment of the structure directly, but by affecting the alignment of the members of the structure, it may affect the alignment of the overall structure.

The **#pragma pack** directive cannot increase the alignment of a member, but rather can decrease the alignment. For example, for a member with data type of short, a **#pragma pack(1)** directive would cause that member to be packed in the structure on a 1-byte boundary, while a **#pragma pack(4)** directive would have no effect.

The **#pragma pack** directive aligns all bit fields in a structure/union on 1-bit boundaries. Example:

```
#pragma pack(2)
struct A{
int a:31;
int b:2;
}x;

int main(){
printf("size of S = %d\n", sizeof(s));
}
```

When compiled and run, the output is:  
size of S = 6

But if you remove the **#pragma pack** directive, you get this output:  
size of S = 8

The **#pragma pack** directive applies only to complete declarations of structures or unions; this excludes forward declarations, in which member lists are not specified. For example, in the following code fragment, the alignment for struct S is 4, since this is the rule in effect when the member list is declared:

```
#pragma pack(1)
struct S;
#pragma pack(4)
struct S { int i, j, k; };
```

A nested structure has the alignment that precedes its declaration, not the alignment of the structure in which it is contained, as shown in the following example:

```
#pragma pack (4)                // 4-byte alignment
    struct nested {
        int x;
        char y;
        int z;
    };

#pragma pack(1)                // 1-byte alignment
struct packedcxx{
    short b;
    struct nested s1;          // 4-byte alignment
};
```

If more than one **#pragma pack** directive appears in a structure defined in an inlined function, the **#pragma pack** directive in effect at the beginning of the structure takes precedence.

## Examples

The following example shows how the **#pragma pack** directive can be used to set the alignment of a structure definition:

```
// header file file.h

#pragma pack(1)

struct jeff{           // this structure is packed
    short bill;        // along 1-byte boundaries
    int *chris;
};
#pragma pack(pop)     // reset to previous alignment rule
// source file anyfile.c

#include "file.h"

struct jeff j;        // uses the alignment specified
                     // by the pragma pack directive
                     // in the header file and is
                     // packed along 1-byte boundaries
```

This example shows how a **#pragma pack** directive can affect the size and mapping of a structure:

```
struct s_t {
    char a;
    int b;
    short c;
    int d;
}S;
```

### Default mapping:

size of s\_t = 16  
offset of a = 0  
offset of b = 4  
offset of c = 8  
offset of d = 12  
alignment of a = 1  
alignment of b = 4  
alignment of c = 2  
alignment of d = 4

### With #pragma pack(1):

size of s\_t = 11  
offset of a = 0  
offset of b = 1  
offset of c = 5  
offset of d = 7  
alignment of a = 1  
alignment of b = 1  
alignment of c = 1  
alignment of d = 1

The following example defines a union uu containing a structure as one of its members, and declares an array of 2 unions of type uu:

```
union uu {
    short a;
    struct {
        char x;
        char y;
        char z;
    };
};
```


```

    } b;
};

union uu nonpacked[2];

```

Since the largest alignment requirement among the union members is that of short a, namely, 2 bytes, one byte of padding is added at the end of each union in the array to enforce this requirement:



The next example uses `#pragma pack(1)` to set the alignment of unions of type uu to 1 byte:

```


#pragma pack(1)

union uu {
    short a;
    struct {
        char x;
        char y;
        char z;
    } b;
};

union uu pack_array[2];

```

Now, each union in the array packed has a length of only 3 bytes, as opposed to the 4 bytes of the previous case:



### Related information

- “-qalign” on page 100
- "Using alignment modifiers" in the *XL C/C++ Optimization and Programming Guide*

### #pragma pass\_by\_value (C++ only)

See “-qoldpassbyvalue (C++ only)” on page 274.

### #pragma priority (C++ only)

See “-qpriority (C++ only)” on page 297.

### #pragma reachable

#### Category

Optimization and tuning

## Purpose

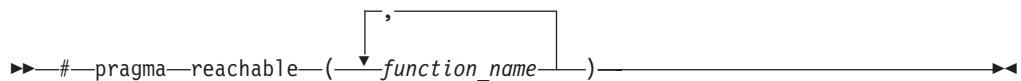
Informs the compiler that the point in the program after a named function can be the target of a branch from some unknown location.

By informing the compiler that the instruction after the specified function can be reached from a point in your program other than the return statement in the named function, the pragma allows for additional opportunities for optimization.

**Note:** The compiler automatically inserts **#pragma reachable** directives for the `setjmp` family of functions (`setjmp`, `_setjmp`, `sigsetjmp`, and `_sigsetjmp`) when you include the `setjmp.h` header file.

## Syntax

```
▶▶ #pragma reachable (function_name) ▶▶
```

A diagram showing the syntax of the #pragma reachable directive. It starts with a double arrow pointing right, followed by the text "#pragma reachable". Then, an opening parenthesis "(" is shown. A line goes from the opening parenthesis to the text "function\_name". A vertical line goes up from "function\_name", then a horizontal line goes right, then a vertical line goes down to a comma ",". Another vertical line goes down from the comma to the closing parenthesis ")", which is followed by another double arrow pointing right.

## Parameters

*function\_name*

The name of a function preceding the instruction which is reachable from a point in the program other than the function's return statement.

## Defaults

Not applicable.

## Related information

- “#pragma leaves” on page 409

## #pragma reg\_killed\_by

### Category

Optimization and tuning

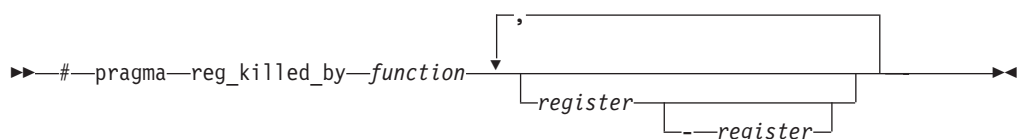
### Purpose

Specifies registers that may be altered by functions specified by **#pragma mc\_func**.

Ordinarily, code generated for functions specified by **#pragma mc\_func** may alter any or all volatile registers available on your system. You can use **#pragma reg\_killed\_by** to explicitly list a specific set of volatile registers to be altered by such functions. Registers not in this list will not be altered.

## Syntax

```
▶▶ #pragma reg_killed_by function register register ▶▶
```

A diagram showing the syntax of the #pragma reg\_killed\_by directive. It starts with a double arrow pointing right, followed by the text "#pragma reg\_killed\_by". Then, a line goes from the text to the text "function". A vertical line goes up from "function", then a horizontal line goes right, then a vertical line goes down to a comma ",". Another vertical line goes down from the comma to the text "register". A horizontal line goes from "register" to the text "register", with a vertical line going down from the second "register" to a second horizontal line below it. This second horizontal line ends with a double arrow pointing right.

## Parameters

### *function*

The name of a function previously defined using the `#pragma mc_func` directive.

### *register*

The symbolic name(s) of either a single register or a range of registers to be altered by the named *function*. The symbolic name must be a valid register name on the target platform. Valid registers are:

#### **cr0, cr1, and cr5 to cr7**

Condition registers

#### **ctr**

Count register

#### **gr0 and gr3 to gr12**

General purpose registers

#### **fp0 to fp13**

Floating-point registers

#### **fs**

Floating-point and status control register

#### **lr**

Link register

#### **vr0 to vr31**

Vector registers (on selected processors only)

#### **xer**

Fixed-point exception register

You can identify a range of registers by providing the symbolic names of both starting and ending registers, separated by a dash.

If no *register* is specified, no volatile registers will be killed by the named *function*.

## Examples

The following example shows how to use `#pragma reg_killed_by` to list a specific set of volatile registers to be used by the function defined by `#pragma mc_func`.

```
int add_logical(int, int);
#pragma mc_func add_logical {"7c632014" "7c630194"}
/*   addc      r3 <- r3, r4      */
/*   addze     r3 <- r3, carry bit */

#pragma reg_killed_by add_logical gr3, xer
/* only gpr3 and the xer are altered by this function */

main() {
    int i,j,k;

    i = 4;
    k = -4;
    j = add_logical(i,k);
    printf("\n\nresult = %d\n\n",j);
}
```

## Related information

- “`#pragma mc_func`” on page 413

## #pragma report (C++ only)

### Category


Listings, messages and compiler information

### Purpose

Controls the generation of diagnostic messages.

The pragma allows you to specify a minimum severity level for a message for it to display, or allows you to enable or disable a specific message regardless of the prevailing report level.

### Syntax



### Defaults

The default report level is Informational (I), which displays messages of all types.

### Parameters

#### level

Indicates that the pragma is set according to the minimum severity level of diagnostic messages to display.

- E** Indicates that only error messages will display. Error messages are of the highest severity. This is equivalent to the `-qflag=e:e` compiler option.
- W** Indicates that warning and error messages will display. This is equivalent to the `-qflag=w:w` compiler option.
- I** Indicates that all diagnostic messages will display: warning, error and informational messages. Informational messages are of the lowest severity. This is equivalent to the `-qflag=i:i` compiler option.

#### enable

Enables the specified `"message_number"`.

#### disable

Disables the specified `"message_number"`.

#### "message\_number"

Represents a message identifier, which consists of a prefix followed by the message number in quotation marks; for example, "CCN1004".

**Note:** You must use quotation marks with `message_number` as in the preceding example "CCN1004".

#### pop

Reverts the report level to that which was previously in effect. If no previous report level has been specified, a warning is issued, and the report level remains unchanged.



## Usage

The pragma takes precedence over **#pragma info** and most compiler options. For example, if you use **#pragma report** to disable a compiler message, that message will not be displayed with any **-qflag** compiler option setting.

## Related information

- “-qflag” on page 153

## #pragma simd\_level

### Category

Optimization and tuning

### Purpose

Controls the compiler code generation of vector instructions for individual loops.

Vector instructions can offer high performance when used with algorithmic-intensive tasks such as multimedia applications. You have the flexibility to control the aggressiveness of autosimdization on a loop-by-loop basis, and might be able to achieve further performance gain with this fine grain control.

The supported levels are from 0 to 10. level(0) indicates performing no autosimdization on the loop that follows the pragma directive. level(10) indicates performing the most aggressive form of autosimdization on the loop. With this pragma directive, you can control the autosimdization behavior on a loop-by-loop basis.

### Syntax

►► #pragma simd\_level (—*n*—) ◀◀

### Parameters

*n* A scalar integer initialization expression, from 0 to 10, specifying the aggressiveness of autosimdization on the loop that follows the pragma directive.

### Usage

A loop with no `simd_level` pragma is set to `simd level 5` by default, if `-qsimd=auto` is in effect.

**#pragma simd\_level(0)** is equivalent to **#pragma nosimd**, where autosimdization is not performed on the loop that follows the pragma directive.

**#pragma simd\_level(10)** instructs the compiler to perform autosimdization on the loop that follows the pragma directive most aggressively, including bypassing cost analysis.

### Rules

The rules of **#pragma simd\_level** directive are listed as follows:

- The `#pragma simd_level` directive has effect only for architectures that support vector instructions and when used with `-qsimd=auto`.
- The `#pragma simd_level` directive applies to `while`, `do while`, and `for` loops.
- The `#pragma simd_level` directive applies only to the loop immediately following it. The directive has no effect on other loops that are nested within the specified loop. It is possible to set different `simd` levels for the inner and outer loops by specifying separate `#pragma simd_level` directives.
- The `#pragma simd_level` directive can be mixed with loop optimization (`-qhot`) and OpenMP directives without requiring any specific optimization level. For more information about `-qhot` and OpenMP directives, see `-qhot` in the *IBM XL C/C++ Compiler Reference* and Using OpenMP directives in the *IBM XL C/C++ Optimization and Programming Guide*.

## Examples

```
...
#pragma simd_level(10)
for (i=1; i<1000; i++) {
/* program code */

} ...
```

## Related information

- 

## #pragma STDC cx\_limited\_range

### Category

Optimization and tuning

### Purpose

Instructs the compiler that complex division and absolute value are only invoked with values such that intermediate calculation will not overflow or lose significance.

### Syntax

```
▶▶ #pragma STDC cx_limited_range [off | on | default] ▶▶
```

### Usage

Using values outside the limited range may generate wrong results, where the limited range is defined such that the "obvious symbolic definition" will not overflow or run out of precision.

The `pragma` is effective from its first occurrence until another `cx_limited_range` `pragma` is encountered, or until the end of the translation unit. When the `pragma` occurs inside a compound statement (including within a nested compound statement), it is effective from its first occurrence until another `cx_limited_range` `pragma` is encountered, or until the end of the compound statement.

## Examples

The following example shows the use of the pragma for complex division:

```
#include <complex.h>

_Complex double a, b, c, d;
void p() {

d = b/c;

{

#pragma STDC CX_LIMITED_RANGE ON

a = b / c;

}

}
```

The following example shows the use of the pragma for complex absolute value:

```
#include <complex.h>

_Complex double cd = 10.10 + 10.10*I;
int p() {

#pragma STDC CX_LIMITED_RANGE ON

double d = cabs(cd);

}
```

## Related information

- "Standard pragmas" in the *XL C/C++ Language Reference*

## #pragma stream\_unroll

### Category

Optimization and tuning

### Purpose



When optimization is enabled, breaks a stream contained in a for loop into multiple streams.

### Syntax

```
▶▶ #pragma stream_unroll [ (—number—) ] ▶▶
```

### Parameters

*number*



A loop unrolling factor.  The value of *number* is a positive integral constant expression.  The value of *number* is a positive scalar integer or compile-time constant initialization expression.

An unroll factor of 1 disables unrolling.

If *number* is not specified, the optimizer determines an appropriate unrolling factor for each nested loop.

## Usage

To enable stream unrolling, you must specify `-qhot` and `-qstrict`, or `-qsmp`, or use optimization level `-O4` or higher. If `-qstrict` is in effect, no stream unrolling takes place.

For stream unrolling to occur, the `#pragma stream_unroll` directive must be the last pragma specified preceding a for loop.  Specifying `#pragma stream_unroll` more than once for the same for loop or combining it with other loop unrolling pragmas (`#pragma unroll`, `#pragma nounroll`, `#pragma unrollandfuse`, `#pragma nounrollandfuse`) results in a warning.  The compiler silently ignores all but the last of multiple loop unrolling pragmas specified on the same for loop.

## Examples

The following example shows how `#pragma stream_unroll` can increase performance.

```
int i, m, n;
int a[1000];
int b[1000];
int c[1000];

....

#pragma stream_unroll(4)
for (i=0; i<n; i++) {
    a[i] = b[i] * c[i];
}
```

The unroll factor of 4 reduces the number of iterations from  $n$  to  $n/4$ , as follows:

```
m = n/4;

for (i=0; i<n/4; i++){
    a[i] = b[i] + c[i];
    a[i+m] = b[i+m] + c[i+m];
    a[i+2*m] = b[i+2*m] + c[i+2*m];
    a[i+3*m] = b[i+3*m] + c[i+3*m];
}
```

The increased number of read and store operations are distributed among a number of streams determined by the compiler, which reduces computation time and increase performance.

### Related information

- “`-qunroll`” on page 363
- “`#pragma unrollandfuse`” on page 435

## #pragma strings

See “`-qro`” on page 305.

## #pragma unroll

See “`-qunroll`” on page 363.

# #pragma unrollandfuse


## Category

Optimization and tuning

## Purpose



Instructs the compiler to attempt an unroll and fuse operation on nested for loops.

## Syntax



## Parameters

*number*

A loop unrolling factor.  The value of *number* is a positive integral constant expression.  The value of *number* is a positive scalar integer or compile-time constant initialization expression.

If *number* is not specified, the optimizer determines an appropriate unrolling factor for each nested loop.

## Usage

The **#pragma unrollandfuse** directive applies only to the outer loops of nested for loops that meet the following conditions:

- There must be only one loop counter variable, one increment point for that variable, and one termination variable. These cannot be altered at any point in the loop nest.
- Loops cannot have multiple entry and exit points. The loop termination must be the only means to exit the loop.
- Dependencies in the loop must not be "backwards-looking". For example, a statement such as  $A[i][j] = A[i - 1][j + 1] + 4$  must not appear within the loop.

For loop unrolling to occur, the **#pragma unrollandfuse** directive must precede a for loop. You must not specify **#pragma unrollandfuse** for the innermost for loop.

You must not specify **#pragma unrollandfuse** more than once, or combine the directive with **#pragma nounrollandfuse**, **#pragma nounroll**, **#pragma unroll**, or **#pragma stream\_unroll** directives for the same for loop.

## Predefined macros

None.

## Examples

In the following example, a **#pragma unrollandfuse** directive replicates and fuses the body of the loop. This reduces the number of cache misses for array b.

```

int i, j;
int a[1000][1000];
int b[1000][1000];
int c[1000][1000];

....

#pragma unrollandfuse(2)
for (i=1; i<1000; i++) {
    for (j=1; j<1000; j++) {
        a[j][i] = b[i][j] * c[j][i];
    }
}

```

The for loop below shows a possible result of applying the **#pragma unrollandfuse(2)** directive to the loop shown above:

```

for (i=1; i<1000; i=i+2) {
    for (j=1; j<1000; j++) {
        a[j][i] = b[i][j] * c[j][i];
        a[j][i+1] = b[i+1][j] * c[j][i+1];
    }
}

```

You can also specify multiple **#pragma unrollandfuse** directives in a nested loop structure.

```

int i, j, k;
int a[1000][1000];
int b[1000][1000];
int c[1000][1000];
int d[1000][1000];
int e[1000][1000];

....

#pragma unrollandfuse(4)
for (i=1; i<1000; i++) {
    #pragma unrollandfuse(2)
    for (j=1; j<1000; j++) {
        for (k=1; k<1000; k++) {
            a[j][i] = b[i][j] * c[j][i] + d[j][k] * e[i][k];
        }
    }
}

```

### Related information

- “-qunroll” on page 363
- “#pragma stream\_unroll” on page 433

## #pragma weak

### Category

Object code control

### Purpose

Prevents the linker from issuing error messages if it encounters a symbol multiply-defined during linking, or if it does not find a definition for a symbol.

The pragma can be used to allow a program to call a user-defined function that has the same name as a library function. By marking the library function definition as "weak", the programmer can reference a "strong" version of the function and cause the linker to accept multiple definitions of a global symbol in the object code. While this pragma is intended for use primarily with functions, it will also work for most data objects.

## Syntax

```
▶▶ #pragma weak name1 [=name2] ▶▶
```

## Parameters

*name1*

A name of a data object or function with external linkage.

*name2*

A name of a data object or function with external linkage.

▶ C++ *name2* must not be a member function. If *name2* is a template function, you must explicitly instantiate the template function.

▶ C++ Names must be specified using their mangled names. To obtain C++ mangled names, compile your source to object files only, using the `-c` compiler option, and use the `nm` operating system command on the resulting object file. You can also use the `c++filt` utility provided by the compiler for a side-by-side listing of source names and mangled names; see "Demangling compiled C++ names" in the *XL C/C++ Optimization and Programming Guide* for details. (See also "Name mangling" in the *XL C/C++ Language Reference* for details on using the `extern "C"` linkage specifier on declarations to prevent name mangling.)

## Usage

There are two forms of the **weak** pragma:

**#pragma weak** *name1*

This form of the pragma marks the definition of the *name1* as "weak" in a given compilation unit. If *name1* is referenced from anywhere in the program, the linker will use the "strong" version of the definition (that is, the definition not marked with **#pragma weak**), if there is one. If there is no strong definition, the linker will use the weak definition; if there are multiple weak definitions, it is unspecified which weak definition the linker will select (typically, it uses the definition found in the first object file specified on the command line during the link step). *name1* must be defined in the same compilation unit as **#pragma weak**.

**#pragma weak** *name1*=*name2*

This form of the pragma creates a weak definition of the *name1* for a given compilation unit, and an alias for *name2*. If *name1* is referenced from anywhere in the program, the linker will use the "strong" version of the definition (that is, the definition not marked with **#pragma weak**), if there is one. If there is no strong definition, the linker will use the weak definition, which resolves to the definition of *name2*. If there are multiple weak definitions, it is unspecified which weak definition the linker will select (typically, it uses the definition found in the first object file specified on the command line during the link step).

*name2* must be defined in the same compilation unit as **#pragma weak**. *name1* may or may not be declared in the same compilation unit as the **#pragma weak**, but must never be defined in the compilation unit. If *name1* is declared in the compilation unit, *name1*'s declaration must be compatible to that of *name2*. For example, if *name2* is a function, *name1* must have the same return and argument types as *name2*.

This pragma should not be used with uninitialized global data, or with shared library data objects that are exported to executables.

## Examples

The following is an example of the **#pragma weak name1** form:

```
// Compilation unit 1:
#include <stdio.h>

void foo();

int main()
{
    foo();
}

// Compilation unit 2:
#include <stdio.h>

#if __cplusplus
#pragma weak foo__Fv
#else
#pragma weak foo
#endif
void foo()
{
    printf("Foo called from compilation unit 2\n");
}

// Compilation unit 3:
#include <stdio.h>

void foo()
{
    printf("Foo called from compilation unit 3\n");
}
```

If all three compilation units are compiled and linked together, the linker will use the strong definition of `foo` in compilation unit 3 for the call to `foo` in compilation unit 1, and the output will be:

```
Foo called from compilation unit 3
```

If only compilation unit 1 and 2 are compiled and linked together, the linker will use the weak definition of `foo` in compilation unit 2, and the output will be:

```
Foo called from compilation unit 2
```

The following is an example of the **#pragma weak name1=name2** form:

```
// Compilation unit 1:
#include <stdio.h>
```



```

void foo();

int main()
{
foo();
}

// Compilation unit 2:

#include <stdio.h>

void foo(); // optional

#ifdef __cplusplus
#pragma weak foo__Fv = foo2__Fv
#else#pragma weak foo = foo2
#endif
void foo2()
{
printf("Hello from foo2!\n");
}

// Compilation unit 3:

#include <stdio.h>

void foo()
{
printf("Hello from foo!\n");
}

```

If all three compilation units are compiled and linked together, the linker will use the strong definition of `foo` in compilation unit 3 for the call to `foo` from compilation unit 1, and the output will be:

```
Hello from foo!
```

If only compilation unit 1 and 2 are compiled and linked together, the linker will use the weak definition of `foo` in compilation unit 2, which is an alias for `foo2`, and the output will be:

```
Hello from foo2!
```

### Related information

- "The weak variable attribute" in the *XL C/C++ Language Reference*
- "The weak function attribute" in the *XL C/C++ Language Reference*
- "#pragma map" on page 411
- "-qweaksymbol" on page 378
- "-qweakexp" on page 377

## Pragma directives for parallel processing

Parallel processing operations are controlled by pragma directives in your program source. The pragmas have effect only when parallelization is enabled with the `-qsmp` compiler option.

You can use IBM SMP or OpenMP directives in C programs, and OpenMP directives in C++ programs. Each has its own usage characteristics.

## #pragma ibm critical (C only)

### Purpose

**Note:** `#pragma ibm critical` has been deprecated and might be removed from the future release. Use “`#pragma omp critical`” on page 457. For more information about deprecated SMP directives, see “Deprecated directives” on page 388.

The **critical** pragma identifies a critical section of program code that must only be run by one process at a time.

### Syntax

```
▶▶ #pragma ibm critical (name) ▶▶
```

where *name* can be used to optionally identify the critical region. Identifiers naming a critical region have external linkage.

### Usage

The compiler reports an error if you try to branch into or out of a critical section. Some situations that will cause an error are:

- A critical section that contains the return statement.
- A critical section that contains goto, continue, or break statements that transfer program flow outside of the critical section.
- A goto statement outside a critical section that transfers program flow to a label defined within a critical section.

## #pragma ibm independent\_calls (C only)

### Description

The **independent\_calls** pragma asserts that specified function calls within the chosen loop have no loop-carried dependencies. This information helps the compiler perform dependency analysis.

### Syntax

```
▶▶ #pragma ibm independent_calls (identifier) ▶▶
```

Where *identifier* is a comma-separated list that represents the name of the functions.

### Usage

*identifier* cannot be the name of a pointer to a function.

If no function identifiers are specified, the compiler assumes that all functions inside the loop are free of carried dependencies.

## #pragma ibm independent\_loop

### Purpose

The `independent_loop` pragma explicitly states that the iterations of the chosen loop are independent, and that the iterations can be executed in parallel.

### Syntax

```
▶▶ #pragma ibm independent_loop _____▶▶  
                                └if exp┘
```

where `exp` represents a scalar expression.

### Usage

If the iterations of a loop are independent, you can put the pragma before the loop block. Then the compiler executes these iterations in parallel. When the `exp` argument is specified, the loop iterations are considered independent only if `exp` evaluates to TRUE at run time.

#### Notes:

- If the iterations of the chosen loop are dependent, the compiler executes the loop iterations sequentially no matter whether you specify the `independent_loop` pragma.
- To have an effect on a loop, you must put the `independent_loop` pragma immediately before this loop. Otherwise, the pragma is ignored.
- If several `independent_loop` pragmas are specified before a loop, only the last one takes effect.

This pragma can be combined with the `omp parallel for` pragma to select a specific parallel process scheduling algorithm. For more information, see “`#pragma omp parallel for`” on page 454.

### Examples

In the following example, the loop iterations are executed in parallel if the value of the argument `k` is larger than 2.

```
int a[1000], b[1000], c[1000];  
int main(int k){  
    if(k>0){  
        #pragma ibm independent_loop if (k>2)  
        for(int i=0; i<900; i++){  
            a[i]=b[i]*c[i];  
        }  
    }  
}
```

## #pragma ibm parallel\_loop (C only)

### Purpose

**Note:** The `parallel_loop` and `schedule` pragmas have been deprecated and might be removed from the future release. Use “`#pragma omp parallel for`” on page 454 with the `schedule` clause. For more information about deprecated SMP directives, see “Deprecated directives” on page 388.

The `parallel_loop` pragma explicitly instructs the compiler to parallelize the chosen loop.

### Syntax

```
▶▶ #pragma ibm parallel_loop [if exp] [schedule (sched-type)] ▶▶
```

where *exp* represents a scalar expression, and *sched-type* represents any scheduling algorithm as valid for the `schedule` directive.

### Usage

When the `if` argument is specified, the loop executes in parallel only if *exp* evaluates to TRUE at run time. Otherwise the loop executes sequentially. The loop will also run sequentially if it is in a critical section.

This pragma can be applied to a wide variety of C loops, and the compiler will try to determine if a loop is countable or not.

Program sections using the `parallel_loop` pragma must be able to produce a correct result in both sequential and parallel mode. For example, loop iterations must be independent before the loop can be parallelized. Explicit parallel programming techniques involving condition synchronization are not permitted.

The compiler will not automatically detect reductions on loops marked with this pragma. To properly parallelize loops with reductions, use:

- `parallel for` and specify reductions explicitly, or,
- `#pragma ibm independent_loop`, which will let the compiler discover the reductions.

This pragma can be combined with the `schedule` pragma to select a specific parallel process scheduling algorithm. For more information, see the description for the “`#pragma ibm schedule (C only)`” on page 443 pragma.

A warning is generated if this pragma is not followed by a countable loop.

### #pragma ibm permutation (C only)

#### Purpose

The `permutation` pragma asserts that on the following loop, different elements of the named arrays are guaranteed to have different values (that is,  $a[i] \neq a[j]$  iff  $i \neq j$ ).

### Syntax

```
▶▶ #pragma ibm permutation ( identifier ) ▶▶
```

where *identifier* represents the name of an array. The *identifier* cannot be a function parameter or the name of a pointer.

## Usage

Pragma must appear immediately before the loop or loop block directive to be affected.

This assertion may enable loop transformations if elements are used to index other arrays. This pragma is useful for programs that deal with sparse data structures.

## #pragma ibm schedule (C only)

### Purpose

**Note:** #pragma ibm schedule and #pragma ibm parallel\_loop have been deprecated and might be removed from the future release. Use “#pragma omp parallel for” on page 454 with the **schedule** clause. For more information about SMP directives, see “Deprecated directives” on page 388.

The **schedule** pragma specifies the scheduling algorithms used for parallel processing.

### Syntax

▶▶ #pragma ibm schedule (*sched-type*) ▶▶

### Parameters

*sched-type* represents one of the following options:

#### affinity

Iterations of a loop are initially divided into local partitions of size **ceiling**(*number\_of\_iterations*/*number\_of\_threads*). Each local partition is then further subdivided into chunks of size **ceiling**(*number\_of\_iterations\_remaining\_in\_partition*/2).

When a thread becomes available, it takes the next chunk from its local partition. If there are no more chunks in the local partition, the thread takes an available chunk from the partition of another thread.

#### affinity,*n*

As above, except that each local partition is subdivided into chunks of size *n*. *n* must be an integral assignment expression of value 1 or greater.

#### dynamic

Iterations of a loop are divided into chunks of size 1.

Chunks are assigned to threads on a first-come, first-serve basis as threads become available. This continues until all work is completed.

#### dynamic,*n*

As above, except that all chunks are set to size *n*. *n* must be an integral assignment expression of value 1 or greater.

#### guided

Chunks are made progressively smaller until a chunk size of one is reached. The first chunk is of size **ceiling**(*number\_of\_iterations*/*number\_of\_threads*). Remaining chunks are of size **ceiling**(*number\_of\_iterations\_remaining*/*number\_of\_threads*).

Chunks are assigned to threads on a first-come, first-serve basis as threads become available. This continues until all work is completed.

**guided,*n***

As above, except the minimum chunk size is set to *n*. *n* must be an integral assignment expression of value 1 or greater.

**runtime**

Scheduling policy is determined at run time.

**static**

Iterations of a loop are divided into chunks of size **ceiling**(*number\_of\_iterations*/*number\_of\_threads*). Each thread is assigned a separate chunk.

This scheduling policy is also known as *block scheduling*.

**static,*n***

Iterations of a loop are divided into chunks of size *n*. Each chunk is assigned to a thread in *round-robin* fashion.

*n* must be an integral assignment expression of value 1 or greater.

**Note:** If *n*=1, iterations of a loop are divided into chunks of size 1 and each chunk is assigned to a thread in *round-robin* fashion. This scheduling policy is also known as *block cyclic scheduling*

**Usage**

Pragma must appear immediately before the loop or loop block directive to be affected.

Scheduling algorithms for parallel processing can be specified using any of the methods shown below. If used, methods higher in the list override entries lower in the list.

- pragma statements
- compiler command line options
- runtime command line options
- runtime default options

Scheduling algorithms can also be specified using the **schedule** argument of the **parallel\_loop** and **independent\_loop** pragma statements. For example, the following sets of statements are equivalent:

```
#pragma ibm parallel_loop
#pragma ibm schedule(sched_type)
<countable for|while|do loop>
and
#pragma ibm parallel_loop(sched_type)
<countable for|while|do loop>
```

If different scheduling types are specified for a given loop, the last one specified is applied.

**#pragma ibm sequential\_loop (C only)****Purpose**

The **sequential\_loop** pragma explicitly instructs the compiler to execute the chosen loop sequentially.

## Syntax

▶▶ #pragma ibm sequential\_loop ▶▶

## Usage

Pragma must appear immediately before the loop or loop block directive to be affected.

This pragma disables automatic parallelization of the chosen loop, and is always respected by the compiler.

## #pragma omp atomic

### Purpose

The **omp atomic** directive allows access of a specific memory location atomically. It ensures that race conditions are avoided through direct control of concurrent threads that might read or write to or from the particular memory location. With the **omp atomic** directive, you can write more efficient concurrent algorithms with fewer locks.

## Syntax

▶▶ #pragma omp atomic { update, read, write, capture } [ expression\_statement, structured\_block ] ▶▶

where *expression\_statement* is an expression statement of scalar type, and *structured\_block* is a structured block of two expression statements.

## Clauses

### update

Updates the value of a variable atomically. Guarantees that only one thread at a time updates the shared variable, avoiding errors from simultaneous writes to the same variable. An **omp atomic** directive without a clause is equivalent to an **omp atomic update**.

**Note:** Atomic updates cannot write arbitrary data to the memory location, but depend on the previous data at the memory location.

### read

Reads the value of a variable atomically. The value of a shared variable can be read safely, avoiding the danger of reading an intermediate value of the variable when it is accessed simultaneously by a concurrent thread.

### write

Writes the value of a variable atomically. The value of a shared variable can be written exclusively to avoid errors from simultaneous writes.

### capture

Updates the value of a variable while capturing the original or final value of the variable atomically.

The *expression\_statement* or *structured\_block* takes one of the following forms, depending on the atomic directive clause:

Directive clause	<i>expression_statement</i>	<i>structured_block</i>
update (equivalent to no clause)	x++; x--; ++x; --x; x binop = expr; x = x binop expr;	
read	v = x;	
write	x = expr;	
capture	v = x++; v = x--; v = ++x; v = --x; v = x binop = expr;	{v = x; x binop = expr;} {v = x; xOP;} {v = x; OPx;} {x binop = expr; v = x;} {xOP; v = x;} {OPx; v = x;} {v = x; x = x binop expr;} {x = x binop expr; v = x;}

where:

*x*, *v* are both lvalue expressions with scalar type.

*expr* is an expression of scalar type that does not reference *x*.

*binop* is one of the following binary operators:

+ \* - / & ^ | << >>

*OP* is one of ++ or --.

**Note:** *binop*, *binop=*, and *OP* are not overloaded operators.

## Usage

Objects that can be updated in parallel and that might be subject to race conditions should be protected with the **omp atomic** directive.

All atomic accesses to the storage locations designated by *x* throughout the program should have a compatible type.

Within an atomic region, multiple syntactic occurrences of *x* must designate the same storage location.

All accesses to a certain storage location throughout a concurrent program must be atomic. A non-atomic access to a memory location might break the expected atomic behavior of all atomic accesses to that storage location.



Neither  $v$  nor  $expr$  can access the storage location that is designated by  $x$ .

Neither  $x$  nor  $expr$  can access the storage location that is designated by  $v$ .

All accesses to the storage location designated by  $x$  are atomic. Evaluations of the expression  $expr$ ,  $v$ ,  $x$  are not atomic.

For atomic capture access, the operation of writing the captured value to the storage location represented by  $v$  is not atomic.

## Examples

### Example 1: Atomic update

```
extern float x[], *p = x, y;

/* Protect against race conditions among multiple updates. */
#pragma omp atomic
x[index[i]] += y;

/* Protect against race conditions with updates through x. */
#pragma omp atomic
p[i] -= 1.0f;
```

### Example 2: Atomic read, write, and update

```
extern int x[10];
extern int f(int);
int temp[10], i;

for(i = 0; i < 10; i++)
{
    #pragma omp atomic read
    temp[i] = x[f(i)];

    #pragma omp atomic write
    x[i] = temp[i]*2;

    #pragma omp atomic update
    x[i] *= 2;
}
```

### Example 3: Atomic capture

```
extern int x[10];
extern int f(int);
int temp[10], i;

for(i = 0; i < 10; i++)
{
    #pragma omp atomic capture
    temp[i] = x[f(i)]++;

    #pragma omp atomic capture
    {
        temp[i] = x[f(i)]; //the two occurrences of x[f(i)] must evaluate to the
        x[f(i)] -= 3; //same memory location, otherwise behavior is undefined.
    }
}
```

## #pragma omp parallel Purpose

The **omp parallel** directive explicitly instructs the compiler to parallelize the chosen block of code.

## Syntax

```
▶▶ #pragma omp parallel clause ...
```

## Parameters

*clause* is any of the following clauses:

### **if** (*exp*)

When the **if** argument is specified, the program code executes in parallel only if the scalar expression represented by *exp* evaluates to a nonzero value at run time. Only one **if** clause can be specified.

### **private** (*list*)

Declares the scope of the data variables in *list* to be private to each thread. Data variables in *list* are separated by commas.

### **firstprivate** (*list*)

Declares the scope of the data variables in *list* to be private to each thread. Each new private object is initialized with the value of the original variable as if there was an implied declaration within the statement block. Data variables in *list* are separated by commas.

### **num\_threads** (*int\_exp*)

The value of *int\_exp* is an integer expression that specifies the number of threads to use for the parallel region. If dynamic adjustment of the number of threads is also enabled, then *int\_exp* specifies the maximum number of threads to be used.

### **shared** (*list*)

Declares the scope of the comma-separated data variables in *list* to be shared across all threads.

### **default** (**shared** | **none**)

Defines the default data scope of variables in each thread. Only one **default** clause can be specified on an **omp parallel** directive.

Specifying **default(shared)** is equivalent to stating each variable in a **shared(list)** clause.

Specifying **default(none)** requires that each data variable visible to the parallelized statement block must be explicitly listed in a data scope clause, with the exception of those variables that are:

- const-qualified,
- specified in an enclosed data scope attribute clause, or,
- used as a loop control variable referenced only by a corresponding **omp for** or **omp parallel for** directive.

### **copyin** (*list*)

For each data variable specified in *list*, the value of the data variable in the master thread is copied to the thread-private copies at the beginning of the parallel region. Data variables in *list* are separated by commas.

Each data variable specified in the **copyin** clause must be a **threadprivate** variable.

### **reduction** (*operator*: *list*)

Performs a reduction on all scalar variables in *list* using the specified *operator*. Reduction variables in *list* are separated by commas.

A private copy of each variable in *list* is created for each thread. At the end of the statement block, the final values of all private copies of the reduction variable are combined in a manner appropriate to the operator, and the result is placed back in the original value of the shared reduction variable. For example, when the max operator is specified, the original reduction variable value combines with the final values of the private copies by using the following expression:

```
original_reduction_variable = original_reduction_variable < private_copy ?
private_copy : original_reduction_variable;
```

For variables specified in the **reduction** clause, they must satisfy the following conditions:

- Must be of a type appropriate to the operator. If the max or min operator is specified, the variables must be one of the following types with or without long, short, signed, or unsigned:
  - `C` `_Bool` `C`
  - `C++` `bool` `C++`
  - `char`
  - `C++` `wchar_t` `C++`
  - `int`
  - `float`
  - `double`
- Must be shared in the enclosing context.
- Must not be const-qualified.
- Must not have pointer type.

## Usage

When a parallel region is encountered, a logical team of threads is formed. Each thread in the team executes all statements within a parallel region except for work-sharing constructs. Work within work-sharing constructs is distributed among the threads in a team.

Loop iterations must be independent before the loop can be parallelized. An implied barrier exists at the end of a parallelized statement block.


Nested parallel regions are always serialized.

## #pragma omp for

### Purpose

The **omp for** directive instructs the compiler to distribute loop iterations within the team of threads that encounters this work-sharing construct.

### Syntax



### Parameters

*clause* is any of the following clauses:

### **collapse** (*n*)

Allows you to parallelize multiple loops in a nest without introducing nested parallelism.

▶▶—COLLAPSE—(—*n*—)—————▶▶

- Only one collapse clause is allowed on a worksharing **for** or **parallel for** pragma.
- The specified number of loops must be present lexically. That is, none of the loops can be in a called subroutine.
- The loops must form a rectangular iteration space and the bounds and stride of each loop must be invariant over all the loops.
- If the loop indices are of different size, the index with the largest size will be used for the collapsed loop.
- The loops must be perfectly nested; that is, there is no intervening code nor any OpenMP pragma between the loops which are collapsed.
- The associated do-loops must be structured blocks. Their execution must not be terminated by an **break** statement.
- If multiple loops are associated to the loop construct, only an iteration of the innermost associated loop may be curtailed by a **continue** statement. If multiple loops are associated to the loop construct, there must be no branches to any of the loop termination statements except for the innermost associated loop.

### **Ordered construct**

During execution of an iteration of a loop or a loop nest within a loop region, the executing thread must not execute more than one ordered region which binds to the same loop region. As a consequence, if multiple loops are associated to the loop construct by a collapse clause, the ordered construct has to be located inside all associated loops.

### **Lastprivate clause**

When a lastprivate clause appears on the pragma that identifies a work-sharing construct, the value of each new list item from the sequentially last iteration of the associated loops, is assigned to the original list item even if a collapse clause is associated with the loop

### **Other SMP and performance pragmas**

**stream\_unroll**, **unroll**, **unrollandfuse**, **nounrollandfuse** pragmas cannot be used for any of the loops associated with the **collapse** clause loop nest.

### **private** (*list*)

Declares the scope of the data variables in *list* to be private to each thread. Data variables in *list* are separated by commas.

### **firstprivate** (*list*)

Declares the scope of the data variables in *list* to be private to each thread. Each new private object is initialized as if there was an implied declaration within the statement block. Data variables in *list* are separated by commas.

### **lastprivate** (*list*)

Declares the scope of the data variables in *list* to be private to each thread. The final value of each variable in *list*, if assigned, will be the value assigned to that variable in the last iteration. Variables not assigned a value will have an indeterminate value. Data variables in *list* are separated by commas.

### **reduction** (*operator: list*)

Performs a reduction on all scalar variables in *list* using the specified *operator*. Reduction variables in *list* are separated by commas.

A private copy of each variable in *list* is created for each thread. At the end of the statement block, the final values of all private copies of the reduction variable are combined in a manner appropriate to the operator, and the result is placed back in the original value of the shared reduction variable. For example, when the max operator is specified, the original reduction variable value combines with the final values of the private copies by using the following expression:

```
original_reduction_variable = original_reduction_variable < private_copy ?  
private_copy : original_reduction_variable;
```

For variables specified in the **reduction** clause, they must satisfy the following conditions:

- Must be of a type appropriate to the operator. If the max or min operator is specified, the variables must be one of the following types with or without long, short, signed, or unsigned:

- `<C> _Bool <C>`
- `<C++> bool <C++>`
- `char`
- `<C++> wchar_t <C++>`
- `int`
- `float`
- `double`

- Must be shared in the enclosing context.
- Must not be const-qualified.
- Must not have pointer type.

### **ordered**

Specify this clause if an ordered construct is present within the dynamic extent of the **omp for** directive.

### **schedule** (*type*)

Specifies how iterations of the **for** loop are divided among available threads. Acceptable values for *type* are:

**auto** With **auto**, scheduling is delegated to the compiler and runtime system. The compiler and runtime system can choose any possible mapping of iterations to threads (including all possible valid schedules) and these may be different in different loops.

### **dynamic**

Iterations<sup>®</sup> of a loop are divided into chunks of size **ceiling**(*number\_of\_iterations*/*number\_of\_threads*).

Chunks are dynamically assigned to active threads on a "first-come, first-do" basis until all work has been assigned.

### **dynamic,*n***

As above, except chunks are set to size *n*. *n* must be an integral assignment expression of value 1 or greater.

### **guided**

Chunks are made progressively smaller until the default minimum chunk size is reached. The first chunk is of size **ceiling**(*number\_of\_iterations*/*number\_of\_threads*). Remaining chunks are of size **ceiling**(*number\_of\_iterations\_left*/*number\_of\_threads*).

The minimum chunk size is 1.

Chunks are assigned to active threads on a "first-come, first-do" basis until all work has been assigned.

**guided,*n***

As above, except the minimum chunk size is set to *n*; *n* must be an integral assignment expression of value 1 or greater.

**runtime**

Scheduling policy is determined at run time. Use the OMP\_SCHEDULE environment variable to set the scheduling type and chunk size.

**static**

Iterations of a loop are divided into chunks of size **ceiling**(*number\_of\_iterations/number\_of\_threads*). Each thread is assigned a separate chunk.

This scheduling policy is also known as *block scheduling*.

**static,*n***

Iterations of a loop are divided into chunks of size *n*. Each chunk is assigned to a thread in *round-robin* fashion.

*n* must be an integral assignment expression of value 1 or greater.

This scheduling policy is also known as *block cyclic scheduling*.

**Note:** if *n*=1, iterations of a loop are divided into chunks of size 1 and each chunk is assigned to a thread in *round-robin* fashion. This scheduling policy is also known as *block cyclic scheduling*.

**nowait**

Use this clause to avoid the implied **barrier** at the end of the **for** directive. This is useful if you have multiple independent work-sharing sections or iterative loops within a given parallel region. Only one **nowait** clause can appear on a given **for** directive.

and where *for\_loop* is a for loop construct with the following canonical shape:

```
for (init_expr; exit_cond; incr_expr)  
  statement
```

where:

<i>init_expr</i>	takes the form:	<i>iv</i> = <i>b</i> <i>integer-type iv</i> = <i>b</i>
<i>exit_cond</i>	takes the form:	<i>iv</i> <= <i>ub</i> <i>iv</i> < <i>ub</i> <i>iv</i> >= <i>ub</i> <i>iv</i> > <i>ub</i>
<i>incr_expr</i>	takes the form:	++ <i>iv</i> <i>iv</i> ++ -- <i>iv</i> <i>iv</i> -- <i>iv</i> += <i>incr</i> <i>iv</i> -= <i>incr</i> <i>iv</i> = <i>iv</i> + <i>incr</i> <i>iv</i> = <i>incr</i> + <i>iv</i> <i>iv</i> = <i>iv</i> - <i>incr</i>

and where:

<i>iv</i>	Iteration variable. The iteration variable must be a signed integer not modified anywhere within the for loop. It is implicitly made private for the duration of the for operation. If not specified as <b>lastprivate</b> , the iteration variable will have an indeterminate value after the operation completes.
<i>b, ub, incr</i>	Loop invariant signed integer expressions. No synchronization is performed when evaluating these expressions and evaluated side effects may result in indeterminate values.

## Usage

This pragma must appear immediately before the loop or loop block directive to be affected.

Program sections using the **omp for** pragma must be able to produce a correct result regardless of which thread executes a particular iteration. Similarly, program correctness must not rely on using a particular scheduling algorithm.

The for loop iteration variable is implicitly made private in scope for the duration of loop execution. This variable must not be modified within the body of the for loop. The value of the increment variable is indeterminate unless the variable is specified as having a data scope of **lastprivate**.

An implicit barrier exists at the end of the for loop unless the **nowait** clause is specified.

### Restrictions:

- The for loop must be a structured block, and must not be terminated by a break statement.
- Values of the loop control expressions must be the same for all iterations of the loop.
- An **omp for** directive can accept only one **schedule** clause.
- The value of *n* (chunk size) must be the same for all threads of a parallel region.

## #pragma omp ordered

### Purpose

The **omp ordered** directive identifies a structured block of code that must be executed in sequential order.

### Syntax

▶▶ #pragma omp ordered ◀◀

## Usage

The **omp ordered** directive must be used as follows:

- It must appear within the extent of a **omp for** or **omp parallel for** construct containing an **ordered** clause.
- It applies to the statement block immediately following it. Statements in that block are executed in the same order in which iterations are executed in a sequential loop.


- An iteration of a loop must not execute the same **omp ordered** directive more than once.
- An iteration of a loop must not execute more than one distinct **omp ordered** directive.

## #pragma omp parallel for

### Purpose

The **omp parallel for** directive effectively combines the **omp parallel** and **omp for** directives. This directive lets you define a parallel region containing a single **for** directive in one step.

### Syntax



### Usage

With the exception of the **nowait** clause, clauses and restrictions described in the **omp parallel** and **omp for** directives also apply to the **omp parallel for** directive.

## #pragma omp section, #pragma omp sections

### Purpose

The **omp sections** directive distributes work among threads bound to a defined parallel region.

### Syntax



### Parameters

*clause* is any of the following clauses:

#### **private** (*list*)

Declares the scope of the data variables in *list* to be private to each thread. Data variables in *list* are separated by commas.

#### **firstprivate** (*list*)

Declares the scope of the data variables in *list* to be private to each thread. Each new private object is initialized as if there was an implied declaration within the statement block. Data variables in *list* are separated by commas.

#### **lastprivate** (*list*)

Declares the scope of the data variables in *list* to be private to each thread. The final value of each variable in *list*, if assigned, will be the value assigned to that variable in the last **section**. Variables not assigned a value will have an indeterminate value. Data variables in *list* are separated by commas.



### **reduction** (*operator: list*)

Performs a reduction on all scalar variables in *list* using the specified *operator*. Reduction variables in *list* are separated by commas.

A private copy of each variable in *list* is created for each thread. At the end of the statement block, the final values of all private copies of the reduction variable are combined in a manner appropriate to the operator, and the result is placed back in the original value of the shared reduction variable. For example, when the max operator is specified, the original reduction variable value combines with the final values of the private copies by using the following expression:

```
original_reduction_variable = original_reduction_variable < private_copy ?  
private_copy : original_reduction_variable;
```

For variables specified in the **reduction** clause, they must satisfy the following conditions:

- Must be of a type appropriate to the operator. If the max or min operator is specified, the variables must be one of the following types with or without long, short, signed, or unsigned:

- `_Bool`
- `bool`
- `char`
- `wchar_t`
- `int`
- `float`
- `double`

- Must be shared in the enclosing context.
- Must not be const-qualified.
- Must not have pointer type.

### **nowait**

Use this clause to avoid the implied **barrier** at the end of the **sections** directive. This is useful if you have multiple independent work-sharing sections within a given parallel region. Only one **nowait** clause can appear on a given **sections** directive.

## Usage


The **omp section** directive is optional for the first program code segment inside the **omp sections** directive. Following segments must be preceded by an **omp section** directive. All **omp section** directives must appear within the lexical construct of the program source code segment associated with the **omp sections** directive.

When program execution reaches a **omp sections** directive, program segments defined by the following **omp section** directive are distributed for parallel execution among available threads. A barrier is implicitly defined at the end of the larger program region associated with the **omp sections** directive unless the **nowait** clause is specified.

## #pragma omp parallel sections Purpose

The **omp parallel sections** directive effectively combines the **omp parallel** and **omp sections** directives. This directive lets you define a parallel region containing a single **sections** directive in one step.

## Syntax



## Usage


All clauses and restrictions described in the **omp parallel** and **omp sections** directives apply to the **omp parallel sections** directive.

## #pragma omp single

### Purpose

The **omp single** directive identifies a section of code that must be run by a single available thread.

## Syntax



## Parameters

*clause* is any of the following:

### **private** (*list*)

Declares the scope of the data variables in *list* to be private to each thread. Data variables in *list* are separated by commas.

A variable in the **private** clause must not also appear in a **copyprivate** clause for the same **omp single** directive.

### **copyprivate** (*list*)

Broadcasts the values of variables specified in *list* from one member of the team to other members. This occurs after the execution of the structured block associated with the **omp single** directive, and before any of the threads leave the barrier at the end of the construct. For all other threads in the team, each variable in the *list* becomes defined with the value of the corresponding variable in the thread that executed the structured block. Data variables in *list* are separated by commas. Usage restrictions for this clause are:

- A variable in the **copyprivate** clause must not also appear in a **private** or **firstprivate** clause for the same **omp single** directive.
- If an **omp single** directive with a **copyprivate** clause is encountered in the dynamic extent of a parallel region, all variables specified in the **copyprivate** clause must be private in the enclosing context.
- Variables specified in **copyprivate** clause within dynamic extent of a parallel region must be private in the enclosing context.
- A variable that is specified in the **copyprivate** clause must have an accessible and unambiguous copy assignment operator.

- The **copyprivate** clause must not be used together with the **nowait** clause.

### **firstprivate** (*list*)

Declares the scope of the data variables in *list* to be private to each thread. Each new private object is initialized as if there was an implied declaration within the statement block. Data variables in *list* are separated by commas.

A variable in the **firstprivate** clause must not also appear in a **copyprivate** clause for the same **omp single** directive.

### **nowait**

Use this clause to avoid the implied **barrier** at the end of the **single** directive. Only one **nowait** clause can appear on a given **single** directive. The **nowait** clause must not be used together with the **copyprivate** clause.

## **Usage**

An implied barrier exists at the end of a parallelized statement block unless the **nowait** clause is specified.

## **#pragma omp master**

### **Purpose**

The **omp master** directive identifies a section of code that must be run only by the master thread.

### **Syntax**

```
▶▶ #pragma omp master ◀◀
```

## **Usage**

Threads other than the master thread will not execute the statement block associated with this construct.

No implied barrier exists on either entry to or exit from the master section.

## **#pragma omp critical**

### **Purpose**

The **omp critical** directive identifies a section of code that must be executed by a single thread at a time.

### **Syntax**

```
▶▶ #pragma omp critical (name) ◀◀
```

where *name* can optionally be used to identify the critical region. Identifiers naming a critical region have external linkage and occupy a namespace distinct from that used by ordinary identifiers.

## Usage

A thread waits at the start of a critical region identified by a given name until no other thread in the program is executing a critical region with that same name. Critical sections not specifically named by **omp critical** directive invocation are mapped to the same unspecified name.

## #pragma omp barrier

### Purpose

The **omp barrier** directive identifies a synchronization point at which threads in a parallel region will wait until all other threads in that section reach the same point. Statement execution past the **omp barrier** point then continues in parallel.

### Syntax

```
▶▶ #pragma omp barrier ▶▶
```

## Usage

The **omp barrier** directive must appear within a block or compound statement. For example:

```
if (x!=0) {  
    #pragma omp barrier    /* valid usage    */  
}  
if (x!=0)  
    #pragma omp barrier    /* invalid usage */
```

## #pragma omp flush

### Purpose

The **omp flush** directive identifies a point at which the compiler ensures that all threads in a parallel region have the same view of specified objects in memory.

### Syntax

```
▶▶ #pragma omp flush [list] ▶▶
```

where *list* is a comma-separated list of variables that will be synchronized.

## Usage

If *list* includes a pointer, the pointer is flushed, not the object being referred to by the pointer. If *list* is not specified, all shared objects are synchronized except those inaccessible with automatic storage duration.

An implied **flush** directive appears in conjunction with the following directives:

- **omp barrier**
- Entry to and exit from **omp critical**.
- Exit from **omp parallel**.

- Exit from **omp for**.
- Exit from **omp sections**.
- Exit from **omp single**.

The **omp flush** directive must appear within a block or compound statement. For example:

```
if (x!=0) {
    #pragma omp flush    /* valid usage    */
}
if (x!=0)
    #pragma omp flush    /* invalid usage */
```

## #pragma omp threadprivate

### Purpose

The **omp threadprivate** directive makes the named file-scope, namespace-scope, or static block-scope variables private to a thread.

### Syntax

```
▶ #pragma omp threadprivate (identifier) ▶
```

where *identifier* is a file-scope, name space-scope or static block-scope variable.

### Usage

Each copy of an **omp threadprivate** data variable is initialized once prior to first use of that copy. If an object is changed before being used to initialize a **threadprivate** data variable, behavior is unspecified.

A thread must not reference another thread's copy of an **omp threadprivate** data variable. References will always be to the master thread's copy of the data variable when executing serial and master regions of the program.

Use of the **omp threadprivate** directive is governed by the following points:

- An **omp threadprivate** directive must appear at file scope outside of any definition or declaration.
- The **omp threadprivate** directive is applicable to static-block scope variables and may appear in lexical blocks to reference those block-scope variables. The directive must appear in the scope of the variable and not in a nested scope, and must precede all references to variables in its list.
- A data variable must be declared with file scope prior to inclusion in an **omp threadprivate** directive *list*.
- An **omp threadprivate** directive and its *list* must lexically precede any reference to a data variable found in that *list*.
- A data variable specified in an **omp threadprivate** directive in one translation unit must also be specified as such in all other translation units in which it is declared.
- Data variables specified in an **omp threadprivate** *list* must not appear in any clause other than the **copyin**, **copyprivate**, **if**, **num\_threads**, and **schedule** clauses.

- The address of a data variable in an **omp threadprivate list** is not an address constant.
- A data variable specified in an **omp threadprivate list** must not have an incomplete or reference type.

## #pragma omp task

### Purpose

The **task** pragma can be used to explicitly define a task.

Use the **task** pragma when you want to identify a block of code to be executed in parallel with the code outside the task region. The **task** pragma can be useful for parallelizing irregular algorithms such as pointer chasing or recursive algorithms. The **task** directive takes effect only if you specify the **-qsmp** compiler option.

### Syntax



### Parameters

The *clause* parameter can be any of the following types of clauses:

#### default (shared | none)

Defines the default data scope of variable in each task. Only one default clause can be specified on an **omp task** directive.

Specifying **default(shared)** is equivalent to stating each variable in a **shared(list)** clause.

Specifying **default(none)** requires that each data variable visible to the construct must be explicitly listed in a data scope clause, with the exception of variables with the following attributes:

- Threadprivate
- Automatic and declared in a scope inside the construct
- Objects with dynamic storage duration
- Static data members
- The loop iteration variables in the associated for-loops for a work-sharing **for** or **parallel for** construct
- Static and declared in a scope inside the construct

#### final (exp)

If you specify a **final** clause and *exp* evaluates to a nonzero value, the generated task is a final task. All task constructs encountered inside a final task create final and included tasks.

You can specify only one **final** clause on the **task** pragma.

#### firstprivate (list)

Declares the scope of the data variables in *list* to be private to each thread. Each new private object is initialized with the value of the original variable as if there was an implied declaration within the statement block. Data variables in *list* are separated by commas.

**if** (*exp*)

When the `if` clause is specified, an undeferred task is generated if the scalar expression *exp* evaluates to a nonzero value. Only one `if` clause can be specified.

**mergeable**

If you specify a `mergeable` clause and the generated task is an undeferred task or included task, a merged task might be generated.

**private** (*list*)

Declares the scope of the data variables in *list* to be private to each thread. Data variables in *list* are separated by commas.

**shared** (*list*)

Declares the scope of the comma-separated data variables in *list* to be shared across all threads.

**untied**

When a task region is suspended, untied tasks can be resumed by any thread in a team. The `untied` clause on a task construct is ignored if either of the following conditions is a nonzero value:

- A `final` clause is specified on the same task construct and the `final` clause expression evaluates to a nonzero value.
- The task is an included task.

## Usage

A final task is a task that makes all its child tasks become final and included tasks. A final task is generated when either of the following conditions is a nonzero value:

- A `final` clause is specified on a task construct and the `final` clause expression evaluates to nonzero value.
- The generated task is a child task of a final task.

An undeferred task is a task whose execution is not deferred with respect to its generating task region. In other words, the generating task region is suspended until the undeferred task has finished running. An undeferred task is generated when an `if` clause is specified on a task construct and the `if` clause expression evaluates to zero.

An included task is a task whose execution is sequentially included in the generating task region. In other words, an included task is undeferred and executed immediately by the encountering thread. An included task is generated when the generated task is a child task of a final task.

A merged task is a task that has the same data environment as that of its generating task region. A merged task might be generated when both the following conditions nonzero values:

- A `mergeable` clause is specified on a task construct.
- The generated task is an undeferred task or an included task.

The `if` clause expression and the `final` clause expression are evaluated outside of the task construct, and the evaluation order is not specified.

Related reference:

"#pragma omp taskwait"

## #pragma omp taskyield

### Purpose

The **omp taskyield** pragma instructs the compiler to suspend the current task in favor of running a different task. The **taskyield** region includes an explicit task scheduling point in the current task region.

### Syntax

▶▶ #pragma omp taskyield ◀◀

## #pragma omp taskwait

### Purpose

Use the **taskwait** pragma to specify a *wait* for child tasks to be completed that are generated by the current task.

### Syntax

▶▶ #pragma omp taskwait ◀◀

Related reference:

"#pragma omp task" on page 460



---

## Chapter 6. Compiler predefined macros

Predefined macros can be used to conditionally compile code for specific compilers, specific versions of compilers, specific environments and/or specific language features.

Predefined macros fall into several categories:

- “General macros”
- “Macros related to the platform” on page 465
- “Macros related to compiler features” on page 465

“Examples of predefined macros” on page 476 show how you can use them in your code.

---

### General macros

The following predefined macros are always predefined by the compiler. Unless noted otherwise, all the following macros are *protected*, which means that the compiler will issue a warning if you try to undefine or redefine them.

Table 44. General predefined macros

Predefined macro name	Description	Predefined value
<code>__BASE_FILE__</code>	Indicates the name of the primary source file.	The fully qualified file name of the primary source file.
<code>__DATE__</code>	Indicates the date that the source file was preprocessed.	A character string containing the date when the source file was preprocessed.
<code>__FILE__</code>	Indicates the name of the preprocessed source file.	A character string containing the name of the preprocessed source file.
<code>__FUNCTION__</code>	Indicates the name of the function currently being compiled.	A character string containing the name of the function currently being compiled.
<code>__LINE__</code>	Indicates the current line number in the source file.	An integer constant containing the line number in the source file.
<code>__SIZE_TYPE__</code>	Indicates the underlying type of <code>size_t</code> on the current platform. Not protected.	unsigned int in 32-bit compilation mode and unsigned long in 64-bit compilation mode.
<code>__TIME__</code>	Indicates the time that the source file was preprocessed.	A character string containing the time when the source file was preprocessed.

Table 44. General predefined macros (continued)

Predefined macro name	Description	Predefined value
<code>__TIMESTAMP__</code>	Indicates the date and time when the source file was last modified. The value changes as the compiler processes any include files that are part of your source program.	A character string literal in the form " <i>Day Mmm dd hh:mm:ss yyyy</i> ", where::  <i>Day</i> Represents the day of the week (Mon, Tue, Wed, Thu, Fri, Sat, or Sun).  <i>Mmm</i> Represents the month in an abbreviated form (Jan, Feb, Mar, Apr, May, Jun, Jul, Aug, Sep, Oct, Nov, or Dec).  <i>dd</i> Represents the day. If the day is less than 10, the first d is a blank character.  <i>hh</i> Represents the hour.  <i>mm</i> Represents the minutes.  <i>ss</i> Represents the seconds.  <i>yyyy</i> Represents the year.

## Macros indicating the XL C/C++ compiler product

Macros related to the XL C/C++ compiler are always predefined, and are protected (the compiler issues a warning if you try to undefine or redefine them). You can use the `-qshowmacros=pre -E` compiler options to view the values of the predefined macros.

Table 45. Compiler product predefined macros




Predefined macro name	Description	Predefined value
 <code>__IBMC__</code>	Indicates the level of the XL C compiler.	An integer in the format <i>VRM</i> , where :  <i>V</i> Represents the version number <i>R</i> Represents the release number <i>M</i> Represents the modification number
 <code>__IBMCP__</code>	Indicates the level of the XL C++ compiler.	An integer in the format <i>VRM</i> , where :  <i>V</i> Represents the version number <i>R</i> Represents the release number <i>M</i> Represents the modification number
 <code>__xlC__</code>	Indicates the level of the XL C compiler.	A string in the format " <i>V.R.M.F</i> ", where:  <i>V</i> Represents the version number <i>R</i> Represents the release number <i>M</i> Represents the modification number <i>F</i> Represents the fix level
<code>__xlC__</code>	Indicates the VR level of the XL C and XL C++ compilers in hexadecimal format. Using the XL C compiler also automatically defines this macro.	A four-digit hexadecimal integer in the format <i>0xVVRR</i> , where:  <i>V</i> Represents the version number <i>R</i> Represents the release number

Table 45. Compiler product predefined macros (continued)

Predefined macro name	Description	Predefined value
<code>__xlC_ver__</code>	Indicates the MF level of the XL C and XL C++ compilers in hexadecimal format. Using the XL C compiler also automatically defines this macro.	<p>A eight-digit hexadecimal integer in the format <code>0x0000MMFF</code>, where:</p> <p><i>M</i> Represents the modification number</p> <p><i>F</i> Represents the fix level</p> <p>For example, in XL C/C++ V12.1, PTF 10.1.0.3, the value of the macro is <code>0x00000003</code>.</p>

## Macros related to the platform

The following predefined macros are provided to facilitate porting applications between platforms. All platform-related predefined macros are unprotected and can be undefined or redefined without warning unless otherwise specified.

Table 46. Platform-related predefined macros

Predefined macro name	Description	Predefined value	Predefined under the following conditions
<code>_BIG_ENDIAN</code> , <code>__BIG_ENDIAN__</code>	Indicates that the platform is big-endian (that is, the most significant byte is stored at the memory location with the lowest address).	1	Always predefined.
<code>__powerpc</code> , <code>__powerpc__</code>	Indicates that the target is a Power architecture.	1	Predefined when the target is a Power architecture.
<code>__PPC</code> , <code>__PPC__</code>	Indicates that the target is a Power architecture.	1	Predefined when the target is a Power architecture.
<code>__unix</code> , <code>__unix__</code>	Indicates that the operating system is a variety of UNIX.	1	Always predefined.

## Macros related to compiler features

Feature-related macros are predefined according to the setting of specific compiler options or pragmas. Unless noted otherwise, all feature-related macros are protected (the compiler will issue a warning if you try to undefine or redefine them).

Feature-related macros are discussed in the following sections:

- “Macros related to compiler option settings”
- “Macros related to architecture settings” on page 469
- “Macros related to language levels” on page 470

## Macros related to compiler option settings

The following macros can be tested for various features, including source input characteristics, output file characteristics, and optimization. All of these macros are predefined by a specific compiler option or suboption, or any invocation or pragma that implies that suboption. If the suboption enabling the feature is not in effect, then the macro is undefined.

Table 47. General option-related predefined macros


Predefined macro name	Description	Predefined value	Predefined when the following compiler option or equivalent pragma is in effect:
<code>__ALTIVEC__</code>	Indicates support for vector data types. (unprotected)	1	<code>-qaltivec</code>
<code>__64BIT__</code>	Indicates that 64-bit compilation mode is in effect.	1	<code>-q64</code>
<code>__CHAR_SIGNED, __CHAR_SIGNED__</code>	Indicates that the default character type is signed char.	1	<code>-qchars=signed</code>
<code>__CHAR_UNSIGNED, __CHAR_UNSIGNED__</code>	Indicates that the default character type is unsigned char.	1	<code>-qchars=unsigned</code>
<code>__DEBUG_ALLOC__</code>	Indicates that debug versions of the standard memory management functions are being used.	1	<code>-qheapdebug</code>
 <code>__EXCEPTIONS</code>	Indicates that C++ exception handling is enabled.	1	<code>-qeh</code>
<code>__IBM_GCC_ASM</code>	Indicates support for GCC inline asm statements.	1	 <code>-qasm=gcc</code> and <code>-qlanglvl=extc99   extc89   extended</code> or <code>-qkeyword=asm</code>  <code>-qasm=gcc</code> and <code>-qlanglvl=extended</code>
		0	 <code>-qnoasm</code> and <code>-qlanglvl=extc99   extc89   extended</code> or <code>-qkeyword=asm</code>  <code>-qnoasm</code> and <code>-qlanglvl=extended</code>

Table 47. General option-related predefined macros (continued)






Predefined macro name	Description	Predefined value	Predefined when the following compiler option or equivalent pragma is in effect:
 <code>__IBM_STDCPP_ASM</code>	Indicates that support for GCC inline asm statements is disabled.	0	<code>-qnoasm=stdcpp</code> or <code>-qlanglvl=strict98   compat366</code>
<code>__IBM_DFP__</code>	Indicates support for decimal floating-point types.	1	<code>-qdfp</code>
<code>__IBM_DFP_SW_EMULATION__</code>	Indicates that decimal floating-point computations are implemented through software emulation rather than in hardware instructions.	1	<code>-qfloat=dfpemulate</code>
 <code>__IBMSMP</code>	Indicates that IBM SMP directives are recognized.	1	<code>-qsmp</code>
<code>__IBM_UTF_LITERAL</code>	Indicates support for UTF-16 and UTF-32 string literals.	1	<code>LANGLVL(EXTENDED)</code>  <code>LANGLVL(EXTENDED0X)</code>
 <code>__IGNERRNO__</code>	Indicates that system calls do not modify <code>errno</code> , thereby enabling certain compiler optimizations.	1	<code>-qignerrno</code>
 <code>__INITAUTO__</code>	Indicates the value to which automatic variables which are not explicitly initialized in the source program are to be initialized.	The two-digit hexadecimal value specified in the <code>-qinitauto</code> compiler option.	<code>-qinitauto=hex value</code>

Table 47. General option-related predefined macros (continued)









Predefined macro name	Description	Predefined value	Predefined when the following compiler option or equivalent pragma is in effect:
 <code>__INITAUTO_W__</code>	Indicates the value to which automatic variables which are not explicitly initialized in the source program are to be initialized.	An eight-digit hexadecimal corresponding to the value specified in the <b>-qinitauto</b> compiler option repeated 4 times.	<code>-qinitauto=hex value</code>
 <code>__LIBANSI__</code>	Indicates that calls to functions whose names match those in the C Standard Library are in fact the C library functions, enabling certain compiler optimizations.	1	<code>-qlibansi</code>
<code>__LONGDOUBLE64</code>	Indicates that the size of a long double type is 64 bits.	1	<b>-qnoldbl128</b>
<code>__LONGDOUBLE128</code>	Indicates that the size of a long double type is 128 bits.	1	<code>-qldbl128</code>
 <code>__OBJECT_MODEL_CLASSIC__</code>	Indicates that the "classic" object model is in effect.	1	<code>-qobjmodel=classic</code>
 <code>__OBJECT_MODEL_IBM__</code>	Indicates that the IBM object is in effect.	1	<code>-qobjmodel=ibm</code>
<code>__OPTIMIZE__</code>	Indicates the level of optimization in effect.	2	<code>-O   -O2</code>
		3	<code>-O3   -O4   -O5</code>
<code>__OPTIMIZE_SIZE__</code>	Indicates that optimization for code size is in effect.	1	<code>-O   -O2   -O3   -O4   -O5</code> and <code>-qcompact</code>

Table 47. General option-related predefined macros (continued)

Predefined macro name	Description	Predefined value	Predefined when the following compiler option or equivalent pragma is in effect:
 <code>__RTTI_DYNAMIC_CAST__</code>	Indicates that runtime type identification information for the <code>dynamic_cast</code> operator is generated.	1	<code>-qrtti</code>   <code>-qrtti=all</code>   <code>dynamiccast</code>
 <code>__RTTI_TYPE_INFO__</code>	Indicates that runtime type identification information for the <code>typeid</code> operator is generated.	1	<code>-qrtti</code>   <code>-qrtti=all</code>   <code>typeid</code>
 <code>__NO_RTTI__</code>	Indicates that runtime type identification information is disabled.	1	<code>-qnortti</code>
 <code>__TEMPINC__</code>	Indicates that the compiler is using the template-implementation file method of resolving template functions.	1	<code>-qtempinc</code>
<code>__VEC__</code>	Indicates support for vector data types.	10205	<code>-qaltivec</code>

## Macros related to architecture settings

The following macros can be tested for target architecture settings. All of these macros are predefined to a value of 1 by a `-qarch` compiler option setting, or any other compiler option that implies that setting. If the `-qarch` suboption enabling the feature is not in effect, then the macro is undefined.

Table 48. `-qarch`-related macros

Macro name	Description	Predefined by the following <code>-qarch</code> suboptions
<code>_ARCH_403</code>	Indicates that the application is targeted to run on the PowerPC 403 processor.	403
<code>_ARCH_604</code>	Indicates that the application is targeted to run on the PowerPC 604 processor.	604
<code>_ARCH_COM</code>	Indicates that the application is targeted to run on any Power processor.	Defined for all <code>-qarch</code> suboptions except <code>auto</code> .

Table 48. **-qarch-related macros** (continued)

Macro name	Description	Predefined by the following <b>-qarch</b> suboptions
<code>_ARCH_PPC</code>	Indicates that the application is targeted to run on any Power processor.	Defined for all <b>-qarch</b> suboptions except <b>auto</b> .
<code>_ARCH_PPC64</code>	Indicates that the application is targeted to run on Power processors with 64-bit support.	<code>ppc64</code>   <code>pwr3</code>   <code>rs64b</code>   <code>rs64c</code>   <code>ppc64gr</code>   <code>ppc64grsq</code>   <code>ppc64v</code>   <code>pwr4</code>   <code>pwr5</code>   <code>pwr5x</code>   <code>pwr6</code>   <code>pwr6e</code>   <code>pwr7</code>   <code>ppc970</code>
<code>_ARCH_PPCGR</code>	Indicates that the application is targeted to run on Power processors with graphics support.	<code>ppcgr</code>   <code>604</code>   <code>pwr3</code>   <code>rs64b</code>   <code>rs64c</code>   <code>ppc64gr</code>   <code>ppc64grsq</code>   <code>ppc64v</code>   <code>pwr4</code>   <code>pwr5</code>   <code>pwr5x</code>   <code>pwr6</code>   <code>pwr6e</code>   <code>pwr7</code>   <code>ppc970</code>
<code>_ARCH_PPC64GR</code>	Indicates that the application is targeted to run on Power processors with 64-bit and graphics support.	<code>pwr3</code>   <code>rs64b</code>   <code>rs64c</code>   <code>ppc64gr</code>   <code>ppc64v</code>   <code>pwr4</code>   <code>pwr5</code>   <code>pwr5x</code>   <code>pwr6</code>   <code>pwr6e</code>   <code>pwr7</code>   <code>ppc970</code>
<code>_ARCH_PPC64GRSQ</code>	Indicates that the application is targeted to run on Power processors with 64-bit, graphics, and square root support.	<code>pwr3</code>   <code>rs64b</code>   <code>rs64c</code>   <code>ppc64grsq</code>   <code>ppc64v</code>   <code>pwr4</code>   <code>pwr5</code>   <code>pwr5x</code>   <code>pwr6</code>   <code>pwr6e</code>   <code>pwr7</code>   <code>ppc970</code>
<code>_ARCH_PPC64V</code>	Indicates that the application is targeted to run on Power processors with 64-bit and vector processing support.	<code>ppc64v</code>   <code>ppc970</code>   <code>pwr6</code>   <code>pwr6e</code>   <code>pwr7</code>
<code>_ARCH_PPC970</code>	Indicates that the application is targeted to run on the PowerPC 970 processor.	<code>ppc970</code>
<code>_ARCH_PWR3</code>	Indicates that the application is targeted to run on POWER3 processors.	<code>pwr3</code>   <code>pwr4</code>   <code>pwr5</code>   <code>pwr5x</code>   <code>pwr6</code>   <code>pwr6e</code>   <code>pwr7</code>   <code>ppc970</code>
<code>_ARCH_PWR4</code>	Indicates that the application is targeted to run on POWER4 processors.	<code>pwr4</code>   <code>pwr5</code>   <code>pwr5x</code>   <code>pwr6</code>   <code>pwr6e</code>   <code>pwr7</code>   <code>ppc970</code>
<code>_ARCH_PWR5</code>	Indicates that the application is targeted to run on POWER5 processors.	<code>pwr5</code>   <code>pwr5x</code>   <code>pwr6</code>   <code>pwr6e</code>   <code>pwr7</code>
<code>_ARCH_PWR5X</code>	Indicates that the application is targeted to run on POWER5+ processors.	<code>pwr5x</code>   <code>pwr6</code>   <code>pwr6e</code>   <code>pwr7</code>
<code>_ARCH_PWR6</code>	Indicates that the application is targeted to run on POWER6 processors.	<code>pwr6</code>   <code>pwr6e</code>   <code>pwr7</code>
<code>_ARCH_PWR6E</code>	Indicates that the application is targeted to run on POWER6 processors running in POWER6 raw mode.	<code>pwr6e</code>
<code>_ARCH_PWR7</code>	Indicates that the application is targeted to run on POWER7 processors.	<code>pwr7</code>
<code>_ARCH_RS64A</code>	Indicates that the application is targeted to run on the RS64I processor.	<code>rs64a</code>
<code>_ARCH_RS64B</code>	Indicates that the application is targeted to run on the RS64II processor.	<code>rs64b</code>
<code>_ARCH_RS64C</code>	Indicates that the application is targeted to run on the RS64III processor.	<code>rs64c</code>

## Macros related to language levels

The following macros can be tested for C99 features, features related to GNU C or C++, and other IBM language extensions. All of these macros are predefined to a value of 1 by a specific language level, represented by a suboption of the **-qlanglvl** compiler option, or any invocation or pragma that implies that suboption. If the



suboption enabling the feature is not in effect, then the macro is undefined. For descriptions of the features related to these macros, see the *XL C/C++ Language Reference*.

Table 49. Predefined macros for language features



















Predefined macro name	Description	Predefined when the following language level is in effect
 <code>__BOOL__</code>	Indicates that the <code>bool</code> keyword is accepted.	Always defined except when <code>-qnokeyword=bool</code> is in effect.
 <code>__C99_BOOL</code>	Indicates support for the <code>_Bool</code> data type.	<code>extc1x</code>   <code>stdc99</code>   <code>extc99</code>   <code>extc89</code>   extended
 <code>__C99_COMPLEX</code>	Indicates that the support for C99 complex types is enabled or that the C99 complex header should be included.	<code>extc1x</code>   <code>stdc99</code>   <code>extc99</code>   <code>extc89</code>   extended
 <code>__C99_COMPLEX_HEADER__</code>	Indicates support for C99-style complex headers.	<code>c99complexheader</code>
 <code>__C99_CPLUSCMT</code>	Indicates support for C++ style comments	<code>extc1x</code>   <code>stdc99</code>   <code>extc99</code>   <code>stdc89</code>   <code>extc89</code>   extended (also <code>-qcplusplus</code> )
<code>__C99_COMPOUND_LITERAL</code>	Indicates support for compound literals.	 <code>extc1x</code>   <code>stdc99</code>   <code>extc99</code>   <code>extc89</code>   extended   extended   <code>extended0x</code>
 <code>__C99_DESIGNATED_INITIALIZER</code>	Indicates support for designated initialization.	<code>extc1x</code>   <code>stdc99</code>   <code>extc99</code>   <code>extc89</code>   extended
 <code>__C99_DUP_TYPE_QUALIFIER</code>	Indicates support for duplicated type qualifiers.	<code>extc1x</code>   <code>stdc99</code>   <code>extc99</code>   <code>extc89</code>   extended
 <code>__C99_EMPTY_MACRO_ARGUMENTS</code>	Indicates support for empty macro arguments.	<code>extc1x</code>   <code>stdc99</code>   <code>extc99</code>   <code>extc89</code>   extended
 <code>__C99_FLEXIBLE_ARRAY_MEMBER</code>	Indicates support for flexible array members.	<code>extc1x</code>   <code>stdc99</code>   <code>extc99</code>   <code>extc89</code>   extended
<code>__C99_FUNC__</code>	Indicates support for the <code>__func__</code> predefined identifier.	 <code>extc1x</code>   <code>stdc99</code>   <code>extc99</code>   <code>extc89</code>   extended   extended   <code>extended0x</code>   <code>c99_func__</code>
<code>__C99_HEX_FLOAT_CONST</code>	Indicates support for hexadecimal floating constants.	 <code>extc1x</code>   <code>stdc99</code>   <code>extc99</code>   <code>extc89</code>   extended   extended   <code>extended0x</code>   <code>c99hexfloat</code>
 <code>__C99_INLINE</code>	Indicates support for the <code>inline</code> function specifier.	<code>extc1x</code>   <code>stdc99</code>   <code>extc99</code> (also <code>-qkeyword=inline</code> )
<code>__C99_LLONG</code>	Indicates support for C99-style long long data types and literals.	 <code>extc1x</code>   <code>stdc99</code>   <code>extc99</code>   extended0x   <code>c99longlong</code>

Table 49. Predefined macros for language features (continued)

Predefined macro name	Description	Predefined when the following language level is in effect
<code>__C99_MACRO_WITH_VA_ARGS</code>	Indicates support for function-like macros with variable arguments.	<p>▶ C extc1x   stdc99   extc99   extc89   extended</p> <p>▶ C++ extended   extended0x   varargmacros</p>
<code>__C99_MAX_LINE_NUMBER</code>	Indicates that the maximum line number is 2147483647.	<p>▶ C extc1x   stdc99   extc99   extc89   extended</p> <p>▶ C++ extended0x   c99preprocessor</p>
▶ C <code>__C99_MIXED_DECL_AND_CODE</code>	Indicates support for mixed declaration and code.	extc1x   stdc99   extc99   extc89   extended
<code>__C99_MIXED_STRING_CONCAT</code>	Indicates support for concatenation of wide string and non-wide string literals.	<p>▶ C extc1x   stdc99   extc99   extc89   extended</p> <p>▶ C++ extended0x   c99preprocessor</p>
▶ C <code>__C99_NON_LVALUE_ARRAY_SUB</code>	Indicates support for non-lvalue subscripts for arrays.	extc1x   stdc99   extc99   extc89   extended
▶ C <code>__C99_NON_CONST_AGGR_INITIALIZER</code>	Indicates support for non-constant aggregate initializers.	extc1x   stdc99   extc99   extc89   extended
<code>__C99_PRAGMA_OPERATOR</code>	Indicates support for the <code>_Pragma</code> operator.	<p>▶ C extc1x   stdc99   extc99   extc89   extended</p> <p>▶ C++ extended extended0x</p>
▶ C <code>__C99_REQUIRE_FUNC_DECL</code>	Indicates that implicit function declaration is not supported.	stdc99
<code>__C99_RESTRICT</code>	Indicates support for the C99 <code>restrict</code> qualifier.	<p>▶ C extc1x   stdc99   extc99 (also <b>-qkeyword=restrict</b>)</p> <p>▶ C++ extended   extended0x (also <b>-qkeyword=restrict</b>)</p>
▶ C <code>__C99_STATIC_ARRAY_SIZE</code>	Indicates support for the <code>static</code> keyword in array parameters to functions.	extc1x   stdc99   extc99   extc89   extended
▶ C <code>__C99_STD_PRAGMAS</code>	Indicates support for standard pragmas.	extc1x   stdc99   extc99   extc89   extended
▶ C <code>__C99_TGMATH</code>	Indicates support for type-generic macros in <code>tgmath.h</code>	extc1x   stdc99   extc99   extc89   extended
<code>__C99_UCN</code>	Indicates support for universal character names.	<p>▶ C extc1x   stdc99   extc99   ucs</p> <p>▶ C++ ucs</p>

Table 49. Predefined macros for language features (continued)
















Predefined macro name	Description	Predefined when the following language level is in effect
 <code>__C99_VAR_LEN_ARRAY</code>	Indicates support for variable length arrays.	<code>extc1x</code>   <code>stdc99</code>   <code>extc99</code>   <code>extc89</code>   <code>extended</code>
 <code>__C99_VARIABLE_LENGTH_ARRAY</code>	Indicates support for variable length arrays.	<code>extended</code>   <code>extended0x</code>   <code>c99vla</code>
<code>__DIGRAPHS__</code>	Indicates support for digraphs.	 <code>extc1x</code>   <code>stdc99</code>   <code>extc99</code>   <code>extc89</code>   <code>extended</code> (also <b>-qdigraph</b> )   <code>extended</code>   <code>extended0x</code>   <code>compat366</code>   <code>strict98</code> (also <b>-qdigraph</b> )
 <code>__EXTENDED__</code>	Indicates that language extensions are supported.	<code>extended</code>
 <code>__IBM_ALIGN</code>	Indicates support for the <code>__align</code> specifier.	Always defined except when <b>-qnokeyword=__alignof</b> is specified
<code>__IBM_ALIGNOF__</code> , <code>__IBM_ALIGNOF__</code>	Indicates support for the <code>__alignof__</code> operator.	 <code>extc1x</code>   <code>extc99</code>   <code>extc89</code>   <code>extended</code>   <code>extended</code>
<code>__IBM_ATTRIBUTES</code>	Indicates support for type, variable, and function attributes.	 <code>extc1x</code>   <code>extc99</code>   <code>extc89</code>   <code>extended</code>   <code>extended</code>   <code>extended0x</code>
<code>__IBM_COMPUTED_GOTO</code>	Indicates support for computed goto statements.	 <code>extc1x</code>   <code>extc99</code>   <code>extc89</code>   <code>extended</code>   <code>extended</code>   <code>extended0x</code>   <code>gnu_computedgoto</code>
<code>__IBM_EXTENSION_KEYWORD</code>	Indicates support for the <code>__extension__</code> keyword.	 <code>extc1x</code>   <code>extc99</code>   <code>extc89</code>   <code>extended</code>   <code>extended</code>   <code>extended0x</code>   <code>compat366</code>   <code>strict98</code>
 <code>__IBM_GCC_INLINE__</code>	Indicates support for the GCC <code>__inline__</code> specifier.	<code>extc1x</code>   <code>extc99</code>   <code>extc89</code>   <code>extended</code>
 <code>__IBM_DOLLAR_IN_ID</code>	Indicates support for dollar signs in identifiers.	<code>extc1x</code>   <code>extc99</code>   <code>extc89</code>   <code>extended</code>
 <code>__IBM_GENERALIZED_LVALUE</code>	Indicates support for generalized lvalues.	<code>extc1x</code>   <code>extc99</code>   <code>extc89</code>   <code>extended</code>

Table 49. Predefined macros for language features (continued)

Predefined macro name	Description	Predefined when the following language level is in effect
<code>__IBM_INCLUDE_NEXT</code>	Indicates support for the <code>#include_next</code> preprocessing directive.	<p>▶ <b>C</b> Always defined</p> <p>▶ <b>C++</b> Always defined except when <code>-qlanglvl=nognu_include_next</code> is in effect.</p>
<code>__IBM_LABEL_VALUE</code>	Indicates support for labels as values.	<p>▶ <b>C</b> extc1x   extc99   extc89   extended</p> <p>▶ <b>C++</b> extended   extended0x   gnu_labelvalue</p>
<code>__IBM_LOCAL_LABEL</code>	Indicates support for local labels.	<p>▶ <b>C</b> extc1x   extc99   extc89   extended</p> <p>▶ <b>C++</b> extended   extended0x   gnu_locallabel</p>
<code>__IBM_MACRO_WITH_VA_ARGS</code>	Indicates support for variadic macro extensions.	<p>▶ <b>C</b> extc1x   extc99   extc89   extended</p> <p>▶ <b>C++</b> extended   extended0x   gnu_varargmacros</p>
▶ <b>C</b> <code>__IBM_NESTED_FUNCTION</code>	Indicates support for nested functions.	extc1x   extc99   extc89   extended
▶ <b>C</b> <code>__IBM_PP_PREDICATE</code>	Indicates support for <code>#assert</code> , <code>#unassert</code> , <code>#cpu</code> , <code>#machine</code> , and <code>#system</code> preprocessing directives.	extc1x   extc99   extc89   extended
▶ <b>C</b> <code>__IBM_PP_WARNING</code>	Indicates support for the <code>#warning</code> preprocessing directive.	extc1x   extc99   extc89   extended
▶ <b>C</b> <code>__IBM_REGISTER_VARS</code>	Indicates support for variables in specified registers.	Always defined.
<code>__IBM__TYPEOF__</code>	Indicates support for the <code>__typeof__</code> or <code>typeof</code> keyword.	<p>▶ <b>C</b> always defined</p> <p>▶ <b>C++</b> extended   extended0x (Also <code>-qkeyword=typeof</code>)</p>
<code>__IBMC_COMPLEX_INIT</code>	Indicates support for the initialization of complex types: <code>float _Complex</code> , <code>double _Complex</code> , and <code>long double _Complex</code> .	extc1x
<code>__IBMC_NORETURN</code>	Indicates support for the <code>_Noreturn</code> function specifier.	<p>▶ <b>C</b> extc89   extc99   extended   extc1x</p> <p>▶ <b>C++</b> extended   extended0x   c1xnoreturn</p>

Table 49. Predefined macros for language features (continued)



























Predefined macro name	Description	Predefined when the following language level is in effect
 <code>__IBMC_STATIC_ASSERT</code>	Indicates support for the static assertions feature.	<code>extc1x</code>
 <code>__IBMCPP_AUTO_TYPEDEDUCTION</code>	Indicates support for the auto type deduction feature.	<code>extended0x</code>   <code>autotypededuction</code>
 <code>__IBMCPP_C99_LONG_LONG</code>	Indicates support for the C99 long long feature.	<code>extended0x</code>   <code>c99longlong</code>
 <code>__IBMCPP_C99_PREPROCESSOR</code>	Indicates support for the C99 preprocessor features adopted in the C++0x standard.	<code>extended0x</code>   <code>c99preprocessor</code>
<code>__IBMCPP_COMPLEX_INIT</code>	Indicates support for the initialization of complex types: <code>float _Complex</code> , <code>double _Complex</code> , and <code>long double _Complex</code> .	<code>extended</code>
 <code>__IBMCPP_CONSTEXPR</code>	Indicates support for the generalized constant expressions feature. <b>Note:</b> In XL C/C++ V12.1, this feature is a partial implementation of what is defined in the C++0x standard.	<code>extended0x</code>   <code>constexpr</code>
 <code>__IBMCPP_DECLTYPE</code>	Indicates support for the <code>decltype</code> feature.	<code>extended0x</code>   <code>decltype</code>
 <code>__IBMCPP_DELEGATING_CTORS</code>	Indicates support for the delegating constructors feature.	<code>extended0x</code>   <code>delegatingctors</code>
 <code>__IBMCPP_EXPLICIT_CONVERSION_OPERATORS</code>	Indicates support for the explicit conversion operators feature.	<code>extended0x</code>   <code>explicitconversionoperators</code>
 <code>__IBMCPP_EXTENDED_FRIEND</code>	Indicates support for the extended friend declarations feature.	<code>extended0x</code>   <code>extendedfriend</code>
 <code>__IBMCPP_EXTERN_TEMPLATE</code>	Indicates support for the explicit instantiation declarations feature.	<code>extended</code>   <code>extended0x</code>   <code>externtemplate</code>
 <code>__IBMCPP_INLINE_NAMESPACE</code>	Indicates support for the inline namespace definitions feature.	<code>extended0x</code>   <code>inlinenamespace</code>
 <code>__IBMCPP_REFERENCE_COLLAPSING</code>	Indicates support for the reference collapsing feature.	<code>extended0x</code>   <code>referencecollapsing</code>
 <code>__IBMCPP_RIGHT_ANGLE_BRACKET</code>	Indicates support for the right angle bracket feature.	<code>extended0x</code>   <code>rightanglebracket</code>
 <code>__IBMCPP_RVALUE_REFERENCES</code>	Indicates support for the rvalue references feature.	<code>extended0x</code>   <code>rvalueresferences</code>
 <code>__IBMCPP_SCOPED_ENUM</code>	Indicates support for the scoped enumeration feature.	<code>extended0x</code>   <code>scopedenum</code>
 <code>__IBMCPP_STATIC_ASSERT</code>	Indicates support for the static assertions feature.	 <code>extended0x</code>   <code>static_assert</code>

Table 49. Predefined macros for language features (continued)

Predefined macro name	Description	Predefined when the following language level is in effect
 <code>__IBMCPP_VARIADIC_TEMPLATES</code>	Indicates support for the variadic templates feature.	<code>extended0x</code>   <code>variadic[templates]</code>
<code>__LONG_LONG</code>	Indicates support for long long data types.	 <code>extc1x</code>   <code>stdc99</code>   <code>extc99</code>   <code>stdc89</code>   <code>extc89</code>   <code>extended</code> (also <code>-qlonglong</code> )   <code>extended0x</code>   <code>c99longlong</code>   <code>extended</code> (also <code>-qlonglong</code> )
 <code>__SAA__</code>	Indicates that only language constructs that support the most recent level of SAA C standards are allowed.	<code>saa</code>
 <code>__SAA_L2__</code>	Indicates that only language constructs that conform to SAA Level 2 C standards are allowed.	<code>saal2</code>
 <code>__STDC__</code>	Indicates that the compiler conforms to the ANSI/ISO C standard.	Predefined to 1 if ANSI/ISO C standard conformance is in effect.
<code>__STDC_HOSTED__</code>	Indicates that the implementation is a hosted implementation of the ANSI/ISO C standard. (That is, the hosted environment has all the facilities of the standard C available).	 <code>extc1x</code>   <code>stdc99</code>   <code>extc99</code>   <code>extended0x</code>
 <code>__STDC_VERSION__</code>	Indicates the version of ANSI/ISO C standard which the compiler conforms to.	The format is <code>yyyymmL</code> . (For example, the format is <code>199901L</code> for C99.)

## Examples of predefined macros

This example illustrates use of the `__FUNCTION__` and the `__C99_FUNC__` macros to test for the availability of the C99 `__func__` identifier to return the current function name:

```
#include <stdio.h>

#if defined(__C99_FUNC__)
#define PRINT_FUNC_NAME() printf (" In function %s \n", __func__);
#elif defined(__FUNCTION__)
#define PRINT_FUNC_NAME() printf (" In function %s \n", __FUNCTION__);
#else
#define PRINT_FUNC_NAME() printf (" Function name unavailable\n");
#endif

void foo(void);


int main(int argc, char **argv)
{
    int k = 1;
    PRINT_FUNC_NAME();
    foo();
}
```

```
    return 0;
}

void foo (void)
{
    PRINT_FUNC_NAME();
    return;
}
```

The output of this example is:

```
In function main
In function foo
```

 This example illustrates use of the `__FUNCTION__` macro in a C++ program with virtual functions.

```
#include <stdio.h>
class X { public: virtual void func() = 0;};

class Y : public X {
    public: void func() { printf("In function %s \n", __FUNCTION__);}
};

int main() {
    Y aaa;
    aaa.func();
}
```

The output of this example is:

```
In function Y::func()
```






---

## Chapter 7. Compiler built-in functions

A built-in function is a coding extension to C and C++ that allows a programmer to use the syntax of C function calls and C variables to access the instruction set of the processor of the compiling machine. IBM Power architectures have special instructions that enable the development of highly optimized applications. Access to some Power instructions cannot be generated using the standard constructs of the C and C++ languages. Other instructions can be generated through standard constructs, but using built-in functions allows exact control of the generated code. Inline assembly language programming, which uses these instructions directly, is not fully supported by XL C/C++ and other compilers. Furthermore, the technique can be time-consuming to implement.

As an alternative to managing hardware registers through assembly language, XL C/C++ built-in functions provide access to the optimized Power instruction set and allow the compiler to optimize the instruction scheduling.

 To call any of the XL C/C++ built-in functions in C++, you must include the header file `builtins.h` in your source code.

The following sections describe the available built-in functions for the AIX platform.

- “Fixed-point built-in functions”
- “Binary floating-point built-in functions” on page 488
- “Decimal floating-point built-in functions” on page 499
- “Synchronization and atomic built-in functions” on page 518
- “Cache-related built-in functions” on page 525
- “Block-related built-in functions” on page 534
- “Miscellaneous built-in functions” on page 603
- “Built-in functions for parallel processing” on page 609

---

### Fixed-point built-in functions

Fixed-point built-in functions are grouped into the following categories:

- “Absolute value functions” on page 480
- “Assert functions” on page 480
- “Count zero functions” on page 481
- “Load functions” on page 483
- “Multiply functions” on page 483
- “Population count functions” on page 484
- “Rotate functions” on page 485
- “Store functions” on page 486
- “Trap functions” on page 487

## Absolute value functions

**\_\_labs, \_\_llabs**

### Purpose

Absolute Value Long, Absolute Value Long Long

Returns the absolute value of the argument.

### Prototype

```
signed long __labs (signed long);
```

```
signed long long __llabs (signed long long);
```

## Assert functions

**\_\_assert1, \_\_assert2**

### Purpose

Generates trap instructions.

### Prototype

```
int __assert1 (int, int, int);
```

```
void __assert2 (int);
```

## Bit permutation functions

**\_\_bpermd**

### Purpose

Byte Permute Doubleword

Returns the result of a bit permutation operation.

### Prototype

```
long long __bpermd (long long bit_selector, long long source);
```

### Usage

Eight bits are returned, each corresponding to a bit within source, and were selected by a byte of bit\_selector. If byte i of bit\_selector is less than 64, the permuted bit i is set to the bit of source specified by byte i of bit\_selector; otherwise the permuted bit i is set to 0. The permuted bits are placed in the least-significant byte of the result value and the remaining bits are filled with 0s.

Valid only when **-qarch** is set to target POWER7 processors in 64-bit mode.

## Comparison functions

### **\_\_cmpb**

#### **Purpose**

Compare Bytes

Compares each of the eight bytes of *source1* with the corresponding byte of *source2*. If byte *i* of *source1* and byte *i* of *source2* are equal, 0xFF is placed in the corresponding byte of the result; otherwise, 0x00 is placed in the corresponding byte of the result.

#### **Prototype**

```
long long __cmpb (long long source1, long long source2);
```

#### **Usage**

Valid only when **-qarch** is set to target POWER6 and POWER7 processors.

## Count zero functions

### **\_\_cntlz4, \_\_cntlz8**

#### **Purpose**

Count Leading Zeros, 4/8-byte integer

#### **Prototype**

```
unsigned int __cntlz4 (unsigned int);
```

```
unsigned int __cntlz8 (unsigned long long);
```

### **\_\_cnttz4, \_\_cnttz8**

#### **Purpose**

Count Trailing Zeros, 4/8-byte integer

#### **Prototype**

```
unsigned int __cnttz4 (unsigned int);
```

```
unsigned int __cnttz8 (unsigned long long);
```

## Division functions

These division functions are valid only when **-qarch** is set to target POWER7 processors.

### **\_\_divde**

#### **Purpose**

Divide Doubleword Extended

Returns the result of a doubleword extended division. The result has a value equal to *dividend/divisor*.

## Prototype

```
long long __divde (long long dividend, long long divisor);
```

## Usage

Valid only when **-qarch** is set to target POWER7 processors in 64-bit mode.

**Note:** If the result of the division is larger than 32 bits or if the divisor is 0, the return value of the function is undefined.

## **\_\_divdeu**

### Purpose

Divide Doubleword Extended Unsigned

Returns the result of a double word extended unsigned division. The result has a value equal to *dividend/divisor*.

## Prototype

```
unsigned long long __divdeu (unsigned long long dividend, unsigned long long divisor);
```

## Usage

Valid only when **-qarch** is set to target POWER7 processors in 64-bit mode.

**Note:** If the result of the division is larger than 32 bits or if the divisor is 0, the return value of the function is undefined.

## **\_\_divwe**

### Purpose

Divide Word Extended

Returns the result of a word extended division. The result has a value equal to *dividend/divisor*.

## Prototype

```
int __divwe(int dividend, int divisor);
```

## Usage

Valid only when **-qarch** is set to target POWER7 processors.

**Note:** If the divisor is 0, the return value of the function is undefined.

## **\_\_divweu**

### Purpose

Divide Word Extended Unsigned

Returns the result of a word extended unsigned division. The result has a value equal to *dividend/divisor*.

## Prototype

```
unsigned int __divweu(unsigned int dividend, unsigned int divisor);
```

## Usage

Valid only when `-qarch` is set to target POWER7 processors.

**Note:** If the divisor is 0, the return value of the function is undefined.

## Load functions

### `__load2r`, `__load4r`

#### Purpose

Load Halfword Byte Reversed, Load Word Byte Reversed

#### Prototype

```
unsigned short __load2r (unsigned short*);
```

```
unsigned int __load4r (unsigned int*);
```

### `__load8r`

#### Purpose

Load with Byte Reversal (8-byte integer)

Performs an eight-byte byte-reversed load from the given address.

#### Prototype

```
unsigned long long __load8r (unsigned long long * address);
```

#### Usage

Valid only when `-qarch` and `-q64` are set to target POWER7 processors.

## Multiply functions

### `__imul_dbl`

#### Purpose

Computes the product of two long integers and stores the result in a pointer.

#### Prototype

```
void __imul_dbl (long, long, long*);
```

### `__mulhd`, `__mulhdu`

#### Purpose

Multiply High Doubleword Signed, Multiply High Doubleword Unsigned

Returns the highorder 64 bits of the 128bit product of the two parameters.

## Prototype

```
long long int __mulhd ( long int, long int);
```

```
unsigned long long int __mulhdu (unsigned long int, unsigned long int);
```

## Usage

Valid only in 64-bit mode.

**\_\_mulhw, \_\_mulhwu**

## Purpose

Multiply High Word Signed, Multiply High Word Unsigned

Returns the highorder 32 bits of the 64bit product of the two parameters.

## Prototype

```
int __mulhw (int, int);
```

```
unsigned int __mulhwu (unsigned int, unsigned int);
```

## Population count functions

**\_\_popcnt4, \_\_popcnt8**

## Purpose

Population Count, 4/8-byte integer

Returns the number of bits set for a 32/64-bit integer.

## Prototype

```
int __popcnt4 (unsigned int);
```

```
int __popcnt8 (unsigned long long);
```

**\_\_popcntb**

## Purpose

Population Count Byte

Counts the 1 bits in each byte of the parameter and places that count into the corresponding byte of the result.

## Prototype

```
unsigned long __popcntb(unsigned long);
```

**\_\_poppar4, \_\_poppar8**

## Purpose

Population Parity, 4/8-byte integer

Checks whether the number of bits set in a 32/64-bit integer is an even or odd number.

### Prototype

```
int __poppar4(unsigned int);  
  
int __poppar8(unsigned long long);
```

### Return value

Returns 1 if the number of bits set in the input parameter is odd. Returns 0 otherwise.

## Rotate functions

### `__rdlam`

#### Purpose

Rotate Double Left and AND with Mask

Rotates the contents of *rs* left *shift* bits, and ANDs the rotated data with the *mask*.

#### Prototype

```
unsigned long long __rdlam (unsigned long long rs, unsigned int shift,  
unsigned long long mask);
```

#### Parameters

*mask*

Must be a constant that represents a contiguous bit field.

### `__rldimi`, `__rlwimi`

#### Purpose

Rotate Left Doubleword Immediate then Mask Insert, Rotate Left Word Immediate then Mask Insert

Rotates *rs* left *shift* bits then inserts *rs* into *is* under bit mask *mask*.

#### Prototype

```
unsigned long long __rldimi (unsigned long long rs, unsigned long long is,  
unsigned int shift, unsigned long long mask);
```

```
unsigned int __rlwimi (unsigned int rs, unsigned int is, unsigned int shift,  
unsigned int mask);
```

#### Parameters

*shift*

A constant value 0 to 63 (`__rldimi`) or 31 (`__rlwimi`).

*mask*

Must be a constant that represents a contiguous bit field.

## **\_\_rlwnm**

### **Purpose**

Rotate Left Word then AND with Mask

Rotates *rs* left *shift* bits, then ANDs *rs* with bit mask *mask*.

### **Prototype**

```
unsigned int __rlwnm (unsigned int rs, unsigned int shift, unsigned int mask);
```

### **Parameters**

*mask*

Must be a constant that represents a contiguous bit field.

## **\_\_rotatel4, \_\_rotatel8**

### **Purpose**

Rotate Left Word, Rotate Left Doubleword

Rotates *rs* left *shift* bits.

### **Prototype**

```
unsigned int __rotatel4 (unsigned int rs, unsigned int shift);
```

```
unsigned long long __rotatel8 (unsigned long long rs, unsigned long long shift);
```

## **Store functions**

## **\_\_store2r, \_\_store4r**

### **Purpose**

Store 2/4-byte Reversal

### **Prototype**

```
void __store2r (unsigned short, unsigned short*);
```

```
void __store4r (unsigned int, unsigned int*);
```

## **\_\_store8r**

### **Purpose**

Store with Byte-Reversal (eight-byte integer)

Takes the loaded eight-byte integer value and performs a byte-reversed store operation.

### **Prototype**

```
void __store8r (unsigned long long source, unsigned long long * address);
```



## Usage

Valid only when `-qarch` is set to target POWER7 processors in 64-bit mode.

## Trap functions

`__tdw`, `__tw`

### Purpose

Trap Doubleword, Trap Word

Compares parameter *a* with parameter *b*. This comparison results in five conditions which are ANDed with a 5-bit constant *TO*. If the result is not 0 the system trap handler is invoked.

### Prototype

```
void __tdw ( long a, long b, unsigned int TO);
```

```
void __tw (int a, int b, unsigned int TO);
```

### Parameters

*TO* A value of 0 to 31 inclusive. Each bit position, if set, indicates one or more of the following possible conditions:

**0 (high-order bit)**

*a* is less than *b*, using signed comparison.

**1** *a* is greater than *b*, using signed comparison.

**2** *a* is equal to *b*

**3** *a* is less than *b*, using unsigned comparison.

**4 (low-order bit)**

*a* is greater than *b*, using unsigned comparison.

## Usage

`__tdw` is valid only in 64-bit mode.

`__trap`, `__trapd`

### Purpose

Trap if the Parameter is not Zero, Trap if the Parameter is not Zero Doubleword

### Prototype

```
void __trap (int);
```

```
void __trapd ( long);
```

## Usage

`__trapd` is valid only in 64-bit mode.

---

## Binary floating-point built-in functions

Binary floating-point built-in functions are grouped into the following categories:

- “Absolute value functions” on page 480
- “Add functions”
- “Conversion functions” on page 489
- “FPSCR functions” on page 491
- “Multiply functions” on page 493
- “Multiply-add/subtract functions” on page 494
- “Reciprocal estimate functions” on page 495
- “Rounding functions” on page 495
- “Select functions” on page 497
- “Square root functions” on page 497
- “Software division functions” on page 497

For decimal floating-point built-in functions, see Decimal floating-point built-in functions.

### Absolute value functions

#### **\_\_fabss**

##### **Purpose**

Floating Absolute Value Single

Returns the absolute value of the argument.

##### **Prototype**

```
float __fabss (float);
```

#### **\_\_fnabs**

##### **Purpose**

Floating Negative Absolute Value, Floating Negative Absolute Value Single

Returns the negative absolute value of the argument.

##### **Prototype**

```
double __fnabs (double);
```

```
float __fnabss (float);
```

### Add functions

#### **\_\_fadd, \_\_fadds**

##### **Purpose**

Floating Add, Floating Add Single

Adds two arguments and returns the result.

### Prototype

```
double __fadd (double, double);
```

```
float __fadds (float, float);
```

## Conversion functions

### **\_\_cplx, \_\_cplx, \_\_cplx**

#### **Purpose**

Converts two real parameters into a single complex value.

#### **Prototype**

```
double _Complex __cplx (double, double);
```

```
float _Complex __cplx (float, float);
```

```
long double _Complex __cplx (long double, long double);
```

### **\_\_fcid**

#### **Purpose**

Floating Convert from Integer Doubleword

Converts a 64-bit signed integer stored in a double to a double-precision floating-point value.

#### **Prototype**

```
double __fcid (double);
```

### **\_\_fcud**

#### **Purpose**

Floating-point Conversion from Unsigned integer Double word

Converts a 64-bit unsigned integer stored in a double into a double-precision floating-point value.

#### **Prototype**

```
double __fcud(double);
```

### **\_\_fctid**

#### **Purpose**

Floating Convert to Integer Doubleword

Converts a double-precision argument to a 64-bit signed integer, using the current rounding mode, and returns the result in a double.

#### **Prototype**

```
double __fctid (double);
```

## **\_\_fctidz**

### **Purpose**

Floating Convert to Integer Doubleword with Rounding towards Zero

Converts a double-precision argument to a 64-bit signed integer, using the rounding mode round-toward-zero, and returns the result in a double.

### **Prototype**

```
double __fctidz (double);
```

## **\_\_fctiw**

### **Purpose**

Floating Convert to Integer Word

Converts a double-precision argument to a 32-bit signed integer, using the current rounding mode, and returns the result in a double.

### **Prototype**

```
double __fctiw (double);
```

## **\_\_fctiwz**

### **Purpose**

Floating Convert to Integer Word with Rounding towards Zero

Converts a double-precision argument to a 32-bit signed integer, using the rounding mode round-toward-zero, and returns the result in a double.

### **Prototype**

```
double __fctiwz (double);
```

## **\_\_fctudz**

### **Purpose**

Floating-point Conversion to Unsigned integer Double word with rounding towards Zero

Converts a floating-point value to unsigned integer double word and rounds to zero.

### **Prototype**

```
double __fctudz(double);
```

### **Result value**

The result is a double number, which is rounded to zero.

## **\_\_fctuwz**

### **Purpose**

Floating-point Conversion to Unsigned integer Word with rounding to Zero

Converts a floating-point number into a 32-bit unsigned integer and rounds to zero.

### **Prototype**

```
double __fctuwz(double);
```

### **Result value**

The result is a double number, which is rounded to zero.

## **FPSCR functions**

### **\_\_mtfsb0**

#### **Purpose**

Move to Floating-Point Status/Control Register (FPSCR) Bit 0

Sets bit *bt* of the FPSCR to 0.

#### **Prototype**

```
void __mtfsb0 (unsigned int bt);
```

#### **Parameters**

*bt* Must be a constant with a value of 0 to 31.

### **\_\_mtfsb1**

#### **Purpose**

Move to FPSCR Bit 1

Sets bit *bt* of the FPSCR to 1.

#### **Prototype**

```
void __mtfsb1 (unsigned int bt);
```

#### **Parameters**

*bt* Must be a constant with a value of 0 to 31.

### **\_\_mtfsf**

#### **Purpose**

Move to FPSCR Fields

Places the contents of *frb* into the FPSCR under control of the field mask specified by *flm*. The field mask *flm* identifies the 4bit fields of the FPSCR affected.

## Prototype

```
void __mtfsf (unsigned int flm, unsigned int frb);
```

## Parameters

*flm*

Must be a constant 8-bit mask.

## **\_\_mtfsfi**

### Purpose

Move to FPSCR Field Immediate

Places the value of *u* into the FPSCR field specified by *bf*.

## Prototype

```
void __mtfsfi (unsigned int bf, unsigned int u);
```

## Parameters

*bf* Must be a constant with a value of 0 to 7.

*u* Must be a constant with a value of 0 to 15.

## **\_\_readflm**

### Purpose

Returns a 64-bit double precision floating point, whose 32 low order bits contain the contents of the FPSCR. The 32 low order bits are bits 32 - 63 counting from the highest order bit.

## Prototype

```
double __readflm (void);
```

## **\_\_setflm**

### Purpose

Takes a double precision floating-point number and places the lower 32 bits in the FPSCR. The 32 low order bits are bits 32 - 63 counting from the highest order bit. Returns the previous contents of the FPSCR.

## Prototype

```
double __setflm (double);
```

## **\_\_setrnd**

### Purpose

Sets the rounding mode.

## Prototype

```
double __setrnd (int mode);
```

## Parameters

The allowable values for *mode* are:

- 0 — round to nearest
- 1 — round to zero
- 2 — round to +infinity
- 3 — round to -infinity

## **\_\_dfp\_set\_rounding\_mode** Purpose

Set Rounding Mode

Sets the current decimal rounding mode.

## Prototype

```
void __dfp_set_rounding_mode (unsigned long rounding_mode);
```

## Parameters

*rounding\_mode*

One of the compile-time constant values (0 to 7) or macros listed in Table 51 on page 510.

## Usage

If you change the rounding mode within a function, you must restore the rounding mode before the function returns.

## **\_\_dfp\_get\_rounding\_mode** Purpose

Get Rounding Mode

Gets the current decimal rounding mode.

## Prototype

```
unsigned long __dfp_get_rounding_mode (void);
```

## Return value

The current rounding mode as one of the values (0 to 7) listed in Table 51 on page 510.

## Multiply functions

### **\_\_fmul, \_\_fmuls** Purpose

Floating Multiply, Floating Multiply Single

Multiplies two arguments and returns the result.

## Prototype

```
double __fmul (double, double);
```

```
float __fmuls (float, float);
```

## Multiply-add/subtract functions

### **\_\_fmadd, \_\_fmadds**

#### **Purpose**

Floating Multiply-Add, Floating Multiply-Add Single

Multiplies the first two arguments, adds the third argument, and returns the result.

#### **Prototype**

```
double __fmadd (double, double, double);
```

```
float __fmadds (float, float, float);
```

### **\_\_fmsub, \_\_fmsubs**

#### **Purpose**

Floating Multiply-Subtract, Floating Multiply-Subtract Single

Multiplies the first two arguments, subtracts the third argument and returns the result.

#### **Prototype**

```
double __fmsub (double, double, double);
```

```
float __fmsubs (float, float, float);
```

### **\_\_fnmadd, \_\_fnmadds**

#### **Purpose**

Floating Negative Multiply-Add, Floating Negative Multiply-Add Single

Multiplies the first two arguments, adds the third argument, and negates the result.

#### **Prototype**

```
double __fnmadd (double, double, double);
```

```
float __fnmadds (float, float, float);
```

### **\_\_fnmsub, \_\_fnmsubs**

#### **Purpose**

Floating Negative Multiply-Subtract

Multiplies the first two arguments, subtracts the third argument, and negates the result.



## Prototype

```
double __fnmsub (double, double, double);
```

```
float __fnmsubs (float, float, float);
```

## Reciprocal estimate functions

See also “Square root functions” on page 497.

### **\_\_fre, \_\_fres**

#### **Purpose**

Floating Reciprocal Estimate, Floating Reciprocal Estimate Single

#### **Prototype**

```
float __fre (double);
```

```
float __fres (float);
```

#### **Usage**

`__fre` is valid only when `-qarch` is set to target POWER5 or later processors.

## Rounding functions

### **\_\_fric**

#### **Purpose**

Floating-point Rounding to Integer with Current rounding mode

Rounds a double-precision floating-point value to integer with the current rounding mode.

#### **Prototype**

```
double __fric(double);
```

### **\_\_frim, \_\_frims**

#### **Purpose**

Floating Round to Integer Minus

Rounds the floating-point argument to an integer using round-to-minus-infinity mode, and returns the value as a floating-point value.

#### **Prototype**

```
double __frim (double);
```

```
float __frims (float);
```

#### **Usage**

Valid only when `-qarch` is set to target POWER5+ or later processors.

## **\_\_frin, \_\_frins**

### **Purpose**

Floating Round to Integer Nearest

Rounds the floating-point argument to an integer using round-to-nearest mode, and returns the value as a floating-point value.

### **Prototype**

```
double __frin (double);
```

```
float __frins (float);
```

### **Usage**

Valid only when **-qarch** is set to target POWER5+ or later processors.

## **\_\_frip, \_\_frips**

### **Purpose**

Floating Round to Integer Plus

Rounds the floating-point argument to an integer using round-to-plus-infinity mode, and returns the value as a floating-point value.

### **Prototype**

```
double __frip (double);
```

```
float __frips (float);
```

### **Usage**

Valid only when **-qarch** is set to target POWER5+ or later processors.

## **\_\_friz, \_\_frizs**

### **Purpose**

Floating Round to Integer Zero

Rounds the floating-point argument to an integer using round-to-zero mode, and returns the value as a floating-point value.

### **Prototype**

```
double __friz (double);
```

```
float __frizs (float);
```

### **Usage**

Valid only when **-qarch** is set to target POWER5+ or later processors.

## Select functions

### **\_\_fsel, \_\_fsels**

#### **Purpose**

Floating Select, Floating Select Single

Returns the second argument if the first argument is greater than or equal to zero; returns the third argument otherwise.

#### **Prototype**

```
double __fsel (double, double, double);
```

```
float __fsels (float, float, float);
```

## Square root functions

### **\_\_frsqrte, \_\_frsqrtes**

#### **Purpose**

Floating Reciprocal Square Root Estimate, Floating Reciprocal Square Root Estimate Single

#### **Prototype**

```
double __frsqrte (double);
```

```
float __frsqrtes (float);
```

#### **Usage**

`__frsqrtes` is valid only when `-qarch` is set to target POWER5+ or later processors.

### **\_\_fsqrt, \_\_fsqrts**

#### **Purpose**

Floating Square Root, Floating Square Root Single

#### **Prototype**

```
double __fsqrt (double);
```

```
float __fsqrts (float);
```

## Software division functions

### **\_\_swdiv, \_\_swdivs**

#### **Purpose**

Software Divide, Software Divide Single

Divides the first argument by the second argument and returns the result.

## Prototype

```
double __swdiv (double, double);
```

```
float __swdivs (float, float);
```

## **\_\_swdiv\_nochk, \_\_swdivs\_nochk**

### Purpose

Software Divide No Check, Software Divide No Check Single

Divides the first argument by the second argument, without performing range checking, and returns the result.

## Prototype

```
double __swdiv_nochk (double a, double b);
```

```
float __swdivs_nochk (float a, float b);
```

## Parameters

- a* Must not equal infinity. When **-qstrict** is in effect, *a* must have an absolute value greater than  $2^{-970}$  and less than infinity.
- b* Must not equal infinity, zero, or denormalized values. When **-qstrict** is in effect, *b* must have an absolute value greater than  $2^{-1022}$  and less than  $2^{1021}$ .

## Return value

The result must not be equal to positive or negative infinity. When **-qstrict** in effect, the result must have an absolute value greater than  $2^{-1021}$  and less than  $2^{1023}$ .

## Usage

This function can provide better performance than the normal divide operator or the `__swdiv` built-in function in situations where division is performed repeatedly in a loop and when arguments are within the permitted ranges.

## Store functions

### **\_\_stfiw**

#### Purpose

Store Floating Point as Integer Word

Stores the contents of the loworder 32 bits of *value*, without conversion, into the word in storage addressed by *addr*.

## Prototype

```
void __stfiw (const int* addr, double value);
```

---

## Decimal floating-point built-in functions


Decimal floating-point (DFP) built-in functions are grouped into the following categories:

- “Absolute value functions”
- “Coefficient functions” on page 500
- “Comparison functions” on page 501
- “Conversion functions” on page 502
- “Exponent functions” on page 507
- “NaN functions” on page 508
- “Register transfer functions” on page 509
- “Rounding functions” on page 510
- “Test functions” on page 512

For binary floating-point built-in functions, see Binary floating-point built-in functions

When `-qarch` is set to `pwr6` or `pwr6e` to target POWER6 processors, `-qfloat=nodfpemulate` becomes the default. This means that DFP hardware instructions are generated. Lower-performance software emulation code is generated only when:

- `-qarch` is not set to `pwr6` or `pwr6e`
- `-qarch` is set to `pwr6` or `pwr6e` and `-qfloat=dfpemulate` is enabled

**Note:**  In the prototypes given in the following sections, the C keyword `_Bool` is used by convention to represent a Boolean type.

### Absolute value functions

Absolute value functions determine the sign of the returned value.

#### `__d64_abs`, `__d128_abs`

##### Purpose

Absolute Value

Returns the absolute value of the parameter.

##### Prototype

```
_Decimal64 __d64_abs (_Decimal64);
```

```
_Decimal128 __d128_abs (_Decimal128);
```

#### `__d64_nabs`, `__d128_nabs`

##### Purpose

Negative Absolute Value

Returns the negative absolute value of the parameter.

## Prototype

```
_Decimal64 __d64_nabs (_Decimal64);  
_Decimal128 __d128_nabs (_Decimal128);
```

## **\_\_d64\_copysign, \_\_d128\_copysign**

### Purpose

Copysign

Returns the absolute value of the first parameter, with the sign of the second parameter.

## Prototype

```
_Decimal64 __d64_copysign (_Decimal64, _Decimal64);  
_Decimal128 __d128_copysign (_Decimal128, _Decimal128);
```

## Coefficient functions

Coefficient functions manipulate the fraction without affecting the exponent and sign, to support decimal-floating point conversion library functions.

## **\_\_d64\_shift\_left, \_\_d128\_shift\_left**

### Purpose

Shift Coefficient Left.

Shifts the coefficient of the parameter left.

## Prototype

```
_Decimal64 __d64_shift_left (_Decimal64, unsigned long digits);  
_Decimal128 __d128_shift_left (_Decimal128, unsigned long digits);
```

## Parameters

*digits*

The number of digits to be shifted left. The shift count must be in the range 0 to 63; otherwise the result is undefined.

## Return value

The sign and exponent are unchanged. The digits are shifted left.

## **\_\_d64\_shift\_right, \_\_d128\_shift\_right**

### Purpose

Shift Coefficient Right.

Shifts the coefficient of the parameter right.

## Prototype

```
_Decimal64 __d64_shift_right (_Decimal64, unsigned long digits);
```

```
_Decimal128 __d128_shift_right (_Decimal128, unsigned long digits);
```

## Parameters

*digits*

The number of digits to be shifted right. The shift count must be in the range 0 to 63; otherwise the result is undefined.

## Return value

The sign and exponent are unchanged. The digits are shifted right.

## Comparison functions

Comparison functions support extended exception handling and exponent comparisons.

### **\_\_d64\_compare\_exponents, \_\_d128\_compare\_exponents**

#### **Purpose**

Compare Exponents

Compares the exponents of two decimal floating-point values.

#### **Prototype**

```
long __d64_compare_exponents (_Decimal64, _Decimal64);
```

```
long __d128_compare_exponents (_Decimal128, _Decimal128);
```

#### **Return value**

Returns the following values:

- Less than 0 if the exponent of the first parameter is less than the exponent of the second parameter.
- 0 if both parameters have the same exponent value or if both are quiet or signaling NaNs (quiet and signaling are considered equal) or both are infinities.
- Greater than 0 if the exponent of the first argument is greater than the exponent of the second argument.
- -2 if one of the two parameters is a quiet or signaling NaN or one of the two parameters is an infinity.

### **\_\_d64\_compare\_signaling, \_\_d128\_compare\_signaling**

#### **Purpose**

Compare Signaling Exception on NaN

Compares two decimal floating-point values and raises an Invalid Operation exception if either is a quiet or signaling NaN.

## Prototype

```
long __d64_compare_signaling (_Decimal64, _Decimal64);  
long __d128_compare_signaling (_Decimal128, _Decimal128);
```

## Return value

Returns the following values:

- Less than 0 if the value of the first parameter is less than the value of the second parameter.
- 0 if both parameters have the same value.
- Greater than 0 if the value of the first parameter is greater than the value of the second parameter.

If either value is a quiet or signalling NaN, an exception is raised. If no exception handler has been enabled to trap the exception, the function returns -2.

## Usage

If either value is a NaN, normal comparisons using the relational operators (`==`, `!=`, `<`, `<=`, `>` and `>=`) always return false, which raises an exception for a signaling NaN but not for a quiet NaN. If you want an exception to be raised when either value is a quiet or signaling NaN, you should use the Compare Signaling Exception on NaN built-in functions instead of a relational operator.

## Conversion functions

Conversion functions execute decimal floating-point conversions. Some override the current rounding mode.

### **`__cbcdtd`**

#### **Purpose**

Convert Binary Coded Decimal To Declets.

The low-order 24 bits of each word of the source contain six, 4-bit BCD fields that are converted to two declets; each set of the two declets is placed into the low-order 20 bits of the corresponding word in the result. The high-order 12 bits in each word of the result are set to 0. If a 4-bit BCD field has a value greater than 9, the results are undefined.

#### **Prototype**

```
long long __cdtbcd (long long);
```

#### **Usage**

Valid only when `-qarch` is set to target POWER7 processors.

### **`__cdtbcd`**

#### **Purpose**

Convert Declets To Binary Coded Decimal.



The low-order 20 bits of each word of the source contain two deplets that are converted to six, 4-bit BCD fields; each set of six, 4-bit BCD fields is placed into the low-order 24 bits of the corresponding word in the result. The high-order 8 bits in each word of the result are set to 0.

### Prototype

```
long long __cdtbcd (long long);
```

### Usage

Valid only when `-qarch` is set to target POWER7 processors.

### `__d64_to_long_long`, `__d128_to_long_long`

#### Purpose

Convert to Integer

Converts a decimal floating-point value to a 64-bit signed binary integer, using the current rounding mode.

### Prototype

```
long long __d64_to_long_long (_Decimal64);
```

```
long long __d128_to_long_long (_Decimal128);
```

### Return value

The input value converted to a long long, using the current rounding mode (not always rounded towards zero as a cast or implicit conversion would be).

### `__d64_to_long_long_rounding`, `__d128_to_long_long_rounding`

#### Purpose

Convert to Integer

Converts a decimal floating-point value to a 64-bit signed binary integer, using a specified rounding mode.

### Prototype

```
long long __d64_to_long_long_rounding (_Decimal64, long rounding_mode);
```

```
long long __d128_to_long_long_rounding (_Decimal128, long rounding_mode);
```

### Parameters

*mode*

One of the compile time constant values or macros defined in Table 51 on page 510.

### Return value

The input value converted to a long long, using the specified rounding mode (not always rounded towards zero as a cast or implicit conversion would be).

## Usage

These functions temporarily override the rounding mode in effect for the current operation.

### **\_\_d64\_to\_signed\_BCD**

#### **Purpose**

Convert to Signed Binary-Coded Decimal

Converts the lower digits of a 64-bit decimal floating-point value to a Signed Packed Format (packed decimal).

#### **Prototype**

```
unsigned long long __d64_to_signed_BCD (_Decimal64, _Bool value);
```

#### **Return value**

Produces 15 decimal digits followed by a decimal sign in a 64-bit result. The leftmost digit is ignored.

Positive values are given the sign 0xF if *value* is true and 0xC if *value* is false.

Negative values are given the sign 0xD.

## Usage

You can use the `__d64_shift_right` function to access the leftmost digit.

### **\_\_d128\_to\_signed\_BCD**

#### **Purpose**

Convert to Signed Binary Coded Decimal.

Converts the lower digits of a 128-bit decimal floating-point value to a Signed Packed Format (packed decimal).

#### **Prototype**

```
void __d128_to_signed_BCD (_Decimal128, _Bool value, unsigned long long *upper, unsigned long long *lower);
```

#### **Parameters**

*upper*

The address of the variable that will hold the upper digits of the result.

*lower*

The address of the variable that will hold the lower digits of the result.

#### **Return value**

Produces 31 decimal digits followed by a decimal sign in a 128-bit result. Digits to the left are ignored. The higher 16 digits are stored in the parameter *upper*. The lower 15 digits plus the sign are stored in the parameter *lower*.

Positive values are given the sign 0xF if *value* is true and 0xC if *value* is false.

Negative values are given the sign 0xD.

### Usage

You can use the `__d128_shift_right` function to access the digits to the left.

### `__d64_to_unsigned_BCD`

#### Purpose

Convert to Unsigned Binary Coded Decimal.

Converts the lower digits of a 64-bit decimal floating-point value to an Unsigned Packed Format.

#### Prototype

```
unsigned long long __d64_to_unsigned_BCD (_Decimal64);
```

#### Return value

Returns 16 decimal digits with no sign in a 64-bit result.

### Usage

You can use the `__d64_shift_right` function to access the digits to the left.

### `__d128_to_unsigned_BCD`

#### Purpose

Convert to Unsigned Binary Coded Decimal.

Converts the lower digits of a 128-bit decimal floating-point value to an Unsigned Packed Format.

#### Prototype

```
void __d128_to_unsigned_BCD (_Decimal128, unsigned long long *upper,  
unsigned long long *lower);
```

#### Parameters

*upper*

The address of the variable that will hold the upper digits of the result.

*lower*

The address of the variable that will hold the lower digits of the result.

#### Return value

Produces 32 decimal digits with no sign in a 128-bit result. Digits to the left are ignored. The higher 16 digits are stored in the parameter *upper*. The lower 16 digits are stored in the parameter *lower*.

## Usage

You can use the `__d128_shift_right` function to access the digits to the left.

### **`__signed_BCD_to_d64`**

#### **Purpose**

Convert from Signed Binary Coded Decimal.

Converts a 64-bit Signed Packed Format (packed decimal - 15 decimal digits followed by a decimal sign) to a 64-bit decimal floating-point value.

#### **Prototype**

```
_Decimal64 __signed_BCD_to_d64 (unsigned long long);
```

#### **Parameters**

The signs 0xA, 0xC, 0xE, and 0xF are treated as positive. The signs 0xB and 0xD are treated as negative.

### **`__signed_BCD_to_d128`**

#### **Purpose**

Convert from Signed Binary Coded Decimal.

Converts a 128-bit Signed Packed Format (packed decimal - 31 decimal digits followed by a decimal sign) to a 128-bit decimal floating-point value.

#### **Prototype**

```
_Decimal128 __signed_BCD_to_d128 ( unsigned long long upper, unsigned long long lower);
```

#### **Parameters**

*upper*

The upper 16 digits of the input value.

*lower*

The lower 15 digits plus the sign of the input value.

#### **Parameters**

The signs 0xA, 0xC, 0xE, and 0xF are treated as positive. The signs 0xB and 0xD are treated as negative.

### **`__unsigned_BCD_to_d64`**

#### **Purpose**

Convert from Unsigned Binary Coded Decimal.

Converts a 64-bit Unsigned Packed Format (16 decimal digits with no sign) to a 64-bit decimal floating-point value.

## Prototype

`_Decimal64 __unsigned_BCD_to_d64` (unsigned long long);

## `__unsigned_BCD_to_d128`

### Purpose

Convert from Unsigned Binary Coded Decimal.

Converts a 128-bit Unsigned Packed Format (32 decimal digits with no sign) to a 128-bit decimal floating-point value.

## Prototype

`_Decimal128 __unsigned_BCD_to_d128` ( unsigned long long *upper*, unsigned long long *lower*);

## Parameters

*upper*

The upper 16 digits of the input value.

*lower*

The lower 16 digits of the input value.

## Exponent functions

Exponent functions extract the exponent from a value or insert an exponent into a value, primarily to support decimal-floating point conversion library functions. They use special values to identify or specify the exponent type.

Table 50. Biased exponents macros and values

Macro	Integer value
DFP_BIASED_EXPONENT_FINITE	0
DFP_BIASED_EXPONENT_INFINITY	-1
DFP_BIASED_EXPONENT_QNAN	-2
DFP_BIASED_EXPONENT_SNAN	-3

## `__d64_biased_exponent, __d128_biased_exponent`

### Purpose

Extract Biased Exponent

Returns the exponent of a decimal floating-point value as an integer.

## Prototype

`long __d64_biased_exponent` (\_Decimal64);

`long __d128_biased_exponent` (\_Decimal128);

## Return value

Returns special values for infinity, quiet NaN, and signalling NaN, as listed in Table 50.

For finite values, the result is DFP\_BIASED\_EXPONENT\_FINITE plus the exponent bias (398 for `_Decimal64`, 6176 for `_Decimal128`) plus the actual exponent.

### **\_\_d64\_insert\_biased\_exponent, \_\_d128\_insert\_biased\_exponent** **Purpose**

Insert Biased Exponent

Replaces the exponent of a decimal floating-point value.

### **Prototype**

```
_Decimal64 __d64_insert_biased_exponent (_Decimal64, long exponent);  
_Decimal128 __d128_insert_biased_exponent (_Decimal128, long exponent);
```

### **Parameters**

*exponent*

The exponent value to be applied to the first parameter. For infinity, quiet NaN and signalling NaN, use one of the compile-time constant values or macros listed in Table 50 on page 507.

For finite values, the result is DFP\_BIASED\_EXPONENT\_FINITE plus the exponent bias (398 for `_Decimal64`, 6176 for `_Decimal128`) plus the corresponding exponent.

## **NaN functions**

NaN functions create quiet or signaling NaNs.

### **\_\_d32\_sNaN, \_\_d64\_sNaN, \_\_d128\_sNaN** **Purpose**

Make Signalling NaN

Creates a signalling NaN of the specified precision, with a positive sign and zero payload.

### **Prototype**

```
_Decimal32 __d32_sNaN (void);  
_Decimal64 __d64_sNaN (void);  
_Decimal128 __d128_sNaN (void);
```

### **\_\_d32\_qNaN, \_\_d64\_qNaN, \_\_d128qNaN** **Purpose**

Make Quiet NaN

Creates a quiet NaN of the specified precision, with a positive sign and zero payload.

## Prototype

```
_Decimal32 __d32_qNaN (void);  
_Decimal64 __d64_qNaN (void);  
_Decimal128 __d128_qNaN (void);
```

## Register transfer functions

Register transfer functions transfer data between general purpose registers and floating-point registers. No conversion occurs. Register transfer functions handle integer data in floating-point registers or floating-point data in general purpose registers. These functions use instructions that are available with **-qarch=pwr6** or **-qarch=pwr6e** only, on a POWER6 running in POWER6e (raw) mode.

### **\_\_gpr\_to\_d64**

#### **Purpose**

Transfer from General Purpose Register to Floating-Point Register

Transfers a value from a general purpose register (64-bit mode) or a general purpose register pair (32-bit mode).

#### **Prototype**

```
_Decimal64 __gpr_to_d64 (long long);
```

### **\_\_gprs\_to\_d128**

#### **Purpose**

Transfer from General Purpose Register to Floating-Point Register.

Transfers a value from a pair of general purpose registers (64-bit mode) or four general purpose registers (32-bit mode).

#### **Prototype**

```
_Decimal128 __gprs_to_d128 (unsigned long long*upper, unsigned long  
long*lower);
```

#### **Parameters**

*upper*

The address of the variable that will hold the upper 64 bits of the result.

*lower*

The address of the variable that will hold the lower 64 bits of the result.

#### **Return value**

The higher 64 bits are stored in the parameter *upper*. The lower 64 bits are stored in the parameter *lower*.

### **\_\_d64\_to\_gpr**

#### **Purpose**

Transfer from Floating-Point Register to General Purpose Register.

Transfers a value from a floating-point register to a general purpose register (64-bit mode) or a general purpose register pair (32-bit mode).

### Prototype

```
long long __d64_to_gpr (_Decimal64);
```

### **\_\_d128\_to\_gprs** Purpose

Transfer from Floating-Point Register to General Purpose Register.

Transfers a value from a pair of floating-point registers to a pair of general purpose registers (64-bit mode) or four general purpose registers (32-bit mode).

### Prototype

```
void __d128_to_gprs (_Decimal128, unsigned long long*upper, unsigned long long*lower);
```

### Parameters

*upper*

The address of the variable that contains the upper 64 bits of the input value.

*lower*

The address of the variable that contains the lower 64 bits of the input value.

## Rounding functions

Rounding functions perform operations such as rounding and truncation of floating-point values.

Table 51. Rounding mode macros and values

Macro	Integer value
DFP_ROUND_TO_NEAREST_WITH_TIES_TO_EVEN	0
DFP_ROUND_TOWARD_ZERO	1
DFP_ROUND_TOWARD_POSITIVE_INFINITY	2
DFP_ROUND_TOWARD_NEGATIVE_INFINITY	3
DFP_ROUND_TO_NEAREST_WITH_TIES_AWAY_FROM_ZERO	4
DFP_ROUND_TO_NEAREST_WITH_TIES_TOWARD_ZERO	5
DFP_ROUND_AWAY_FROM_ZERO	6
DFP_ROUND_TO_PREPARE_FOR_SHORTER_PRECISION	7
DFP_ROUND_USING_CURRENT_MODE <sup>1</sup>	8

### Note:

1. This value is valid only for functions that override the current rounding mode; it is not valid for `__dfp_set_rounding_mode` and can not be returned by `__dfp_get_rounding_mode`.

### **\_\_d64\_integral, \_\_d128\_integral**

#### Purpose

Round to Integral



Rounds a decimal floating-point value to an integer, allowing an Inexact exception to be raised.

### Prototype

```
_Decimal64 __d64_integral (_Decimal64);
```

```
_Decimal128 __d128_integral (_Decimal128);
```

### Return value

The integer is returned in decimal floating-point format, rounded using the current rounding mode. Digits after the decimal point are discarded.

### **\_\_d64\_integral\_no\_inexact, \_\_d128\_integral\_no\_inexact**

#### Purpose

Round to Integral

Rounds a decimal floating-point value to an integer, suppressing any Inexact exception from being raised.

### Prototype

```
_Decimal64 __d64_integral_no_inexact (_Decimal64);
```

```
_Decimal128 __d128_integral_no_inexact (_Decimal128);
```

### Return value

The integer is returned in decimal floating-point format, rounded using the current rounding mode. Digits after the decimal point are discarded.

### **\_\_d64\_quantize, \_\_d128\_quantize**

#### Purpose

Quantize

Returns the arithmetic value of the first parameter, with the exponent adjusted to match the second parameter, using a specified rounding mode.

### Prototype

```
_Decimal64 __d64_quantize (_Decimal64, _Decimal64, long rounding_mode);
```

```
_Decimal128 __d128_quantize (_Decimal128, _Decimal128, long  
rounding_mode);
```

### Parameters

*rounding\_mode*

One of the compile-time constant values or macros defined in Table 51 on page 510.

## Usage

These functions temporarily override the rounding mode in effect for the current operation.

### `__d64_reround`, `__d128_reround`

#### Purpose

Reround

Complete rounding of a partially rounded value, avoiding double rounding which causes errors.

#### Prototype

```
_Decimal64 __d64_reround (_Decimal64, unsigned long number_of_digits,  
unsigned long rounding_mode);
```

```
_Decimal128 __d128_reround (_Decimal128, unsigned long number_of_digits,  
unsigned long rounding_mode);
```

#### Parameters

*number\_of\_digits*

The number of digits to round to, from 1 to 15 for `__d64_reround` and from 1 to 33 for `__d128_reround`.

*rounding\_mode*

One of the compile-time constant values or macros defined in Table 51 on page 510.

## Usage

These functions temporarily override the rounding mode in effect for the current operation. The value to be rerounded should have been previously rounded using mode `DFP_ROUND_TO_PREPARE_FOR_SHORTER_PRECISION` or 7 to ensure correct rounding.

## Test functions

Test functions allow extended exception handling of invalid results or categorization of input values, primarily to support math library functions.

Those functions that begin with `__d64_is` or `__d128_is` will not raise an exception, even for signaling NaNs.

Table 52. Test data class mask macros and values

Macro	Integer value
<code>DFP_PPC_DATA_CLASS_ZERO</code>	0x20
<code>DFP_PPC_DATA_CLASS_SUBNORMAL</code>	0x10
<code>DFP_PPC_DATA_CLASS_NORMAL</code>	0x08
<code>DFP_PPC_DATA_CLASS_INFINITY</code>	0x04
<code>DFP_PPC_DATA_CLASS_QUIET_NAN</code>	0x02
<code>DFP_PPC_DATA_CLASS_SIGNALING_NAN</code>	0x01

Table 53. Test data group mask macros and values

Macro	Integer value
DFP_PPC_DATA_GROUP_SAFE_ZERO	0x20
DFP_PPC_DATA_GROUP_ZERO_WITH_EXTREME_EXPONENT	0x10
DFP_PPC_DATA_GROUP_NONZERO_WITH_EXTREME_EXPONENT	0x08
DFP_PPC_DATA_GROUP_SAFE_NONZERO	0x04
DFP_PPC_DATA_GROUP_NONZERO_LEFTMOST_DIGIT_NONEXTREME_EXPONENT	0x02
DFP_PPC_DATA_GROUP_SPECIAL	0x01

Table 54. Test data class and group result macros and values

Macro	Integer value
DFP_PPC_DATA_POSITIVE_NO_MATCH	0x00
DFP_PPC_DATA_POSITIVE_MATCH	0x02
DFP_PPC_DATA_NEGATIVE_NO_MATCH	0x08
DFP_PPC_DATA_NEGATIVE_MATCH	0x0A

Table 55. Test data class and group result mask macros and values

Macro	Integer value
DFP_PPC_DATA_NEGATIVE_MASK	0x08
DFP_PPC_DATA_MATCH_MASK	0x02

## **\_\_d64\_same\_quantum, \_\_d128\_same\_quantum**

### **Purpose**

Same Quantum

Returns true if two values have the same quantum

### **Prototype**

```
_Bool __d64_same_quantum (_Decimal64, _Decimal64);
```

```
_Bool __d128_same_quantum (_Decimal28, _Decimal128);
```

## **\_\_d64\_issigned, \_\_d128\_issigned**

### **Purpose**

Is Signed

Returns true if the parameter is negative, negative zero, negative infinity, or negative NaN.

### **Prototype**

```
_Bool __d64_issigned (_Decimal64);
```

```
_Bool __d128_issigned (_Decimal128);
```

## **\_\_d64\_isnormal, \_\_d128\_isnormal**

### **Purpose**

Is Normal

Returns true if the parameter is in the normal range (that is, not a subnormal, infinity or NaN) and not zero.

### **Prototype**

```
_Bool __d64_isnormal (_Decimal64);
```

```
_Bool __d128_isnormal (_Decimal128);
```

## **\_\_d64\_isfinite, \_\_d128\_isfinite**

### **Purpose**

Is Finite

Returns true if the parameter is not positive or negative infinity and not a quiet or signaling NaN.

### **Prototype**

```
_Bool __d64_isfinite (_Decimal64);
```

```
_Bool __d128_isfinite (_Decimal128);
```

## **\_\_d64\_iszero, \_\_d128\_iszero**

### **Purpose**

Is Zero

Returns true if the parameter is positive or negative zero.

### **Prototype**

```
_Bool __d64_iszero (_Decimal64);
```

```
_Bool __d128_iszero (_Decimal128);
```

## **\_\_d64\_issubnormal, \_\_d128\_issubnormal**

### **Purpose**

Is Subnormal

Returns true if the parameter is a subnormal.

### **Prototype**

```
_Bool __d64_issubnormal (_Decimal64);
```

```
_Bool __d128_issubnormal (_Decimal128);
```

## **\_\_d64\_isinf, \_\_d128\_isinf**

### **Purpose**

Is Infinity

Returns true if the parameter is positive or negative infinity.

### **Prototype**

```
_Bool __d64_isinf (_Decimal64);
```

```
_Bool __d128_isinf (_Decimal128);
```

## **\_\_d64\_isnan, \_\_d128\_isnan**

### **Purpose**

Is NaN

Returns true if the parameter is a positive or negative quiet or signaling NaN.

### **Prototype**

```
_Bool __d64_isnan (_Decimal64);
```

```
_Bool __d128_isnan (_Decimal128);
```

## **\_\_d64\_issignaling, \_\_d128\_issignaling**

### **Purpose**

Is Signaling NaN

Returns true if the parameter is a positive or negative signaling NaN.

### **Prototype**

```
_Bool __d64_issignaling (_Decimal64);
```

```
_Bool __d128_issignaling (_Decimal128);
```

## **\_\_d64\_test\_data\_class, \_\_d128\_test\_data\_class**

### **Purpose**

Test Data Class

Reports if a value is a zero, subnormal, normal, infinity, quiet NaN or signaling NaN, and if the value is positive or negative.

### **Prototype**

```
long __d64_test_data_class (_Decimal64, unsigned long mask);
```

```
long __d128_test_data_class (_Decimal128, unsigned long mask);
```

## Parameters

*mask*

One of the values or macros defined in Table 52 on page 512 or several ORed together. The parameter must be a compile time constant expression.

## Return value

One of the values listed in Table 54 on page 513.

## Usage

You can use an appropriate mask to check combinations of values at the same time. Use the masks listed in Table 52 on page 512 to check input values. Use the masks listed in Table 55 on page 513 to check result values.

## **\_\_d64\_test\_data\_group, \_\_d128\_test\_data\_group**

### **Purpose**

Test Data Group

Reports if a value is a safe zero, a zero with an extreme exponent, a subnormal, a safe nonzero, a normal with no leading zero, or an infinity or NaN and if the value is positive or negative. Safe means leading zero digits and a non-extreme exponent. A subnormal can appear as either an extreme nonzero or safe nonzero. The exact meaning of some masks depends on the particular CPU model.

### **Prototype**

```
long __d64_test_data_group (_Decimal64, unsigned long mask);
```

```
long __d128_test_data_group (_Decimal128, unsigned long mask);
```

## Parameters

*mask*

One of the values or macros defined in Table 53 on page 513 or several ORed together. The parameter must be a compile time constant expression.

## Return value

One of the values listed in Table 54 on page 513.

## Usage

You can use an appropriate mask to check combinations of values at the same time. Use the masks listed in Table 53 on page 513 to check input values. Use the masks listed in Table 55 on page 513 to check result values.

## **\_\_d64\_test\_significance, \_\_d128\_test\_significance**

### **Purpose**

Test Significance

Checks whether a decimal floating-point value has a specified number of digits of significance.

## Prototype

```
long __d64_test_significance (_Decimal64, unsigned long digits);  
long __d128_test_significance (_Decimal128, unsigned long digits);
```

## Parameters

### *digits*

The number of digits of significance to be tested for. *digits* must be in the range 0 to 63; otherwise the result is undefined. If it is 0, all values including zero will be considered to have more significant digits, if it is not 0, a zero value will be considered to have fewer significant digits.

## Return value

Returns the following values:

- Less than 0 if the number of digits of significance of the first parameter is less than the second parameter.
- 0 if the number of digits of significance is the same as the second parameter.
- Greater than 0 if the number of digits of significance of the first parameter is greater than that of the second parameter or *digits* is 0.
- -2 if either parameter is a quiet or signaling NaN or positive or negative infinity.

For these functions, the number of significant digits of the value 0 is considered to be zero.

## Miscellaneous functions

This section lists the miscellaneous decimal floating-point built-in functions.

### **\_\_addg6s**

#### **Purpose**

Add and Generate Sixes

Adds *source1* to *source2* and produces 16 carry bits, one for each carry out of decimal position *n* (bit position 4*xn*).

The result is a doubleword composed of the 16 carry bits. The doubleword consists of a decimal six (0b0110) in every decimal digit position for which the corresponding carry bit is 0, and a zero (0b0000) in every position for which the corresponding carry bit is 1.

#### **Prototype**

```
long long __addg6s (long long source1, long long source2);
```

#### **Usage**

Valid only when **-qarch** is set to target POWER7 processors in 64-bit mode.

---

## Synchronization and atomic built-in functions

Synchronization and atomic built-in functions are grouped into the following categories:

- “Check lock functions”
- “Clear lock functions” on page 519
- “Compare and swap functions” on page 520
- “Fetch functions” on page 521
- “Load functions” on page 523
- “Store functions” on page 523
- “Synchronization functions” on page 524

### Check lock functions

`__check_lock_mp`, `__check_lockd_mp`

#### Purpose

Check Lock on Multiprocessor Systems, Check Lock Doubleword on Multiprocessor Systems

Conditionally updates a single word or doubleword variable atomically.

#### Prototype

```
unsigned int __check_lock_mp (const int* addr, int old_value, int new_value);
```

```
unsigned int __check_lockd_mp (const long* addr, long old_value, long new_value);
```

#### Parameters

*addr*

The address of the variable to be updated. Must be aligned on a 4-byte boundary for a single word or on an 8-byte boundary for a doubleword.

*old\_value*

The old value to be checked against the current value in *addr*.

*new\_value*

The new value to be conditionally assigned to the variable in *addr*,

#### Return value

Returns false (0) if the value in *addr* was equal to *old\_value* and has been set to the *new\_value*. Returns true (1) if the value in *addr* was not equal to *old\_value* and has been left unchanged.

#### Usage

`__check_lockd_mp` is valid only in 64-bit mode.



## **\_\_check\_lock\_up, \_\_check\_lockd\_up**

### **Purpose**

Check Lock on Uniprocessor Systems, Check Lock Doubleword on Uniprocessor Systems

Conditionally updates a single word or doubleword variable atomically.

### **Prototype**

```
unsigned int __check_lock_up (const int* addr, int old_value, int new_value);
```

```
unsigned int __check_lockd_up (const long* addr, long old_value, long new_value);
```

### **Parameters**

*addr*

The address of the variable to be updated. Must be aligned on a 4-byte boundary for a single word and on an 8-byte boundary for a doubleword.

*old\_value*

The old value to be checked against the current value in *addr*.

*new\_value*

The new value to be conditionally assigned to the variable in *addr*,

### **Return value**

Returns false (0) if the value in *addr* was equal to *old\_value* and has been set to the new value. Returns true (1) if the value in *addr* was not equal to *old\_value* and has been left unchanged.

### **Usage**

`__check_lockd_up` is valid only in 64-bit mode.

## **Clear lock functions**

### **\_\_clear\_lock\_mp, \_\_clear\_lockd\_mp**

#### **Purpose**

Clear Lock on Multiprocessor Systems, Clear Lock Doubleword on Multiprocessor Systems

Atomic store of the *value* into the variable at the address *addr*.

#### **Prototype**

```
void __clear_lock_mp (const int* addr, int value);
```

```
void __clear_lockd_mp (const long* addr, long value);
```

#### **Parameters**

*addr*

The address of the variable to be updated. Must be aligned on a 4-byte boundary for a single word and on an 8-byte boundary for a doubleword.

*value*

The new value to be assigned to the variable in *addr*,

### Usage

`__clear_lockd_mp` is only valid in 64-bit mode.

### `__clear_lock_up`, `__clear_lockd_up`

#### Purpose

Clear Lock on Uniprocessor Systems, Clear Lock Doubleword on Uniprocessor Systems

Atomic store of the *value* into the variable at the address *addr*.

#### Prototype

```
void __clear_lock_up (const int* addr, int value);
```

```
void __clear_lockd_up (const long* addr, long value);
```

#### Parameters

*addr*

The address of the variable to be updated. Must be aligned on a 4-byte boundary for a single word and on an 8-byte boundary for a doubleword.

*value*

The new value to be assigned to the variable in *addr*.

### Usage

`__clear_lockd_up` is only valid in 64-bit mode.

## Compare and swap functions

### `__compare_and_swap`, `__compare_and_swaplp`

#### Purpose

Conditionally updates a single word or doubleword variable atomically.

#### Prototype

```
int __compare_and_swap (volatile int* addr, int* old_val_addr, int new_val);
```

```
int __compare_and_swaplp (volatile long* addr, long* old_val_addr, long new_val);
```

#### Parameters

*addr*

The address of the variable to be copied. Must be aligned on a 4-byte boundary for a single word and on an 8-byte boundary for a doubleword.

*old\_val\_addr*

The memory location into which the value in *addr* is to be copied.

*new\_val*

The value to be conditionally assigned to the variable in *addr*,

### Return value

Returns true (1) if the value in *addr* was equal to *old\_value* and has been set to the new value. Returns false (0) if the value in *addr* was not equal to *old\_value* and has been left unchanged. In either case, the contents of the memory location specified by *addr* are copied into the memory location specified by *old\_val\_addr*.

### Usage

The `__compare_and_swap` function is useful when a single word value must be updated only if it has not been changed since it was last read. If you use `__compare_and_swap` as a locking primitive, insert a call to the `__i_sync` built-in function at the start of any critical sections.

`__compare_and_swaplp` is valid only in 64-bit mode.

## Fetch functions

### `__fetch_and_and`, `__fetch_and_andlp`

#### Purpose

Clears bits in the word or doubleword specified by *addr* by AND-ing that value with the value specified by *val*, in a single atomic operation, and returns the original value of *addr*.

#### Prototype

```
unsigned int __fetch_and_and (volatile unsigned int* addr, unsigned int val);
```

```
unsigned long __fetch_and_andlp (volatile unsigned long* addr, unsigned long val);
```

#### Parameters

*addr*

The address of the variable to be ANDed. Must be aligned on a 4-byte boundary for a single word and on an 8-byte boundary for a doubleword.

*value*

The value by which the value in *addr* is to be ANDed.

#### Usage

This operation is useful when a variable containing bit flags is shared between several threads or processes.

`__fetch_and_andlp` is valid only in 64-bit mode.

### `__fetch_and_or`, `__fetch_and_orlp`

#### Purpose

Sets bits in the word or doubleword specified by *addr* by OR-ing that value with the value specified *val*, in a single atomic operation, and returns the original value of *addr*.

## Prototype

```
unsigned int __fetch_and_or (volatile unsigned int* addr, unsigned int val);
```

```
unsigned long __fetch_and_orlp (volatile unsigned long* addr, unsigned long val);
```

## Parameters

*addr*

The address of the variable to be ORed. Must be aligned on a 4-byte boundary for a single word and on an 8-byte boundary for a doubleword.

*value*

The value by which the value in *addr* is to be ORed.

## Usage

This operation is useful when a variable containing bit flags is shared between several threads or processes.

`__fetch_and_orlp` is valid only in 64-bit mode.

## `__fetch_and_swap`, `__fetch_and_swaplp`

### Purpose

Sets the word or doubleword specified by *addr* to the value of *val* and returns the original value of *addr*, in a single atomic operation.

## Prototype

```
unsigned int __fetch_and_swap (volatile unsigned int* addr, unsigned int val);
```

```
unsigned long __fetch_and_swaplp (volatile unsigned long* addr, unsigned long val);
```

## Parameters

*addr*

The address of the variable to be updated. Must be aligned on a 4-byte boundary for a single word and on an 8-byte boundary for a doubleword.

*value*

The value which is to be assigned to *addr*.

## Usage

This operation is useful when a variable is shared between several threads or processes, and one thread needs to update the value of the variable without losing the value that was originally stored in the location.

`__fetch_and_swaplp` is valid only in 64-bit mode.

## Load functions

### `__ldarx`, `__lwarx`

#### Purpose

Load Doubleword and Reserve Indexed, Load Word and Reserve Indexed

Loads the value from the memory location specified by *addr* and returns the result. For `__lwarx`, in 64-bit mode, the compiler returns the sign-extended result.

#### Prototype

```
long __ldarx (volatile long* addr);
```

```
int __lwarx (volatile int* addr);
```

#### Parameters

*addr*

The address of the value to be loaded. Must be aligned on a 4-byte boundary for a single word and on an 8-byte boundary for a doubleword.

#### Usage

This function can be used with a subsequent `__stdcx` (or `__stwcx`) built-in function to implement a read-modify-write on a specified memory location. The two built-in functions work together to ensure that if the store is successfully performed, no other processor or mechanism can modify the target doubleword between the time the `__ldarx` function is executed and the time the `__stdcx` function completes. This has the same effect as inserting `__fence` built-in functions before and after the `__ldarx` built-in function and can inhibit compiler optimization of surrounding code (see “`__alignx`” on page 603 for a description of the `__fence` built-in function).

`__ldarx` is valid only in 64-bit mode.

## Store functions

### `__stdcx`, `__stwcx`

#### Purpose

Store Doubleword Conditional Indexed, Store Word Conditional Indexed

Stores the value specified by *val* into the memory location specified by *addr*.

#### Prototype

```
int __stdcx(volatile long* addr, long val);
```

```
int __stwcx(volatile int* addr, int val);
```

#### Parameters

*addr*

The address of the variable to be updated. Must be aligned on a 4-byte boundary for a single word and on an 8-byte boundary for a doubleword.

*value*

The value which is to be assigned to *addr*.

### Return value

Returns 1 if the update of *addr* is successful and 0 if it is unsuccessful.

### Usage

This function can be used with a preceding `__ldarx` (or `__ldarx`) built-in function to implement a read-modify-write on a specified memory location. The two built-in functions work together to ensure that if the store is successfully performed, no other processor or mechanism can modify the target doubleword between the time the `__ldarx` function is executed and the time the `__stdcx` function completes. This has the same effect as inserting `__fence` built-in functions before and after the `__stdcx` built-in function and can inhibit compiler optimization of surrounding code.

`__stdcx` is valid only in 64-bit mode.

## Synchronization functions

### `__eieio`, `__iospace_eieio`

#### Purpose

Enforce In-order Execution of Input/Output

Ensures that all I/O storage access instructions preceding the call to `__eieio` complete in main memory before I/O storage access instructions following the function call can execute.

#### Prototype

```
void __eieio (void);
```

```
void __iospace_eieio (void);
```

#### Usage

This function is useful for managing shared data instructions where the execution order of load/store access is significant. The function can provide the necessary functionality for controlling I/O stores without the cost to performance that can occur with other synchronization instructions.

### `__isync`

#### Purpose

Instruction Synchronize

Waits for all previous instructions to complete and then discards any prefetched instructions, causing subsequent instructions to be fetched (or refetched) and executed in the context established by previous instructions.

#### Prototype

```
void __isync (void);
```

## **\_\_lwsync, \_\_iospace\_lwsync**

### **Purpose**

Load Word Synchronize

Ensures that all instructions preceding the call to `__lwsync` complete before any subsequent store instructions can be executed on the processor that executed the function. Also, it ensures that all load instructions preceding the call to `__lwsync` complete before any subsequent load instructions can be executed on the processor that executed the function. This allows you to synchronize between multiple processors with minimal performance impact, as `__lwsync` does not wait for confirmation from each processor.

### **Prototype**

```
void __lwsync (void);  
  
void __iospace_lwsync (void);
```

## **\_\_sync, \_\_iospace\_sync**

### **Purpose**

Synchronize

Ensures that all instructions preceding the function the call to `__sync` complete before any instructions following the function call can execute.

### **Prototype**

```
void __sync (void);  
  
void __iospace_sync (void);
```

---

## **Cache-related built-in functions**

Cache-related built-in functions are grouped into the following categories:

- “Data cache functions”
- “Prefetch built-in functions” on page 527

### **Data cache functions**

#### **\_\_dcbf**

##### **Purpose**

Data Cache Block Flush

Copies the contents of a modified block from the data cache to main memory and flushes the copy from the data cache.

##### **Prototype**

```
void __dcbf(const void* addr);
```

## **\_\_dcbfl** **Purpose**

Data Cache Block Flush Line

Flushes the cache line at the specified address from the L1 data cache.

## **Prototype**

```
void __dcbfl (const void* addr );
```

## **Usage**

The target storage block is preserved in the L2 cache.

Valid only when **-qarch** is set to target POWER6 processors

## **\_\_dcbflp** **Purpose**

Data Cache Block Flush Line Primary

Flushes the cache line at address from the primary data cache of a single processor.

## **Prototype**

```
void __dcbflp(const void* address);
```

## **Usage**

Valid only when **-qarch** is set to target POWER7 processors.

## **\_\_dcbst** **Purpose**

Data Cache Block Store

Copies the contents of a modified block from the data cache to main memory.

## **Prototype**

```
void __dcbst(const void* addr);
```

## **\_\_dcbt** **Purpose**

Data Cache Block Touch

Loads the block of memory containing the specified address into the L1 data cache.

## **Prototype**

```
void __dcbt (void* addr);
```



## **\_\_dcbtst**

### **Purpose**

Data Cache Block Touch for Store

Fetches the block of memory containing the specified address into the data cache.

### **Prototype**

```
void __dcbtst(void* addr);
```

## **\_\_dcbz**

### **Purpose**

Data Cache Block set to Zero

Sets a cache line containing the specified address in the data cache to zero (0).

### **Prototype**

```
void __dcbz (void* addr);
```

## **Prefetch built-in functions**

## **\_\_dcbtstt**

### **Purpose**

Store Transient Touch provides a hint that describes a block that the program may perform a store access to. The block is likely to be transient, that is, the time interval during which the program accesses the unit is likely to be short.

### **Prototype**

```
void __dcbtstt (void * address);
```

### **Usage**

Valid only when **-qarch** is set to target POWER7 processors.

## **\_\_dcbtt**

### **Purpose**

Data Cache Block Touch Transient

Load Transient Touch provides a hint that describes a block that the program might perform a load access to. The block is likely to be transient, that is, the time interval during which the program accesses the unit is likely to be short.

### **Prototype**

```
void __dcbtt (void * address);
```

### **Usage**

Valid only when **-qarch** is set to target POWER7 processors.

## **\_\_partial\_dcbt**

### **Purpose**

Partial Data Cache Block Touch

Loads half of the cache line that contains the specified address into the L3 data cache.

### **Prototype**

```
void __partial_dcbt (void * address);
```

### **Usage**

Valid only when **-qarch** is set to target POWER7 processors.

## **\_\_prefetch\_by\_load**

### **Purpose**

Touches a memory location by using an explicit load.

### **Prototype**

```
void __prefetch_by_load (const void*);
```

## **\_\_prefetch\_by\_stream**

### **Purpose**

Touches consecutive memory locations by using an explicit stream.

### **Prototype**

```
void __prefetch_by_stream (const int, const void*);
```

## **\_\_protected\_stream\_count**

### **Purpose**

Sets the number of cache lines for a specific limited-length protected stream.

### **Prototype**

```
void __protected_stream_count (unsigned int unit_cnt, unsigned int stream_ID);
```

### **Parameters**

*unit\_cnt*

The number of cache lines. Must be an integer with a value of 0 to 1023.

*stream\_ID*

An integer with a value 0-7 on POWER5 processors, a value 0 to 15 on POWER6 processors, and a value 0 to 11 on POWER7 processors.

### **Usage**

Valid only when **-qarch** is set to target POWER5, POWER6, or POWER7 processors.

## **\_\_protected\_stream\_count\_depth**

### **Purpose**

Sets the number of cache lines and the prefetch depth for a specific limited-length protected stream.

### **Prototype**

```
void __protected_stream_count_depth (unsigned int unit_cnt, unsigned int prefetch_depth, unsigned int stream_ID);
```

### **Parameters**

*unit\_cnt*

The number of cache lines. Must be an integer with a value of 0 to 1023.

*prefetch\_depth*

A relative, qualitative value which sets the steady-state *fetch-ahead* distance of the prefetches for a stream. The fetch-ahead distance is the number of lines being prefetched in advance of the line from which data is currently being loaded, or to which data is currently being stored. Valid values are as follows:

- 0 The default defined in the Data Stream Control Register.
- 1 None.
- 2 Shallowest.
- 3 Shallow.
- 4 Medium.
- 5 Deep.
- 6 Deeper.
- 7 Deepest.

*stream\_ID*

An integer with a value 0 to 15 on POWER6 processors, and a value 0 to 11 on or POWER7 processors.

### **Usage**

Valid only when **-qarch** is set to target POWER6 or POWER7 processors.

## **\_\_protected\_stream\_go**

### **Purpose**

Starts prefetching all limited-length protected streams.

### **Prototype**

```
void __protected_stream_go (void);
```

### **Usage**

Valid only when **-qarch** is set to target POWER5, POWER6, or POWER7 processors.

## **\_\_protected\_stream\_set, \_\_protected\_unlimited\_stream\_set, \_\_protected\_unlimited\_stream\_set\_go**

### **Purpose**

Establishes a limited- or unlimited-length protected stream which fetches from either incremental (forward) or decremental (backward) memory addresses. The stream is protected from being replaced by any hardware detected streams.

### **Prototype**

```
void __protected_stream_set (unsigned int direction, const void* addr,  
unsigned int stream_ID);
```

```
void _protected_unlimited_stream_set (unsigned int direction, const void* addr,  
unsigned int ID);
```

```
void __protected_unlimited_stream_set_go (unsigned int direction, const void*  
addr, unsigned int stream_ID);
```

### **Parameters**

*direction*

An integer with a value of 1 (forward) or 3 (backward).

*addr*

The beginning of the cache line.

*stream\_ID*

An integer with a value 0-7 on POWER4 and POWER5 processors, a value 0 to 15 on POWER6 processors, and a value 0 to 11 on POWER7 processors.

### **Usage**

Valid only when **-qarch** is set to target POWER5, POWER6, or POWER7 processors.

## **\_\_protected\_stream\_stride**

### **Purpose**

Sets the word-offset of the first unit of the stream *address\_offset*, and stride in word size for protected load or store stream with identifier *stream\_id*

### **Prototype**

```
void __protected_stream_stride (unsigned int address_offset, unsigned int stride,  
unsigned int stream_id);
```

### **Parameters**

*address\_offset*

The address of the first unit of the prefetch variable.

*stride*

This is the distance in the number of words of two consecutive elements of the prefetch stream.

*stream\_id*

An integer with a value 0 to 11 on POWER7 processors.

## Usage

Valid only when `-qarch` is set to target POWER7 processors.

### `__protected_stream_stop`

#### Purpose

Stops prefetching a protected stream.

#### Prototype

```
void __protected_stream_stop (unsigned int stream ID);
```

#### Parameters

*stream\_id*

An integer with a value 0-7 on POWER5 processors, a value 0 to 15 on POWER6 processors, and a value 0 to 11 on POWER7 processors.

## Usage

Valid only when `-qarch` is set to target POWER5, POWER6, or POWER7 processors.

### `__protected_stream_stop_all`

#### Purpose

Stops prefetching all protected streams.

#### Prototype

```
void __protected_stream_stop_all (void);
```

## Usage

Valid only when `-qarch` is set to target POWER5, POWER6, or POWER7 processors.

### `__protected_store_stream_set,` `__protected_unlimited_store_stream_set`

#### Purpose

Establishes a limited- or unlimited-length protected store stream which fetches from either incremental (forward) or decremental (backward) memory addresses. The stream is protected from being replaced by any hardware detected streams.

#### Prototype

```
void _protected_store_stream_set (unsigned int direction, const void* addr,  
unsigned int stream_ID );
```

```
void _protected_unlimited_store_stream_set (unsigned int direction, const  
void* addr, unsigned int stream_ID);
```

## Parameters

*direction*

An integer with a value of 1 (forward) or 3 (backward).

*addr*

The beginning of the cache line.

*stream\_ID*

An integer with a value 0 to 15 on POWER6 processors, and a value 0 to 11 on POWER7 processors.

## Usage

Valid only when `-qarch` is set to target POWER6 and POWER7 processors.

## `__transient_protected_stream_count_depth`

### Purpose

Sets the number of cache lines *unit\_cnt* and the prefetch depth *prefetch\_depth* for the limited length protected load or store stream with identifier *stream\_id*. The term "transient" indicates that the time interval during which the program accesses the stream's memory is likely to be short, so the processor can remove it from the cache earlier.

### Prototype

```
void __transient_protected_stream_count_depth (unsigned int unit_cnt,  
unsigned int prefetch_depth, unsigned int stream_id);
```

## Parameters

*unit\_cnt*

The number of cache lines. Must be an integer with a value of 0 to 1023.

*prefetch\_depth*

A relative, qualitative value which sets the steady-state *fetch-ahead* distance of the prefetches for a stream. The fetch-ahead distance is the number of lines being prefetched in advance of the line from which data is currently being loaded, or to which data is currently being stored. Valid values are as follows:

- 0** The default defined in the Data Stream Control Register.
- 1** None.
- 2** Shallowest.
- 3** Shallow.
- 4** Medium.
- 5** Deep.
- 6** Deeper.
- 7** Deepest.

*stream\_id*

An integer with a value 0 to 11 on POWER7 processors.

## Usage

Valid only when `-qarch` is set to target POWER7 processors.

## **\_\_transient\_unlimited\_protected\_stream\_depth**

### **Purpose**

Sets the prefetch depth *prefetch\_depth* for the unlimited length protected load or store stream with identifier *stream\_id*. The stream is likely to be transient, that is, the time interval during which the program accesses the unit is likely to be short.

### **Prototype**

```
void __transient_unlimited_protected_stream_depth (unsigned int
    prefetch_depth, unsigned int stream_id);
```

### **Parameters**

*prefetch\_depth*

A relative, qualitative value which sets the steady-state *fetch-ahead* distance of the prefetches for a stream. The fetch-ahead distance is the number of lines being prefetched in advance of the line from which data is currently being loaded, or to which data is currently being stored. Valid values are as follows:

- 0** The default defined in the Data Stream Control Register.
- 1** None.
- 2** Shallowest.
- 3** Shallow.
- 4** Medium.
- 5** Deep.
- 6** Deeper.
- 7** Deepest.

*stream\_id*

An integer with a value 0 to 11 on POWER7 processors.

### **Usage**

Valid only when `-qarch` is set to target POWER7 processors.

## **\_\_unlimited\_protected\_stream\_depth**

### **Purpose**

Sets the prefetch depth *prefetch\_depth* for the unlimited length protected load or store stream with identifier *stream\_id*.

### **Prototype**

```
void __unlimited_protected_stream_depth (unsigned int prefetch_depth,
    unsigned int stream_id);
```

### **Parameter**

*prefetch\_depth*

A relative, qualitative value which sets the steady-state *fetch-ahead* distance of the prefetches for a stream. The fetch-ahead distance is the number of lines being prefetched in advance of the line from which data is currently being loaded, or to which data is currently being stored. Valid values are as follows:

- 0 The default defined in the Data Stream Control Register.
- 1 None.
- 2 Shallowest.
- 3 Shallow.
- 4 Medium.
- 5 Deep.
- 6 Deeper.
- 7 Deepest.

*stream\_id*

An integer with a value 0 to 15 on POWER6 processors, and a value 0 to 11 on POWER7 processors.

### Usage

Valid only when `-qarch` is set to target POWER6 or POWER7 processors.

## Block-related built-in functions

### `__bcopy`

#### Purpose

Copies *n* bytes from *src* to *dest*. The result is correct even when both areas overlap.

#### Prototype

```
void __bcopy(const void* src, void* dest, size_t n);
```

#### Parameters

*src*

The source address of data to be copied.

*dest*

The destination address of data to be copied

*n*

The size of the data.

### `__bzero`

#### Purpose

Sets the first *n* bytes of the byte area starting at *s* to zero.

#### Prototype

```
void bzero(void* s, size_t n);
```

#### Parameters

*n*

The size of the data.

*s*

The starting address in the byte area.



---

## Vector built-in functions

Individual elements of vectors can be accessed by using the Vector Multimedia eXtension (VMX) or the Vector Scalar eXtension (VSX) built-in functions. This section provides an alphabetical reference to the VMX and the VSX built-in functions. You can use these functions to manipulate vectors.

You must specify appropriate compiler options for your architecture when you use the built-in functions. Built-in functions that use or return a **vector unsigned long long**, **vector signed long long**, **vector bool long long**, or **vector double** type require an architecture that supports the VSX instruction set extensions, such as POWER7. You must specify **-qarch=pwr7** when you use these types.

This section uses pseudocode description to represent function syntax, as shown below:

```
d=func_name(a, b, c)
```

In the description,

- d represents the return value of the function.
- a, b, and c represent the arguments of the function.
- func\_name is the name of the function.

For example, the syntax for the function `vector double vec_xld2(int, double*);` is represented by `d=vec_xld2(a, b)`.

**Note:** This section only describes the IBM specific vector built-in functions and the AltiVec built-in functions with IBM extensions. For information about the other AltiVec built-in functions, see the AltiVec Application Programming Interface specification.

### vec\_abs

#### Purpose

Returns a vector containing the absolute values of the contents of the given vector.

#### Syntax

```
d=vec_abs(a)
```

#### Result and argument types

The following table describes the types of the returned value and the function arguments.

d	a
vector signed char	vector signed char
vector signed short	vector signed short
vector signed int	vector signed int
vector float	vector float
vector double	vector double

## Result value

The value of each element of the result is the absolute value of the corresponding element of a.

## vec\_add

### Purpose

Returns a vector containing the sums of each set of corresponding elements of the given vectors.

This function emulates the operation on long long vectors.

### Syntax

```
d=vec_add(a, b)
```

### Result and argument types

The following table describes the types of the returned value and the function arguments.

d	a	b
vector signed char	vector signed char	vector signed char
vector unsigned char	vector unsigned char	vector unsigned char
vector signed short	vector signed short	vector signed short
vector unsigned short	vector unsigned short	vector unsigned short
vector signed int	vector signed int	vector signed int
vector unsigned int	vector unsigned int	vector unsigned int
vector signed long long	vector signed long long	vector signed long long
vector unsigned long long	vector unsigned long long	vector unsigned long long
vector float	vector float	vector float
vector double	vector double	vector double

## Result value

The value of each element of the result is the sum of the corresponding elements of a and b. For integer vectors and unsigned vectors, the arithmetic is modular.

## vec\_all\_eq

### Purpose

Tests whether all sets of corresponding elements of the given vectors are equal.

### Syntax

```
d=vec_all_eq(a, b)
```

### Result and argument types

The following table describes the types of the returned value and the function arguments.

<b>d</b>	<b>a</b>	<b>b</b>
int	vector bool char	vector bool char
		vector signed char
		vector unsigned char
	vector signed char	vector bool char
		vector signed char
	vector unsigned char	vector bool char
		vector unsigned char
	vector bool short	vector bool short
		vector signed short
		vector unsigned short
	vector signed short	vector bool short
		vector signed short
	vector unsigned short	vector bool short
		vector unsigned short
	vector bool int	vector bool int
		vector signed int
		vector unsigned int
	vector signed int	vector bool int
		vector signed int
	vector unsigned int	vector bool int
		vector unsigned int
	vector bool long long	vector bool long long
		vector signed long long
		vector unsigned long long
	vector signed long long	vector bool long long
		vector signed long long
	vector unsigned long long	vector bool long long
		vector unsigned long long
	vector float	vector float
	vector double	vector double

## Result value

The result is 1 if each element of a is equal to the corresponding element of b. Otherwise, the result is 0.

## vec\_all\_ge

### Purpose

Tests whether all elements of the first argument are greater than or equal to the corresponding elements of the second argument.

## Syntax

d=vec\_all\_ge(a, b)

## Result and argument types

The following table describes the types of the returned value and the function arguments.

d	a	b
int	vector bool char	vector signed char
		vector unsigned char
	vector signed char	vector bool char
		vector signed char
	vector unsigned char	vector bool char
		vector unsigned char
	vector bool short	vector signed short
		vector unsigned short
	vector signed short	vector bool short
		vector signed short
	vector unsigned short	vector bool short
		vector unsigned short
	vector bool int	vector signed int
		vector unsigned int
	vector signed int	vector bool int
		vector signed int
	vector unsigned int	vector bool int
		vector unsigned int
	vector bool long long	vector signed long long
		vector unsigned long long
vector signed long long	vector bool long long	
	vector signed long long	
vector unsigned long long	vector bool long long	
	vector unsigned long long	
vector float	vector float	
vector double	vector double	

## Result value

The result is 1 if all elements of a are greater than or equal to the corresponding elements of b. Otherwise, the result is 0.

## vec\_all\_gt

### Purpose

Tests whether all elements of the first argument are greater than the corresponding elements of the second argument.

## Syntax

`d=vec_all_gt(a, b)`

## Result and argument types

The following table describes the types of the returned value and the function arguments.

d	a	b
int	vector bool char	vector signed char
		vector unsigned char
	vector signed char	vector bool char
		vector signed char
	vector unsigned char	vector bool char
		vector unsigned char
	vector bool short	vector signed short
		vector unsigned short
	vector signed short	vector bool short
		vector signed short
	vector unsigned short	vector bool short
		vector unsigned short
	vector bool int	vector signed int
		vector unsigned int
	vector signed int	vector bool int
		vector signed int
	vector unsigned int	vector bool int
		vector unsigned int
	vector bool long long	vector signed long long
		vector unsigned long long
vector signed long long	vector bool long long	
	vector signed long long	
vector unsigned long long	vector bool long long	
	vector unsigned long long	
vector float	vector float	
vector double	vector double	

## Result value

The result is 1 if all elements of a are greater than the corresponding elements of b. Otherwise, the result is 0.

## vec\_all\_le

### Purpose

Tests whether all elements of the first argument are less than or equal to the corresponding elements of the second argument.

## Syntax

d=vec\_all\_le(a, b)

## Result and argument types

The following table describes the types of the returned value and the function arguments.

d	a	b
int	vector bool char	vector signed char
		vector unsigned char
	vector signed char	vector bool char
		vector signed char
	vector unsigned char	vector bool char
		vector unsigned char
	vector bool short	vector signed short
		vector unsigned short
	vector signed short	vector bool short
		vector signed short
	vector unsigned short	vector bool short
		vector unsigned short
	vector bool int	vector signed int
		vector unsigned int
	vector signed int	vector bool int
		vector signed int
	vector unsigned int	vector bool int
		vector unsigned int
	vector bool long long	vector signed long long
		vector unsigned long long
vector signed long long	vector bool long long	
	vector signed long long	
vector unsigned long long	vector bool long long	
	vector unsigned long long	
vector float	vector float	
vector double	vector double	

## Result value

The result is 1 if all elements of a are less than or equal to the corresponding elements of b. Otherwise, the result is 0.

## vec\_all\_lt

### Purpose

Tests whether all elements of the first argument are less than the corresponding elements of the second argument.

## Syntax

`d=vec_all_lt(a, b)`

## Result and argument types

The following table describes the types of the returned value and the function arguments.

d	a	b
int	vector bool char	vector signed char
		vector unsigned char
	vector signed char	vector bool char
		vector signed char
	vector unsigned char	vector bool char
		vector unsigned char
	vector bool short	vector signed short
		vector unsigned short
	vector signed short	vector bool short
		vector signed short
	vector unsigned short	vector bool short
		vector unsigned short
	vector bool int	vector signed int
		vector unsigned int
	vector signed int	vector bool int
		vector signed int
	vector unsigned int	vector bool int
		vector unsigned int
	vector bool long long	vector signed long long
		vector unsigned long long
vector signed long long	vector bool long long	
	vector signed long long	
vector unsigned long long	vector bool long long	
	vector unsigned long long	
vector float	vector float	
vector double	vector double	

## Result value

The result is 1 if all elements of a are less than the corresponding elements of b. Otherwise, the result is 0.

## vec\_all\_nan

### Purpose

Tests whether each element of the given vector is a NaN.

## Syntax

`d=vec_all_nan(a)`

## Result and argument types

The following table describes the types of the returned value and the function arguments.

d	a
int	vector float
	vector double

## Result value

The result is 1 if each element of a is a NaN. Otherwise, the result is 0.

## vec\_all\_ne

### Purpose

Tests whether all sets of corresponding elements of the given vectors are not equal.

### Syntax

`d=vec_all_ne(a, b)`

## Result and argument types

The following table describes the types of the returned value and the function arguments.



<b>d</b>	<b>a</b>	<b>b</b>
int	vector bool char	vector signed char
		vector unsigned char
	vector signed char	vector bool char
		vector signed char
	vector unsigned char	vector bool char
		vector unsigned char
	vector bool short	vector signed short
		vector unsigned short
	vector signed short	vector bool short
		vector signed short
	vector unsigned short	vector bool short
		vector unsigned short
	vector bool int	vector signed int
		vector unsigned int
	vector signed int	vector bool int
		vector signed int
	vector unsigned int	vector bool int
		vector unsigned int
	vector bool long long	vector signed long long
		vector unsigned long long
vector signed long long	vector bool long long	
	vector signed long long	
vector unsigned long long	vector bool long long	
	vector unsigned long long	
vector float	vector float	
vector double	vector double	

## Result value

The result is 1 if each element of a is not equal to the corresponding element of b. Otherwise, the result is 0.

## vec\_all\_nge

### Purpose

Tests whether each element of the first argument is not greater than or equal to the corresponding element of the second argument.

### Syntax

```
d=vec_all_nge(a, b)
```

## Result and argument types

The following table describes the types of the returned value and the function arguments.

d	a	b
int	vector float	vector float
	vector double	vector double

## Result value

The result is 1 if each element of a is not greater than or equal to the corresponding element of b. Otherwise, the result is 0.

## vec\_all\_ngt

### Purpose

Tests whether each element of the first argument is not greater than the corresponding element of the second argument.

### Syntax

```
d=vec_all_ngt(a, b)
```

## Result and argument types

The following table describes the types of the returned value and the function arguments.

d	a	b
int	vector float	vector float
	vector double	vector double

## Result value

The result is 1 if each element of a is not greater than the corresponding element of b. Otherwise, the result is 0.

## vec\_all\_nle

### Purpose

Tests whether each element of the first argument is not less than or equal to the corresponding element of the second argument.

### Syntax

```
d=vec_all_nle(a, b)
```

## Result and argument types

The following table describes the types of the returned value and the function arguments.

d	a	b
int	vector float	vector float
	vector double	vector double

### Result value

The result is 1 if each element of a is not less than or equal to the corresponding element of b. Otherwise, the result is 0.

## vec\_all\_nlt

### Purpose

Tests whether each element of the first argument is not less than the corresponding element of the second argument.

### Syntax

d=vec\_all\_nlt(a, b)

### Result and argument types

The following table describes the types of the returned value and the function arguments.

d	a	b
int	vector float	vector float
	vector double	vector double

### Result value

The result is 1 if each element of a is not less than the corresponding element of b. Otherwise, the result is 0.

## vec\_all\_numeric

### Purpose

Tests whether each element of the given vector is numeric (not a NaN).

### Syntax

d=vec\_all\_numeric(a)

### Result and argument types

The following table describes the types of the returned value and the function arguments.

d	a
int	vector float
	vector double

## Result value

The result is 1 if each element of a is numeric (not a NaN). Otherwise, the result is 0.

## vec\_and

### Purpose

Performs a bitwise AND of the given vectors.

### Syntax

```
d=vec_and(a, b)
```

### Result and argument types

The following table describes the types of the returned value and the function arguments.

d	a	b
vector bool char	vector bool char	vector bool char
vector signed char	vector bool char	vector signed char
	vector signed char	vector signed char
		vector bool char
vector unsigned char	vector bool char	vector unsigned char
	vector unsigned char	vector unsigned char
		vector bool char
vector bool short	vector bool short	vector vector bool short
vector signed short	vector bool short	vector signed short
	vector signed short	vector signed short
		vector bool short
vector unsigned short	vector bool short	vector unsigned short
	vector unsigned short	vector unsigned short
		vector bool short
vector bool int	vector bool int	vector bool int
vector signed int	vector bool int	vector signed int
	vector signed int	vector signed int
		vector bool int
vector unsigned int	vector bool int	vector unsigned int
	vector unsigned int	vector unsigned int
		vector bool int
vector bool long long	vector bool long long	vector bool long long
vector signed long long	vector bool long long	vector signed long long
	vector signed long long	vector signed long long
		vector bool long long

<b>d</b>	<b>a</b>	<b>b</b>
vector unsigned long long	vector bool long long	vector unsigned long long
	vector unsigned long long	vector unsigned long long
		vector bool long long
vector float	vector bool int	vector float
	vector float	vector bool int
		vector float
vector double	vector bool long long	vector double
	vector double	vector double
		vector bool long long

## vec\_andc

### Purpose

Performs a bitwise AND of the first argument and the bitwise complement of the second argument.

### Syntax

`d=vec_andc(a, b)`

### Result and argument types

The following table describes the types of the returned value and the function arguments.

<b>d</b>	<b>a</b>	<b>b</b>
vector bool char	vector bool char	vector bool char
vector signed char	vector bool char	vector signed char
	vector signed char	vector signed char
		vector bool char
vector unsigned char	vector bool char	vector unsigned char
	vector unsigned char	vector unsigned char
		vector bool char
vector bool short	vector bool short	vector bool short
vector signed short	vector bool short	vector signed short
	vector signed short	vector signed short
		vector bool short
vector unsigned short	vector bool short	vector unsigned short
	vector unsigned short	vector unsigned short
		vector bool short
vector bool int	vector bool int	vector bool int
vector signed int	vector bool int	vector signed int
	vector signed int	vector signed int
		vector bool int

<b>d</b>	<b>a</b>	<b>b</b>
vector unsigned int	vector bool int	vector unsigned int
	vector unsigned int	vector unsigned int
		vector bool int
vector bool long long	vector bool long long	vector bool long long
vector signed long long	vector bool long long	vector signed long long
	vector signed long long	vector signed long long
		vector bool long long
vector unsigned long long	vector bool long long	vector unsigned long long
	vector unsigned long long	vector unsigned long long
		vector bool long long
vector float	vector bool int	vector float
	vector float	vector bool int
		vector float
vector double	vector bool long long	vector double
	vector double	vector bool long long
		vector double

## Result value

The result is the bitwise AND of a with the bitwise complement of b.

## vec\_any\_eq

### Purpose

Tests whether any set of corresponding elements of the given vectors are equal.

### Syntax

```
d=vec_any_eq(a, b)
```

### Result and argument types

The following table describes the types of the returned value and the function arguments.

<b>d</b>	<b>a</b>	<b>b</b>
int	vector bool char	vector bool char
		vector signed char
		vector unsigned char
	vector signed char	vector bool char
		vector signed char
	vector unsigned char	vector bool char
		vector unsigned char
	vector bool short	vector bool short
		vector signed short
		vector unsigned short
	vector signed short	vector bool short
		vector signed short
	vector unsigned short	vector bool short
		vector unsigned short
	vector bool int	vector bool int
		vector signed int
		vector unsigned int
	vector signed int	vector bool int
		vector signed int
	vector unsigned int	vector bool int
		vector unsigned int
	vector bool long long	vector bool long long
		vector signed long long
		vector unsigned long long
	vector signed long long	vector bool long long
		vector signed long long
	vector unsigned long long	vector bool long long
		vector unsigned long long
	vector float	vector float
	vector double	vector double

### Result value

The result is 1 if any element of a is equal to the corresponding element of b. Otherwise, the result is 0.

### vec\_any\_ge

#### Purpose

Tests whether any element of the first argument is greater than or equal to the corresponding element of the second argument.

## Syntax

d=vec\_any\_ge(a, b)

## Result and argument types

The following table describes the types of the returned value and the function arguments.

d	a	b
int	vector bool char	vector signed char
		vector unsigned char
	vector signed char	vector bool char
		vector signed char
	vector unsigned char	vector bool char
		vector unsigned char
	vector bool short	vector signed short
		vector unsigned short
	vector signed short	vector signed short
		vector bool short
	vector unsigned short	vector bool short
		vector unsigned short
	vector bool int	vector signed int
		vector unsigned int
	vector signed int	vector bool int
		vector signed int
	vector unsigned int	vector bool int
		vector unsigned int
	vector bool long long	vector signed long long
		vector unsigned long long
vector signed long long	vector bool long long	
	vector signed long long	
vector unsigned long long	vector bool long long	
	vector unsigned long long	
vector float	vector float	
vector double	vector double	

## Result value

The result is 1 if any element of a is greater than or equal to the corresponding element of b. Otherwise, the result is 0.

## vec\_any\_gt

### Purpose

Tests whether any element of the first argument is greater than the corresponding element of the second argument.



## Syntax

d=vec\_any\_gt(a, b)

## Result and argument types

The following table describes the types of the returned value and the function arguments.

d	a	b
int	vector bool char	vector signed char
		vector unsigned char
	vector signed char	vector bool char
		vector signed char
	vector unsigned char	vector bool char
		vector unsigned char
	vector bool short	vector signed short
		vector unsigned short
	vector signed short	vector signed short
		vector bool short
	vector unsigned short	vector bool short
		vector unsigned short
	vector bool int	vector signed int
		vector unsigned int
	vector signed int	vector bool int
		vector signed int
	vector unsigned int	vector bool int
		vector unsigned int
	vector bool long long	vector signed long long
		vector unsigned long long
vector signed long long	vector bool long long	
	vector signed long long	
vector unsigned long long	vector bool long long	
	vector unsigned long long	
vector float	vector float	
vector double	vector double	

## Result value

The result is 1 if any element of a is greater than the corresponding element of b. Otherwise, the result is 0.

## vec\_any\_le

### Purpose

Tests whether any element of the first argument is less than or equal to the corresponding element of the second argument.

## Syntax

d=vec\_any\_lt(a, b)

## Result and argument types

The following table describes the types of the returned value and the function arguments.

d	a	b
int	vector bool char	vector signed char
		vector unsigned char
	vector signed char	vector bool char
		vector signed char
	vector unsigned char	vector bool char
		vector unsigned char
	vector bool short	vector signed short
		vector unsigned short
	vector signed short	vector signed short
		vector bool short
	vector unsigned short	vector bool short
		vector unsigned short
	vector bool int	vector signed int
		vector unsigned int
	vector signed int	vector bool int
		vector signed int
	vector unsigned int	vector bool int
		vector unsigned int
	vector bool long long	vector signed long long
		vector unsigned long long
vector signed long long	vector bool long long	
	vector signed long long	
vector unsigned long long	vector bool long long	
	vector unsigned long long	
vector float	vector float	
vector double	vector double	

## Result value

The result is 1 if any element of a is less than or equal to the corresponding element of b. Otherwise, the result is 0.

## vec\_any\_lt

### Purpose

Tests whether any element of the first argument is less than the corresponding element of the second argument.

## Syntax

d=vec\_any\_lt(a, b)

## Result and argument types

The following table describes the types of the returned value and the function arguments.

d	a	b
int	vector bool char	vector signed char
		vector unsigned char
	vector signed char	vector bool char
		vector signed char
	vector unsigned char	vector bool char
		vector unsigned char
	vector bool short	vector signed short
		vector unsigned short
	vector signed short	vector signed short
		vector bool short
	vector unsigned short	vector bool short
		vector unsigned short
	vector bool int	vector signed int
		vector unsigned int
	vector signed int	vector bool int
		vector signed int
	vector unsigned int	vector bool int
		vector unsigned int
	vector bool long long	vector signed long long
		vector unsigned long long
vector signed long long	vector bool long long	
	vector signed long long	
vector unsigned long long	vector bool long long	
	vector unsigned long long	
vector float	vector float	
vector double	vector double	

## Result value

The result is 1 if any element of a is less than the corresponding element of b. Otherwise, the result is 0.

## vec\_any\_nan

### Purpose

Tests whether any element of the given vector is a NaN.

## Syntax

`d=vec_any_nan(a)`

## Result and argument types

The following table describes the types of the returned value and the function arguments.

d	a
int	vector float
	vector double

## Result value

The result is 1 if any element of a is a NaN. Otherwise, the result is 0.

## vec\_any\_ne

### Purpose

Tests whether any set of corresponding elements of the given vectors are not equal.

### Syntax

`d=vec_any_ne(a, b)`

## Result and argument types

The following table describes the types of the returned value and the function arguments.

<b>d</b>	<b>a</b>	<b>b</b>
int	vector bool char	vector bool char
		vector signed char
		vector unsigned char
	vector signed char	vector bool char
		vector signed char
	vector unsigned char	vector bool char
		vector unsigned char
	vector bool short	vector bool short
		vector signed short
		vector unsigned short
	vector signed short	vector bool short
		vector signed short
	vector unsigned short	vector bool short
		vector unsigned short
	vector bool int	vector bool int
		vector signed int
		vector unsigned int
	vector signed int	vector bool int
		vector signed int
	vector unsigned int	vector bool int
		vector unsigned int
	vector bool long long	vector bool long long
		vector signed long long
		vector unsigned long long
	vector signed long long	vector bool long long
		vector signed long long
	vector unsigned long long	vector bool long long
		vector unsigned long long
	vector float	vector float
	vector double	vector double

### Result value

The result is 1 if any element of a is not equal to the corresponding element of b. Otherwise, the result is 0.

### vec\_any\_nge

#### Purpose

Tests whether any element of the first argument is not greater than or equal to the corresponding element of the second argument.

## Syntax

`d=vec_any_nge(a, b)`

## Result and argument types

The following table describes the types of the returned value and the function arguments.

d	a	b
int	vector float	vector float
	vector double	vector double

## Result value

The result is 1 if any element of a is not greater than or equal to the corresponding element of b. Otherwise, the result is 0.

## vec\_any\_ngt

### Purpose

Tests whether any element of the first argument is not greater than the corresponding element of the second argument.

## Syntax

`d=vec_any_ngt(a, b)`

## Result and argument types

The following table describes the types of the returned value and the function arguments.

d	a	b
int	vector float	vector float
	vector double	vector double

## Result value

The result is 1 if any element of a is not greater than the corresponding element of b. Otherwise, the result is 0.

## vec\_any\_nle

### Purpose

Tests whether any element of the first argument is not less than or equal to the corresponding element of the second argument.

## Syntax

`d=vec_any_nle(a, b)`

## Result and argument types

The following table describes the types of the returned value and the function arguments.

d	a	b
int	vector float	vector float
	vector double	vector double

## Result value

The result is 1 if any element of a is not less than or equal to the corresponding element of b. Otherwise, the result is 0.

## vec\_any\_nlt

### Purpose

Tests whether any element of the first argument is not less than the corresponding element of the second argument.

### Syntax

`d=vec_any_nlt(a, b)`

## Result and argument types

The following table describes the types of the returned value and the function arguments.

d	a	b
int	vector float	vector float
	vector double	vector double

## Result value

The result is 1 if any element of a is not less than the corresponding element of b. Otherwise, the result is 0.

## vec\_any\_numeric

### Purpose

Tests whether any element of the given vector is numeric (not a NaN).

### Syntax

`d=vec_any_numeric(a)`

## Result and argument types

The following table describes the types of the returned value and the function arguments.

d	a
int	vector float
	vector double

## Result value

The result is 1 if any element of a is numeric (not a NaN). Otherwise, the result is 0.

## vec\_ceil

### Purpose

Returns a vector containing the smallest representable floating-point integral values greater than or equal to the values of the corresponding elements of the given vector.

**Note:** `vec_ceil` is another name for `vec_roundp`. For details, see “`vec_roundp`” on page 580.

## vec\_cmpeq

### Purpose

Returns a vector containing the results of comparing each set of corresponding elements of the given vectors for equality.

This function emulates the operation on long long vectors.

### Syntax

`d=vec_cmpeq(a, b)`

### Result and argument types

The following table describes the types of the returned value and the function arguments.

d	a	b
vector bool char	vector bool char	vector bool char
	vector signed char	vector signed char
	vector unsigned char	vector unsigned char
vector bool short	vector bool short	vector bool short
	vector signed short	vector signed short
	vector unsigned short	vector unsigned short
vector bool int	vector bool int	vector bool int
	vector signed int	vector signed int
	vector unsigned int	vector unsigned int
	vector float	vector float



<b>d</b>	<b>a</b>	<b>b</b>
vector bool long long	vector bool long long	vector bool long long
	vector signed long long	vector signed long long
	vector unsigned long long	vector unsigned long long
	vector double	vector double

## Result value

For each element of the result, the value of each bit is 1 if the corresponding elements of a and b are equal. Otherwise, the value of each bit is 0.

## vec\_cmpge

### Purpose

Returns a vector containing the results of a greater-than-or-equal-to comparison between each set of corresponding elements of the given vectors.

### Syntax

`d=vec_cmpge(a, b)`

### Result and argument types

The following table describes the types of the returned value and the function arguments.

<b>d</b>	<b>a</b>	<b>b</b>
vector bool char	vector signed char	vector signed char
	vector unsigned char	vector unsigned char
vector bool short	vector signed short	vector signed short
	vector unsigned short	vector unsigned short
vector bool int	vector signed int	vector signed int
	vector unsigned int	vector unsigned int
	vector float	vector float
vector bool long long	vector signed long long	vector signed long long
	vector unsigned long long	vector unsigned long long
	vector double	vector double

## Result value

For each element of the result, the value of each bit is 1 if the value of the corresponding element of a is greater than or equal to the value of the corresponding element of b. Otherwise, the value of each bit is 0.

## vec\_cmpgt

### Purpose

Returns a vector containing the results of a greater-than comparison between each set of corresponding elements of the given vectors.

This function emulates the operation on long long vectors.

### Syntax

```
d=vec_cmpgt(a, b)
```

### Result and argument types

The following table describes the types of the returned value and the function arguments.

d	a	b
vector bool char	vector signed char	vector signed char
	vector unsigned char	vector unsigned char
vector bool short	vector signed short	vector signed short
	vector unsigned short	vector unsigned short
vector bool int	vector signed int	vector signed int
	vector unsigned int	vector unsigned int
	vector float	vector float
vector bool long long	vector signed long long	vector signed long long
	vector unsigned long long	vector unsigned long long
	vector double	vector double

### Result value

For each element of the result, the value of each bit is 1 if the value of the corresponding element of a is greater than the value of the corresponding element of b. Otherwise, the value of each bit is 0.

## vec\_cmple

### Purpose

Returns a vector containing the results of a less-than-or-equal-to comparison between each set of corresponding elements of the given vectors.

### Syntax

```
d=vec_cmple(a, b)
```

### Result and argument types

The following table describes the types of the returned value and the function arguments.

<b>d</b>	<b>a</b>	<b>b</b>
vector bool char	vector signed char	vector signed char
	vector unsigned char	vector unsigned char
vector bool short	vector signed short	vector signed short
	vector unsigned short	vector unsigned short
vector bool int	vector signed int	vector signed int
	vector unsigned int	vector unsigned int
	vector float	vector float
vector bool long long	vector signed long long	vector signed long long
	vector unsigned long long	vector unsigned long long
	vector double	vector double

## Result value

For each element of the result, the value of each bit is 1 if the value of the corresponding element of a is less than or equal to the value of the corresponding element of b. Otherwise, the value of each bit is 0.

## vec\_cmplt

### Purpose

Returns a vector containing the results of a less-than comparison between each set of corresponding elements of the given vectors.

This operation emulates the operation on long long vectors.

### Syntax

`d=vec_cmplt(a, b)`

### Result and argument types

The following table describes the types of the returned value and the function arguments.

<b>d</b>	<b>a</b>	<b>b</b>
vector bool char	vector signed char	vector signed char
	vector unsigned char	vector unsigned char
vector bool short	vector signed short	vector signed short
	vector unsigned short	vector unsigned short
vector bool int	vector signed int	vector signed int
	vector unsigned int	vector unsigned int
	vector float	vector float
vector bool long long	vector signed long long	vector signed long long
	vector unsigned long long	vector unsigned long long
	vector double	vector double

## Result value

For each element of the result, the value of each bit is 1 if the value of the corresponding element of *a* is less than the value of the corresponding element of *b*. Otherwise, the value of each bit is 0.

## vec\_cpsgn

### Purpose

Returns a vector by copying the sign of the elements in vector *a* to the sign of the corresponding elements in vector *b*.

This function requires the POWER7 architecture.

### Syntax

`d=vec_cpsgn(a, b)`

### Result and argument types

The following table describes the types of the returned value and the function arguments.

<b>d</b>	<b>a</b>	<b>b</b>
vector float	vector float	vector float
vector double	vector double	vector double

## vec\_ctd

### Purpose

Converts the type of each element in *a* from integer to floating-point single precision and divides the result by 2 to the power of *b*.

### Syntax

`d=vec_ctd(a, b)`

### Result and argument types

The following table describes the types of the returned value and the function arguments.

<b>d</b>	<b>a</b>	<b>b</b>
vector double	vector signed int	0–31
	vector unsigned int	
	vector signed long long	
	vector unsigned long long	

## vec\_ctf

### Purpose

Converts a vector of fixed-point numbers into a vector of floating-point numbers.

## Syntax

`d=vec_ctf(a, b)`

## Result and argument types

The following table describes the types of the returned value and the function arguments.

d	a	b
vector float	vector signed int	0–31
	vector unsigned int	
	vector signed long long	
	vector unsigned long long	

## Result value

The value of each element of the result is the closest floating-point estimate of the value of the corresponding element of a divided by 2 to the power of b.

## vec\_cts

### Purpose

Converts a vector of floating-point numbers into a vector of signed fixed-point numbers.

### Syntax

`d=vec_cts(a, b)`

## Result and argument types

The following table describes the types of the returned value and the function arguments.

d	a	b
vector signed int	vector float	0–31
	vector double	

## Result value

The value of each element of the result is the saturated value obtained by multiplying the corresponding element of a by 2 to the power of b.

## vec\_ctsl

### Purpose

Multiplies each element in a by 2 to the power of b and converts the result into an integer.

**Note:** This function does not use elements 1 and 3 of a when a is a double vector.

## Syntax

`d=vec_ctsl(a, b)`

## Result and argument types

The following table describes the types of the returned value and the function arguments.

d	a	b
vector signed long long	vector float	0–31
	vector double	

## vec\_ctu

### Purpose

Converts a vector of floating-point numbers into a vector of unsigned fixed-point numbers.

**Note:** Elements 1 and 3 of the result vector are undefined when a is a double vector.

### Syntax

`d=vec_ctu(a, b)`

## Result and argument types

The following table describes the types of the returned value and the function arguments.

d	a	b
vector unsigned int	vector float	0–31
	vector double	

## Result value

The value of each element of the result is the saturated value obtained by multiplying the corresponding element of a by 2 to the power of b.

## vec\_ctul

### Purpose

Multiplies each element in a by 2 to the power of b and converts the result into an unsigned type.

### Syntax

`d=vec_ctul(a, b)`

## Result and argument types

The following table describes the types of the returned value and the function arguments.

d	a	b
vector unsigned long long	vector float	0–31
	vector double	

## Result value

This function does not use elements 1 and 3 of a when a is a float vector.

## vec\_cvf

### Purpose

Converts a single-precision floating-point vector to a double-precision floating-point vector or converts a double-precision floating-point vector to a single-precision floating-point vector.

### Syntax

`d=vec_cvf(a)`

### Result and argument types

The following table describes the types of the returned value and the function arguments.

d	a
vector float	vector double
vector double	vector float

## Result value

When this function converts from vector float to vector double, it converts the types of elements 0 and 2 in the vector.

When this function converts from vector double to vector float, the types of element 1 and 3 in the result vector are undefined.

## vec\_div

### Purpose

Divides the elements in vector a by the corresponding elements in vector b and then assigns the result to corresponding elements in the result vector.

This function emulates the operation on integer vectors. It requires the POWER7 architecture.

### Syntax

`d=vec_div(a, b)`

### Result and argument types

The following table describes the types of the returned value and the function arguments.

d	a	b
vector signed char	vector signed char	vector signed char
vector unsigned char	vector unsigned char	vector unsigned char
vector signed short	vector signed short	vector signed short
vector unsigned short	vector unsigned short	vector unsigned short
vector signed int	vector signed int	vector signed int
vector unsigned int	vector unsigned int	vector unsigned int
vector signed long long	vector signed long long	vector signed long long
vector unsigned long long	vector unsigned long long	vector unsigned long long
vector float	vector float	vector float
vector double	vector double	vector double

## vec\_extract

### Purpose

Returns the value of element a from the vector b.

### Syntax

```
d=vec_extract(a, b)
```

### Result and argument types

The following table describes the types of the returned value and the function arguments.

d	a	b
signed char	vector signed char	signed int
unsigned char	vector unsigned char	
	vector bool char	
signed short	vector signed short	
unsigned short	vector unsigned short	
	vector bool short	
signed int	vector signed int	
unsigned int	vector unsigned int	
	vector bool int	
signed long long	vector signed long long	
unsigned long long	vector unsigned long long	
	vector bool long long	
float	vector float	
double	vector double	

### Result value

This function uses the modulo arithmetic on b to determine the element number. For example, if b is out of range, the compiler uses b modulo the number of



elements in the vector to determine the element position.

## vec\_floor

### Purpose

Returns a vector containing the largest representable floating-point integral values less than or equal to the values of the corresponding elements of the given vector.

**Note:** `vec_floor` is another name for `vec_roundm`. For details, see “`vec_roundm`” on page 580.

## vec\_insert

### Purpose

Returns a copy of the vector `b` with the value of its element `c` replaced by `a`.

### Syntax

`d=vec_insert(a, b, c)`

### Result and argument types

The following table describes the types of the returned value and the function arguments.

d	a	b	c
vector signed char	signed char	vector signed char	signed int
vector unsigned char	unsigned char	vector bool char	
		vector unsigned char	
vector signed short	signed short	vector signed short	
vector unsigned short	unsigned short	vector bool short	
		vector unsigned short	
vector signed int	signed int	vector signed int	
vector unsigned int	unsigned int	vector bool int	
		vector unsigned int	
vector signed long long	signed long long	vector signed long long	
vector unsigned long long	unsigned long long	vector bool long long	
		vector unsigned long long	
vector float	float	vector float	
vector double	double	vector double	

### Result value

This function uses the modulo arithmetic on `c` to determine the element number. For example, if `c` is out of range, the compiler uses `c` modulo the number of elements in the vector to determine the element position.

## vec\_madd

### Purpose

Returns a vector containing the results of performing a fused multiply-add operation for each corresponding set of elements of the given vectors.

### Syntax

```
d=vec_madd(a, b, c)
```

### Result and argument types

The following table describes the types of the returned value and the function arguments.

d	a	b	c
vector float	vector float	vector float	vector float
vector double	vector double	vector double	vector double

### Result value

The value of each element of the result is the product of the values of the corresponding elements of a and b, added to the value of the corresponding element of c.

## vec\_max

### Purpose

Returns a vector containing the maximum value from each set of corresponding elements of the given vectors.

### Syntax

```
d=vec_max(a, b)
```

### Result and argument types

The following table describes the types of the returned value and the function arguments.

d	a	b
vector signed char	vector bool char	vector signed char
	vector signed char	vector signed char
		vector bool char
vector unsigned char	vector bool char	vector unsigned char
	vector unsigned char	vector unsigned char
		vector bool char
vector signed short	vector bool short	vector signed short
	vector signed short	vector signed short
		vector bool short

<b>d</b>	<b>a</b>	<b>b</b>
vector unsigned short	vector bool short	vector unsigned short
	vector unsigned short	vector unsigned short
		vector bool short
vector signed int	vector bool int	vector signed int
	vector signed int	vector signed int
		vector bool int
vector unsigned int	vector bool int	vector unsigned int
	vector unsigned int	vector unsigned int
		vector bool int
vector float	vector float	vector float
vector double	vector double	vector double

## Result value

The value of each element of the result is the maximum of the values of the corresponding elements of a and b.

## vec\_mergeh

### Purpose

Merges the most significant halves of two vectors.

### Syntax

`d=vec_mergeh(a, b)`

### Result and argument types

The following table describes the types of the returned value and the function arguments.

<b>d</b>	<b>a</b>	<b>b</b>
vector bool char	vector bool char	vector bool char
vector signed char	vector signed char	vector signed char
vector unsigned char	vector unsigned char	vector unsigned char
vector bool short	vector bool short	vector bool short
vector signed short	vector signed short	vector signed short
vector unsigned short	vector unsigned short	vector unsigned short
vector bool int	vector bool int	vector bool int
vector signed int	vector signed int	vector signed int
vector unsigned int	vector unsigned int	vector unsigned int
vector bool long long	vector bool long long	vector bool long long
vector signed long long	vector signed long long	vector signed long long
vector unsigned long long	vector unsigned long long	vector unsigned long long
vector float	vector float	vector float
vector double	vector double	vector double

## Result value

Assume that the elements of each vector are numbered beginning with 0. The even-numbered elements of the result are taken, in order, from the elements in the most significant 8 bytes of a. The odd-numbered elements of the result are taken, in order, from the elements in the most significant 8 bytes of b.

## vec\_mergel

### Purpose

Merges the least significant halves of two vectors.

### Syntax

```
d=vec_mergel(a, b)
```

### Result and argument types

The following table describes the types of the returned value and the function arguments.

d	a	b
vector bool char	vector bool char	vector bool char
vector signed char	vector signed char	vector signed char
vector unsigned char	vector unsigned char	vector unsigned char
vector bool short	vector bool short	vector bool short
vector signed short	vector signed short	vector signed short
vector unsigned short	vector unsigned short	vector unsigned short
vector bool int	vector bool int	vector bool int
vector signed int	vector signed int	vector signed int
vector unsigned int	vector unsigned int	vector unsigned int
vector bool long long	vector bool long long	vector bool long long
vector signed long long	vector signed long long	vector signed long long
vector unsigned long long	vector unsigned long long	vector unsigned long long
vector float	vector float	vector float
vector double	vector double	vector double

## Result value

Assume that the elements of each vector are numbered beginning with 0. The even-numbered elements of the result are taken, in order, from the elements in the least significant 8 bytes of a. The odd-numbered elements of the result are taken, in order, from the elements in the least significant 8 bytes of b.

## vec\_min

### Purpose

Returns a vector containing the minimum value from each set of corresponding elements of the given vectors.

## Syntax

d=vec\_min(a, b)

## Result and argument types

The following table describes the types of the returned value and the function arguments.

d	a	b
vector signed char	vector bool char	vector signed char
	vector signed char	vector signed char
		vector bool char
vector unsigned char	vector bool char	vector unsigned char
	vector unsigned char	vector unsigned char
		vector bool char
vector signed short	vector bool short	vector signed short
	vector signed short	vector signed short
		vector bool short
vector unsigned short	vector bool short	vector unsigned short
	vector unsigned short	vector unsigned short
		vector bool short
vector signed int	vector bool int	vector signed int
	vector signed int	vector signed int
		vector bool int
vector unsigned int	vector bool int	vector unsigned int
	vector unsigned int	vector unsigned int
		vector bool int
vector float	vector float	vector float
vector double	vector double	vector double

## Result value

The value of each element of the result is the minimum of the values of the corresponding elements of a and b.

## vec\_msub

### Purpose

Returns a vector containing the results of performing a multiply-subtract operation using the given vectors.

This function requires the POWER7 architecture.

### Syntax

d=vec\_msub(a, b, c)

## Result and argument types

The following table describes the types of the returned value and the function arguments.

d	a	b	c
vector float	vector float	vector float	vector float
vector double	vector double	vector double	vector double

## Result value

This function multiplies each element in a by the corresponding element in b and then subtracts the corresponding element in c from the result.

## vec\_mul

### Purpose

Returns a vector containing the results of performing a multiply operation using the given vectors.

This function emulates the operation on integer vectors. It requires the POWER7 architecture.

### Syntax

```
d=vec_mul(a, b)
```

## Result and argument types

The following table describes the types of the returned value and the function arguments.

d	a	b
vector signed char	vector signed char	vector signed char
vector unsigned char	vector unsigned char	vector unsigned char
vector signed short	vector signed short	vector signed short
vector unsigned short	vector unsigned short	vector unsigned short
vector signed int	vector signed int	vector signed int
vector unsigned int	vector unsigned int	vector unsigned int
vector signed long long	vector signed long long	vector signed long long
vector unsigned long long	vector unsigned long long	vector unsigned long long
vector float	vector float	vector float
vector double	vector double	vector double

## Result value

This function multiplies corresponding elements in the given vectors and then assigns the result to corresponding elements in the result vector.

## vec\_nabs

### Purpose

Returns a vector containing the results of performing a negative-absolute operation using the given vector.

This function requires the POWER7 architecture.

### Syntax

d=vec\_nabs(a)

### Result and argument types

The following table describes the types of the returned value and the function arguments.

d	a
vector float	vector float
vector double	vector double

### Result value

This function computes the absolute value of each element in the given vector and then assigns the negated value of the result to the corresponding elements in the result vector.

## vec\_neg

### Purpose

Returns a vector containing the negated value of the corresponding elements in the given vector.

**Note:** For vector signed long long, this function emulates the operation. This function requires the POWER7 architecture.

### Syntax

d=vec\_neg(a)

### Result and argument types

The following table describes the types of the returned value and the function arguments.

d	a
vector signed char	vector signed char
vector signed short	vector signed short
vector signed int	vector signed int
vector signed long long	vector signed long long
vector float	vector float

<b>d</b>	<b>a</b>
vector double	vector double

## Result value

This function multiplies the value of each element in the given vector by -1.0 and then assigns the result to the corresponding elements in the result vector.

## vec\_nmadd

### Purpose

Returns a vector containing the results of performing a negative multiply-add operation on the given vectors.

This function requires the POWER7 architecture.

### Syntax

`d=vec_nmadd(a, b, c)`

### Result and argument types

The following table describes the types of the returned value and the function arguments.

<b>d</b>	<b>a</b>	<b>b</b>	<b>c</b>
vector double	vector double	vector double	vector double
vector float	vector float	vector float	vector float

## Result value

The value of each element of the result is the product of the corresponding elements of a and b, added to the corresponding elements of c, and then multiplied by -1.0.

## vec\_nmsub

### Purpose

Returns a vector containing the results of performing a negative multiply-subtract operation on the given vectors.

### Syntax

`d=vec_nmsub(a, b, c)`

### Result and argument types

The following table describes the types of the returned value and the function arguments.

<b>d</b>	<b>a</b>	<b>b</b>	<b>c</b>
vector float	vector float	vector float	vector float



d	a	b	c
vector double	vector double	vector double	vector double

## Result value

The value of each element of the result is the product of the corresponding elements of a and b, subtracted from the corresponding element of c.

## vec\_nor

### Purpose

Performs a bitwise NOR of the given vectors.

### Syntax

d=vec\_nor(a, b)

### Result and argument types

The following table describes the types of the returned value and the function arguments.

d	a	b
vector bool char	vector bool char	vector bool char
vector signed char	vector bool char	vector signed char
	vector signed char	vector signed char
		vector bool char
vector unsigned char	vector bool char	vector unsigned char
	vector unsigned char	vector unsigned char
		vector bool char
vector bool short	vector bool short	vector vector bool short
vector signed short	vector bool short	vector signed short
	vector signed short	vector signed short
		vector bool short
vector unsigned short	vector bool short	vector unsigned short
	vector unsigned short	vector unsigned short
		vector bool short
vector bool int	vector bool int	vector bool int
vector signed int	vector bool int	vector signed int
	vector signed int	vector signed int
		vector bool int
vector unsigned int	vector bool int	vector unsigned int
	vector unsigned int	vector unsigned int
		vector bool int
vector bool long long	vector bool long long	vector bool long long
vector signed long long	vector signed long long	vector signed long long

<b>d</b>	<b>a</b>	<b>b</b>
vector unsigned long long	vector unsigned long long	vector unsigned long long
vector float	vector bool int	vector float
	vector float	vector bool int
vector double	vector double	vector double

## Result value

The result is the bitwise NOR of a and b.

## vec\_or

### Purpose

Performs a bitwise OR of the given vectors.

### Syntax

`d=vec_or(a, b)`

### Result and argument types

The following table describes the types of the returned value and the function arguments.

<b>d</b>	<b>a</b>	<b>b</b>
vector bool char	vector bool char	vector bool char
vector signed char	vector bool char	vector signed char
	vector signed char	vector signed char
		vector bool char
vector unsigned char	vector bool char	vector unsigned char
	vector unsigned char	vector unsigned char
		vector bool char
vector bool short	vector bool short	vector vector bool short
vector signed short	vector bool short	vector signed short
	vector signed short	vector signed short
		vector bool short
vector unsigned short	vector bool short	vector unsigned short
	vector unsigned short	vector unsigned short
		vector bool short
vector bool int	vector bool int	vector bool int
vector signed int	vector bool int	vector signed int
	vector signed int	vector signed int
		vector bool int
vector unsigned int	vector bool int	vector unsigned int
	vector unsigned int	vector unsigned int
		vector bool int

<b>d</b>	<b>a</b>	<b>b</b>
vector bool long long	vector bool long long	vector bool long long
vector signed long long	vector bool long long	vector signed long long
	vector signed long long	vector signed long long
vector unsigned long long	vector bool long long	vector bool long long
	vector unsigned long long	vector unsigned long long
		vector bool long long
vector float	vector bool int	vector float
	vector float	vector bool int
		vector float
vector double	vector bool long long	vector double
	vector double	vector bool long long
		vector double

## Result value

The result is the bitwise OR of a and b.

## vec\_permi

### Purpose

Returns a vector by permuting and combining the two eight-byte-long vector elements in a and b based on the value of c.

### Syntax

d=vec\_permi(a, b, c)

### Result and argument types

The following table describes the types of the returned value and the function arguments.

<b>d</b>	<b>a</b>	<b>b</b>	<b>c</b>
vector bool long long	vector bool long long	vector bool long long	0-3
vector signed long long	vector signed long long	vector signed long long	
vector unsigned long long	vector unsigned long long	vector unsigned long long	
vector double	vector double	vector double	

## Result value

If we use a[0] and a[1] to represent the first and second eight-byte-long elements in a, and use b[0] and b[1] for elements in b, then this function determines the elements in the result vector based on the binary value of c. This is illustrated as follows:

- 00 - a[0], b[0]

- 01 - a[0], b[1]
- 10 - a[1], b[0]
- 11 - a[1], b[1]

## vec\_promote

### Purpose

Returns a vector with a in element position b.

### Syntax

d=vec\_promote(a, b)

### Result and argument types

The following table describes the types of the returned value and the function arguments.

d	a	b
vector signed char	signed char	signed int
vector unsigned char	unsigned char	
vector signed short	signed short	
vector unsigned short	unsigned short	
vector signed int	signed int	
vector unsigned int	unsigned int	
vector signed long long	signed long long	
vector unsigned long long	unsigned long long	
vector float	float	
vector double	double	

### Result value

The result is a vector with a in element position b. This function uses modulo arithmetic on b to determine the element number. For example, if b is out of range, the compiler uses b modulo the number of elements in the vector to determine the element position. The other elements of the vector are undefined.

## vec\_re

### Purpose

Returns a vector containing estimates of the reciprocals of the corresponding elements of the given vector.

### Syntax

d=vec\_re(a)

## Result and argument types

The following table describes the types of the returned value and the function arguments.

<b>d</b>	<b>a</b>
vector float	vector float
vector double	vector double

## Result value

Each element of the result contains the estimated value of the reciprocal of the corresponding element of a.

## vec\_round

### Purpose

Returns a vector containing the rounded values of the corresponding elements of the given vector.

### Syntax

d=vec\_round(a)

## Result and argument types

The following table describes the types of the returned value and the function arguments.

<b>d</b>	<b>a</b>
vector float	vector float
vector double	vector double

## Result value

Each element of the result contains the value of the corresponding element of a, rounded to the nearest representable floating-point integer, using IEEE round-to-nearest rounding.

**Note:** This function might not follow the strict operation definition of the resolution of a tie during a round when you specify the `-qstrict=nooperationprecision` compiler option.

## vec\_roundc

### Purpose

Returns a vector by rounding every single-precision or double-precision floating-point element in the given vector to integer.

This function uses the current rounding mode. It requires the POWER7 architecture.

## Syntax

`d=vec_roundc(a)`

## Result and argument types

The following table describes the types of the returned value and the function arguments.

d	a
vector float	vector float
vector double	vector double

## vec\_roundm

### Purpose

Returns a vector containing the largest representable floating-point integer values less than or equal to the values of the corresponding elements of the given vector.

**Note:** `vec_roundm` is another name for `vec_floor`. For details, see “`vec_floor`” on page 567.

## Syntax

`d=vec_roundm(a)`

## Result and argument types

The following table describes the types of the returned value and the function arguments.

d	a
vector float	vector float
vector double	vector double

## vec\_roundp

### Purpose

Returns a vector containing the smallest representable floating-point integer values greater than or equal to the values of the corresponding elements of the given vector.

**Note:** `vec_roundp` is another name for `vec_ceil`. For details, see “`vec_ceil`” on page 558.

## Syntax

`d=vec_roundp(a)`

## Result and argument types

The following table describes the types of the returned value and the function arguments.

<b>d</b>	<b>a</b>
vector float	vector float
vector double	vector double

## **vec\_roundz**

### **Purpose**

Returns a vector containing the truncated values of the corresponding elements of the given vector.

**Note:** `vec_roundz` is another name for `vec_trunc`. For details, see “`vec_trunc`” on page 588.

### **Syntax**

`d=vec_roundz(a)`

### **Result and argument types**

The following table describes the types of the returned value and the function arguments.

<b>d</b>	<b>a</b>
vector float	vector float
vector double	vector double

### **Result value**

Each element of the result contains the value of the corresponding element of `a`, truncated to an integral value.

## **vec\_rsqrte**

### **Purpose**

Returns a vector containing estimates of the reciprocal square roots of the corresponding elements of the given vector.

### **Syntax**

`d=vec_rsqrte(a)`

### **Result and argument types**

The following table describes the types of the returned value and the function arguments.

<b>d</b>	<b>a</b>
vector float	vector float
vector double	vector double

## Result value

Each element of the result contains the estimated value of the reciprocal square root of the corresponding element of a.

## vec\_sel

### Purpose

Returns a vector containing the value of either a or b depending on the value of c.

### Syntax

```
d=vec_sel(a, b, c)
```

### Result and argument types

The following table describes the types of the returned value and the function arguments.

d	a	b	c
vector bool char	vector bool char	vector bool char	vector bool char
			vector unsigned char
vector signed char	vector signed char	vector signed char	vector bool char
			vector unsigned char
vector unsigned char	vector unsigned char	vector unsigned char	vector bool char
			vector unsigned char
vector bool short	vector bool short	vector bool short	vector bool short
			vector unsigned short
vector signed short	vector signed short	vector signed short	vector bool shot
			vector unsigned short
vector unsigned short	vector unsigned short	vector unsigned short	vector bool short
			vector unsigned short
vector bool int	vector bool int	vector bool int	vector bool int
			vector unsigned int
vector signed int	vector signed int	vector signed int	vector bool int
			vector unsigned int
vector unsigned int	vector unsigned int	vector unsigned int	vector bool int
			vector unsigned int
vector bool long long	vector bool long long	vector bool long long	vector bool long long
			vector unsigned long long
vector signed long long	vector signed long long	vector signed long long	vector bool long long
			vector unsigned long long
vector unsigned long long	vector unsigned long long	vector unsigned long long	vector bool long long
			vector unsigned long long



d	a	b	c
vector float	vector float	vector float	vector bool int
			vector unsigned int
vector double	vector double	vector double	vector bool long long
			vector unsigned long long

## Result value

Each bit of the result vector has the value of the corresponding bit of a if the corresponding bit of c is 0, or the value of the corresponding bit of b otherwise.

## vec\_sl

### Purpose

Performs a left shift for each element of a vector.

### Syntax

`d=vec_sl(a, b)`

### Result and argument types

The following table describes the types of the returned value and the function arguments.

d	a	b
vector signed char	vector signed char	vector unsigned char
vector unsigned char	vector unsigned char	vector unsigned char
vector signed short	vector signed short	vector unsigned short
vector unsigned short	vector unsigned short	vector unsigned short
vector signed int	vector signed int	vector unsigned int
vector unsigned int	vector unsigned int	vector unsigned int
vector signed long long	vector signed long long	vector unsigned long long
vector unsigned long long	vector unsigned long long	vector unsigned long long

## Result value

Each element of the result vector is the result of left shifting the corresponding element of a by the number of bits specified by the value of the corresponding element of b, modulo the number of bits in the element. The bits that are shifted out are replaced by zeroes.

## vec\_sldw

### Purpose

Shift Left Double by Word Immediate

Returns a vector by concatenating a and b, and then left-shifting the result vector by multiples of 4 bytes. c specifies the offset for the shifting operation.

### Syntax

d=vec\_sldw(a, b, c)

### Result and argument types

The following table describes the types of the returned value and the function arguments.

d	a	b	c
vector bool char	vector bool char	vector bool char	0-3
vector signed char	vector signed char	vector signed char	
vector unsigned char	vector unsigned char	vector unsigned char	
vector bool short	vector bool short	vector bool short	
vector signed short	vector signed short	vector signed short	
vector unsigned short	vector unsigned short	vector unsigned short	
vector bool int	vector bool int	vector bool int	
vector signed int	vector signed int	vector signed int	
vector unsigned int	vector unsigned int	vector unsigned int	
vector bool long long	vector bool long long	vector bool long long	
vector signed long long	vector signed long long	vector signed long long	
vector unsigned long long	vector unsigned long long	vector unsigned long long	
vector float	vector float	vector float	
vector double	vector double	vector double	

### Result value

After left-shifting the concatenated a and b by multiples of 4 bytes specified by c, the function takes the four leftmost 4-byte values and forms the result vector.

## vec\_splat

### Purpose

Returns a vector that has all of its elements set to a given value.

### Syntax

d=vec\_splat(a, b)

### Result and argument types

The following table describes the types of the returned value and the function arguments.

d	a	b
vector bool char	vector bool char	0 - 15

<b>d</b>	<b>a</b>	<b>b</b>
vector signed char	vector signed char	0 - 15
vector unsigned char	vector unsigned char	0 - 15
vector bool short	vector bool short	0 - 7
vector signed short	vector signed short	0 - 7
vector unsigned short	vector unsigned short	0 - 7
vector bool int	vector bool int	0-3
vector signed int	vector signed int	0-3
vector unsigned int	vector unsigned int	0-3
vector bool long long	vector bool long long	0-1
vector signed long long	vector signed long long	0-1
vector unsigned long long	vector unsigned long long	0-1
vector float	vector float	0-3
vector double	vector double	0-1

## Result value

The value of each element of the result is the value of the element of *a* specified by *b*.

## vec\_splats

### Purpose

Returns a vector of which the value of each element is set to *a*.

### Syntax

`d=vec_splats(a)`

### Result and argument types

The following table describes the types of the returned value and the function arguments.

<b>d</b>	<b>a</b>
vector signed char	signed char
vector unsigned char	unsigned char
vector signed short	signed short
vector unsigned short	unsigned short
vector signed int	signed int
vector unsigned int	unsigned int
vector signed long long	signed long long
vector unsigned long long	unsigned long long
vector float	float
vector double	double

## vec\_sqrt

### Purpose

Returns a vector containing the square root of each element in the given vector.

This function requires the POWER7 architecture.

### Syntax

```
d=vec_sqrt(a)
```

### Result and argument types

The following table describes the types of the returned value and the function arguments.

d	a
vector float	vector float
vector double	vector double

## vec\_sr

### Purpose

Performs a right shift for each element of a vector.

### Syntax

```
d=vec_sr(a, b)
```

### Result and argument types

The following table describes the types of the returned value and the function arguments.

d	a	b
vector signed char	vector signed char	vector unsigned char
vector unsigned char	vector unsigned char	vector unsigned char
vector signed short	vector signed short	vector unsigned short
vector unsigned short	vector unsigned short	vector unsigned short
vector signed int	vector signed int	vector unsigned int
vector unsigned int	vector unsigned int	vector unsigned int
vector signed long long	vector signed long long	vector unsigned long long
vector unsigned long long	vector unsigned long long	vector unsigned long long

### Result value

Each element of the result vector is the result of right shifting the corresponding element of a by the number of bits specified by the value of the corresponding element of b, modulo the number of bits in the element. The bits that are shifted out are replaced by zeroes.

## vec\_sra

### Purpose

Performs an algebraic right shift for each element of a vector.

### Syntax

```
d=vec_sra(a, b)
```

### Result and argument types

The following table describes the types of the returned value and the function arguments.

d	a	b
vector signed char	vector signed char	vector unsigned char
vector unsigned char	vector unsigned char	vector unsigned char
vector signed short	vector signed short	vector unsigned short
vector unsigned short	vector unsigned short	vector unsigned short
vector signed int	vector signed int	vector unsigned int
vector unsigned int	vector unsigned int	vector unsigned int
vector signed long long	vector signed long long	vector unsigned long long
vector unsigned long long	vector unsigned long long	vector unsigned long long

### Result value

Each element of the result vector is the result of algebraically right shifting the corresponding element of a by the number of bits specified by the value of the corresponding element of b, modulo the number of bits in the element. The bits that are shifted out are replaced by copies of the most significant bit of the element of a.

## vec\_sub

### Purpose

Returns a vector containing the result of subtracting each element of b from the corresponding element of a.

This function emulates the operation on long long vectors.

### Syntax

```
d=vec_sub(a, b)
```

### Result and argument types

The following table describes the types of the returned value and the function arguments.

d	a	b
vector signed char	vector signed char	vector signed char
vector unsigned char	vector unsigned char	vector unsigned char

d	a	b
vector signed short	vector signed short	vector signed short
vector unsigned short	vector unsigned short	vector unsigned short
vector signed int	vector signed int	vector signed int
vector unsigned int	vector unsigned int	vector unsigned int
vector signed long long	vector signed long long	vector signed long long
vector unsigned long long	vector unsigned long long	vector unsigned long long
vector float	vector float	vector float
vector double	vector double	vector double

## Result value

The value of each element of the result is the result of subtracting the value of the corresponding element of *b* from the value of the corresponding element of *a*. The arithmetic is modular for integer vectors.

## vec\_trunc

### Purpose

Returns a vector containing the truncated values of the corresponding elements of the given vector.

**Note:** `vec_trunc` is another name for `vec_roundz`. For details, see “`vec_roundz`” on page 581.

## vec\_xld2

### Purpose

Loads a 16-byte vector from two 8-byte elements at the memory address specified by the displacement *a* and the pointer *b*.

This function requires the POWER7 architecture.

### Syntax

```
d=vec_xld2(a, b)
```

### Result and argument types

The following table describes the types of the returned value and the function arguments.

**Note:** The type for operand *a* in the following table is: `int` in 32-bit mode, and `long` in 64-bit mode.

d	a	b
vector signed char	int	signed char *
	long	
vector unsigned char	int	unsigned char *
	long	

<b>d</b>	<b>a</b>	<b>b</b>
vector signed short	int	signed short *
	long	
vector unsigned short	int	unsigned short *
	long	
vector signed int	int	signed int *
	long	
vector unsigned int	int	unsigned int *
	long	
vector signed long long	int	signed long long *
	long	
vector unsigned long long	int	unsigned long long *
	long	
vector float	int	float *
	long	
vector double	int	double *
	long	

## Result value

This function adds the displacement and the pointer R-value to obtain the address for the load operation. It does not truncate the affected address to a multiple of 16 bytes.

## vec\_xlds

### Purpose

Loads an 8-byte element from the memory address specified by the displacement *a* and the pointer *b* and then splats it onto a vector.

This function requires the POWER7 architecture.

### Syntax

`d=vec_xlds(a, b)`

### Result and argument types

The following table describes the types of the returned value and the function arguments.

**Note:** The type for operand *a* in the following table is: `int` in 32-bit mode, and `long` in 64-bit mode.

<b>d</b>	<b>a</b>	<b>b</b>
vector signed long long	int	signed long long *
	long	

d	a	b
vector unsigned long long	int	unsigned long long *
	long	
vector double	int	double *
	long	

## Result value

This function adds the displacement and the pointer R-value to obtain the address for the load operation. It does not truncate the affected address to a multiple of 16 bytes.

## vec\_xlw4

### Purpose

Loads a 16-byte vector from four 4-byte elements at the memory address specified by the displacement a and the pointer b.

This function requires the POWER7 architecture.

### Syntax

d=vec\_xlw4(a, b)

### Result and argument types

The following table describes the types of the returned value and the function arguments.

**Note:** The type for operand a in the following table is: int in 32-bit mode, and long in 64-bit mode.

d	a	b
vector signed char	int	signed char *
	long	
vector unsigned char	int	unsigned char *
	long	
vector signed short	int	signed short *
	long	
vector unsigned short	int	unsigned short *
	long	
vector signed int	int	signed int *
	long	
vector unsigned int	int	unsigned int *
	long	
vector float	int	float *
	long	



## Result value

This function adds the displacement and the pointer R-value to obtain the address for the load operation. It does not truncate the affected address to a multiple of 16 bytes.

## vec\_xor

### Purpose

Performs a bitwise XOR of the given vectors.

### Syntax

`d=vec_xor(a, b)`

### Result and argument types

The following table describes the types of the returned value and the function arguments.

<b>d</b>	<b>a</b>	<b>b</b>
vector bool char	vector bool char	vector bool char
vector signed char	vector bool char	vector signed char
	vector signed char	vector signed char
		vector bool char
vector unsigned char	vector bool char	vector unsigned char
	vector unsigned char	vector unsigned char
		vector bool char
vector bool short	vector bool short	vector vector bool short
vector signed short	vector bool short	vector signed short
	vector signed short	vector signed short
		vector bool short
vector unsigned short	vector bool short	vector unsigned short
	vector unsigned short	vector unsigned short
		vector bool short
vector bool int	vector bool int	vector bool int
vector signed int	vector bool int	vector signed int
	vector signed int	vector signed int
		vector bool int
vector unsigned int	vector bool int	vector unsigned int
	vector unsigned int	vector unsigned int
		vector bool int
vector bool long long	vector bool long long	vector bool long long
vector signed long long	vector bool long long	vector signed long long
	vector signed long long	vector signed long long
		vector bool long long

<b>d</b>	<b>a</b>	<b>b</b>
vector unsigned long long	vector bool long long	vector unsigned long long
	vector unsigned long long	vector unsigned long long
		vector bool long long
vector float	vector bool int	vector float
	vector float	vector bool int
		vector float
vector double	vector bool long long	vector double
	vector double	vector bool long long
		vector double

### Result value

The result is the bitwise XOR of a and b.

## vec\_xstd2

### Purpose

Puts a 16-byte vector a as two 8-byte elements to the memory address specified by the displacement b and the pointer c.

This function requires the POWER7 architecture.

### Syntax

d=vec\_xstd2(a, b, c)

### Result and argument types

The following table describes the types of the returned value and the function arguments.

**Note:** The type for operand a in the following table is: int in 32-bit mode, and long in 64-bit mode.

d	a	b	c
void	vector signed char	int	signed char *
		long	
	vector unsigned char	int	unsigned char *
		long	
	vector signed short	int	signed short *
		long	
	vector unsigned short	int	unsigned short *
		long	
	vector signed int	int	signed int *
		long	
	vector unsigned int	int	unsigned int *
		long	
	vector signed long long	int	signed long long *
		long	
vector unsigned long long	int	unsigned long long *	
	long		
vector float	int	float *	
	long		
vector double	int	double *	
	long		
vector pixel	int	signed short *	
		unsigned short *	
	long	signed short *	
		unsigned short *	

## Result value

This function adds the displacement and the pointer R-value to obtain the address for the store operation. It does not truncate the affected address to a multiple of 16 bytes.

## vec\_xstw4

### Purpose

Puts a 16-byte vector a to four 4-byte elements at the memory address specified by the displacement b and the pointer c.

This function requires the POWER7 architecture.

### Syntax

d=vec\_xstw4(a, b, c)

## Result and argument types

The following table describes the types of the returned value and the function arguments.

**Note:** The type for operand *b* in the following table is: `int` in 32-bit mode, and `long` in 64-bit mode.

d	a	b	c
void	vector signed char	int	signed char *
		long	
	vector unsigned char	int	unsigned char *
		long	
	vector signed short	int	signed short *
		long	
	vector unsigned short	int	unsigned short *
		long	
	vector signed int	int	signed int *
		long	
	vector unsigned int	int	unsigned int *
		long	
	vector float	int	float *
		long	
vector pixel	int		signed short *
			unsigned short *
	long		signed short *
			unsigned short *

## Result value

This function adds the displacement and the pointer R-value to obtain the address for the store operation. It does not truncate the affected address to a multiple of 16 bytes.

---

## GCC atomic memory access built-in functions

This section provides reference information for atomic memory access built-in functions whose behavior corresponds to that provided by GNU Compiler Collection (GCC). In a program with multiple threads, you can use these functions to atomically and safely modify data in one thread without interference from other threads.

These built-in functions manipulate data atomically, regardless of how many processors are installed in the host machine. The data types of the parameters that are supported in these functions include Boolean type and any integral scalar type of 1, 2, 4 or 8 bytes, listed as follows.

Table 56. Supported data types


 <code>bool</code>	 <code>_Bool</code>
---	--

Table 56. Supported data types (continued)

signed char	unsigned char
short int	unsigned short int
int	unsigned int
long int	unsigned long int
long long int <b>1</b>	unsigned long long int <b>1</b>
<b>1</b> Restriction: This type is supported only on 64-bit platforms.	

In the prototype of each function, the ellipsis (...) represents an optional list of parameters. XL C/C++ ignores these optional parameters and protects all globally accessible variables.

The GCC atomic memory access built-in functions are grouped into the following categories.

## Atomic lock, release, and synchronize functions

### `__sync_lock_test_and_set`

#### Purpose

This function atomically assigns the value of `__v` to the variable that `__p` points to.

An acquire memory barrier is created when this function is invoked.

#### Prototype

```
T __sync_lock_test_and_set (T* __p, T __v, ...);
```

where *T* is one of the data types listed in Supported data types.

#### Parameters

`__p`  
The pointer of the variable that is to be set.

`__v`  
The value to set to the variable that `__p` points to.

#### Return value

The function returns the initial value of the variable that `__p` points to.

### `__sync_lock_release`

#### Purpose

This function releases the lock acquired by the `__sync_lock_test_and_set` function, and assigns the value of zero to the variable that `__p` points to.

A release memory barrier is created when this function is invoked.

## Prototype

```
void __sync_lock_release (T* __p, ...);
```

where  $T$  is one of the data types listed in Supported data types.

## Parameters

$\_p$   
The pointer of the variable that is to be set.

## `__sync_synchronize`

### Purpose

This function synchronizes data in all threads.

A full memory barrier is created when this function is invoked.

## Prototype

```
__sync_synchronize (...);
```

## Atomic fetch and operation functions

## `__sync_fetch_and_and`

### Purpose

This function performs an atomic bitwise AND operation on the variable  $\_v$  with the variable that  $\_p$  points to. The result is stored in the address that is specified by  $\_p$ .

A full memory barrier is created when this function is invoked.

## Prototype

```
T __sync_fetch_and_and (T* __p, T __v, ...);
```

where  $T$  is one of the data types listed in Supported data types.

## Parameters

$\_p$   
The pointer of a variable on which the bitwise AND operation is to be performed. The value of this variable is to be changed to the result of the operation.

$\_v$   
The variable with which the bitwise AND operation is to be performed.

## Return value

The function returns the initial value of the variable that  $\_p$  points to.

## **\_\_sync\_fetch\_and\_nand**

### **Purpose**

This function performs an atomic bitwise NAND operation on the variable `__v` with the variable that `__p` points to. The result is stored in the address that is specified by `__p`.

A full memory barrier is created when this function is invoked.

### **Prototype**

```
T __sync_fetch_and_nand (T* __p, T __v, ...);
```

where *T* is one of the data types listed in Supported data types.

### **Parameters**

`__p`  
The pointer of a variable on which the bitwise NAND operation is to be performed. The value of this variable is to be changed to the result of the operation.

`__v`  
The variable with which the bitwise NAND operation is to be performed.

### **Return value**

The function returns the initial value of the variable that `__p` points to.

## **\_\_sync\_fetch\_and\_or**

### **Purpose**

This function performs an atomic bitwise inclusive OR operation on the variable `__v` with the variable that `__p` points to. The result is stored in the address that is specified by `__p`.

A full memory barrier is created when this function is invoked.

### **Prototype**

```
T __sync_fetch_and_or (T* __p, T __v, ...);
```

where *T* is one of the data types listed in Supported data types.

### **Parameters**

`__p`  
The pointer of a variable on which the bitwise inclusive OR operation is to be performed. The value of this variable is to be changed to the result of the operation.

`__v`  
The variable with which the bitwise inclusive OR operation is to be performed.

### **Return value**

The function returns the initial value of the variable that `__p` points to.

## **\_\_sync\_fetch\_and\_xor**

### **Purpose**

This function performs an atomic bitwise exclusive OR operation on the variable `__v` with the variable that `__p` points to. The result is stored in the address that is specified by `__p`.

A full memory barrier is created when this function is invoked.

### **Prototype**

```
T __sync_fetch_and_xor (T* __p, T __v, ...);
```

where *T* is one of the data types listed in Supported data types.

### **Parameters**

`__p`

The pointer of a variable on which the bitwise exclusive OR operation is to be performed. The value of this variable is to be changed to the result of the operation.

`__v`

The variable with which the bitwise exclusive OR operation is to be performed.

### **Return value**

The function returns the initial value of the variable that `__p` points to.

## **\_\_sync\_fetch\_and\_add**

### **Purpose**

This function atomically adds the value of `__v` to the variable that `__p` points to. The result is stored in the address that is specified by `__p`.

A full memory barrier is created when this function is invoked.

### **Prototype**

```
T __sync_fetch_and_add (T* __p, T __v, ...);
```

where *T* is one of the data types listed in Supported data types.

### **Parameters**

`__p`

The pointer of a variable to which `__v` is to be added. The value of this variable is to be changed to the result of the add operation.

`__v`

The variable whose value is to be added to the variable that `__p` points to.

### **Return value**

The function returns the initial value of the variable that `__p` points to.



## **\_\_sync\_fetch\_and\_sub**

### **Purpose**

This function atomically subtracts the value of `__v` from the variable that `__p` points to. The result is stored in the address that is specified by `__p`.

A full memory barrier is created when this function is invoked.

### **Prototype**

```
T __sync_fetch_and_sub (T* __p, T __v, ...);
```

where *T* is one of the data types listed in Supported data types.

### **Parameters**

`__p`  
The pointer of a variable from which `__v` is to be subtracted. The value of this variable is to be changed to the result of the sub operation.

`__v`  
The variable whose value is to be subtracted from the variable that `__p` points to.

### **Return value**

The function returns the initial value of the variable that `__p` points to.

## **Atomic operation and fetch functions**

### **\_\_sync\_and\_and\_fetch**

#### **Purpose**

This function performs an atomic bitwise AND operation on the variable `__v` with the variable that `__p` points to. The result is stored in the address that is specified by `__p`.

A full memory barrier is created when this function is invoked.

#### **Prototype**

```
T __sync_and_and_fetch (T* __p, T __v, ...);
```

where *T* is one of the data types listed in Supported data types.

#### **Parameters**

`__p`  
The pointer of a variable on which the bitwise AND operation is to be performed. The value of this variable is to be changed to the result of the operation.

`__v`  
The variable with which the bitwise AND operation is to be performed.

#### **Return value**

The function returns the new value of the variable that `__p` points to.

## **\_\_sync\_nand\_and\_fetch**

### **Purpose**

This function performs an atomic bitwise NAND operation on the variable `__v` with the variable that `__p` points to. The result is stored in the address that is specified by `__p`.

A full memory barrier is created when this function is invoked.

### **Prototype**

```
T __sync_nand_and_fetch (T* __p, T __v, ...);
```

where *T* is one of the data types listed in Supported data types.

### **Parameters**

`__p`  
The pointer of a variable on which the bitwise NAND operation is to be performed. The value of this variable is to be changed to the result of the operation.

`__v`  
The variable with which the bitwise NAND operation is to be performed.

### **Return value**

The function returns the new value of the variable that `__p` points to.

## **\_\_sync\_or\_and\_fetch**

### **Purpose**

This function performs an atomic bitwise inclusive OR operation on the variable `__v` with variable that `__p` points to. The result is stored in the address that is specified by `__p`.

A full memory barrier is created when this function is invoked.

### **Prototype**

```
T __sync_or_and_fetch (T* __p, T __v, ...);
```

where *T* is one of the data types listed in Supported data types.

### **Parameters**

`__p`  
The pointer of a variable on which the bitwise inclusive OR operation is to be performed. The value of this variable is to be changed to the result of the operation.

`__v`  
The variable with which the bitwise inclusive OR operation is to be performed.

### **Return value**

The function returns the new value of the variable that `__p` points to.

## **\_\_sync\_xor\_and\_fetch**

### **Purpose**

This function performs an atomic bitwise exclusive OR operation on the variable `__v` with the variable that `__p` points to. The result is stored in the address that is specified by `__p`.

A full memory barrier is created when this function is invoked.

### **Prototype**

```
T __sync_xor_and_fetch (T* __p, T __v, ...);
```

where *T* is one of the data types listed in Supported data types.

### **Parameters**

`__p`

The pointer of the variable on which the bitwise exclusive OR operation is to be performed. The value of this variable is to be changed to the result of the operation.

`__v`

The variable with which the bitwise exclusive OR operation is to be performed.

### **Return value**

The function returns the new value of the variable that `__p` points to.

## **\_\_sync\_add\_and\_fetch**

### **Purpose**

This function atomically adds the value of `__v` to the variable that `__p` points to. The result is stored in the address that is specified by `__p`.

A full memory barrier is created when this function is invoked.

### **Prototype**

```
T __sync_add_and_fetch (T* __p, T __v, ...);
```

where *T* is one of the data types listed in Supported data types.

### **Parameters**

`__p`

The pointer of a variable to which `__v` is to be added. The value of this variable is to be changed to the result of the add operation.

`__v`

The variable whose value is to be added to the variable that `__p` points to.

### **Return value**

The function returns the new value of the variable that `__p` points to.

## **`__sync_sub_and_fetch`**

### **Purpose**

This function atomically subtracts the value of `__v` from the variable that `__p` points to. The result is stored in the address that is specified by `__p`.

A full memory barrier is created when this function is invoked.

### **Prototype**

```
T __sync_sub_and_fetch (T* __p, T __v, ...);
```

where *T* is one of the data types listed in Supported data types.

### **Parameters**

`__p`  
The pointer of a variable from which `__v` is to be subtracted. The value of this variable is to be changed to the result of the sub operation.

`__v`  
The variable whose value is to be subtracted from the variable that `__p` points to.

### **Return value**

The function returns the new value of the variable that `__p` points to.

## **Atomic compare and swap functions**

### **`__sync_val_compare_and_swap`**

#### **Purpose**

This function compares the value of `__compVal` to the value of the variable that `__p` points to. If they are equal, the value of `__exchVal` is stored in the address that is specified by `__p`; otherwise, no operation is performed.

A full memory barrier is created when this function is invoked.

#### **Prototype**

```
T __sync_val_compare_and_swap (T* __p, T __compVal, T __exchVal, ...);
```

where *T* is one of the data types listed in Supported data types.

#### **Parameters**

`__p`  
The pointer to a variable whose value is to be compared with.

`__compVal`  
The value to be compared with the value of the variable that `__p` points to.

`__exchVal`  
The value to be stored in the address that `__p` points to.

## Return value

The function returns the initial value of the variable that `__p` points to.

## `__sync_bool_compare_and_swap`

### Purpose

This function compares the value of `__compVal` with the value of the variable that `__p` points to. If they are equal, the value of `__exchVal` is stored in the address that is specified by `__p`; otherwise, no operation is performed.

A full memory barrier is created when this function is invoked.

### Prototype

```
bool __sync_bool_compare_and_swap (T* __p, T __compVal, T __exchVal, ...);
```

where *T* is one of the data types listed in Supported data types.

### Parameters

`__p`  
The pointer to a variable whose value is to be compared with.

`__compVal`  
The value to be compared with the value of the variable that `__p` points to.

`__exchVal`  
The value to be stored in the address that `__p` points to.

### Return value

If the value of `__compVal` and the value of the variable that `__p` points to are equal, the function returns true; otherwise, it returns false.

---

## Miscellaneous built-in functions

Miscellaneous functions are grouped into the following categories:

- “Optimization-related functions”
- “Move to/from register functions” on page 604
- “Memory-related functions” on page 607

## Optimization-related functions

### `__alignx`

#### Purpose

Allows for optimizations such as automatic vectorization by informing the compiler that the data pointed to by *pointer* is aligned at a known compile-time offset.

#### Prototype

```
void __alignx (int alignment, const void* pointer);
```

## Parameters

*alignment*

Must be a constant integer with a value greater than zero and of a power of two.

## **\_\_builtin\_expect**

### Purpose

Indicates that an expression is likely to evaluate to a specified value. The compiler may use this knowledge to direct optimizations.

### Prototype

```
long __builtin_expect (long expression, long value);
```

## Parameters

*expression*

Should be an integral-type expression.

*value*

Must be a constant literal.

## Usage

If the *expression* does not actually evaluate at run time to the predicted value, performance may suffer. Therefore, this built-in function should be used with caution.

## **\_\_fence**

### Purpose

Acts as a barrier to compiler optimizations that involve code motion, or reordering of machine instructions. Compiler optimizations will not move machine instructions past the location of the `__fence` call.

### Prototype

```
void __fence (void);
```

## Examples

This function is useful to guarantee the ordering of instructions in the object code generated by the compiler when optimization is enabled.

## Move to/from register functions

## **\_\_mftb**

### Purpose

Move from Time Base

In 32-bit compilation mode, returns the lower word of the time base register. In 64-bit mode, returns the entire doubleword of the time base register.

## Prototype

```
unsigned long __mftb (void);
```

## Usage

In 32-bit mode, this function can be used in conjunction with the `__mftbu` built-in function to read the entire time base register. In 64-bit mode, the entire doubleword of the time base register is returned, so separate use of `__mftbu` is unnecessary.

It is recommended that you insert the `__fence` built-in function before and after the `__mftb` built-in function.

## `__mftbu`

### Purpose

Move from Time Base Upper

Returns the upper word of the time base register.

## Prototype

```
unsigned int __mftbu (void);
```

## Usage

In 32-bit mode you can use this function in conjunction with the `__mftb` built-in function to read the entire time base register.

It is recommended that you insert the `__fence` built-in function before and after the `__mftbu` built-in function.

## `__mfmsr`

### Purpose

Move from Machine State Register

Moves the contents of the machine state register (MSR) into bits 32 to 63 of the designated general-purpose register.

## Prototype

```
unsigned long __mfmsr (void);
```

## Usage

Execution of this instruction is privileged and restricted to supervisor mode only.

## `__mfspr`

### Purpose

Move from Special-Purpose Register

Returns the value of given special purpose register.

## Prototype

```
unsigned long __mfspr (const int registerNumber);
```

## Parameters

*registerNumber*

The number of the special purpose register whose value is to be returned. The *registerNumber* must be known at compile time.

## **\_\_mtmsr** Purpose

Move to Machine State Register

Moves the contents of bits 32 to 62 of the designated GPR into the MSR.

## Prototype

```
void __mtmsr (unsigned long value);
```

## Parameters

*value*

The bitwise OR result of bits 48 and 49 of *value* is placed into MSR<sub>48</sub>. The bitwise OR result of bits 58 and 49 of *value* is placed into MSR<sub>58</sub>. The bitwise OR result of bits 59 and 49 of *value* is placed into MSR<sub>59</sub>. Bits 32:47, 49:50, 52:57, and 60:62 of *value* are placed into the corresponding bits of the MSR.

## Usage

Execution of this instruction is privileged and restricted to supervisor mode only.

## **\_\_mtspr** Purpose

Move to Special-Purpose Register

Sets the value of a special purpose register.

## Prototype

```
void __mtspr (const int registerNumber, unsigned long value);
```

## Parameters

*registerNumber*

The number of the special purpose register whose value is to be set. The *registerNumber* must be known at compile time.

*value*

Must be known at compile time.

## Related information

- “Register transfer functions” on page 509



## Memory-related functions

### **\_\_alloca**

#### **Purpose**

Allocates space for an object. The allocated space is put on the stack and freed when the calling function returns.

#### **Prototype**

```
void* __alloca (size_t size)
```

#### **Parameters**

*size*

An integer representing the amount of space to be allocated, measured in bytes.

### **\_\_builtin\_frame\_address, \_\_builtin\_return\_address**

#### **Purpose**

Returns the address of the stack frame, or return address, of the current function, or of one of its callers.

#### **Prototype**

```
void* __builtin_frame_address (unsigned int level);
```

```
void* __builtin_return_address (unsigned int level);
```

#### **Parameters**

*level*

A constant literal indicating the number of frames to scan up the call stack. The *level* must range from 0 to 63. A value of 0 returns the frame or return address of the current function, a value of 1 returns the frame or return address of the caller of the current function and so on.

#### **Return value**

Returns 0 when the top of the stack is reached. Optimizations such as inlining may affect the expected return value by introducing extra stack frames or fewer stack frames than expected. If a function is inlined, the frame or return address corresponds to that of the function that is returned to.

### **\_\_mem\_delay**

#### **Purpose**

The **\_\_mem\_delay** built-in function specifies how many delay cycles there are for specific loads. These specific loads are delinquent loads with a long memory access latency because of cache misses.

When you specify which load is delinquent the compiler takes that information and carries out optimizations such as data prefetching. In addition, when you run **-qprefetch=assistthread**, the compiler uses the delinquent load information to perform analysis and generate prefetching assist threads. For more information, see “-qprefetch” on page 293.

## Prototype

```
void* __mem_delay (const void *address, const unsigned int cycles);
```

## Parameters

*address*

The address of the data to be loaded or stored.

*cycles*

A compile time constant, typically either L1 miss latency or L2 miss latency.

## Usage

The `__mem_delay` built-in function is placed immediately before a statement that contains a specified memory reference.

## Examples

Here is how you generate code using assist threads with `__mem_delay`:

Initial code:

```
int y[64], x[1089], w[1024];

void foo(void){
    int i, j;
    for (i = 0; i &1; 64; i++) {
        for (j = 0; j < 1024; j++) {

            /* what to prefetch? y[i]; inserted by the user */
            __mem_delay(&y[i], 10);
            y[i] = y[i] + x[i + j] * w[j];
            x[i + j + 1] = y[i] * 2;
        }
    }
}
```

Assist thread generated code:

```
void foo@clone(unsigned thread_id, unsigned version)

{ if (!1) goto lab_1;

/* version control to synchronize assist and main thread */
if (version == @2version0) goto lab_5;

goto lab_1;

lab_5:

@CIV1 = 0;

do { /* id=1 guarded */ /* ~2 */

if (!1) goto lab_3;

@CIV0 = 0;

do { /* id=2 guarded */ /* ~4 */

/* region = 0 */

/* __dcbt call generated to prefetch y[i] access */
__dcbt(((char *)&y + (4)*(@CIV1)))
```

```

@CIV0 = @CIV0 + 1;
} while ((unsigned) @CIV0 < 1024u); /* ~4 */

lab_3:
@CIV1 = @CIV1 + 1;
} while ((unsigned) @CIV1 < 64u); /* ~2 */

lab_1:

return;
}

```

### Related information

- “-qprefetch” on page 293

---

## Built-in functions for parallel processing

Use these built-in functions to obtain information about the parallel environment:

- “IBM SMP built-in functions”
- Chapter 8, “OpenMP runtime functions for parallel processing,” on page 611

## IBM SMP built-in functions

### \_\_parthds (C only)

#### Purpose

Returns the value of the **parthds** runtime option.

#### Prototype

```
int __parthds (void);
```

#### Return value

If the **parthds** option is not explicitly set, returns the default value set by the runtime library. If the **-qsmp** compiler option was not specified during program compilation, returns 1 regardless of runtime options selected.

### \_\_usrthds (C only)

#### Purpose

Returns the value of the **usrthds** runtime option.

#### Prototype

```
int __usrthds (void);
```

#### Return value

If the **usrthds** option is not explicitly set, or the **-qsmp** compiler option is not specified during program compilation, returns 0 regardless of runtime options selected.



---

## Chapter 8. OpenMP runtime functions for parallel processing

Function definitions for the `omp_` functions can be found in the `omp.h` header file.

For complete information about OpenMP runtime library functions, refer to the OpenMP Application Program Interface specification in [www.openmp.org](http://www.openmp.org).

### Related information

- “Environment variables for parallel processing” on page 27

---

### `omp_get_max_active_levels`

#### Purpose

Retrieves the value of the *max-active-levels-var* internal control variable that determines the maximum number of nested active parallel regions. *max-active-levels-var* can be set with the `OMP_MAX_ACTIVE_LEVELS` environment variable or the `omp_set_max_active_levels` function.

#### Prototype

```
int omp_get_max_active_levels(void);
```

---

### `omp_set_max_active_levels`

#### Purpose

Sets the value of the *max-active-levels-var* internal control variable to the value in the argument. If the number of parallel levels requested exceeds the number of the supported level of parallelism, the value of *max-active-levels-var* is set to the number of parallel levels supported by the runtime. If the number of parallel levels requested is not a positive integer, this routine call is ignored. When the nested parallelism is turned off, this routine has no effects and the value of *max-active-levels-var* remains 1. *max-active-levels-var* can also be set with the `OMP_MAX_ACTIVE_LEVELS` environment variable. To retrieve the value for *max-active-levels-var*, use the `omp_get_max_active_levels` function.

#### Prototype

```
void omp_set_max_active_levels(int max_levels);
```

---

### `omp_get_schedule`

#### Purpose

Returns the *run-sched-var* internal control variable of the team that is processing the parallel region. The argument *kind* returns the type of schedule that will be used. *modifier* represents the chunk size that is set for applicable schedule types. *run-sched-var* can be set with the `OMP_SCHEDULE` environment variable or the `omp_set_schedule` function.

## Prototype

```
int omp_get_schedule(omp_sched_t * kind, int * modifier );
```

## Parameters

*kind*

The value returned for *kind* is one of the schedule types affinity, auto, dynamic, guided, runtime, or static.

*modifier*

For the schedule type dynamic, guided, or static, modifier is the chunk size that is set. For the schedule type auto, modifier has no meaning.

### Related reference:

“omp\_set\_schedule”

### Related information:

“OMP\_SCHEDULE” on page 38

---

## omp\_set\_schedule

### Purpose

Sets the value of the *run-sched-var* internal control variable. Use **omp\_set\_schedule** if you want to set the schedule type separately from the *OMP\_SCHEDULE* environment variable.

### Prototype

```
void omp_set_schedule (omp_sched_t kind, int modifier);
```

### Parameters

*kind*

Must be one of the schedule types affinity, auto, dynamic, guided, runtime, or static.

*modifier*

For the schedule type dynamic, guided, or static, modifier is the chunk size that you want to set. Generally it is a positive integer. If the value is less than one, the default will be used. For the schedule type auto, modifier has no meaning.

### Related reference:

“omp\_get\_schedule” on page 611

### Related information:

“OMP\_SCHEDULE” on page 38

---

## omp\_get\_thread\_limit

### Purpose

Retrieves the maximum number of OpenMP threads stored in the *thread-limit-var* internal control variable that are available to the program. *thread-limit-var* can be set with the *OMP\_THREAD\_LIMIT* environment variable.

### Prototype

```
int omp_get_thread_limit(void);
```

---

## omp\_get\_level

### Purpose

Use **omp\_get\_level** to return the number of active and inactive nested parallel regions that the generating task is executing in. This does not include the implicit parallel region.

### Prototype

```
int omp_get_level(void);
```

---

## omp\_get\_ancestor\_thread\_num

### Purpose

Use **omp\_get\_ancestor\_thread\_num** to return the thread number in the current level of the ancestor that is at the specified nested level.

**omp\_get\_ancestor\_thread\_num** returns -1 if the nested level is not within the range of 0 and the current thread's nested level as returned by **omp\_get\_level**.

### Prototype

```
int omp_get_ancestor_thread_num(int level);
```

---

## omp\_get\_team\_size

### Purpose

Use **omp\_get\_team\_size** to return the thread team size that the ancestor belongs to. **omp\_get\_team\_size** returns -1 if the nested level is not within the range of 0 and the current thread's nested level as returned by **omp\_get\_level**.

### Prototype

```
int omp_get_team_size(int level);
```

---

## omp\_get\_active\_level

### Purpose

Use **omp\_get\_active\_level** to return the number of active parallel regions that are nested.

### Prototype

```
int omp_get_active_level(void);
```

---

## omp\_get\_num\_threads

### Purpose

Returns the number of threads currently in the team executing the parallel region from which it is called.

## Prototype

```
int omp_get_num_threads (void);
```

---

## omp\_set\_num\_threads

### Purpose

Overrides the setting of the OMP\_NUM\_THREADS environment variable, and specifies the number of threads to use for a subsequent parallel region by setting the first value of *num\_list* for OMP\_NUM\_THREADS.

### Prototype

```
void omp_set_num_threads (int num_threads);
```

### Parameters

*num\_threads*

Must be a positive integer.

### Usage

If the *num\_threads* clause is present, then for the parallel region it is applied to, it supersedes the number of threads requested by this function or the OMP\_NUM\_THREADS environment variable. Subsequent parallel regions are not affected by it.

---

## omp\_get\_max\_threads

### Purpose

Returns the first value of *num\_list* for the OMP\_NUM\_THREADS environment variable. This value is the maximum number of threads that can be used to form a new team if a parallel region without a **num\_threads** clause is encountered.

### Prototype

```
int omp_get_max_threads (void);
```

---

## omp\_get\_thread\_num

### Purpose

Returns the thread number, within its team, of the thread executing the function.

### Prototype

```
int omp_get_thread_num (void);
```

### Return value

The thread number lies between 0 and `omp_get_num_threads()-1`, inclusive. The master thread of the team is thread 0.



---

## **omp\_get\_num\_procs**

### **Purpose**

Returns the maximum number of processors that could be assigned to the program.

### **Prototype**

```
int omp_get_num_procs (void);
```

---

## **omp\_in\_final**

### **Purpose**

Returns a nonzero integer value if the function is called in a final task region; otherwise, it returns 0.

### **Prototype**

```
int omp_in_final(void);
```

---

## **omp\_in\_parallel**

### **Purpose**

Returns non-zero if it is called within the dynamic extent of a parallel region executing in parallel; otherwise, returns 0.

### **Prototype**

```
int omp_in_parallel (void);
```

---

## **omp\_set\_dynamic**

### **Purpose**

Enables or disables dynamic adjustment of the number of threads available for execution of parallel regions.

### **Prototype**

```
void omp_set_dynamic (int dynamic_threads);
```

---

## **omp\_get\_dynamic**

### **Purpose**

Returns non-zero if dynamic thread adjustments enabled and returns 0 otherwise.

### **Prototype**

```
int omp_get_dynamic (void);
```

---

## omp\_set\_nested

### Purpose

Enables or disables nested parallelism.

### Prototype

```
void omp_set_nested (int);
```

### Return value

In the current implementation, nested parallel regions are always serialized. As a result, has no effect.

---

## omp\_get\_nested

### Purpose

Returns non-zero if nested parallelism is enabled and 0 if it is disabled.

### Prototype

```
int omp_get_nested (void);
```

### Return value

In the current implementation, nested parallel regions are always serialized. As a result, always returns 0.

---

## omp\_init\_lock, omp\_init\_nest\_lock

### Purpose

Initializes the lock associated with the parameter *lock* for use in subsequent calls.

### Prototype

```
void omp_init_lock (omp_lock_t *lock);  
void omp_init_nest_lock (omp_nest_lock_t *lock);
```

---

## omp\_destroy\_lock, omp\_destroy\_nest\_lock

### Purpose

Ensures that the specified lock variable *lock* is uninitialized.

### Prototype

```
void omp_destroy_lock (omp_lock_t *lock);  
void omp_destroy_nest_lock (omp_nest_lock_t *lock);
```

---

## omp\_set\_lock, omp\_set\_nest\_lock

### Purpose

Blocks the thread executing the function until the specified lock is available and then sets the lock.

### Prototype

```
void omp_set_lock (omp_lock_t * lock);  
void omp_set_nest_lock (omp_nest_lock_t * lock);
```

### Usage

A simple lock is available if it is unlocked. A nestable lock is available if it is unlocked or if it is already owned by the thread executing the function.

---

## omp\_unset\_lock, omp\_unset\_nest\_lock

### Purpose

Releases ownership of a lock.

### Prototype

```
void omp_unset_lock (omp_lock_t * lock);  
void omp_unset_nest_lock (omp_nest_lock_t * lock);
```

---

## omp\_test\_lock, omp\_test\_nest\_lock

### Purpose

Attempts to set a lock but does not block execution of the thread.

### Prototype

```
int omp_test_lock (omp_lock_t * lock);  
int omp_test_nest_lock (omp_nest_lock_t * lock);
```

---

## omp\_get\_wtime

### Purpose

Returns the time elapsed from a fixed starting time.

### Prototype

```
double omp_get_wtime (void);
```

### Usage

The value of the fixed starting time is determined at the start of the current program, and remains constant throughout program execution.

---

## **omp\_get\_wtick**

### **Purpose**

Returns the number of seconds between clock ticks.

### **Prototype**

```
double omp_get_wtick (void);
```

### **Usage**

The value of the fixed starting time is determined at the start of the current program, and remains constant throughout program execution.

---

## Notices

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

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

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

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

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

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

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

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

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

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

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

Lab Director  
IBM Canada Ltd. Laboratory  
8200 Warden Avenue  
Markham, Ontario L6G 1C7  
Canada

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

#### COPYRIGHT LICENSE:

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

Each copy or any portion of these sample programs or any derivative work, must include a copyright notice as follows:

© (your company name) (year). Portions of this code are derived from IBM Corp. Sample Programs. © Copyright IBM Corp. 1998, 2012. All rights reserved.

---

## Trademarks and service marks

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

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, other countries, or both.

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

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

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

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

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





# Index

## Special characters

- qassert compiler option 113
- qdbgfnt compiler option 135
- qfdpr compiler option 152
- qflttrap compiler option 159
- qfunctrace 166
- qinline 197
- qlibmpi 241
- qlistfnt compiler option 245
- qnofunctrace 166
- qnoinline 197
- qoptdebug compiler option 275
- qreport compiler option 302
- qsaveopt compiler option 311
- qskipsrc compiler option 318
- qsmp compiler option 320
- qstackprotect compiler option 329
- qunique compiler option 362
- qversion compiler option 370
- #pragma nofunctrace 166, 418

## A

- alias 98
  - qalias compiler option 98
  - pragma disjoint 395
- alignment 100
  - qalign compiler option 100
  - pragma align 100
  - pragma pack 424
- alter program semantics 333
- appending macro definitions, preprocessed output 315
- architecture 9, 106
  - q32 compiler option 96
  - q64 compiler option 96
  - qarch compiler option 106
  - qcache compiler option 121
  - qtune compiler option 357
  - architecture combination 358
  - macros 469
- arrays
  - padding 178
- auto
  - qlanglvl compiler option
  - qlanglvl=autotypededuction 215

## B

- backward 16
- Backward compatibility issues 16
- basic example, described xiii
- built-in functions 479
  - block-related 534
  - cache-related 525
  - fixed-point 479
  - floating-point 488, 499
    - binary 488
    - decimal 499
  - for parallel processing 609

- built-in functions (*continued*)
  - GCC atomic memory access 594
  - miscellaneous 603
  - synchronization and atomic 518

## C

- C++0x
  - qlanglvl compiler options
    - qlanglvl=autotypededuction 215
    - qlanglvl=c99longlong 215
    - qlanglvl=c99preprocessor 215
    - qlanglvl=decltype 215
    - qlanglvl=delegatingctors 215
    - qlanglvl=extended0x 215
    - qlanglvl=extendedfriend 215
    - qlanglvl=extendedintegersafe 215
    - qlanglvl=externtemplate 215
    - qlanglvl=inlinenamespace 215
    - qlanglvl=referencecollapsing 215
    - qlanglvl=rvalueresferences 215
    - qlanglvl=static\_assert 215
    - qlanglvl=variadic[templates] 215
  - qwarn0x compiler option 374
- C99 long long
  - qlanglvl compiler option
  - qlanglvl=c99longlong 215
- C99 preprocessor
  - qlanglvl compiler option
  - qlanglvl=c99preprocessor 215
- cleanpdf command 286
- compatibility 16
  - qoldpassbyvalue compiler option 274
  - compatibility
    - options for compatibility 91
- compiler options 5
  - architecture-specific 9
  - performance optimization 87
  - resolving conflicts 8
  - specifying compiler options 5
    - command line 6
    - configuration file 7
    - source files 8
  - summary of command line options 75
- configuration 40
  - custom configuration files 40
  - gxc and gxc++ options 43
  - specifying compiler options 7
- configuration file 150
- constructor
  - delegating constructors
    - qlanglvl=delegatingctors 215
- control of implicit timestamps 351
- control of transformations 333

## D

- data types 105
  - qaltivec compiler option 105
- debug optimized code 275
- decltype
  - qlanglvl compiler option
  - qlanglvl=decltype 215
- delegating constructors
  - qlanglvl compiler option
  - qlanglvl=delegatingctors 215
- Dynamic Probe Class Library
  - qdpcl compiler option 140
- dynamic profiling environment variable 33

## E

- environment variable 25
  - environment variables 26
  - scheduling algorithm environment variable 38
  - XLSMPOPTS environment variable 27
- environment variables
  - runtime
    - XLSMPOPTS 28
- error checking and debugging 83
  - g compiler option 169
  - qcheck compiler option 125
  - qheapdebug compiler option 177
  - qlinedebug compiler option 242
- exception handling 419
  - for floating point 159
- explicit instantiation declarations
  - qlanglvl compiler option
  - qlanglvl=externtemplate 215
- extended friend declarations
  - qlanglvl compiler option
  - qlanglvl=extendedfriend 215

## F

- fini 401
- floating-point
  - exceptions 159
- function declarator
  - trailing return type
    - qlanglvl=autotypededuction 215
- function trace 166
- functrace 166

## G

- GCC options 11
- gxc and gxc++ utilities 11

## H

- high order transformation 178

## I

- implicit timestamps, control of 351
- init pragma 407
- inlining 197
- interprocedural analysis (IPA) 199
- invocations 1
  - compiler or components 1
  - preprocessor 12
  - selecting 1
  - syntax 2

## L

- language level
  - extended0x 215
- language standards 215
- large pages 238
- lib\*.a library files 213
- lib\*.so library files 213
- libraries
  - libraries
    - redistributable 15
  - XL C/C++ 15
- linker 14
  - G compiler option 172
  - invoking 14
- linking 14
  - brtl compiler option 118
  - G compiler option 172
  - options that control linking 91
  - order of linking 14
- listing 20, 313
  - qattr compiler option 114
  - qlist compiler option 243
  - qlistopt compiler option 248
  - qsource compiler option 324
  - qxref compiler option 379
  - options that control listings and messages 85

## M

- machines, compiling for different types 106
- macro definitions, preprocessed output 315
- macros 463
  - related to architecture 469
  - related to compiler options 465
  - related to language features 470
  - related to the compiler 464
  - related to the platform 465
- maf suboption of -qfloat 336
- mergepdf 286
- mpi 241
- MPI 241

## N

- name mangling
  - pragma 415
- namespace
  - qlanglvl compiler option
  - qlanglvl=inlinenamespace 215
- nofunctrace 418

NOFUNCTRACE 166

## O

- object model 272
  - qobjmodel compiler option 272
  - pragma object\_model 272
- object output, implicit timestamps 351
- OMP\_DYNAMIC environment variable 33
- OMP\_NESTED environment variable 34
- OMP\_NUM\_THREADS environment variable 34
- OMP\_PROC\_BIND environment variable 36
- OMP\_SCHEDULE environment variable 38
- OMP\_STACKSIZE environment variable 38
- OMP\_WAIT\_POLICY environment variable 39
- OpenMP 33
  - OpenMP environment variables 33
- operator\_new pragma 419
- optimization 87
  - O compiler option 269
  - qalias compiler option 98
  - qoptimize compiler option 269
  - controlling, using option\_override pragma 422
  - loop optimization 87
  - qhot compiler option 178
  - qstrict\_induction compiler option 337
  - options for performance optimization 87

## P

- parallel processing 33
  - built-in functions 609
  - OpenMP environment variables 33
  - parallel processing pragmas 439
  - pragma directives 439
  - setting parallel processing environment variables 27
- performance 87
  - O compiler option 269
  - qalias compiler option 98
  - qoptimize compiler option 269
- platform, compiling for a specific type 106
- pragma nofunctrace 418
- pragmas
  - fini 401
  - init 407
  - namemanglingrule 415
  - operator\_new 419
  - priority 297
  - report 430
- priority pragma 297
- procedure trace 166
- profile-directed feedback (PDF) 283
  - qpdf1 compiler option 283
  - qpdf2 compiler option 283
- profiling 279

profiling (*continued*)

- qdpcl compiler option 140
- qpdf1 compiler option 283
- qpdf2 compiler option 283
- qshowpdf compiler option 316
- environment variable 33

## R

- report
  - pragma 430
- resetpdf command 286
- rrm suboption of -qfloat 336

## S

- scoped enumerations
  - qlanglvl compiler option
  - qlanglvl=scopedenum 215
- shared objects 260
  - b compiler option 115
  - qmksrobj 260
- shared-memory parallelism (SMP) 28
  - qsmp compiler option 320
  - environment variables 28
- showpdf 286
- SIGTRAP signal 159
- skipsrc
  - skipsrc 318
- stackprotect
  - stackprotect 329
- static assertions
  - qlanglvl compiler option
  - qlanglvl=extc1x 215
  - qlanglvl=static\_assert 215

## T

- target machine, compiling for 106
- templates 345
  - qlanglvl compiler option
  - qlanglvl=externtemplate 215
  - qlanglvl=variadic[templates] 215
- qtempinc compiler option 345
- qtemplaterecompile compiler option 347
- qtemplateregistry compiler option 348
- qtempmax compiler option 349
- qtmplinst compiler option 353
- qtmplparse compiler option 354
- pragma define 394
- pragma do\_not\_instantiate 397
- pragma implementation 406
- pragma instantiate 394
- threads, wait policy 39
- trace 166
- trailing return type
  - qlanglvl compiler option
  - qlanglvl=autotypededuction 215
- transformations, control of 333
- tuning 357
  - qarch compiler option 357
  - qtune compiler option 357

type specifier  
  auto  
    -qlanglvl=autotypededuction 215  
  decltype(expression)  
    -qlanglvl=decltype 215

## V

variadic templates  
  -qlanglvl compiler options  
    -qlanglvl=extendedintegersafe 215  
    -qlanglvl=variadic[templates] 215  
vector built-in functions  
  vec\_abs 535  
  vec\_add 536  
  vec\_and 546  
  vec\_andc 547  
  vec\_ceil 558  
  vec\_cmpeq 558  
  vec\_cmpgt 560  
  vec\_cmplt 561  
  vec\_cpsgn 562  
  vec\_extract 566  
  vec\_floor 567  
  vec\_insert 567  
  vec\_madd 568  
  vec\_mul 572  
  vec\_nabs 573  
  vec\_neg 573  
  vec\_nor 575  
  vec\_round 579  
  vec\_sldw 583  
  vec\_splat 584  
  vec\_splats 585  
  vec\_trunc 588  
vector data types 105  
  -qaltivec compiler option 105  
vector processing 316  
  -qaltivec compiler option 105  
virtual function table (VFT)  
  pragma hashome 401, 408

## W

waiting threads, handling 39

## X

XLSMPOPTS environment variable 28







Product Number: 5765-J02; 5725-C72

Printed in USA

SC14-7330-00

