

Intel® Fortran Compiler User and Reference Guides

Document number: 304970-005US

Disclaimer and Legal Information

INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL(R) PRODUCTS. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. EXCEPT AS PROVIDED IN INTEL'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, INTEL ASSUMES NO LIABILITY WHATSOEVER, AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO SALE AND/OR USE OF INTEL PRODUCTS INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT.

UNLESS OTHERWISE AGREED IN WRITING BY INTEL, THE INTEL PRODUCTS ARE NOT DESIGNED NOR INTENDED FOR ANY APPLICATION IN WHICH THE FAILURE OF THE INTEL PRODUCT COULD CREATE A SITUATION WHERE PERSONAL INJURY OR DEATH MAY OCCUR.

Intel may make changes to specifications and product descriptions at any time, without notice. Designers must not rely on the absence or characteristics of any features or instructions marked "reserved" or "undefined." Intel reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them. The information here is subject to change without notice. Do not finalize a design with this information. The products described in this document may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

Contact your local Intel sales office or your distributor to obtain the latest specifications and before placing your product order.

Copies of documents which have an order number and are referenced in this document, or other Intel literature, may be obtained by calling 1-800-548-4725, or by visiting Intel's Web Site.

Intel processor numbers are not a measure of performance. Processor numbers differentiate features within each processor family, not across different processor families. See http://www.intel.com/products/processor_number for details.

This document contains information on products in the design phase of development.

BunnyPeople, Celeron, Celeron Inside, Centrino, Centrino Atom, Centrino Atom Inside, Centrino Inside, Centrino logo, Core Inside, FlashFile, i960, InstantIP, Intel, Intel logo, Intel386, Intel486, IntelDX2, IntelDX4, IntelSX2, Intel Atom, Intel Atom Inside, Intel Core, Intel Inside, Intel Inside logo, Intel. Leap ahead., Intel. Leap ahead. logo, Intel NetBurst, Intel NetMerge, Intel NetStructure, Intel SingleDriver, Intel SpeedStep, Intel StrataFlash, Intel Viiv, Intel vPro, Intel XScale, Itanium, Itanium Inside, MCS, MMX, Oplus, OverDrive, PDCharm, Pentium, Pentium Inside, skool, Sound Mark, The Journey Inside, Viiv Inside, vPro Inside, VTune, Xeon, and Xeon Inside are trademarks of Intel Corporation in the U.S. and other countries.

* Other names and brands may be claimed as the property of others.

Copyright (C) 1996-2008, Intel Corporation. All rights reserved.

Portions Copyright (C) 2001, Hewlett-Packard Development Company, L.P.

Table of Contents

Intel(R) Fortran Compiler User and Reference Guides	1
Intel® Fortran Compiler User and Reference Guides.....	1
Disclaimer and Legal Information	1
Welcome to the Intel® Fortran Compiler	2
See Also	2
Conventions	3
Notational Conventions.....	3
Platform Labels.....	7
Introduction to the Intel® Fortran Compiler	8
Product Website and Support.....	8
System Requirements	9
FLEXlm* Electronic Licensing.....	9
Related Publications.....	9
Tutorial information	9
Standard and Specification Documents.....	10
Associated Intel Documents	11
Optimization and Vectorization Terminology and Technology	11
Additional Training on the Intel® Fortran Compiler.....	12
Compiler Options.....	13
Overview: Compiler Options	13
New Options	14
Deprecated and Removed Compiler Options	34
Alphabetical Compiler Options.....	40
Quick Reference Guides and Cross References	699
Related Options	795
Floating-point Operations	805
Overview: Floating-point Operations.....	805
Floating-point Options Quick Reference	806
Understanding Floating-point Operations	811

Tuning Performance	822
Handling Floating-point Exceptions	826
Understanding IEEE Floating-point Standard	847

Intel(R) Fortran Compiler User and Reference Guides

Intel® Fortran Compiler User and Reference Guides

Document number: 304970-005US

[Start Here](#)

www.intel.com

[Disclaimer and Legal Information](#)

Disclaimer and Legal Information

INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL(R) PRODUCTS. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. EXCEPT AS PROVIDED IN INTEL'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, INTEL ASSUMES NO LIABILITY WHATSOEVER, AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO SALE AND/OR USE OF INTEL PRODUCTS INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT.

UNLESS OTHERWISE AGREED IN WRITING BY INTEL, THE INTEL PRODUCTS ARE NOT DESIGNED NOR INTENDED FOR ANY APPLICATION IN WHICH THE FAILURE OF THE INTEL PRODUCT COULD CREATE A SITUATION WHERE PERSONAL INJURY OR DEATH MAY OCCUR.

Intel may make changes to specifications and product descriptions at any time, without notice. Designers must not rely on the absence or characteristics of any features or instructions marked "reserved" or "undefined." Intel reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them. The information here is

subject to change without notice. Do not finalize a design with this information. The products described in this document may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

Contact your local Intel sales office or your distributor to obtain the latest specifications and before placing your product order.

Copies of documents which have an order number and are referenced in this document, or other Intel literature, may be obtained by calling 1-800-548-4725, or by visiting Intel's Web Site.

Intel processor numbers are not a measure of performance. Processor numbers differentiate features within each processor family, not across different processor families. See http://www.intel.com/products/processor_number for details.

This document contains information on products in the design phase of development.

BunnyPeople, Celeron, Celeron Inside, Centrino, Centrino Atom, Centrino Atom Inside, Centrino Inside, Centrino logo, Core Inside, FlashFile, i960, InstantIP, Intel, Intel logo, Intel386, Intel486, IntelDX2, IntelDX4, IntelSX2, Intel Atom, Intel Atom Inside, Intel Core, Intel Inside, Intel Inside logo, Intel. Leap ahead., Intel. Leap ahead. logo, Intel NetBurst, Intel NetMerge, Intel NetStructure, Intel SingleDriver, Intel SpeedStep, Intel StrataFlash, Intel Viiv, Intel vPro, Intel XScale, Itanium, Itanium Inside, MCS, MMX, Oplus, OverDrive, PDCharm, Pentium, Pentium Inside, skool, Sound Mark, The Journey Inside, Viiv Inside, vPro Inside, VTune, Xeon, and Xeon Inside are trademarks of Intel Corporation in the U.S. and other countries.

* Other names and brands may be claimed as the property of others.

Copyright (C) 1996-2008, Intel Corporation. All rights reserved.

Portions Copyright (C) 2001, Hewlett-Packard Development Company, L.P.

Welcome to the Intel® Fortran Compiler

The Intel® Fortran Compiler lets you build and optimize Fortran applications for the Linux* OS (operating system).

See Also

- Introduction
- Building Applications
- Compiler Options
- Optimizing Applications
- Floating-point Operations
- Language Reference

For details on getting started with the Intel Fortran Compiler, see:

- Getting Started
- Invoking the Compiler from the Command Line

Conventions

Information in this documentation applies to all supported operating systems and architectures unless otherwise specified.

This documentation uses the following conventions:

- [Notational Conventions](#)
- [Platform Labels](#)

Notational Conventions

THIS TYPE	Indicates statements, data types, directives, and other language keywords. Examples of statement keywords are WRITE, INTEGER, DO, and OPEN.
<i>this type</i>	Indicates command-line or option arguments, new terms, or emphasized text. Most new terms are defined in the Glossary.
<code>This type</code>	Indicates a code example.
This type	Indicates what you type as input.
This type	Indicates menu names, menu items, button names, dialog window names, and other user-interface items.
File>Open	Menu names and menu items joined by a greater than (>) sign indicate a sequence of actions. For example, "Click File>Open " indicates that in the File menu, click Open to perform this action.

{value value}	Indicates a choice of items or values. You can usually only choose one of the values in the braces.
[<i>item</i>]	Indicates items that are optional. Brackets are also used in code examples to show arrays.
<i>item</i> [, <i>item</i>]...	Indicates that the item preceding the ellipsis (three dots) can be repeated. In some code examples, a horizontal ellipsis means that not all of the statements are shown.
Windows* OS Windows operating system	These terms refer to all supported Microsoft* Windows* operating systems.
Linux* OS Linux operating system	These terms refer to all supported Linux* operating systems.
Mac OS* X Mac OS X operating system	These terms refer to Intel®-based systems running the Mac OS* X operating system.
Microsoft Windows XP*	An asterisk at the end of a word or name indicates it is a third-party product trademark.
compiler option	This term refers to Windows* OS options, Linux* OS options, or MAC OS* X options that can be used on the compiler command line.

Conventions Used in *Compiler Options*

/option or -option	A slash before an option name indicates the option is available on Windows OS. A dash before an option name indicates the option is available on Linux OS* and Mac OS* X systems. For example:
-----------------------	--

Windows option: `/fast`

Linux and Mac OS X option: `-fast`

Note: If an option is available on Windows* OS, Linux* OS, and Mac OS* X systems, no slash or dash appears in the general description of the option. The slash and dash will only appear where the option syntax is described.

`/option:argument` Indicates that an option requires a argument (parameter).
 or
`-option argument` For example, you must specify an argument for the following options:

Windows OS option: `/Qdiag-error-limit:n`

Linux OS and Mac OS X option: `-diag-error-limit n`

`/option:keyword` Indicates that an option requires one of the *keyword*
 or
 values.

`-option keyword`

`/option[:keyword]` Indicates that the option can be used alone or with an
 or
 optional keyword.

`-option [keyword]`

`option[n]` Indicates that the option can be used alone or with an
 optional value; for example, in `/Qunroll[:n]` or `-unroll[n]`, the *n* can be omitted or a valid value can be specified for *n*.

`option[-]` Indicates that a trailing hyphen disables the option; for
 example, `/Qglobal_hoist-` disables the Windows OS
 option `/Qglobal_hoist`.

`[no]option` or
`[no-]option` Indicates that "no" or "no-" preceding an option disables
 the option. For example:

In the Windows OS option `/[no]traceback`,
`/traceback` enables the option, while `/notraceback`

disables it.

In the Linux OS and Mac OS X option `-[no-]global_hoist`, `-global_hoist` enables the option, while `-no-global_hoist` disables it.

In some options, the "no" appears later in the option name; for example, `-fno-alias` disables the `-falias` option.

Conventions Used in *Language Reference*

This color Indicates extensions to the Fortran 95 Standard. These extensions may or may not be implemented by other compilers that conform to the language standard.

Intel Fortran This term refers to the name of the common compiler language supported by the Intel® Visual Fortran Compiler and the Intel® Fortran Compiler. For more information on these compilers, see <http://developer.intel.com/software/products/>.

Fortran This term refers to language information that is common to ANSI FORTRAN 77, ANSI/ISO Fortran 95 and 90, and Intel Fortran.

Fortran 95/90 This term refers to language information that is common to ANSI FORTRAN 77, ANSI/ISO Fortran 95, ANSI/ISO Fortran 90, and Intel Fortran.

Fortran 95 This term refers to language features specific to ANSI/ISO Fortran 95.

integer This term refers to the `INTEGER(KIND=1)`, `INTEGER(KIND=2)`, `INTEGER (INTEGER(KIND=4))`, and `INTEGER(KIND=8)` data types as a group.

real This term refers to the `REAL (REAL(KIND=4))`, `DOUBLE PRECISION (REAL(KIND=8))`, and `REAL(KIND=16)` data types as a group.

- REAL** This term refers to the default data type of objects declared to be REAL. REAL is equivalent to REAL(KIND=4), unless a compiler option specifies otherwise.
- complex** This term refers to the COMPLEX (COMPLEX(KIND=4)), **DOUBLE COMPLEX** (COMPLEX(KIND=8)), and COMPLEX(KIND=16) data types as a group.
- logical** This term refers to the LOGICAL(KIND=1), LOGICAL(KIND=2), LOGICAL (LOGICAL(KIND=4)), and LOGICAL(KIND=8) data types as a group.
- Compatibility** This term introduces a list of the projects or libraries that are compatible with the library routine.
- < Tab >** This symbol indicates a nonprinting tab character.
- ^** This symbol indicates a nonprinting blank character.

Platform Labels

A platform is a combination of operating system and central processing unit (CPU) that provides a distinct environment in which to use a product (in this case, a computer language). An example of a platform is Microsoft* Windows* XP on processors using IA-32 architecture.

In this documentation, information applies to all supported platforms unless it is otherwise labeled for a specific platform (or platforms).

These labels may be used to identify specific platforms:

- L*X** Applies to Linux* OS on processors using IA-32 architecture, Intel® 64 architecture, and IA-64 architecture.
- L*X32** Applies to Linux* OS on processors using IA-32 architecture and Intel® 64 architecture.
- L*X64** Applies to Linux OS on processors using IA-64 architecture.
- M*X** Applies to Apple* Mac OS* X on processors using IA-32 architecture and Intel® 64 architecture.

M*X32 Applies to Apple* Mac OS* X on processors using IA-32 architecture.

M*X64 Applies to Apple* Mac OS* X on processors using Intel® 64 architecture.

W*32 Applies to Microsoft Windows* 2000, Windows XP, and Windows Server 2003 on processors using IA-32 architecture and Intel® 64 architecture. For a complete list of supported Windows* operating systems, see your Release Notes.

W*64 Applies to Microsoft Windows* XP operating systems on IA-64 architecture.

i32 Applies to 32-bit operating systems on IA-32 architecture.

i64em Applies to 32-bit operating systems on Intel® 64 architecture.

i64 Applies to 64-bit operating systems on IA-64 architecture.

Introduction to the Intel® Fortran Compiler

The Intel® Fortran Compiler can generate code for IA-32, Intel® 64, or IA-64 applications on any Intel®-based Linux* system. IA-32 applications (32-bit) can run on all Intel®-based Linux systems. Intel® 64 applications and IA-64 applications can only run on Intel® 64-based or IA-64-based Linux systems. For more information about the compiler features and other components, see your *Release Notes*.

This documentation assumes that you are familiar with the Fortran programming language and with your processor's architecture. You should also be familiar with the host computer's operating system.

Product Website and Support

For general information on support for Intel software products, visit the Intel web site <http://developer.intel.com/software/products/>

At this site, you will find comprehensive product information, including:

- Links to each product, where you will find technical information such as white papers and articles
- Links to user forums

- Links to news and events

To find technical support information, to register your product, or to contact Intel, please visit: <http://www.intel.com/software/products/support/>

For additional information, see the Technical Support section of your Release Notes.

System Requirements

For detailed information on system requirements, see the Release Notes.

FLEXlm* Electronic Licensing

The Intel® Fortran Compiler uses Macrovision*'s FLEXlm* licensing technology.

The compiler requires a valid license file in the `licenses` directory in the installation path. The default directory is `/opt/intel/licenses`.

License files have a file extension of `.lic`.

For information on how to install and use the Intel® License Manager for FLEXlm to configure a license server for systems using counted licenses, see *Using the Intel® License Manager for FLEXlm* (flex Ug.pdf)*.

Related Publications

Tutorial information

The following commercially published documents provide reference or tutorial information on Fortran 2003, Fortran 95, and Fortran 90:

- *Fortran 95/2003 for Scientists & Engineers* by S. Chapman; published by McGraw-Hill Science/Engineering/Math, ISBN 0073191574.
- *Fortran 95/2003 Explained* by M. Metcalf, J. Reid, and M. Cohen; published by Oxford University Press, ISBN 0-19-852693-8.
- *Compaq Visual Fortran* by N. Lawrence; published by Digital Press* (Butterworth-Heinemann), ISBN 1-55558-249-4.
- *Digital Visual Fortran Programmers Guide* by M. Etzel and K. Dickinson; published by Digital Press (Butterworth-Heinemann), ISBN 1-55558-218-4
- *Fortran 90 Explained* by M. Metcalf and J. Reid; published by Oxford University Press, ISBN 0-19-853772-7.
- *Fortran 90/95 Explained* by M. Metcalf and J. Reid; published by Oxford University Press, ISBN 0-19-851888-9.

- *Fortran 90/95 for Scientists and Engineers* by S. Chapman; published by McGraw-Hill, ISBN 0-07-011938-4.
- *Fortran 90 Handbook* by J. Adams, W. Brainerd, J. Martin, B. Smith, and J. Wagener; published by Intertext Publications (McGraw-Hill), ISBN 0-07-000406-4.
- *Fortran 90 Programming* by T. Ellis, I. Philips, and T. Lahey; published by Addison-Wesley, ISBN 0201-54446-6.
- *Introduction to Fortran 90/95* by Stephen J. Chapman; published by McGraw-Hill, ISBN 0-07-011969-4.
- *User's guide to Fortran 90, Second Edition* by W. Brainerd, C. Goldberg, and J. Adams; published by Unicomp, ISBN 0-07-000248-7.

Intel does not endorse these books or recommend them over other books on the same subjects.

Standard and Specification Documents

The following copyrighted standard and specification documents provide descriptions of many of the features found in Intel® Fortran:

- American National Standard Programming Language FORTRAN, ANSI X3.9-1978
- American National Standard Programming Language Fortran 90, ANSI X3.198-1992
This Standard is equivalent to: International Standards Organization Programming Language Fortran, ISO/IEC 1539:1991 (E).
- American National Standard Programming Language Fortran 95, ANSI X3J3/96-007
This Standard is equivalent to: International Standards Organization Programming Language Fortran, ISO/IEC 1539-1:1997 (E).
- International Standards Organization Information Technology - Programming Languages - Fortran, ISO/IEC 1539-1:2004 (E)
This is the Fortran 2003 Standard.
- High Performance Fortran Language Specification, Version 1.1, Technical Report CRPC-TR-92225

- OpenMP Fortran Application Program Interface, Version 1.1, November 1999
- OpenMP Fortran Application Program Interface, Version 2.0, November 2000

Associated Intel Documents

The following Intel documents provide additional information about the Intel® Fortran Compiler, Intel® architecture, Intel® processors, or tools:

- *Using the Intel® License Manager for FLEXlm**
- *Intel® 64 and IA-32 Architectures Software Developer's Manual Volume 1: Basic Architecture*, Intel Corporation
- *Intel® 64 and IA-32 Architectures Software Developer's Manual Volume 2A: Instruction Set Reference, A-M*, Intel Corporation
- *Intel® 64 and IA-32 Architectures Software Developer's Manual Volume 2B: Instruction Set Reference, N-Z*, Intel Corporation
- *Intel® 64 and IA-32 Architectures Software Developer's Manual Volume 3A: System Programming Guide*, Intel Corporation
- *Intel® 64 and IA-32 Architectures Software Developer's Manual Volume 3B: System Programming Guide*, Intel Corporation
- *Intel® 64 and IA-32 Architectures Optimization Reference Manual*
- *Intel® Itanium® Architecture Software Developer's Manual - Volume 1: Application Architecture, Revision 2.2*
- *Intel® Itanium® Architecture Software Developer's Manual - Volume 2: System Architecture, Revision 2.2*
- *Intel® Itanium® Architecture Software Developer's Manual - Volume 3: Instruction Set Reference, Revision 2.2*
- *Intel® Processor Identification with the CPUID Instruction*, Intel Corporation, doc. number 241618
- *IA-64 Architecture Assembler User's Guide*
- *IA-64 Architecture Assembly Language Reference Guide*

Most Intel documents can be found at the Intel web site

<http://developer.intel.com/software/products/>

Optimization and Vectorization Terminology and Technology

The following documents provide details on basic optimization and vectorization terminology and technology:

- *Intel® Architecture Optimization Reference Manual*
- *Dependence Analysis*, Utpal Banerjee (A Book Series on Loop Transformations for Restructuring Compilers). Kluwer Academic Publishers. 1997.
- *The Structure of Computers and Computation: Volume I*, David J. Kuck. John Wiley and Sons, New York, 1978.
- *Loop Transformations for Restructuring Compilers: The Foundations*, Utpal Banerjee (A Book Series on Loop Transformations for Restructuring Compilers). Kluwer Academic Publishers. 1993.
- *Loop parallelization*, Utpal Banerjee (A Book Series on Loop Transformations for Restructuring Compilers). Kluwer Academic Publishers. 1994.
- *High Performance Compilers for Parallel Computers*, Michael J. Wolfe. Addison-Wesley, Redwood City. 1996.
- *Supercompilers for Parallel and Vector Computers*, H. Zima. ACM Press, New York, 1990.
- *An Auto-vectorizing Compiler for the Intel® Architecture*, Aart Bik, Paul Grey, Milind Girkar, and Xinmin Tian. Submitted for publication
- *Efficient Exploitation of Parallelism on Pentium® III and Pentium® 4 Processor-Based Systems*, Aart Bik, Milind Girkar, Paul Grey, and Xinmin Tian.
- *The Software Vectorization Handbook. Applying Multimedia Extensions for Maximum Performance*, A.J.C. Bik. Intel Press, June, 2004.
- *Multi-Core Programming: Increasing Performance through Software Multithreading*, Shameem Akhter and Jason Roberts. Intel Press, April, 2006.

Additional Training on the Intel® Fortran Compiler

For additional training on the Intel Fortran Compiler, choose a course in the Intel® Software College - Course Catalog at

<http://shale.intel.com/SoftwareCollege/CourseCatalog.asp>

For additional technical product information including white papers about Intel compilers, open the page associated with your product at

<http://developer.intel.com/software/products/>

Compiler Options

Overview: Compiler Options

This document provides details on all current Linux*, Mac OS* X, and Windows* compiler options.

It provides the following information:

- [New options](#)
This topic lists new compiler options in this release.
- [Deprecated](#)
This topic lists deprecated and removed compiler options for this release. Some deprecated options show suggested replacement options.
- [Alphabetical Compiler Options](#)
This topic is the main source in the documentation set for general information on all compiler options. Options are described in alphabetical order. The [Overview](#) describes what information appears in each compiler option description.
- [Quick Reference Guide and Cross Reference](#)
This topic contains tables summarizing compiler options. The tables show the option name, a short description of the option, the default setting for the option, and the equivalent option on the operating system, if any.
- [Related Options](#)
This topic lists related options that can be used under certain conditions.

In this guide, compiler options are available on all supported operating systems and architectures unless otherwise identified.

For further information on compiler options, see Building Applications and Optimizing Applications.

Functional Groupings of Compiler Options

To see functional groupings of compiler options, specify a functional category for option `help` on the command line. For example, to see a list of options that affect diagnostic messages displayed by the compiler, enter one of the following commands:

```
-help diagnostics      ! Linux and Mac OS X systems
/help diagnostics     ! Windows systems
```

For details on the categories you can specify, see [help](#).

New Options

This topic lists the options that provide new functionality in this release.

Some compiler options are only available on certain systems, as indicated by these labels:

Label	Meaning
i32	The option is available on systems using IA-32 architecture.
i64em	The option is available on systems using Intel® 64 architecture.
i64	The option is available on systems using IA-64 architecture.

If no label appears, the option is available on all supported systems.

If "only" appears in the label, the option is only available on the identified system.

For more details on the options, refer to the [Alphabetical Compiler Options](#) section.

For information on conventions used in this table, see [Conventions](#).

New compiler options are listed in tables below:

- The first table lists new options that are available on Windows* systems.
- The second table lists new options that are available on Linux* and Mac OS* X systems. If an option is only available on one of these operating systems, it is labeled.

Windows* Options	Description	Default
/arch:IA32 (i32 only)	Generates code that will run on any Pentium or later	OFF

Windows* Options	Description	Default
	processor.	
/arch:SSE3 (i32, i64em)	Optimizes for Intel® Streaming SIMD Extensions 3 (Intel® SSE3).	OFF
/arch:SSSE3 (i32, i64em)	Optimizes for Intel® Supplemental Streaming SIMD Extensions 3 (Intel® SSSE3).	OFF
/arch:SSE4.1 (i32, i64em)	Optimizes for Intel® Streaming SIMD Extensions 4 Vectorizing Compiler and Media Accelerators.	OFF
/GS (i32, i64em)	Determines whether the compiler generates code that detects some buffer overruns.	/GS-
/QaxSSE2 (i32, i64em)	Can generate Intel® SSE2 and SSE instructions for Intel processors, and it can optimize for Intel® Pentium® 4 processors, Intel®	OFF

Windows* Options	Description	Default
	Pentium® M processors, and Intel® Xeon® processors with Intel® SSE2.	
/QaxSSE3 (i32, i64em)	Can generate Intel® SSE3, SSE2, and SSE instructions for Intel processors and it can optimize for processors based on Intel® Core™ microarchitecture and Intel NetBurst® microarchitecture.	OFF
/QaxSSSE3 (i32, i64em)	Can generate Intel® SSSE3, SSE3, SSE2, and SSE instructions for Intel processors and it can optimize for the Intel® Core™2 Duo processor family.	OFF
/QaxSSE4.1 (i32, i64em)	Can generate Intel® SSE4 Vectorizing Compiler and Media Accelerator instructions for Intel processors. Can	OFF

Windows* Options	Description	Default
	generate Intel® SSSE3, SSE3, SSE2, and SSE instructions and it can optimize for Intel® 45nm Hi-k next generation Intel® Core™ microarchitecture.	
/QaxSSE4.2 (i32, i64em)	Can generate Intel® SSE4 Efficient Accelerated String and Text Processing instructions supported by Intel® Core™ i7 processors. Can generate Intel® SSE4 Vectorizing Compiler and Media Accelerator, Intel® SSSE3, SSE3, SSE2, and SSE instructions and it can optimize for the Intel® Core™ processor family.	OFF
/Qdiag-error-limit:n	Specifies the maximum number of errors allowed before compilation stops.	n=30

Windows* Options	Description	Default
<code>/Qdiag-once:id[,id, ...]</code>	Tells the compiler to issue one or more diagnostic messages only once.	OFF
<code>/Qfast-transcendentals</code>	Enables the compiler to replace calls to transcendental functions with faster but less precise implementations.	OFF
<code>/Qfma</code> (i64 only)	Enables the combining of floating-point multiplies and add/subtract operations.	ON
<code>/Qfp-relaxed</code> (i64 only)	Enables use of faster but slightly less accurate code sequences for math functions.	OFF
<code>/Qinstruction:[no]movbe</code> (i32, i64em)	Determines whether MOVBE instructions are generated for Intel processors.	ON
<code>/Qopenmp-link:library</code>	Controls whether the compiler links to static or dynamic OpenMP run-time libraries.	<code>/Qopenmp-link:dynamic</code>

Windows* Options	Description	Default
<code>/Qopenmp- threadprivate:type</code>	Lets you specify an OpenMP* threadprivate implementation.	<code>/Qopenmp- threadprivate:legacy</code>
<code>/Qopt-block-factor:n</code>	Lets you specify a loop blocking factor.	OFF
<code>/Qopt-jump- tables:keyword</code>	Enables or disables generation of jump tables for switch statements.	<code>/Qopt-jump- tables:default</code>
<code>/Qopt-loadpair (i64 only)</code>	Enables loadpair optimization.	<code>/Qopt-loadpair-</code>
<code>/Qopt-mod-versioning (i64 only)</code>	Enables versioning of modulo operations for certain types of operands.	<code>/Qopt-mod- versioning-</code>
<code>/Qopt-prefetch-initial- values (i64 only)</code>	Enables or disables prefetches that are issued before a loop is entered.	<code>/Qopt-prefetch- initial-values</code>
<code>/Qopt-prefetch-issue- excl-hint (i64 only)</code>	Determines whether the compiler issues prefetches for stores with exclusive hint.	<code>/Qopt-prefetch- issue-excl-hint-</code>
<code>/Qopt-prefetch-next- iteration (i64 only)</code>	Enables or disables prefetches for a memory access in the	<code>/Qopt-prefetch-next- iteration</code>

Windows* Options	Description	Default
	next iteration of a loop.	
<code>/Qopt-subscript-in-range</code> (i32, i64em)	Determines whether the compiler assumes no overflows in the intermediate computation of subscript expressions in loops.	<code>/Qopt-subscript-in-range-</code>
<code>/Qprof-data-order</code>	Enables or disables data ordering if profiling information is enabled.	<code>/Qprof-data-order</code>
<code>/Qprof-func-order</code>	Enables or disables function ordering if profiling information is enabled.	<code>/Qprof-func-order</code>
<code>/Qprof-hotness-threshold</code>	Lets you set the hotness threshold for function grouping and function ordering.	OFF
<code>/Qprof-src-dir</code>	Determines whether directory information of the source file under compilation is considered when looking up profile data records.	<code>/Qprof-src-dir</code>

Windows* Options	Description	Default
<code>/Qprof-src-root</code>	Lets you use relative directory paths when looking up profile data and specifies a directory as the base.	OFF
<code>/Qprof-src-root-cwd</code>	Lets you use relative directory paths when looking up profile data and specifies the current working directory as the base.	OFF
<code>/Qtcollect-filter</code>	Lets you enable or disable the instrumentation of specified functions.	OFF
<code>/Quse-msasm-symbols</code> (i32, i64em)	Tells the compiler to use a dollar sign ("\$") when producing symbol names.	OFF
<code>/Qvc9</code> (i32, i64em)	Specifies compatibility varies with Microsoft* Visual Studio 2008.	
<code>/Qvec</code> (i32, i64em)	Enables or disables vectorization and transformations enabled for vectorization.	<code>/Qvec</code>

Windows* Options	Description	Default
/QxHost (i32, i64em)	Can generate specialized code paths for the highest instruction set and processor available on the compilation host.	OFF
/QxSSE2 (i32, i64em)	Can generate Intel® SSE2 and SSE instructions for Intel processors, and it can optimize for Intel® Pentium® 4 processors, Intel® Pentium® M processors, and Intel® Xeon® processors with Intel® SSE2.	ON
/QxSSE3 (i32, i64em)	Can generate Intel® SSE3, SSE2, and SSE instructions for Intel processors, and it can optimize for processors based on Intel® Core™ microarchitecture and Intel NetBurst® microarchitecture.	OFF

Windows* Options	Description	Default
/QxSSE3_ATOM (i32, i64em)	Can generate MOVBE instructions for Intel processors and it can optimize for the Intel® Atom™ processor and Intel® Centrino® Atom™ Processor Technology.	OFF
/QxSSSE3 (i32, i64em)	Can generate Intel® SSSE3, SSE3, SSE2, and SSE instructions for Intel processors and it can optimize for the Intel® Core™2 Duo processor family.	OFF
/QxSSE4.1 (i32, i64em)	Can generate Intel® SSE4 Vectorizing Compiler and Media Accelerator instructions for Intel processors. Can generate Intel® SSSE3, SSE3, SSE2, and SSE instructions and it can optimize for Intel® 45nm Hi-k next generation Intel® Core™	OFF

Windows* Options	Description	Default
	microarchitecture.	
/QxSSE4.2 (i32, i64em)	Can generate Intel® SSE4 Efficient Accelerated String and Text Processing instructions supported by Intel® Core™ i7 processors. Can generate Intel® SSE4 Vectorizing Compiler and Media Accelerator, Intel® SSSE3, SSE3, SSE2, and SSE instructions and it can optimize for the Intel® Core™ processor family.	OFF

Linux* and Mac OS* X Options	Description	Default
-axSSE2 (i32, i64em)	Can generate Intel® SSE2 and SSE instructions for Intel processors, and it can optimize for Intel® Pentium® 4 processors, Intel® Pentium® M processors, and Intel®	OFF

Linux* and Mac OS* X Options	Description	Default
	Xeon® processors with Intel® SSE2.	
-axSSE3 (i32, i64em)	Can generate Intel® SSE3, SSE2, and SSE instructions for Intel processors, and it can optimize for processors based on Intel® Core microarchitecture and Intel NetBurst® microarchitecture.	OFF
-axSSSE3 (i32, i64em)	Can generate Intel® SSSE3, SSE3, SSE2, and SSE instructions for Intel processors and it can optimize for the Intel® Core™2 Duo processor family.	OFF
-axSSE4.1 (i32, i64em)	Can generate Intel® SSE4 Vectorizing Compiler and Media Accelerator instructions for Intel processors. Can generate Intel® SSSE3, SSE3, SSE2, and SSE instructions and it can optimize for	OFF

Linux* and Mac OS* X Options	Description	Default
	Intel® 45nm Hi-k next generation Intel® Core™ microarchitecture.	
-axSSE4.2 (i32, i64em)	Can generate Intel® SSE4 Efficient Accelerated String and Text Processing instructions supported by Intel® Core™ i7 processors. Can generate Intel® SSE4 Vectorizing Compiler and Media Accelerator, Intel® SSSE3, SSE3, SSE2, and SSE instructions and it can optimize for the Intel® Core™ processor family.	OFF
-diag-error-limit n	Specifies the maximum n=30 number of errors allowed before compilation stops.	
-diag-once id[,id,...]	Tells the compiler to issue one or more diagnostic messages	OFF

Linux* and Mac OS* X Options	Description	Default
	only once.	
<code>-falign-stack</code> (i32 only)	Tells the compiler the stack alignment to use on entry to routines.	<code>-falign-stack=default</code>
<code>-fast-transcendentals</code>	Enables the compiler to replace calls to transcendental functions with faster but less precise implementation.	OFF
<code>-finline</code>	Tells the compiler to inline functions declared with <code>cDEC\$ ATTRIBUTES FORCEINLINE</code> .	<code>-fno-inline</code>
<code>-fma</code> (i64 only; Linux only)	Enables the combining of floating-point multiplies and add/subtract operations.	ON
<code>-fp-relaxed</code> (i64 only; Linux only)	Enables use of faster but slightly less accurate code sequences for math functions.	OFF
<code>-fpie</code> (Linux only)	Tells the compiler to	OFF

Linux* and Mac OS* X Options	Description	Default
	generate position-independent code to link into executables.	
-fstack-protector (i32, i64em)	Determines whether the compiler generates code that detects some buffer overruns. Same as option <code>-fstack-security-check</code> .	-fno-stack-protector
-m32, -m64 (i32, i64em)	Tells the compiler to generate code for a specific architecture.	OFF
-mia32 (i32 only)	Generates code that will run on any Pentium or later processor.	OFF
- minstruction=[no]movbe (i32, i64em)	Determines whether MOVBE instructions are generated for Intel processors.	ON
-mssse3 (i32, i64em)	Generates code for Intel® Supplemental Streaming SIMD Extensions 3 (Intel® SSSE3).	Linux systems: OFF Mac OS X systems using Intel® 64 architecture: ON
-msse4.1 (i32, i64em)	Generates code for Intel® Streaming SIMD	OFF

Linux* and Mac OS* X Options	Description	Default
	Extensions 4 Vectorizing Compiler and Media Accelerators.	
-openmp-link library	Controls whether the compiler links to static or dynamic OpenMP run-time libraries.	-openmp-link dynamic
-openmp- threadprivate=type (Linux only)	Lets you specify an OpenMP* threadprivate implementation.	-openmp- threadprivate=legacy
-opt-block-factor=n	Lets you specify a loop blocking factor.	OFF
-opt-jump- tables=keyword	Enables or disables generation of jump tables for switch statements.	-opt-jump- tables=default
-opt-loadpair (i64 only; Linux only)	Enables loadpair optimization.	-no-opt-loadpair
-opt-mod-versioning (i64 only; Linux only)	Enables versioning of modulo operations for certain types of operands.	-no-opt-mod- versioning
-opt-prefetch-initial- values (i64 only; Linux only)	Enables or disables prefetches that are issued before a loop is	-opt-prefetch- initial-values

Linux* and Mac OS* X Options	Description	Default
	entered.	
<code>-opt-prefetch-issue-excl-hint</code> (i64 only)	Determines whether the compiler issues prefetches for stores with exclusive hint.	<code>-no-opt-prefetch-issue-excl-hint</code>
<code>-opt-prefetch-next-iteration</code> (i64 only; Linux only)	Enables or disables prefetches for a memory access in the next iteration of a loop.	<code>-opt-prefetch-next-iteration</code>
<code>-opt-subscript-in-range</code> (i32, i64em)	Determines whether the compiler assumes no overflows in the intermediate computation of subscript expressions in loops.	<code>-no-opt-subscript-in-range</code>
<code>-pie</code> (Linux only)	Produces a position-independent executable on processors that support it.	OFF
<code>-prof-data-order</code> (Linux only)	Enables or disables data ordering if profiling information is enabled.	<code>-no-prof-data-order</code>
<code>-prof-func-groups</code> (i32, i64em; Linux only)	Enables or disables function grouping if	<code>-no-prof-func-groups</code>

Linux* and Mac OS* X Options	Description	Default
	profiling information is enabled.	
<code>-prof-func-order</code> (Linux only)	Enables or disables function ordering if profiling information is enabled.	<code>-no-prof-func-order</code>
<code>-prof-hotness-threshold</code> (Linux only)	Lets you set the hotness threshold for function grouping and function ordering.	OFF
<code>-prof-src-root</code>	Lets you use relative directory paths when looking up profile data and specifies a directory as the base.	OFF
<code>-prof-src-root-cwd</code>	Lets you use relative directory paths when looking up profile data and specifies the current working directory as the base.	OFF
<code>-staticlib</code> (i32, i64em; Mac OS X only)	Invokes the <code>libtool</code> command to generate static libraries.	OFF
<code>-tcollect-filter</code> (Linux only)	Lets you enable or disable the	OFF

Linux* and Mac OS* X Options	Description	Default
	instrumentation of specified functions.	
-vec (i32, i64em)	Enables or disables vectorization and transformations enabled for vectorization.	-vec
-xHost (i32, i64em)	Can generate instructions for the highest instruction set and processor available on the compilation host.	OFF
-xSSE2 (i32, i64em)	Can generate Intel® SSE2 and SSE instructions for Intel processors, and it can optimize for Intel® Pentium® 4 processors, Intel® Pentium® M processors, and Intel® Xeon® processors with Intel® SSE2.	Linux systems:ON Mac OS X systems: OFF
-xSSE3 (i32, i64em)	Can generate Intel® SSE3, SSE2, and SSE instructions for Intel	Linux systems:OFF Mac OS X systems using IA-32 architecture: ON

Linux* and Mac OS* X Options	Description	Default
	processors and it can optimize for processors based on Intel® Core™ microarchitecture and Intel NetBurst® microarchitecture.	
-xSSE3_ATOM (i32, i64em)	Can generate MOVBE instructions for Intel processors and it can optimize for the Intel® Atom™ processor and Intel® Centrino® Atom™ Processor Technology.	OFF
-xSSSE3 (i32, i64em)	Can generate Intel® SSSE3, SSE3, SSE2, and SSE instructions for Intel processors and it can optimize for the Intel® Core™2 Duo processor family.	Linux systems:OFF Mac OS X systems using Intel® 64 architecture: ON
-xSSE4.1 (i32, i64em)	Can generate Intel® SSE4 Vectorizing Compiler and Media Accelerator instructions for Intel processors. Can generate Intel®	OFF

Linux* and Mac OS* X Options	Description	Default
	SSSE3, SSE3, SSE2, and SSE instructions and it can optimize for Intel® 45nm Hi-k next generation Intel® Core™ microarchitecture.	
-xSSE4.2 (i32, i64em)	Can generate Intel® SSE4 Efficient Accelerated String and Text Processing instructions supported by Intel® Core™ i7 processors. Can generate Intel® SSE4 Vectorizing Compiler and Media Accelerator, Intel® SSSE3, SSE3, SSE2, and SSE instructions and it can optimize for the Intel® Core™ processor family.	OFF

Deprecated and Removed Compiler Options

This topic lists deprecated and removed compiler options and suggests replacement options, if any are available.

Deprecated Options

Occasionally, compiler options are marked as "deprecated." Deprecated options are still supported in the current release, but are planned to be unsupported in future releases.

The following options are deprecated in this release of the compiler:

Linux* and Mac OS* X Options Suggested Replacement

-axK	None
-axN	-axSSE2
-axP	-axSSE3
-axS	-axSSE4.1
-axT	-axSSSE3
-axW	-msse2
-func-groups	-prof-func-groups
-i-dynamic	-shared-intel
-i-static	-static-intel
-inline-debug-info	-debug
-IPF-flt-eval-method0	-fp-model source
-IPF-fltacc	-fp-model precise
-no-IPF-fltacc	-fp-model fast
-IPF-fma	-fma
-IPF-fp-relaxed	-fp-relaxed
-march=pentiumii	None
-march=pentiumiii	-march=pentium3
-mcpu	-mtune
-mp	-fp-model
-Ob	-inline-level

Linux* and Mac OS* X Options Suggested Replacement

<code>-openmp-lib legacy</code>	None
<code>-openmpP</code>	<code>-openmp</code>
<code>-openmpS</code>	<code>-openmp-stubs</code>
<code>-prefetch</code>	<code>-opt-prefetch</code>
<code>-prof-genx</code>	<code>-prof-gen=srcpos</code>
<code>-use-asm</code>	None
<code>-use-pch</code>	<code>-pch-use</code>
<code>-xK</code>	<code>-mia32</code>
<code>-xN</code>	<code>-xSSE2</code>
<code>-xO</code>	<code>-msse3</code>
<code>-xP</code>	<code>-xSSE3</code>
<code>-xS</code>	<code>-xSSE4.1</code>
<code>-xT</code>	<code>-xSSSE3</code>
<code>-xW</code>	<code>-msse2</code>

Windows* Options Suggested Replacement

<code>/4Nb</code>	<code>/check:none</code>
<code>/4Yb</code>	<code>/check:all</code>
<code>/debug:partial</code>	None
<code>/Fm</code>	<code>/map</code>
<code>/G5</code>	None
<code>/G6 (or /GB)</code>	None
<code>/G7</code>	None
<code>/Ge</code>	<code>/Gs0</code>

Windows* Options	Suggested Replacement
/ML and/MLd	None
/Op	/fp
/QaxK	None
/QaxN	/QaxSSE2
/QaxP	/QaxSSE3
/QaxS	/QaxSSE4.1
/QaxT	/QaxSSSE3
/QaxW	/arch:SSE2
/Qinline-debug-info	None
/QIPF-flt-eval-method0	/fp:source
/QIPF-fltacc	/fp:precise
/QIPF-fltacc-	/fp:fast
/QIPF-fma	/Qfma
/QIPF-fp-relaxed	/Qfp-relaxed
/Qopenmp-lib:legacy	None
/Qprefetch	/Qopt-prefetch
/Qprof-genx	/Qprof-gen=srcpos
/Quse-asm	None
/QxK	None
/QxN	/QxSSE2
/QxO	/arch:SSE3
/QxP	/QxSSE3
/QxS	/QxSSE4.1

Windows* Options	Suggested Replacement
------------------	-----------------------

/QxT	/QxSSSE3
/QxW	/arch:SSE2
/Zd	/debug:minimal

Deprecated options are not limited to this list.

Removed Options

Some compiler options are no longer supported and have been removed. If you use one of these options, the compiler issues a warning, ignores the option, and then proceeds with compilation.

This version of the compiler no longer supports the following compiler options:

Linux* and Mac OS* X Options	Suggested Replacement
------------------------------	-----------------------

-axB	-axSSE2
-axi	None
-axM	None
-cxxlib-gcc [=dir]	-cxxlib [=dir]
-cxxlib-icc	None
-F	-preprocess-only or -P
-fp	-fno-omit-frame-pointer
-fpstkchk	-fp-stack-check
-IPF-fp-speculation	-fp-speculation
-ipo-obj (and -ipo_obj)	None
-Kpic, -KPIC	-fpic
-mtune=itanium	None
-nobss-init	-no-bss-init

Linux* and Mac OS* X Options	Suggested Replacement
<code>-opt-report-level</code>	<code>-opt-report</code>
<code>-prof-format-32</code>	None
<code>-prof-gen-sampling</code>	None
<code>-qp</code>	<code>-p</code>
<code>-shared-libcxa</code>	<code>-shared-libgcc</code>
<code>-ssp</code>	None
<code>-static-libcxa</code>	<code>-static-libgcc</code>
<code>-syntax</code>	<code>-syntax-only</code> or <code>-fsyntax-only</code>
<code>-tpp1</code>	None
<code>-tpp2</code>	<code>-mtune=itanium2</code>
<code>-tpp5</code>	None
<code>-tpp6</code>	None
<code>-tpp7</code>	<code>-mtune=pentium4</code>
<code>-xB</code>	<code>-xSSE2</code>
<code>-xi</code>	None
<code>-xM</code>	None

Windows* Options	Suggested Replacement
<code>/4ccD</code> (and <code>/4ccd</code>)	None
<code>/G1</code>	None
<code>/QaxB</code>	<code>/QaxSSE2</code>
<code>/Qaxi</code>	None
<code>/QaxM</code>	None
<code>/Qfpstkchk</code>	<code>/Qfp-stack-check</code>

Windows* Options	Suggested Replacement
/QIPF-fp-speculation	/Qfp-speculation
/Qipo-obj (and /Qipo_obj)	None
/Qopt-report-level	/Qopt-report
/Qprof-format-32	None
/Qprof-gen-sampling	None
/Qssp	None
/Qvc6	None
/Qvc7	None
/QxB	/QxSSE2
/Qxi	None
/QxM	None

Removed options are not limited to these lists.

Alphabetical Compiler Options

Compiler Option Descriptions and General Rules

This section describes all the current Linux*, Mac OS* X, and Windows* compiler options in alphabetical order.

Option Descriptions

Each option description contains the following information:

- A short description of the option.
- IDE Equivalent

This shows information related to the integrated development environment (IDE) Property Pages on Windows, Linux, and Mac OS X systems. It shows on which Property Page the option appears, and under what category it's listed. The Windows IDE is Microsoft* Visual Studio* .NET; the Linux IDE is Eclipse*; the Mac OS X IDE is Xcode*. If the option has no IDE equivalent, it

will specify "None". Note that in this release, there is no IDE support for Fortran on Linux.

- Architectures

This shows the architectures where the option is valid. Possible architectures are:

- IA-32 architecture
- Intel® 64 architecture
- IA-64 architecture

- Syntax

This shows the syntax on Linux and Mac OS X systems and the syntax on Windows systems. If the option has no syntax on one of these systems, that is, the option is not valid on a particular system, it will specify "None".

- Arguments

This shows any arguments (parameters) that are related to the option. If the option has no arguments, it will specify "None".

- Default

This shows the default setting for the option.

- Description

This shows the full description of the option. It may also include further information on any applicable arguments.

- Alternate Options

These are options that are synonyms of the described option. If there are no alternate options, it will specify "None".

Many options have an older spelling where underscores ("_") instead of hyphens ("-") connect the main option names. The older spelling is a valid alternate option name.

Some option descriptions may also have the following:

- Example

This shows a short example that includes the option

- See Also

This shows where you can get further information on the option or related options.

General Rules for Compiler Options

You cannot combine options with a single dash (Linux and Mac OS X) or slash (Windows). For example:

- On Linux and Mac OS X systems: This is incorrect: `-wc`; this is correct: `-w -c`
- On Windows systems: This is incorrect: `/wc`; this is correct: `/w /c`

All Linux and Mac OS X compiler options are case sensitive. Many Windows options are case sensitive. Some options have different meanings depending on their case; for example, option "c" prevents linking, but option "C" checks for certain conditions at run time.

Options specified on the command line apply to all files named on the command line.

Options can take arguments in the form of file names, strings, letters, or numbers. If a string includes spaces, the string must be enclosed in quotation marks. For example:

- On Linux and Mac OS X systems, `-dynamic-linker mylink` (file name) or `-Umacro3` (string)
- On Windows systems, `/Famyfile.s` (file name) or `/V"version 5.0"` (string)

Compiler options can appear in any order.

On Windows systems, all compiler options must precede `/link` options, if any, on the command line.

Unless you specify certain options, the command line will both compile and link the files you specify.

You can abbreviate some option names, entering as many characters as are needed to uniquely identify the option.

Certain options accept one or more keyword arguments following the option name. For example, the `arch` option accepts several keywords.

To specify multiple keywords, you typically specify the option multiple times. However, there are exceptions; for example, the following are valid: `-axNB` (Linux) or `/QaxNB` (Windows).



Note

On Windows systems, you can sometimes use a comma to separate keywords. For example, the following is valid:

```
ifort /warn:usage,declarations test.f90
```

On these systems, you can use an equals sign (=) instead of the colon:

```
ifort /warn=usage,declarations test.f90
```

Compiler options remain in effect for the whole compilation unless overridden by a compiler directive.

To disable an option, specify the negative form of the option.

On Windows systems, you can also disable one or more options by specifying option `/Od` last on the command line.



Note

On Windows systems, the `/Od` option is part of a mutually-exclusive group of options that includes `/Od`, `/O1`, `/O2`, `/O3`, and `/Ox`. The last of any of these options specified on the command line will override the previous options from this group.

If there are enabling and disabling versions of an option on the command line, the last one on the command line takes precedence.

Lists and Functional Groupings of Compiler Options

To see a list of all the compiler options, specify option `help` on the command line.

To see functional groupings of compiler options, specify a functional category for option `help`. For example, to see a list of options that affect diagnostic messages displayed by the compiler, enter one of the following commands:

```
-help diagnostics      ! Linux and Mac OS X systems  
/help diagnostics     ! Windows systems
```

For details on the categories you can specify, see [help](#).

1

See [onetrip](#).

4I2, 4I4, 4I8

See [integer-size](#).

4L72, 4L80, 4L132

See [extend-source](#).

4Na, 4Ya

See [automatic](#).

4Naltparam, 4Yaltparam

See [altparam](#).

4Nb,4Yb

See [check](#).

4Nd,4Yd

See [warn](#).

4Nf

See [fixed](#).

4Nportlib, 4Yportlib

Determines whether the compiler links to the library of portability routines.

IDE Equivalent

Windows: **Libraries > Use Portlib Library**

Linux: None

Mac OS X: None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: None

Windows: /4Nportlib
 /4Yportlib

Arguments

None

Default

/4Yportlib The library of portability routines is linked during compilation.

Description

Option /4Yportlib causes the compiler to link to the library of portability routines. This also includes Intel's functions for Microsoft* compatibility.

Option /4Nportlib prevents the compiler from linking to the library of portability routines.

Alternate Options

None

See Also

Building Applications: Portability Routines

4Ns,4Ys

See [stand](#).

4R8,4R16

See [real-size](#).

4Yf

See [free](#).

4Nportlib, 4Yportlib

Determines whether the compiler links to the library of portability routines.

IDE Equivalent

Windows: **Libraries > Use Portlib Library**

Linux: None

Mac OS X: None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: None

Windows: /4Nportlib

 /4Yportlib

Arguments

None

Default

/4Yportlib The library of portability routines is linked during compilation.

Description

Option `/4Yportlib` causes the compiler to link to the library of portability routines. This also includes Intel's functions for Microsoft* compatibility. Option `/4Nportlib` prevents the compiler from linking to the library of portability routines.

Alternate Options

None

See Also

Building Applications: Portability Routines

66

See [f66](#).

72,80,132

See [extend-source](#).

align

Tells the compiler how to align certain data items.

IDE Equivalent

Windows: **Data > Structure Member Alignment** (`/align:recnbyte`)

Data > Common Element Alignment (`/align:[no]commons,`
`/align:[no]dcommons`)

Data > SEQUENCE Types Obey Alignment Rules (`/align:[no]sequence`)

Linux: None

Mac OS X: **Data > Structure Member Alignment** (`-align`
`rec<1,2,4,8,16>byte`)

Data > Common Element Alignment (`-align [no]commons,`
`/align:[no]dcommons`)

Data > SEQUENCE Types Obey Alignment Rules (`-align [no]sequence`)

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-align [keyword]`
`-noalign`

Windows: `/align[:keyword]`
`/noalign`

Arguments

keyword Specifies the data items to align. Possible values are:

<code>none</code>	Prevents padding bytes anywhere in common blocks and structures.
<code>[no]commons</code>	Affects alignment of common block entities.
<code>[no]dcommons</code>	Affects alignment of common block entities.
<code>[no]records</code>	Affects alignment of derived-type components and fields of record structures.
<code>recnbyte</code>	Specifies a size boundary for derived-type components and fields of record structures.
<code>[no]sequence</code>	Affects alignment of sequenced derived-type components.
<code>all</code>	Adds padding bytes whenever possible to data items in common blocks and structures.

Default

<code>nocommons</code>	Adds no padding bytes for alignment of common blocks.
<code>nodcommmons</code>	Adds no padding bytes for alignment of common blocks.
<code>records</code>	Aligns derived-type components and record structure fields on default natural boundaries.
<code>nosequence</code>	Causes derived-type components declared with the <code>SEQUENCE</code> statement to be packed, regardless of current alignment rules set by the user.

By default, no padding is added to common blocks but padding is added to structures.

Description

This option specifies the alignment to use for certain data items. The compiler adds padding bytes to perform the alignment.

Option	Description
<code>align none</code>	Tells the compiler not to add padding bytes anywhere in common blocks or structures. This is the same as specifying <code>noalign</code> .
<code>align commons</code>	Aligns all common block entities on natural boundaries up to 4 bytes, by adding padding bytes as needed. The <code>align nocommons</code> option adds no padding to common blocks. In this case, unaligned data can occur unless the order of data items specified in the <code>COMMON</code> statement places the largest numeric data item first, followed by the next largest numeric data (and so on), followed by any character data.
<code>align</code>	Aligns all common block entities on natural boundaries up to 8

Option	Description
<code>dcommons</code>	<p>bytes, by adding padding bytes as needed.</p> <p>This option is useful for applications that use common blocks, unless your application has no unaligned data or, if the application might have unaligned data, all data items are four bytes or smaller. For applications that use common blocks where all data items are four bytes or smaller, you can specify <code>/align:commons</code> instead of <code>/align:dcommons</code>.</p> <p>The <code>align nodcommons</code> option adds no padding to common blocks.</p> <p>On Windows systems, if you specify the <code>/stand:f90</code> or <code>/stand:f95</code> option, <code>/align:dcommons</code> is ignored.</p> <p>On Linux and Mac OS X systems, if you specify any <code>-std</code> option or the <code>-stand f90</code> or <code>-stand f95</code> option, <code>-align dcommons</code> is ignored.</p>
<code>align norecords</code>	<p>Aligns components of derived types and fields within record structures on arbitrary byte boundaries with no padding.</p> <p>The <code>align records</code> option requests that multiple data items in record structures and derived-type structures without the <code>SEQUENCE</code> statement be naturally aligned, by adding padding as needed.</p>
<code>align recnbyte</code>	<p>Aligns components of derived types and fields within record structures on the smaller of the size boundary specified (<i>n</i>) or the boundary that will naturally align them. <i>n</i> can be 1, 2, 4, 8, or 16.</p> <p>When you specify this option, each structure member after the first is stored on either the size of the member type or <i>n</i>-byte boundaries, whichever is smaller. For example, to specify 2 bytes as the packing boundary (or alignment constraint) for all structures and unions in the file <code>prog1.f</code>, use the following</p>

Option	Description
	<p>command:</p> <pre data-bbox="467 323 1243 352">ifort {-align rec2byte /align:rec2byte} prog1.f</pre> <p>This option does not affect whether common blocks are naturally aligned or packed.</p>
align sequence	<p>Aligns components of a derived type declared with the SEQUENCE statement (sequenced components) according to the alignment rules that are currently in use. The default alignment rules are to align unsequenced components on natural boundaries.</p> <p>The align nosequence option requests that sequenced components be packed regardless of any other alignment rules. Note that align none implies align nosequence.</p> <p>If you specify an option for standards checking, /align:sequence is ignored.</p>
align all	<p>Tells the compiler to add padding bytes whenever possible to obtain the natural alignment of data items in common blocks, derived types, and record structures. Specifies align nocommons, align dcommons, align records, align nosequence. This is the same as specifying align with no <i>keyword</i>.</p>

Alternate Options

align	Linux and Mac OS X: -noalign
none	Windows: /noalign
align records	Linux and Mac OS X: -align rec16byte, -Zp16 Windows: /align:rec16byte, /Zp16
align norecords	Linux and Mac OS X: -Zp1, -align rec1byte

Windows: /Zp1, /align:rec1byte

align Linux and Mac OS X: -Zp{1|2|4|8|16}

recnbyte Windows: /Zp{1|2|4|8|16}

align all Linux and Mac OS X: -align commons -align
dcommons -align records -align nosequence

Windows:

/align:nocommons,dcommons,records,nosequence

See Also

Optimizing Applications: Setting Data Type and Alignment

allow

Determines whether the compiler allows certain behaviors.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: -allow *keyword*

Windows: /allow:*keyword*

Arguments

keyword Specifies the behaviors to allow or disallow. Possible values are:

[no]fpp_comments Determines how the fpp
preprocessor treats Fortran end-
of-line comments in preprocessor

directive lines.

Default

`fpp_comments` The compiler recognizes Fortran-style end-of-line comments in preprocessor lines.

Description

This option determines whether the compiler allows certain behaviors.

Option	Description
<code>allow nofpp_comments</code>	Tells the compiler to disallow Fortran-style end-of-line comments on preprocessor lines. Comment indicators have no special meaning.

Alternate Options

None

Example

Consider the following:

```
#define MAX_ELEMENTS 100 ! Maximum number of elements
```

By default, the compiler recognizes Fortran-style end-of-line comments on preprocessor lines. Therefore, the line above defines `MAX_ELEMENTS` to be "100" and the rest of the line is ignored. If `allow nofpp_comments` is specified, Fortran comment conventions are not used and the comment indicator "!" has no special meaning. So, in the above example, "! Maximum number of elements" is interpreted as part of the value for the `MAX_ELEMENTS` definition.

Option `allow nofpp_comments` can be useful when you want to have a Fortran directive as a define value; for example:

```
#define dline(routname) !dec$ attributes alias:"__routname":: routname
```

altparam

Allows alternate syntax (without parentheses) for PARAMETER statements.

IDE Equivalent

Windows: **Language > Enable Alternate PARAMETER Syntax**

Linux: None

Mac OS X: **Language > Enable Alternate PARAMETER Syntax**

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-altparam`
`-noaltparam`

Windows: `/altparam`
`/noaltparam`

Arguments

None

Default

`altparam` The alternate syntax for PARAMETER statements is allowed.

Description

This option specifies that the alternate syntax for PARAMETER statements is allowed. The alternate syntax is:

```
PARAMETER c = expr [, c = expr] ...
```

This statement assigns a name to a constant (as does the standard PARAMETER statement), but there are no parentheses surrounding the assignment list.

In this alternative statement, the form of the constant, rather than implicit or explicit typing of the name, determines the data type of the variable.

Alternate Options

`altparam` Linux and Mac OS X: `-dps`
Windows: `/Qdps, /4Yaltparam`

`noaltparam` Linux and Mac OS X: `-nodps`
Windows: `/Qdps-, /4Naltparam`

ansi-alias, Qansi-alias

Tells the compiler to assume that the program adheres to Fortran Standard type aliasability rules.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-ansi-alias`
`-no-ansi-alias`

Windows: `Qansi-alias`
`Qansi-alias-`

Arguments

None

Default

`-ansi-alias` Programs adhere to Fortran Standard type aliasability rules.
or
`/Qansi-`

`alias`

Description

This option tells the compiler to assume that the program adheres to type aliasability rules defined in the Fortran Standard.

For example, an object of type real cannot be accessed as an integer. For information on the rules for data types and data type constants, see "Data Types, Constants, and Variables" in the Language Reference.

This option directs the compiler to assume the following:

- Arrays are not accessed out of arrays' bounds.
- Pointers are not cast to non-pointer types and vice-versa.
- References to objects of two different scalar types cannot alias. For example, an object of type integer cannot alias with an object of type real or an object of type real cannot alias with an object of type double precision.

If your program adheres to the Fortran Standard type aliasability rules, this option enables the compiler to optimize more aggressively. If it doesn't adhere to these rules, then you should disable the option with `-no-ansi-alias` (Linux and Mac OS X) or `/Qansi-alias-` (Windows) so the compiler does not generate incorrect code.

Alternate Options

None

arch

Tells the compiler to generate optimized code specialized for the processor that executes your program.

IDE Equivalent

Windows: **Code Generation > Enable Enhanced Instruction Set**

Linux: None

Mac OS X: None

Architectures

IA-32, Intel® 64 architectures

Syntax

Linux and Mac OS X: `-archprocessor`

Windows: `/arch:processor`

Arguments

processor

Is the processor type. Possible values are:

IA32	Generates code that will run on any Pentium or later processor. Disables any default extended instruction settings, and any previously set extended instruction settings. This value is only available on Linux and Windows systems using IA-32 architecture.
SSE	This is the same as specifying IA32.
SSE2	Generates code for Intel® Streaming SIMD Extensions 2 (Intel® SSE2). This value is only available on Linux and Windows systems.
SSE3	Generates code for Intel® Streaming SIMD Extensions 3 (Intel® SSE3).
SSSE3	Generates code for Intel®

Supplemental Streaming SIMD
Extensions 3 (Intel® SSSE3).

SSE4.1

Generates code for Intel®
Streaming SIMD Extensions 4
Vectorizing Compiler and
Media Accelerators.

Default

Windows For more information on the default values, see Arguments
and Linux above.

systems:

SSE2

Mac OS X

systems

using IA-32

architecture:

SSE3

Mac OS X

systems

using Intel®

64

architecture:

SSSE3

Description

This option tells the compiler to generate optimized code specialized for the processor that executes your program.

Code generated with the values IA32, SSE, SSE2 or SSE3 should execute on any compatible non-Intel processor with support for the corresponding instruction set.

For compatibility with Compaq* Visual Fortran, the compiler allows the following keyword values. However, you should use the suggested replacements.

Compatibility Value	Suggested Replacement
pn1	-mia32 or /arch:IA32
pn2	-mia32 or /arch:IA32
pn3	-mia32 or /arch:IA32
pn4	-msse2 or /arch:SSE2

Alternate Options

Linux and Mac OS X: -m

Windows: /architecture

See Also

[x, Qx](#) compiler option

[ax, Qax](#) compiler option

[m](#) compiler option

architecture

See [arch](#).

asmattr

Specifies the contents of an assembly listing file.

IDE Equivalent

Windows: **Output > Assembler Output**

Linux: None

Mac OS X: None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: None

Windows: `/asmattr:keyword`
 `/noasmattr`

Arguments

keyword Specifies the contents of the assembly listing file. Possible values are:

<code>none</code>	Produces no assembly listing.
<code>machine</code>	Produces an assembly listing with machine code.
<code>source</code>	Produces an assembly listing with source code.
<code>all</code>	Produces an assembly listing with machine code and source code.

Default

`/noasmattr` No assembly listing is produced.

Description

This option specifies what information, in addition to the assembly code, should be generated in the assembly listing file.

To use this option, you must also specify option `/asmfile`, which causes an assembly listing to be generated.

Option	Description
<code>/asmattr:none</code>	Produces no assembly listing. This is the same as

Option	Description
	specifying <code>/noasmattr</code> .
<code>/asmattr:machine</code>	<p>Produces an assembly listing with machine code.</p> <p>The assembly listing file shows the hex machine instructions at the beginning of each line of assembly code. The file cannot be assembled; the filename is the name of the source file with an extension of <code>.cod</code>.</p>
<code>/asmattr:source</code>	<p>Produces an assembly listing with source code.</p> <p>The assembly listing file shows the source code as interspersed comments.</p> <p>Note that if you use alternate option <code>-fsource-asm</code>, you must also specify the <code>-S</code> option.</p>
<code>/asmattr:all</code>	<p>Produces an assembly listing with machine code and source code.</p> <p>The assembly listing file shows the source code as interspersed comments and shows the hex machine instructions at the beginning of each line of assembly code. This file cannot be assembled.</p>

Alternate Options

<code>/asmattr:none</code>	<p>Linux and Mac OS X: None</p> <p>Windows: <code>/noasmattr</code></p>
<code>/asmattr:machine</code>	<p>Linux and Mac OS X: <code>-fcode-asm</code></p> <p>Windows: <code>/FAC</code></p>
<code>/asmattr:source</code>	<p>Linux and Mac OS X: <code>-fsource-asm</code></p> <p>Windows: <code>/FAS</code></p>
<code>/asmattr:all</code>	<p>Linux and Mac OS X: None</p> <p>Windows: <code>/FACs</code></p>

See Also

[asmfile](#) compiler option

asmfile

Specifies that an assembly listing file should be generated.

IDE Equivalent

Windows: **Output > ASM Listing Name**

Linux: None

Mac OS X: None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: None

Windows: /*asmfile[:file | dir]*
 /*noasmfile*

Arguments

file Is the name of the assembly listing file.

dir Is the directory where the file should be placed. It
 can include *file*.

Default

/*noasmfile* No assembly listing file is produced.

Description

This option specifies that an assembly listing file should be generated (optionally named *file*).

If *file* is not specified, the filename will be the name of the source file with an extension of *.asm*; the file is placed in the current directory.

Alternate Options

Linux and Mac OS X: `-s`

Windows: `/Fa`

See Also

[s](#) compiler option

assume

Tells the compiler to make certain assumptions.

IDE Equivalent

Windows: **Compatibility > Treat Backslash as Normal Character in Strings**

`(/assume:[no]bscc)`

Data > Assume Dummy Arguments Share Memory Locations

`(/assume:[no]dummy_aliases)`

Data > Constant Actual Arguments Can Be Changed

`(/assume:[no]protect_constants)`

Data > Use Bytes as RECL=Unit for Unformatted Files

`(/assume:[no]byterecl)`

Floating Point > Enable IEEE Minus Zero Support `(/assume:[no]minus0)`

Optimization > I/O Buffering `(/assume:[no]buffered_io)`

Preprocessor > Default Include and Use Path

`(/assume:[no]source_include)`

Preprocessor > OpenMP Conditional Compilation `(/assume:[no]cc_omp)`

External Procedures > Append Underscore to External Names

`(/assume:[no]underscore)`

Linux: None

Mac OS X: **Optimization > I/O Buffering** `(-assume [no]buffered_io)`

Preprocessor > OpenMP Conditional Compilation (-assume [no]cc_omp)

Preprocessor > Default Include and Use Path (-assume [no]source_include)

Compatibility > Treat Backslash as Normal Character in Strings (-assume [no]bscc)

Data > Assume Dummy Arguments Share Memory Locations (-assume [no]dummy_aliases)

Data > Constant Actual Arguments Can Be Changed (-assume [no]protect_constants)

Data > Use Bytes as RECL=Unit for Unformatted Files (-assume [no]byterecl)

Floating Point > Enable IEEE Minus Zero Support (-assume [no]minus0)

External Procedures > Append Underscore to External Names (-assume [no]underscore)

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: -assume *keyword*

Windows: /assume:*keyword*

Arguments

keyword Specifies the assumptions to be made. Possible values are:

none	Disables all assume options.
[no]bscc	Determines whether the backslash character is treated as a C-style control

	character syntax in character literals.
<code>[no]buffered_io</code>	Determines whether data is immediately written to disk or accumulated in a buffer.
<code>[no]byterecl</code>	Determines whether units for the OPEN statement RECL specifier (record length) value in unformatted files are in bytes or longwords (four-byte units).
<code>[no]cc_omp</code>	Determines whether conditional compilation as defined by the OpenMP Fortran API is enabled or disabled.
<code>[no]dummy_aliases</code>	Determines whether the compiler assumes that dummy arguments to procedures share memory locations with other dummy arguments or with COMMON variables that are assigned.
<code>[no]minus0</code>	Determines whether the compiler uses Fortran 95 or Fortran 90/77 standard

	<p>semantics in the SIGN intrinsic when treating -0.0 and +0.0 as 0.0, and how it writes the value on formatted output.</p>
<code>[no]old_boz</code>	<p>Determines whether the binary, octal, and hexadecimal constant arguments in intrinsic functions INT, REAL, DBLE, and CMPLX are treated as signed integer constants.</p>
<code>[no]old_unit_star</code>	<p>Determines whether READs or WRITEs to UNIT=* go to stdin or stdout, respectively.</p>
<code>[no]old_xor</code>	<p>Determines whether .XOR. is defined by the compiler as an intrinsic operator.</p>
<code>[no]protect_constants</code>	<p>Determines whether a constant actual argument or a copy of it is passed to a called routine.</p>
<code>[no]protect_parens</code>	<p>Determines whether the optimizer honors parentheses in REAL and COMPLEX expression evaluations by not</p>

	reassociating operations.
<code>[no]realloc_lhs</code>	Determines whether allocatable objects on the left-hand side of an assignment are treated according to Fortran 95/90 rules or Fortran 2003 rules.
<code>[no]source_include</code>	Determines whether the compiler searches for USE modules and INCLUDE files in the default directory or in the directory where the source file is located.
<code>[no]std_mod_proc_name</code>	Determines whether the names of module procedures are allowed to conflict with user external symbol names.
<code>[no]underscore</code>	Determines whether the compiler appends an underscore character to external user-defined names.
<code>[no]2underscores</code> (Linux and Mac OS X)	Determines whether the compiler appends two underscore characters to external user-defined names.

`[no]writeable-strings` Determines whether character constants go into non-read-only memory.

Default

<code>nobsc</code>	The backslash character is treated as a normal character in character literals.
<code>nobuffered_io</code>	Data in the internal buffer is immediately written (flushed) to disk (OPEN specifier <code>BUFFERED='NO'</code>). If you set the <code>FORT_BUFFERED</code> environment variable to true, the default is <code>assume buffered_io</code> .
<code>nobyterecl</code>	Units for OPEN statement RECL values with unformatted files are in four-byte (longword) units.
<code>nocc_omp</code>	Conditional compilation as defined by the OpenMP Fortran API is disabled unless option <code>-openmp</code> (Linux) or <code>/Qopenmp</code> (Windows) is specified. If compiler option <code>-openmp</code> (Linux and Mac OS X) or <code>/Qopenmp</code> (Windows) is specified, the default is <code>assume cc_omp</code> .
<code>nodummy_aliases</code>	Dummy arguments to procedures do not share memory locations with other dummy arguments or with variables shared through use association, host association, or common block use.

<code>nominus0</code>	The compiler uses Fortran 90/77 standard semantics in the SIGN intrinsic to treat -0.0 and +0.0 as 0.0, and writes a value of 0.0 with no sign on formatted output.
<code>noold_boz</code>	The binary, octal, and hexadecimal constant arguments in intrinsic functions INT, REAL, DBLE, and CMPLX are treated as bit strings that represent a value of the data type of the intrinsic, that is, the bits are not converted.
<code>old_unit_star</code>	The READs or WRITEs to UNIT=* go to stdin or stdout, respectively, even if UNIT=5 or 6 has been connected to another file.
<code>old_xor</code>	Intrinsic operator .XOR. is defined by the compiler.
<code>protect_constants</code>	A constant actual argument is passed to a called routine. Any attempt to modify it results in an error.
<code>noprotect_parens</code>	The optimizer reorders REAL and COMPLEX expressions without regard for parentheses if it produces faster executing code.
<code>norealloc_lhs</code>	The compiler uses Fortran 95/90 rules when interpreting assignment statements. The left-hand side is assumed to be allocated with the correct shape to hold the right-hand side. If it is not, incorrect behavior will occur.
<code>source_include</code>	The compiler searches for USE modules and INCLUDE files in the directory where the source file is located.

`nostd_mod_proc_name` The compiler allows the names of module procedures to conflict with user external symbol names.

Windows: On Windows systems, the compiler does not append an underscore character to external user-defined names. On Linux and Mac OS X systems, the compiler appends an underscore character to external user-defined names.

`nunderscore`

Linux and Mac OS X:

`underscore`

`no2underscores` The compiler does not append two underscore characters to external user-defined names that contain an embedded underscore.
(Linux and Mac OS X)

`nowriteable-strings` The compiler puts character constants into read-only memory.

Description

This option specifies assumptions to be made by the compiler.

Option	Description
<code>assume none</code>	Disables all the assume options.
<code>assume bsc</code>	Tells the compiler to treat the backslash character (\) as a C-style control (escape) character syntax in character literals. The "bsc" keyword means "BackSlashControlCharacters."
<code>assume buffered_io</code>	Tells the compiler to accumulate records in a buffer. This sets the default for opening sequential output files to BUFFERED='YES', which also occurs if the FORT_BUFFERED run-time environment variable is specified. When this option is specified, the internal buffer is

Option	Description
	<p>filled, possibly by many record output statements (WRITE), before it is written to disk by the Fortran run-time system. If a file is opened for direct access, I/O buffering is ignored.</p> <p>Using buffered writes usually makes disk I/O more efficient by writing larger blocks of data to the disk less often. However, if you request buffered writes, records not yet written to disk may be lost in the event of a system failure.</p> <p>The OPEN statement BUFFERED specifier applies to a specific logical unit. In contrast, the <code>assume [no]buffered_io</code> option and the FORT_BUFFERED environment variable apply to all Fortran units.</p>
<code>assume byterecl</code>	<p>Specifies that the units for the OPEN statement RECL specifier (record length) value are in bytes for unformatted data files, not longwords (four-byte units). For formatted files, the RECL value is always in bytes. If a file is open for unformatted data and <code>assume byterecl</code> is specified, INQUIRE returns RECL in bytes; otherwise, it returns RECL in longwords. An INQUIRE returns RECL in bytes if the unit is not open.</p>
<code>assume cc_omp</code>	<p>Enables conditional compilation as defined by the OpenMP Fortran API. That is, when "!\$space" appears in free-form source or "c\$spaces" appears in column 1 of fixed-form source, the rest of the line is accepted as a Fortran line.</p>
<code>assume dummy_aliases</code>	<p>Tells the compiler that dummy (formal) arguments to procedures share memory locations with other dummy</p>

Option	Description
	<p>arguments (aliases) or with variables shared through use association, host association, or common block use.</p> <p>Specify the option when you compile the called subprogram. The program semantics involved with dummy aliasing do not strictly obey the Fortran 95/90 standards and they slow performance, so you get better run-time performance if you do not use this option.</p> <p>However, if a program depends on dummy aliasing and you do not specify this option, the run-time behavior of the program will be unpredictable. In such programs, the results will depend on the exact optimizations that are performed. In some cases, normal results will occur, but in other cases, results will differ because the values used in computations involving the offending aliases will differ.</p>
assume minus0	Tells the compiler to use Fortran 95 standard semantics for the treatment of the IEEE* floating value -0.0 in the SIGN intrinsic, which distinguishes the difference between -0.0 and +0.0, and to write a value of -0.0 with a negative sign on formatted output.
assume old_boz	Tells the compiler that the binary, octal, and hexadecimal constant arguments in intrinsic functions INT, REAL, DBLE, and CMPLX should be treated as signed integer constants.
assume noold_unit_star	Tells the compiler that READs or WRITEs to UNIT=* go to whatever file UNIT=5 or 6 is connected.

Option	Description
<code>assume noold_xor</code>	Prevents the compiler from defining <code>.XOR.</code> as an intrinsic operator. This lets you use <code>.XOR.</code> as a user-defined operator. This is a Fortran 2003 feature.
<code>assume noprotect_constants</code>	Tells the compiler to pass a copy of a constant actual argument. This copy can be modified by the called routine, even though the Fortran standard prohibits such modification. The calling routine does not see any modification to the constant.
<code>assume protect_parens</code>	<p>Tells the optimizer to honor parentheses in REAL and COMPLEX expression evaluations by not reassociating operations. For example, $(A+B)+C$ would not be evaluated as $A+(B+C)$.</p> <p>If <code>assume noprotect_parens</code> is specified, $(A+B)+C$ would be treated the same as $A+B+C$ and could be evaluated as $A+(B+C)$ if it produced faster executing code.</p> <p>Such reassociation could produce different results depending on the sizes and precision of the arguments. For example, in $(A+B)+C$, if B and C had opposite signs and were very large in magnitude compared to A, $A+B$ could result in the value as B; adding C would result in 0.0. With reassociation, $B+C$ would be 0.0; adding A would result in a non-zero value.</p>
<code>assume realloc_lhs</code>	Tells the compiler that when the left-hand side of an assignment is an allocatable object, it should be reallocated to the shape of the right-hand side of the assignment before the assignment occurs. This is the Fortran 2003 definition. This feature may cause extra

Option	Description
assume	overhead at run time.
nosource_include	Tells the compiler to search the default directory for module files specified by a USE statement or source files specified by an INCLUDE statement. This option affects fpp preprocessor behavior and the USE statement.
assume std_mod_proc_name	Tells the compiler to revise the names of module procedures so they do not conflict with user external symbol names. For example, procedure <code>proc</code> in module <code>m</code> would be named <code>m_MP_proc</code> . The Fortran 2003 Standard requires that module procedure names not conflict with other external symbols. By default, procedure <code>proc</code> in module <code>m</code> would be named <code>m_mp_proc</code> , which could conflict with a user-defined external name <code>m_mp_proc</code> .
assume underscore	Tells the compiler to append an underscore character to external user-defined names: the main program name, named common blocks, BLOCK DATA blocks, global data names in MODULEs, and names implicitly or explicitly declared EXTERNAL. The name of a blank (unnamed) common block remains <code>_BLNK__</code> , and Fortran intrinsic names are not affected.
assume 2underscores (Linux and Mac OS X)	Tells the compiler to append two underscore characters to external user-defined names that contain an embedded underscore: the main program name, named common blocks, BLOCK DATA blocks, global data names in MODULEs, and names implicitly or explicitly declared EXTERNAL. The name of a blank

Option	Description
	<p>(unnamed) common block remains <code>_BLNK_</code>, and Fortran intrinsic names are not affected.</p> <p>This option does not affect external names that do not contain an embedded underscore. By default, the compiler only appends one underscore to those names. For example, if you specify <code>assume 2underscores</code> for external names <code>my_program</code> and <code>myprogram</code>, <code>my_program</code> becomes <code>my_program_</code>, but <code>myprogram</code> becomes <code>myprogram_</code>.</p>
<code>assume writeable-strings</code>	Tells the compiler to put character constants into non-read-only memory.

Alternate Options

<code>assume nobsc</code>	Linux and Mac OS X: <code>-nbs</code> Windows: <code>/nbs</code>
<code>assume dummy_aliases</code>	Linux and Mac OS X: <code>-common-args</code> Windows: <code>/Qcommon-args</code>
<code>assume underscore</code>	Linux and Mac OS X: <code>-us</code> Windows: <code>/us</code>
<code>assume nounderscore</code>	Linux and Mac OS X: <code>-nus</code> Windows: None

auto, QautoSee [automatic](#).**auto-scalar, Qauto-scalar**

Causes scalar variables of intrinsic types INTEGER, REAL, COMPLEX, and LOGICAL that do not have the SAVE attribute to be allocated to the run-time stack.

IDE Equivalent

Windows: **Data > Local Variable Storage** (/Qsave, /Qauto, /Qauto_scalar)

Linux: None

Mac OS X: None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-auto-scalar`

Windows: `/Qauto-scalar`

Arguments

None

Default

`-auto-scalar` Scalar variables of intrinsic types INTEGER, REAL, COMPLEX, and LOGICAL that do not have the SAVE attribute or are allocated to the run-time stack. Note that if option `/Qauto-recursive`, `-openmp` (Linux and Mac OS X), or `/Qopenmp-scalar` (Windows) is specified, the default is `automatic`.

Description

This option causes allocation of scalar variables of intrinsic types INTEGER, REAL, COMPLEX, and LOGICAL to the run-time stack. It is as if they were declared with the AUTOMATIC attribute.

It does not affect variables that have the SAVE attribute (which include initialized locals) or that appear in an EQUIVALENCE statement or in a common block. This option may provide a performance gain for your program, but if your program depends on variables having the same value as the last time the routine was invoked, your program may not function properly. Variables that need to retain their values across subroutine calls should appear in a SAVE statement. You cannot specify option `save`, `auto`, or `automatic` with this option.

**Note**

On Windows NT* systems, there is a performance penalty for addressing a stack frame that is too large. This penalty may be incurred with `/automatic`, `/auto`, or `/Qauto` because arrays are allocated on the stack along with scalars. However, with `/Qauto-scalar`, you would have to have more than 32K bytes of local scalar variables before you incurred the performance penalty. `/Qauto-scalar` enables the compiler to make better choices about which variables should be kept in registers during program execution.

Alternate Options

None

See Also[auto](#) compiler option[save](#) compiler option**autodouble, Qautodouble**See [real-size](#).**automatic**

Causes all local, non-SAVED variables to be allocated to the run-time stack.

IDE Equivalent

Windows: **Data > Local Variable Storage**

Linux: None

Mac OS X: None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-automatic`
`-noautomatic`

Windows: `/automatic`
`/noautomatic`

Arguments

None

Default

`-auto-` Scalar variables of intrinsic types INTEGER, REAL,
`scalar` COMPLEX, and LOGICAL are allocated to the run-time stack.
`or` Note that if one of the following options are specified, the
`/Qauto-` default is `automatic:recursive`, `-openmp` (Linux and Mac
`scalar` OS X), or `/Qopenmp` (Windows).

Description

This option places local variables (scalars and arrays of all types), except those declared as SAVE, on the run-time stack. It is as if the variables were declared with the AUTOMATIC attribute.

It does not affect variables that have the SAVE attribute or ALLOCATABLE attribute, or variables that appear in an EQUIVALENCE statement or in a common block.

This option may provide a performance gain for your program, but if your program depends on variables having the same value as the last time the routine was invoked, your program may not function properly.

If you want to cause variables to be placed in static memory, specify option `-save` (Linux and Mac OS X) or `/Qsave` (Windows). If you want only scalar variables of certain intrinsic types to be placed on the run-time stack, specify option `auto-scalar`.



On Windows NT* systems, there is a performance penalty for addressing a stack frame that is too large. This penalty may be incurred with `/automatic`, `/auto`, or `/Qauto` because arrays are allocated on the stack along with scalars. However, with `/Qauto-scalar`, you would have to have more than 32K bytes of local scalar variables before you incurred the performance penalty. `/Qauto-scalar` enables the compiler to make better choices about which variables should be kept in registers during program execution.

Alternate Options

<code>automatic</code>	Linux and Mac OS X: <code>-auto</code> Windows: <code>/auto</code> , <code>/Qauto</code> , <code>/4Ya</code>
<code>noautomatic</code>	Linux and Mac OS X: <code>-save</code> , <code>-noauto</code> Windows: <code>/Qsave</code> , <code>/noauto</code> , <code>/4Na</code>

See Also

[auto-scalar](#) compiler option

[save, Qsave](#) compiler option

ax, Qax

Tells the compiler to generate multiple, processor-specific auto-dispatch code paths for Intel processors if there is a performance benefit.

IDE Equivalent

Windows: **Code Generation > Add Processor-Optimized Code Path Optimization > Generate Alternate Code Paths**

Linux: None

Mac OS X: **Code Generation > Add Processor-Optimized Code Path**

Architectures

IA-32, Intel® 64 architectures

Syntax

Linux and Mac OS X: *-axprocessor*

Windows: */Qaxprocessor*

Arguments

processor

Indicates the processor for which code is generated. The following descriptions refer to Intel® Streaming SIMD Extensions (Intel® SSE) and Supplemental Streaming SIMD Extensions (Intel® SSSE). Possible values are:

SSE4.2 Can generate Intel® SSE4 Efficient Accelerated String and Text Processing instructions supported by Intel® Core™ i7 processors. Can generate Intel® SSE4 Vectorizing Compiler and Media Accelerator, Intel® SSSE3,

- SSE3, SSE2, and SSE instructions and it can optimize for the Intel® Core™ processor family.
- SSE4 . 1 Can generate Intel® SSE4 Vectorizing Compiler and Media Accelerator instructions for Intel processors. Can generate Intel® SSSE3, SSE3, SSE2, and SSE instructions and it can optimize for Intel® 45nm Hi-k next generation Intel® Core™ microarchitecture. This replaces value S, which is [deprecated](#).
- SSSE3 Can generate Intel® SSSE3, SSE3, SSE2, and SSE instructions for Intel processors and it can optimize for the Intel® Core™2 Duo processor family. This replaces value T, which is [deprecated](#).
- SSE3 Can generate Intel® SSE3, SSE2, and SSE instructions for Intel processors and it can optimize for processors based on Intel® Core™ microarchitecture and Intel

NetBurst® microarchitecture.

This replaces value P, which is [deprecated](#).

SSE2

Can generate Intel® SSE2 and SSE instructions for Intel processors, and it can optimize for Intel® Pentium® 4 processors, Intel® Pentium® M processors, and Intel® Xeon® processors with Intel® SSE2.

This value is not available on Mac OS X systems. This replaces value N, which is [deprecated](#).

Default

OFF No auto-dispatch code is generated. Processor-specific code is generated and is controlled by the setting of compiler option `-m` (Linux), compiler option `/arch` (Windows), or compiler option `-x` (Mac OS* X).

Description

This option tells the compiler to generate multiple, processor-specific auto-dispatch code paths for Intel processors if there is a performance benefit. It also generates a baseline code path. The baseline code is usually slower than the specialized code.

The baseline code path is determined by the architecture specified by the `-x` (Linux and Mac OS X) or `/Qx` (Windows) option. While there are defaults for the `-x` or `/Qx` option that depend on the operating system being used, you can specify an architecture for the baseline code that is higher or lower than the

default. The specified architecture becomes the effective minimum architecture for the baseline code path.

If you specify both the `-ax` and `-x` options (Linux and Mac OS X) or the `/Qax` and `/Qx` options (Windows), the baseline code will only execute on processors compatible with the processor type specified by the `-x` or `/Qx` option.

This option tells the compiler to find opportunities to generate separate versions of functions that take advantage of features of the specified Intel® processor. If the compiler finds such an opportunity, it first checks whether generating a processor-specific version of a function is likely to result in a performance gain. If this is the case, the compiler generates both a processor-specific version of a function and a baseline version of the function. At run time, one of the versions is chosen to execute, depending on the Intel processor in use. In this way, the program can benefit from performance gains on more advanced Intel processors, while still working properly on older processors.

You can use more than one of the processor values by combining them. For example, you can specify `-axSSE4.1,SSSE3` (Linux and Mac OS X) or `/QaxSSE4.1,SSSE3` (Windows). You cannot combine the old style, deprecated options and the new options. For example, you cannot specify `-axSSE4.1,T` (Linux and Mac OS X) or `/QaxSSE4.1,T` (Windows).

Previous values `W` and `K` are deprecated. The details on replacements are as follows:

- Mac OS X systems: On these systems, there is no exact replacement for `W` or `K`. You can upgrade to the default option `-msse3` (IA-32 architecture) or option `-mssse3` (Intel® 64 architecture).
- Windows and Linux systems: The replacement for `W` is `-msse2` (Linux) or `/arch:SSE2` (Windows). There is no exact replacement for `K`. However, on Windows systems, `/QaxK` is interpreted as `/arch:IA32`; on Linux systems, `-axK` is interpreted as `-mia32`. You can also do one of the following:

Intel® Fortran Compiler User and Reference Guides

- Upgrade to option `-msse2` (Linux) or option `/arch:SSE2` (Windows). This will produce one code path that is specialized for Intel® SSE2. It will not run on earlier processors
- Specify the two option combination `-mia32 -axSSE2` (Linux) or `/arch:IA32 /QaxSSE2` (Windows). This combination will produce an executable that runs on any processor with IA-32 architecture but with an additional specialized Intel® SSE2 code path.

The `-ax` and `/Qax` options enable additional optimizations not enabled with option `-m` or option `/arch`.

Alternate Options

None

See Also

[x, Qx](#) compiler option

[m](#) compiler option

[arch](#) compiler option

B

Specifies a directory that can be used to find include files, libraries, and executables.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-Bdir`

Windows: None

Arguments

dir Is the directory to be used. If necessary, the compiler adds a directory separator character at the end of *dir*.

Default

OFF The compiler looks for files in the directories specified in your PATH environment variable.

Description

This option specifies a directory that can be used to find include files, libraries, and executables.

The compiler uses *dir* as a prefix.

For include files, the *dir* is converted to `-I/dir/include`. This command is added to the front of the includes passed to the preprocessor.

For libraries, the *dir* is converted to `-L/dir`. This command is added to the front of the standard `-L` inclusions before system libraries are added.

For executables, if *dir* contains the name of a tool, such as `ld` or `as`, the compiler will use it instead of those found in the default directories.

The compiler looks for include files in *dir* /include while library files are looked for in *dir*.

Another way to get the behavior of this option is to use the environment variable `GCC_EXEC_PREFIX`.

Alternate Options

None

Bdynamic

Enables dynamic linking of libraries at run time.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux: `-Bdynamic`

Mac OS X: `None`

Windows: `None`

Arguments

None

Default

OFF Limited dynamic linking occurs.

Description

This option enables dynamic linking of libraries at run time. Smaller executables are created than with static linking.

This option is placed in the linker command line corresponding to its location on the user command line. It controls the linking behavior of any library that is passed using the command line.

All libraries on the command line following option `-Bdynamic` are linked dynamically until the end of the command line or until a `-Bstatic` option is encountered. The `-Bstatic` option enables static linking of libraries.

Alternate Options

None

See Also

[Bstatic](#) compiler option

bigobj

Increases the number of sections that an object file can contain.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: None

Windows: /bigobj

Arguments

None

Default

OFF An object file can hold up to 65,536 (2^{16}) addressable sections.

Description

This option increases the number of sections that an object file can contain. It increases the address capacity to 4,294,967,296 (2^{32}).

An .obj file produced with this option can only be effectively passed to a linker that shipped in Microsoft Visual C++ 2005 or later. Linkers shipped with earlier versions of the product cannot read .obj files of this size.

This option may be helpful for .obj files that can hold more sections, such as machine generated code.

Alternate Options

None

bintext

Places a text string into the object file (.obj) being generated by the compiler.

IDE Equivalent

Windows: **Code Generation > Object Text String**

Linux: None

Mac OS X: None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: None

Windows: /*bintext:string*
 /*nobintext*

Arguments

string is the text string to go into the object file.

Default

/nobintext No text string is placed in the object file.

Description

This option places a text string into the object file (.obj) being generated by the compiler. The string also gets propagated into the executable file.

For example, this option is useful if you want to place a version number or copyright information into the object and executable.

If the string contains a space or tab, the string must be enclosed by double quotation marks ("). A backslash (\) must precede any double quotation marks contained within the string.

If the command line contains multiple */bintext* options, the last (rightmost) one is used.

Alternate Options

Linux and Mac OS X: None

Windows: */Vstring*

Bstatic

Enables static linking of a user's library.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux: `-Bstatic`

Mac OS X: `None`

Windows: `None`

Arguments

None

Default

OFF Default static linking occurs.

Description

This option enables static linking of a user's library.

This option is placed in the linker command line corresponding to its location on the user command line. It controls the linking behavior of any library that is passed using the command line.

All libraries on the command line following option `-Bstatic` are linked statically until the end of the command line or until a `-Bdynamic` option is encountered.

The `-Bdynamic` option enables dynamic linking of libraries.

Alternate Options

None

See Also

[Bdynamic](#) compiler option

c

Prevents linking.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-c`

Windows: `/c`

Arguments

None

Default

OFF Linking is performed.

Description

This option prevents linking. Compilation stops after the object file is generated. The compiler generates an object file for each Fortran source file.

Alternate Options

Linux and Mac OS X: None

Windows: `/compile-only`, `/nolink`

C

See [check](#).

CB

See [check](#).

ccdefault

Specifies the type of carriage control used when a file is displayed at a terminal screen.

IDE Equivalent

Windows: **Run-time > Default Output Carriage Control**

Linux: None

Mac OS X: **Run-time > Default Output Carriage Control**

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-ccdefault keyword`

Windows: `/ccdefault:keyword`

Arguments

keyword Specifies the carriage-control setting to use. Possible values

are:

<code>none</code>	Tells the compiler to use no carriage control processing.
<code>default</code>	Tells the compiler to use the default carriage-control setting.

<code>fortran</code>	Tells the compiler to use normal Fortran interpretation of the first character. For example, the character 0 causes output of a blank line before a record.
<code>list</code>	Tells the compiler to output one line feed between records.

Default

`ccdefault` The compiler uses the default carriage control setting.
`default`

Description

This option specifies the type of carriage control used when a file is displayed at a terminal screen (units 6 and *). It provides the same functionality as using the CARRIAGECONTROL specifier in an OPEN statement.

The default carriage-control setting can be affected by the `vms` option. If `vms` is specified with `ccdefault default`, carriage control defaults to normal Fortran interpretation (`ccdefault fortran`) if the file is formatted and the unit is connected to a terminal. If `novms` (the default) is specified with `ccdefault default`, carriage control defaults to list (`ccdefault list`).

Alternate Options

None

check

Checks for certain conditions at run time.

IDE Equivalent

Windows:

Run-time > Runtime Error Checking (/nocheck, /check:all, or /xcheck:none)

Run-time > Check Array and String Bounds (/check:[no]bounds)

Run-time > Check Uninitialized Variables (/check:[no]uninit)

Run-time > Check Edit Descriptor Data Type (/check:[no]format)

Run-time > Check Edit Descriptor Data Size

(/check:[no]output_conversion)

Run-time > Check For Actual Arguments Using Temporary Storage

(/check:[no]arg_temp_created)

Run-time > Check For Null Pointers and Allocatable Array References

(/check:[no]pointers)

Linux: None

Mac OS X: Run-time > Runtime Error Checking (-check all, -check none)

Run-time > Check Array and String Bounds (-check [no]bounds)

Run-time > Check Edit Descriptor Data Type (-check [no]format)

Run-time > Check Edit Descriptor Data Size (-check

[no]output_conversion)

Run-time > Check For Actual Arguments Using Temporary Storage (-check

[no]arg_temp_created)

Run-time > Check for Uninitialized Variables (-check [no]uninit)

Run-time > Check For Null Pointers and Allocatable Array References

(/check:[no]pointers)

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: -check [*keyword*]
-nocheck

Windows: /check[:*keyword*]

/nocheck

Arguments

keyword Specifies the conditions to check. Possible values are:

<code>none</code>	Disables all check options.
<code>[no]arg_temp_created</code>	Determines whether checking occurs for actual arguments before routine calls.
<code>[no]bounds</code>	Determines whether checking occurs for array subscript and character substring expressions.
<code>[no]format</code>	Determines whether checking occurs for the data type of an item being formatted for output.
<code>[no]output_conversion</code>	Determines whether checking occurs for the fit of data items within a designated format descriptor field.
<code>[no]pointers</code>	Determines whether checking occurs for certain disassociated or uninitialized pointers or unallocated allocatable objects.
<code>[no]uninit</code>	Determines whether

checking occurs for uninitialized variables.

`all`

Enables all check options.

Default

`nocheck` No checking is performed for run-time failures. Note that if option `vms` is specified, the defaults are `check format` and `check output_conversion`.

Description

This option checks for certain conditions at run time.

Option	Description
<code>check none</code>	Disables all check options (same as <code>nocheck</code>).
<code>check arg_temp_created</code>	Enables run-time checking on whether actual arguments are copied into temporary storage before routine calls. If a copy is made at run-time, an informative message is displayed.
<code>check bounds</code>	<p>Enables compile-time and run-time checking for array subscript and character substring expressions. An error is reported if the expression is outside the dimension of the array or the length of the string.</p> <p>For array bounds, each individual dimension is checked. Array bounds checking is not performed for arrays that are dummy arguments in which the last dimension bound is specified as <code>*</code> or when both upper and lower dimensions are 1.</p> <p>Once the program is debugged, omit this option to reduce executable program size and slightly improve run-time</p>

Option	Description
check format	<p>performance.</p> <p>Issues the run-time FORVARMIS fatal error when the data type of an item being formatted for output does not match the format descriptor being used (for example, a REAL*4 item formatted with an I edit descriptor). With check noformat, the data item is formatted using the specified descriptor unless the length of the item cannot accommodate the descriptor (for example, it is still an error to pass an INTEGER*2 item to an E edit descriptor).</p>
check output_conversion	<p>Issues the run-time OUTCONERR continuable error message when a data item is too large to fit in a designated format descriptor field without loss of significant digits. Format truncation occurs, the field is filled with asterisks (*), and execution continues.</p>
check pointers	<p>Enables run-time checking for disassociated or uninitialized Fortran pointers, unallocated allocatable objects, and integer pointers that are uninitialized.</p>
check uninit	<p>Enables run-time checking for uninitialized variables. If a variable is read before it is written, a run-time error routine will be called. Only local scalar variables of intrinsic type INTEGER, REAL, COMPLEX, and LOGICAL without the SAVE attribute are checked.</p>
check all	<p>Enables all check options. This is the same as specifying check with no keyword.</p>

To get more detailed location information about where an error occurred, use option `traceback`.

Alternate Options

check none	Linux and Mac OS X: -nocheck Windows: /nocheck, /4Nb
check bounds	Linux and Mac OS X: -CB Windows: /CB
check uninit	Linux and Mac OS X: -CU Windows: /RTCu, /CU
check all	Linux and Mac OS X: -check, -C Windows: /check, /4Yb, /C

See Also

[traceback](#) compiler option

cm

See [warn](#).

common-args, Qcommon-args

See [assume](#).

compile-only

See [c](#).

complex-limited-range, Qcomplex-limited-range

Determines whether the use of basic algebraic expansions of some arithmetic operations involving data of type COMPLEX is enabled.

IDE Equivalent

Windows: **Floating point > Limit COMPLEX Range**

Linux: None

Mac OS X: **Floating point > Limit COMPLEX Range**

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-complex-limited-range`
`-no-complex-limited-range`

Windows: `/Qcomplex-limited-range`
`/Qcomplex-limited-range-`

Arguments

None

Default

`-no-complex-limited-range` Basic algebraic expansions of some arithmetic operations involving data of type `COMPLEX` are disabled.
`or/Qcomplex-limited-range-`

Description

This option determines whether the use of basic algebraic expansions of some arithmetic operations involving data of type `COMPLEX` is enabled.

When the option is enabled, this can cause performance improvements in programs that use a lot of `COMPLEX` arithmetic. However, values at the extremes of the exponent range may not compute correctly.

Alternate Options

None

convert

Specifies the format of unformatted files containing numeric data.

IDE Equivalent

Windows: **Compatibility > Unformatted File Conversion**

Linux: None

Mac OS X: **Compatibility > Unformatted File Conversion**

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-convert keyword`

Windows: `/convert:keyword`

Arguments

keyword Specifies the format for the unformatted numeric data.

Possible values are:

<code>native</code>	Specifies that unformatted data should not be converted.
<code>big_endian</code>	Specifies that the format will be big endian for integer data and big endian IEEE floating-point for real and complex data.
<code>cray</code>	Specifies that the format will be big endian for integer data and CRAY* floating-point for real and complex data.

<code>fdx</code> (Linux, Mac OS X)	Specifies that the format will be little endian for integer data, and VAX processor floating-point format F_floating, D_floating, and X_floating for real and complex data.
<code>fgx</code> (Linux, Mac OS X)	Specifies that the format will be little endian for integer data, and VAX processor floating-point format F_floating, G_floating, and X_floating for real and complex data.
<code>ibm</code>	Specifies that the format will be big endian for integer data and IBM* System\370 floating-point format for real and complex data.
<code>little_endian</code>	Specifies that the format will be little endian for integer data and little endian IEEE floating-point for real and complex data.
<code>vaxd</code>	Specifies that the format will be little endian for integer data, and VAX* processor floating-point format F_floating, D_floating, and H_floating for real and complex data.
<code>vaxg</code>	Specifies that the format will be little endian for integer data, and VAX processor floating-point format F_floating, G_floating, and H_floating for real and complex data.

Default

`convert` No conversion is performed on unformatted files containing
`native` numeric data.

Description

This option specifies the format of unformatted files containing numeric data.

Option	Description
<code>convert</code> <code>native</code>	Specifies that unformatted data should not be converted.
<code>convert</code> <code>big_endian</code>	Specifies that the format will be big endian for INTEGER*1, INTEGER*2, INTEGER*4, or INTEGER*8, and big endian IEEE floating-point for REAL*4, REAL*8, REAL*16, COMPLEX*8, COMPLEX*16, or COMPLEX*32.
<code>convert</code> <code>cray</code>	Specifies that the format will be big endian for INTEGER*1, INTEGER*2, INTEGER*4, or INTEGER*8, and CRAY* floating-point for REAL*8 or COMPLEX*16.
<code>convert</code> <code>fdx</code>	Specifies that the format will be little endian for INTEGER*1, INTEGER*2, INTEGER*4, or INTEGER*8, and VAX processor floating-point format F_floating for REAL*4 or COMPLEX*8, D_floating for REAL*8 or COMPLEX*16, and X_floating for REAL*16 or COMPLEX*32.
<code>convert</code> <code>fgx</code>	Specifies that the format will be little endian for INTEGER*1, INTEGER*2, INTEGER*4, or INTEGER*8, and VAX processor floating-point format F_floating for REAL*4 or COMPLEX*8, G_floating for REAL*8 or COMPLEX*16, and X_floating for REAL*16 or COMPLEX*32.
<code>convert</code> <code>ibm</code>	Specifies that the format will be big endian for INTEGER*1, INTEGER*2, or INTEGER*4, and IBM* System\370 floating-point format for REAL*4 or COMPLEX*8 (IBM short 4) and

Option	Description
	REAL*8 or COMPLEX*16 (IBM long 8).
<code>convert</code> <code>little_endian</code>	Specifies that the format will be little endian for INTEGER*1, INTEGER*2, INTEGER*4, or INTEGER*8 and little endian IEEE floating-point for REAL*4, REAL*8, REAL*16, COMPLEX*8, COMPLEX*16, or COMPLEX*32.
<code>convert vaxd</code>	Specifies that the format will be little endian for INTEGER*1, INTEGER*2, INTEGER*4, or INTEGER*8, and VAX processor floating-point format F_floating for REAL*4 or COMPLEX*8, D_floating for REAL*8 or COMPLEX*16, and H_floating for REAL*16 or COMPLEX*32.
<code>convert vaxg</code>	Specifies that the format will be little endian for INTEGER*1, INTEGER*2, INTEGER*4, or INTEGER*8, and VAX processor floating-point format F_floating for REAL*4 or COMPLEX*8, G_floating for REAL*8 or COMPLEX*16, and H_floating for REAL*16 or COMPLEX*32.

Alternate Options

None

cpp, QcppSee [fpp, Qfpp](#).**CU**See [check](#).**cxxlib**

Determines whether the compile links using the C++ run-time libraries provided by gcc.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-cxxlib[=dir]`
`-cxxlib-nostd`
`-no-cxxlib`

Windows: None

Arguments

dir Is an optional top-level location for the gcc binaries and libraries.

Default

`-no-cxxlib` The compiler uses the default run-time libraries and does not link to any additional C++ run-time libraries.

Description

This option determines whether the compiler links using the C++ run-time libraries provided by gcc.

Option `-cxxlib-nostd` prevents the compiler from linking with the standard C++ library. It is only useful for mixed-language applications.

Alternate Options

None

See Also

D

Defines a symbol name that can be associated with an optional value.

IDE Equivalent

Windows: **General > Preprocessor Definitions**

Preprocessor > Preprocessor Definitions

Preprocessor > Preprocessor Definitions to FPP only

Linux: None

Mac OS X: **Preprocessor > Preprocessor Definitions**

Preprocessor > Preprocessor Definitions to FPP only

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-Dname [=value]`

Windows: `/Dname [=value]`

Arguments

name Is the name of the symbol.

value Is an optional integer or an optional character string delimited by double quotes; for example, `Dname=string`.

Default

`noD` Only default symbols or macros are defined.

Description

Defines a symbol name that can be associated with an optional value.

This definition is used during preprocessing.

If a *value* is not specified, *name* is defined as "1".

If you want to specify more than one definition, you must use separate `D` options.

If you specify `noD`, all preprocessor definitions apply only to fpp and not to Intel® Fortran conditional compilation directives. To use this option, you must also specify option `fpp`.



Caution

On Linux and Mac OS X systems, if you are not specifying a *value*, do not use `D` for *name*, because it will conflict with the `-DD` option.

Alternate Options

<code>D</code>	Linux and Mac OS X: None Windows: <code>/define:name[=value]</code>
<code>noD</code>	Linux and Mac OS X: <code>-nodefine</code> Windows: <code>/nodefine</code>

See Also

Building Applications: Predefined Preprocessor Symbols

d-lines, Qd-lines

Compiles debug statements.

IDE Equivalent

Windows: **Language > Compile Lines With D in Column 1**

Linux: None

Mac OS X: **Language > Compile Lines With D in Column 1**

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-d-lines`
`-nod-lines`

Windows: `/d-lines`
`/nod-lines`
`/Qd-lines`

Arguments

None

Default

`nod-` Debug lines are treated as comment lines.
`lines`

Description

This option compiles debug statements. It specifies that lines in fixed-format files that contain a D in column 1 (debug statements) should be treated as source code.

Alternate Options

Linux and Mac OS X: `-DD`

Windows: None

dbglibs

Tells the linker to search for unresolved references in a debug run-time library.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: None

Windows: /dbglibs
 /nodbglibs**Arguments**

None

Default`/nodbglibs` The linker does not search for unresolved references in a debug run-time library.**Description**

This option tells the linker to search for unresolved references in a debug run-time library.

The following table shows which options to specify for a debug run-time library:

Type of Library	Options Required	Alternate Option
Debug single-threaded	<code>/libs:static</code> <code>/dbglibs</code>	<code>/MLd</code> (this is a deprecated option)
Debug multithreaded	<code>/libs:static</code> <code>/threads</code> <code>/dbglibs</code>	<code>/MTd</code>
Multithreaded debug DLLs	<code>/libs:dll</code> <code>/threads</code> <code>/dbglibs</code>	<code>/MDd</code>
Debug Fortran QuickWin multi-thread applications	<code>/libs:qwin</code> <code>/dbglibs</code>	None

Type of Library	Options Required	Alternate Option
Debug Fortran standard graphics (QuickWin single-thread) applications	<code>/libs:qwins</code> <code>/dbglibs</code>	None

Alternate Options

None

See Also

Building Applications: Specifying Consistent Library Types; Programming with Mixed Languages Overview

DD

See [dlines](#), [Qdlines](#).

debug (Linux* OS and Mac OS* X)

Enables or disables generation of debugging information.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-debug[keyword]`

Windows: None

Arguments

keyword Is the type of debugging information to be generated. Possible values are:

Intel(R) Fortran Compiler User and Reference Guides

<code>none</code>	Disables generation of debugging information.
<code>full or all</code>	Generates complete debugging information.
<code>minimal</code>	Generates line number information for debugging.
<code>[no]emit_column</code>	Determines whether the compiler generates column number information for debugging.
<code>[no]inline- debug-info</code>	Determines whether the compiler generates enhanced debug information for inlined code.
<code>[no]semantic- stepping</code>	Determines whether the compiler generates enhanced debug information useful for breakpoints and stepping.
<code>[no]variable- locations</code>	Determines whether the compiler generates enhanced debug information useful in finding scalar local variables.
<code>extended</code>	Sets keyword values <code>semantic-stepping</code>

and `variable-`
`locations`.

For information on the non-default settings for these keywords, see the Description section.

Default

`-debug none` No debugging information is generated.

Description

This option enables or disables generation of debugging information.

Note that if you turn debugging on, optimization is turned off.

Keywords `semantic-stepping`, `inline-debug-info`, `variable-locations`, and `extended` can be used in combination with each other. If conflicting keywords are used in combination, the last one specified on the command line has precedence.

Option	Description
<code>-debug none</code>	Disables generation of debugging information.
<code>-debug full</code> or <code>-debug all</code>	Generates complete debugging information. It is the same as specifying <code>-debug</code> with no keyword.
<code>-debug minimal</code>	Generates line number information for debugging.
<code>-debug emit_column</code>	Generates column number information for debugging.
<code>-debug inline-debug-info</code>	Generates enhanced debug information for inlined code. It provides more information to debuggers for function call traceback.
<code>-debug semantic-stepping</code>	Generates enhanced debug information useful for breakpoints and stepping. It tells the debugger to stop only at machine instructions that achieve the final effect

Option	Description
	<p>of a source statement.</p> <p>For example, in the case of an assignment statement, this might be a store instruction that assigns a value to a program variable; for a function call, it might be the machine instruction that executes the call. Other instructions generated for those source statements are not displayed during stepping.</p> <p>This option has no impact unless optimizations have also been enabled.</p>
-debug variable-locations	<p>Generates enhanced debug information useful in finding scalar local variables. It uses a feature of the Dwarf object module known as "location lists".</p> <p>This feature allows the run-time locations of local scalar variables to be specified more accurately; that is, whether, at a given position in the code, a variable value is found in memory or a machine register.</p>
-debug extended	<p>Sets keyword values <code>semantic-stepping</code> and <code>variable-locations</code>. It also tells the compiler to include column numbers in the line information.</p>

On Linux* systems, debuggers read debug information from executable images. As a result, information is written to object files and then added to the executable by the linker. On Mac OS* X systems, debuggers read debug information from object files. As a result, the executables don't contain any debug information. Therefore, if you want to be able to debug on these systems, you must retain the object files.

Alternate Options

For `debug full`, - Linux and Mac OS X: `-g`

Intel® Fortran Compiler User and Reference Guides

debug all, or - debug	Windows: /debug:full, /debug:all, or /debug
For -debug inline-debug- info	Linux and Mac OS X: -inline-debug-info (this is a deprecated option) Windows: None

See Also

[debug \(Windows*\)](#) compiler option

Building Applications: Debugging Overview

debug (Windows* OS)

Enables or disables generation of debugging information.

IDE Equivalent

Windows: **General > Debug Information Format** (/debug:minimal, /debug:full)

Linux: None

Mac OS X: None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: None

Windows: /debug[:*keyword*]
 /nodebug

Arguments

keyword Is the type of debugging information to be generated. Possible values are:

none	Generates no symbol table
------	---------------------------

	information.
<code>full</code> or <code>all</code>	Generates complete debugging information.
<code>minimal</code>	Generates line numbers and minimal debugging information.
<code>partial</code>	Deprecated . Generates global symbol table information needed for linking.

For information on the non-default settings for these keywords, see the Description section.

Default

<code>/debug:none</code>	This is the default on the command line and for a release configuration in the IDE.
<code>/debug:full</code>	This is the default for a debug configuration in the IDE.

Description

This option enables or disables generation of debugging information. It is passed to the linker.

Note that if you turn debugging on, optimization is turned off.

If conflicting keywords are used in combination, the last one specified on the command line has precedence.

Option	Description
<code>/debug:none</code>	Disables generation of debugging information. It is the same as specifying <code>/nodebug</code> .
<code>/debug:full</code> or <code>/debug:all</code>	Generates complete debugging

Option	Description
/debug:minimal	<p>information. It produces symbol table information needed for full symbolic debugging of unoptimized code and global symbol information needed for linking. It is the same as specifying /debug with no keyword. If you specify /debug:full for an application that makes calls to C library routines and you need to debug calls into the C library, you should also specify /dbglibs to request that the appropriate C debug library be linked against.</p> <p>Generates line number information for debugging. It produces global symbol information needed for linking, but not local symbol table information needed for debugging.</p>
/debug:partial	<p>Generates global symbol table information needed for linking, but not local symbol table information needed for debugging. This option is deprecated and is not available in the IDE.</p>

Alternate Options

For /debug:minimal Linux and Mac OS X: None
 Windows: /zd (this is a [deprecated](#) option)

For /debug:full Linux and Mac OS X: None
 Windows: /zi, /z7

or

/debug

See Also

[dbglibs](#) compiler option

[debug \(Linux* and Mac OS* X\)](#) compiler option

Building Applications: Debugging Fortran Programs

debug-parameters

Tells the compiler to generate debug information for PARAMETERS used in a program.

IDE Equivalent

Windows: **Debugging > Information for PARAMETER Constants**

Linux: None

Mac OS X: **Debug > Information for PARAMETER Constants**

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-debug-parameters [keyword]`

`-nodebug-parameters`

Windows: `/debug-parameters[:keyword]`

`/nodebug-parameters`

Arguments

keyword Are the PARAMETERS to generate debug information for.

Possible values are:

`none` Generates no debug information for

	any PARAMETERS used in the program. This is the same as specifying <code>nodebug-parameters</code> .
<code>used</code>	Generates debug information for only PARAMETERS that have actually been referenced in the program. This is the default if you do not specify a <i>keyword</i> .
<code>all</code>	Generates debug information for all PARAMETERS defined in the program.

Default

`nodebug-parameters` The compiler generates no debug information for any PARAMETERS used in the program. This is the same as specifying *keyword* `none`.

Description

This option tells the compiler to generate debug information for PARAMETERS used in a program.

Note that if a `.mod` file contains PARAMETERS, debug information is only generated for the PARAMETERS that have actually been referenced in the program, even if you specify *keyword* `all`.

Alternate Options

None

define

See [D](#).

diag, Qdiag

Controls the display of diagnostic information.

IDE Equivalent

Windows: **Diagnostics > Disable Specific Diagnostics** (/Qdiag-disable id)

Diagnostics > Level of Static Analysis (/Qdiag-enable[:sv1,sv2, sv3])

Linux: None

Mac OS X: **Diagnostics > Disable Specific Diagnostics** (-diag-disable id)

Diagnostics > Level of Static Analysis (-diag-enable [sv1,sv2, sv3])

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-diag-type diag-list`

Windows: `/Qdiag-type:diag-list`

Arguments

<i>type</i>	Is an action to perform on diagnostics. Possible values are:
<code>enable</code>	Enables a diagnostic message or a group of messages.
<code>disable</code>	Disables a diagnostic message or a group of messages.
<code>error</code>	Tells the compiler to change diagnostics to errors.
<code>warning</code>	Tells the compiler to change diagnostics to warnings.

	<code>remark</code>	Tells the compiler to change diagnostics to remarks (comments).
<i>diag-list</i>		Is a diagnostic group or ID value. Possible values are:
	<code>driver</code>	Specifies diagnostic messages issued by the compiler driver.
	<code>vec</code>	Specifies diagnostic messages issued by the vectorizer.
	<code>par</code>	Specifies diagnostic messages issued by the auto-parallelizer (parallel optimizer).
	<code>sv[n]</code>	Specifies diagnostic messages issued by the Static Verifier. <i>n</i> can be any of the following: 1, 2, 3. For more details on these values, see below.
	<code>warn</code>	Specifies diagnostic messages that have a "warning" severity level.
	<code>error</code>	Specifies diagnostic messages that have an "error" severity level.
	<code>remark</code>	Specifies diagnostic

- messages that are remarks or comments.
- `cpu-dispatch` Specifies the CPU dispatch remarks for diagnostic messages. These remarks are enabled by default. This diagnostic group is only available on IA-32 architecture and Intel® 64 architecture.
- `id[,id, ...]` Specifies the ID number of one or more messages. If you specify more than one message number, they must be separated by commas. There can be no intervening white space between each id.
- `tag[,tag, ...]` Specifies the mnemonic name of one or more messages. If you specify more than one mnemonic name, they must be separated by commas. There can be no intervening white space between each tag.

Default

OFF The compiler issues certain diagnostic messages by default.

Description

This option controls the display of diagnostic information. Diagnostic messages are output to `stderr` unless compiler option `-diag-file` (Linux and Mac OS X) or `/Qdiag-file` (Windows) is specified.

When *diag-list* value "warn" is used with the Static Verifier (sv) diagnostics, the following behavior occurs:

- Option `-diag-enable warn` (Linux and Mac OS X) and `/Qdiag-enable:warn` (Windows) enable all Static Verifier diagnostics except those that have an "error" severity level. They enable all Static Verifier warnings, cautions, and remarks.
- Option `-diag-disable warn` (Linux and Mac OS X) and `/Qdiag-disable:warn` (Windows) disable all Static Verifier diagnostics except those that have an "error" severity level. They suppress all Static Verifier warnings, cautions, and remarks.

The following table shows more information on values you can specify for *diag-list* item *sv*.

<i>diag-list</i> Item	Description						
<i>sv</i> [<i>n</i>]	The value of <i>n</i> for Static Verifier messages can be any of the following: <table border="0"> <tr> <td style="padding-left: 20px;">1</td> <td>Produces the diagnostics with severity level set to all critical errors.</td> </tr> <tr> <td style="padding-left: 20px;">2</td> <td>Produces the diagnostics with severity level set to all errors. This is the default if <i>n</i> is not specified.</td> </tr> <tr> <td style="padding-left: 20px;">3</td> <td>Produces the diagnostics with severity level set to all errors and warnings.</td> </tr> </table>	1	Produces the diagnostics with severity level set to all critical errors.	2	Produces the diagnostics with severity level set to all errors. This is the default if <i>n</i> is not specified.	3	Produces the diagnostics with severity level set to all errors and warnings.
1	Produces the diagnostics with severity level set to all critical errors.						
2	Produces the diagnostics with severity level set to all errors. This is the default if <i>n</i> is not specified.						
3	Produces the diagnostics with severity level set to all errors and warnings.						

To control the diagnostic information reported by the vectorizer, use the `-vec-report` (Linux and Mac OS X) or `/Qvec-report` (Windows) option.

To control the diagnostic information reported by the auto-parallelizer, use the `-par-report` (Linux and Mac OS X) or `/Qpar-report` (Windows) option.

Alternate Options

enable vec	Linux and Mac OS X: <code>-vec-report</code> Windows: <code>/Qvec-report</code>
disable vec	Linux and Mac OS X: <code>-vec-report0</code> Windows: <code>/Qvec-report0</code>
enable par	Linux and Mac OS X: <code>-par-report</code> Windows: <code>/Qpar-report</code>
disable par	Linux and Mac OS X: <code>-par-report0</code> Windows: <code>/Qpar-report0</code>

Example

The following example shows how to enable diagnostic IDs 117, 230 and 450:

```
-diag-enable 117,230,450      ! Linux and Mac OS X systems
/Qdiag-enable:117,230,450    ! Windows systems
```

The following example shows how to change vectorizer diagnostic messages to warnings:

```
-diag-enable vec -diag-warning vec      ! Linux and Mac OS X systems
/Qdiag-enable:vec /Qdiag-warning:vec    ! Windows systems
```

Note that you need to enable the vectorizer diagnostics before you can change them to warnings.

The following example shows how to disable all auto-parallelizer diagnostic messages:

```
-diag-disable par          ! Linux and Mac OS X systems
/Qdiag-disable:par        ! Windows systems
```

The following example shows how to produce Static Verifier diagnostic messages for all critical errors:

```
-diag-enable sv1          ! Linux and Mac OS X systems
/Qdiag-enable:sv1        ! Windows system
```

The following example shows how to cause Static Verifier diagnostics (and default diagnostics) to be sent to a file:

```
-diag-enable sv -diag-file=stat_ver_msg ! Linux and Mac OS X systems
/Qdiag-enable:sv /Qdiag-file:stat_ver_msg ! Windows systems
```

Note that you need to enable the Static Verifier diagnostics before you can send them to a file. In this case, the diagnostics are sent to file `stat_ver_msg.diag`. If a

file name is not specified, the diagnostics are sent to name-of-the-first-source-file.diag.

The following example shows how to change all diagnostic warnings and remarks to errors:

```
-diag-error warn,remark      ! Linux and Mac OS X systems  
/Qdiag-error:warn,remark    ! Windows systems
```

See Also

[diag-dump, Qdiag-dump](#) compiler option

[diag-id-numbers, Qdiag-id-numbers](#) compiler option

[diag-file, Qdiag-file](#) compiler option

[par-report, Qpar-report](#) compiler option

[vec-report, Qvec-report](#) compiler option

diag-dump, Qdiag-dump

Tells the compiler to print all enabled diagnostic messages and stop compilation.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-diag-dump`

Windows: `/Qdiag-dump`

Arguments

None

Default

OFF The compiler issues certain diagnostic messages by default.

Description

This option tells the compiler to print all enabled diagnostic messages and stop compilation. The diagnostic messages are output to `stdout`.

This option prints the enabled diagnostics from all possible diagnostics that the compiler can issue, including any default diagnostics.

If `-diag-enable diag-list` (Linux and Mac OS X) or `/Qdiag-enable diag-list` (Windows) is specified, the print out will include the *diag-list* diagnostics.

Alternate Options

None

Example

The following example adds vectorizer diagnostic messages to the printout of default diagnostics:

```
-diag-enable vec -diag-dump           ! Linux and Mac OS X systems
/Qdiag-enable:vec /Qdiag-dump       ! Windows systems
```

See Also

[diag, Qdiag](#) compiler option

diag-enable sv-include, Qdiag-enable sv-include

Tells the Static Verifier to analyze include files and source files when issuing diagnostic messages.

IDE Equivalent

Windows: **Diagnostics > Analyze Include Files**

Linux: None

Mac OS X: **Diagnostics > Analyze Include Files**

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-diag-enable sv-include`

Windows: `/Qdiag-enable sv-include`

Arguments

None

Default

OFF The compiler issues certain diagnostic messages by default. If the Static Verifier is enabled, include files are not analyzed by default.

Description

This option tells the Static Verifier to analyze include files and source files when issuing diagnostic messages. Normally, when Static Verifier diagnostics are enabled, only source files are analyzed.

To use this option, you must also specify `-diag-enable sv` (Linux and Mac OS X) or `/Qdiag-enable:sv` (Windows) to enable the Static Verifier diagnostics.

Alternate Options

None

Example

The following example shows how to cause include files to be analyzed as well as source files:

```
-diag-enable sv -diag-enable sv-include      ! Linux and Mac OS systems  
/Qdiag-enable:sv /Qdiag-enable:sv-include  ! Windows systems
```

In the above example, the first compiler option enables Static Verifier messages.

The second compiler option causes include files referred to by the source file to be analyzed also.

See Also

[diag, Qdiag](#) compiler option

diag-error-limit, Qdiag-error-limit

Specifies the maximum number of errors allowed before compilation stops.

IDE Equivalent

Windows: **Compilation Diagnostics > Error Limit**

Linux: None

Mac OS X: **Compiler Diagnostics > Error Limit**

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-diag-error-limitn`
`-no-diag-error-limit`

Windows: `/Qdiag-error-limit:n`
`/Qdiag-error-limit-`

Arguments

n Is the maximum number of error-level or fatal-level compiler errors allowed.

Default

30 A maximum of 30 error-level and fatal-level messages are allowed.

Description

This option specifies the maximum number of errors allowed before compilation stops. It indicates the maximum number of error-level or fatal-level compiler errors allowed for a file specified on the command line.

If you specify `-no-diag-error-limit` (Linux and Mac OS X) or `/Qdiag-error-limit-` (Windows) on the command line, there is no limit on the number of errors that are allowed.

If the maximum number of errors is reached, a warning message is issued and the next file (if any) on the command line is compiled.

Alternate Options

Linux and Mac OS X: `-error-limit` and `-noerror-limit`

Windows: `/error-limit` and `/noerror-limit`

diag-file, Qdiag-file

Causes the results of diagnostic analysis to be output to a file.

IDE Equivalent

Windows: **Diagnostics > Diagnostics File**

Linux: None

Mac OS X: **Diagnostics > Diagnostics File**

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-diag-file[=file]`

Windows: `/Qdiag-file[:file]`

Arguments

file Is the name of the file for output.

Default

OFF Diagnostic messages are output to stderr.

Description

This option causes the results of diagnostic analysis to be output to a file. The file is placed in the current working directory.

If *file* is specified, the name of the file is *file.diag*. The file can include a file extension; for example, if *file.ext* is specified, the name of the file is *file.ext*.

If *file* is not specified, the name of the file is *name-of-the-first-source-file.diag*. This is also the name of the file if the name specified for file conflicts with a source file name provided in the command line.



If you specify `-diag-file` (Linux and Mac OS X) or `/Qdiag-file` (Windows) and you also specify `-diag-file-append` (Linux and Mac OS X) or `/Qdiag-file-append` (Windows), the last option specified on the command line takes precedence.

Alternate Options

None

Example

The following example shows how to cause diagnostic analysis to be output to a file named *my_diagnostics.diag*:

```
-diag-file=my_diagnostics      ! Linux and Mac OS X systems
/Qdiag-file:my_diagnostics    ! Windows systems
```

See Also

[diag, Qdiag](#) compiler option

[diag-file-append, Qdiag-file-append](#) compiler option

diag-file-append, Qdiag-file-append

Causes the results of diagnostic analysis to be appended to a file.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-diag-file-append[=file]`

Windows: `/Qdiag-file-append[:file]`

Arguments

file Is the name of the file to be appended to. It can include a path.

Default

OFF Diagnostic messages are output to `stderr`.

Description

This option causes the results of diagnostic analysis to be appended to a file. If you do not specify a path, the driver will look for *file* in the current working directory.

If *file* is not found, then a new file with that name is created in the current working directory. If the name specified for file conflicts with a source file name provided in the command line, the name of the file is *name-of-the-first-source-file.diag*.



Note

If you specify `-diag-file-append` (Linux and Mac OS X) or `/Qdiag-file-append` (Windows) and you also specify `-diag-file` (Linux and Mac OS X) or `/Qdiag-file` (Windows), the last option specified on the command line takes precedence.

Alternate Options

None

Example

The following example shows how to cause diagnostic analysis to be appended to a file named `my_diagnostics.txt`:

```
-diag-file-append=my_diagnostics.txt      ! Linux and Mac OS X systems
/Qdiag-file-append:my_diagnostics.txt    ! Windows systems
```

See Also

[diag, Qdiag](#) compiler option

[diag-file, Qdiag-file](#) compiler option

diag-id-numbers, Qdiag-id-numbers

Determines whether the compiler displays diagnostic messages by using their ID number values.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-diag-id-numbers`
`-no-diag-id-numbers`

Windows: `/Qdiag-id-numbers`
`/Qdiag-id-numbers-`

Arguments

None

Default

`-diag-id-numbers` The compiler displays diagnostic messages by using their ID number values.

or /Qdiag-id-
numbers

Description

This option determines whether the compiler displays diagnostic messages by using their ID number values. If you specify `-no-diag-id-numbers` (Linux and Mac OS X) or `/Qdiag-id-numbers-` (Windows), mnemonic names are output for driver diagnostics only.

Alternate Options

None

See Also

[diag, Qdiag](#) compiler option

diag-once, Qdiag-once

Tells the compiler to issue one or more diagnostic messages only once.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-diag-onceid[,id, ...]`

Windows: `/Qdiag-once:id[,id, ...]`

Arguments

id Is the ID number of the diagnostic message. If you specify more than one message number, they must be separated by commas.

There can be no intervening white space between each *id*.

Default

OFF The compiler issues certain diagnostic messages by default.

Description

This option tells the compiler to issue one or more diagnostic messages only once.

Alternate Options

None

dll

Specifies that a program should be linked as a dynamic-link (DLL) library.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: None

Windows: /dll

Arguments

None

Default

OFF The program is not linked as a dynamic-link (DLL) library.

Description

This option specifies that a program should be linked as a dynamic-link (DLL) library instead of an executable (.exe) file. It overrides any previous specification of run-time routines to be used and enables the `/libs:dll` option.

If you use this option with the `/libs:qwin` or `/libs:qwins` option, the compiler issues a warning.

Alternate Options

Linux and Mac OS X: None

Windows: `/LD`

double-size

Specifies the default KIND for DOUBLE PRECISION and DOUBLE COMPLEX variables.

IDE Equivalent

Windows: **Data > Default Double Precision KIND**

Linux: None

Mac OS X: **Data > Default Double Precision KIND**

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-double-size size`

Windows: `/double-size:size`

Arguments

size Specifies the default KIND for DOUBLE PRECISION and DOUBLE COMPLEX declarations, constants, functions, and intrinsics. Possible values are: 64 (KIND=8) or 128 (KIND=16).

Default

64 DOUBLE PRECISION variables are defined as REAL*8 and
 DOUBLE COMPLEX variables are defined as COMPLEX*16.

Description

This option defines the default KIND for DOUBLE PRECISION and DOUBLE COMPLEX declarations, constants, functions, and intrinsics.

Option	Description
double-size 64	Defines DOUBLE PRECISION declarations, constants, functions, and intrinsics as REAL(KIND=8) (REAL*8) and defines DOUBLE COMPLEX declarations, functions, and intrinsics as COMPLEX(KIND=8) (COMPLEX*16).
double-size 128	Defines DOUBLE PRECISION declarations, constants, functions, and intrinsics as REAL(KIND=16) (REAL*16) and defines DOUBLE COMPLEX declarations, functions, and intrinsics as COMPLEX(KIND=16) (COMPLEX*32).

Alternate Options

None

dps, Qdps

See [altparam](#).

dryrun

Specifies that driver tool commands should be shown but not executed.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-dryrun`

Windows: None

Arguments

None

Default

OFF No tool commands are shown, but they are executed.

Description

This option specifies that driver tool commands should be shown but not executed.

Alternate Options

None

See Also

[v](#) compiler option

dumpmachine

Displays the target machine and operating system configuration.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-dumpmachine`

Windows: None

Arguments

None

Default

OFF The compiler does not display target machine or operating system information.

Description

This option displays the target machine and operating system configuration. No compilation is performed.

Alternate Options

None

dynamic-linker

Specifies a dynamic linker other than the default.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux: `-dynamic-linker file`

Mac OS X: None

Windows: None

Arguments

file Is the name of the dynamic linker to be used.

Default

OFF The default dynamic linker is used.

Description

This option lets you specify a dynamic linker other than the default.

Alternate Options

None

dynamiclib

Invokes the `libtool` command to generate dynamic libraries.

IDE Equivalent

None

Architectures

IA-32, Intel® 64 architectures

Syntax

Linux: None

Mac OS X: `-dynamiclib`

Windows: None

Arguments

None

Default

OFF The compiler produces an executable.

Description

This option invokes the `libtool` command to generate dynamic libraries. When passed this option, the compiler uses the `libtool` command to produce a dynamic library instead of an executable when linking. To build static libraries, you should specify option `-staticlib` or `libtool -static <objects>`.

Alternate Options

None

See Also

[staticlib](#) compiler option

dyncom, Qdyncom

Enables dynamic allocation of common blocks at run time.

IDE Equivalent

Windows: **Data > Dynamic Common Blocks**

Linux: None

Mac OS X: **Data > Dynamic Common Blocks**

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-dyncom "common1,common2,..."`

Windows: `/Qdyncom "common1,common2,..."`

Arguments

`common1,common2,...` Are the names of the common blocks to be dynamically allocated. The list of names must be within quotes.

Default

OFF Common blocks are not dynamically allocated at run time.

Description

This option enables dynamic allocation of the specified common blocks at run time. For example, to enable dynamic allocation of common blocks a, b, and c at run time, use this syntax:

```
/Qdyncom "a,b,c"      ! on Windows systems  
-dyncom "a,b,c"      ! on Linux and Mac OS X systems
```

The following are some limitations that you should be aware of when using this option:

- An entity in a dynamic common cannot be initialized in a DATA statement.
- Only named common blocks can be designated as dynamic COMMON.
- An entity in a dynamic common block must not be used in an EQUIVALENCE expression with an entity in a static common block or a DATA-initialized variable.

Alternate Options

None

See Also

Building Applications: Allocating Common Blocks

E

Causes the preprocessor to send output to `stdout`.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: -E

Windows: /E

Arguments

None

Default

OFF Preprocessed source files are output to the compiler.

Description

This option causes the preprocessor to send output to `stdout`. Compilation stops when the files have been preprocessed.

When you specify this option, the compiler's preprocessor expands your source module and writes the result to `stdout`. The preprocessed source contains `#line` directives, which the compiler uses to determine the source file and line number.

Alternate Options

None

e90, e95, e03

See [warn](#).

EP

Causes the preprocessor to send output to `stdout`, omitting `#line` directives.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-EP`

Windows: `/EP`

Arguments

None

Default

OFF Preprocessed source files are output to the compiler.

Description

This option causes the preprocessor to send output to `stdout`, omitting `#line` directives.

If you also specify option `preprocess-only`, option `P`, or option `F`, the preprocessor will write the results (without `#line` directives) to a file instead of `stdout`.

Alternate Options

None

error-limit

See [diag-error-limit](#), [Qdiag-error-limit](#).

exe

Specifies the name for a built program or dynamic-link library.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: None

Windows: `/exe:{file | dir}`

Arguments

file Is the name for the built program or dynamic-link library.

dir Is the directory where the built program or dynamic-link library should be placed. It can include *file*.

Default

OFF The name of the file is the name of the first source file on the command line with file extension `.exe`, so `file.f` becomes `file.exe`.

Description

This option specifies the name for a built program (`.EXE`) or a dynamic-link library (`.DLL`).

You can use this option to specify an alternate name for an executable file. This is especially useful when compiling and linking a set of input files. You can use the option to give the resulting file a name other than that of the first input file (source or object) on the command line.

You can use this option to specify an alternate name for an executable file. This is especially useful when compiling and linking a set of input files. You can use the option to give the resulting file a name other than that of the first input file (source or object) on the command line.

Alternate Options

Linux and Mac OS X: `-o`

Windows: `/Fe`

Example

The following example creates a dynamic-link library file named file.dll (note that you can use /LD in place of /dll):

```
ifort /dll /exe:file.dll a.f
```

In the following example (which uses the alternate option /Fe), the command produces an executable file named outfile.exe as a result of compiling and linking three files: one object file and two Fortran source files.

```
prompt>ifort /Feoutfile.exe file1.obj file2.for file3.for
```

By default, this command produces an executable file named file1.exe.

See Also

[o](#) compiler option

extend-source

Specifies the length of the statement field in a fixed-form source file.

IDE Equivalent

Windows: **Language > Fixed Form Line Length**

Linux: None

Mac OS X: **Language > Fixed Form Line Length**

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-extend-source [size]`

`-noextend-source`

Windows: `/extend-source[:size]`

`/noextend-source`

Arguments

size Is the length of the statement field in a fixed-form source file.

Possible values are: 72, 80, or 132.

Default

- 72 If you do not specify this option or you specify `noextend-source`, the statement field ends at column 72.
- 132 If you specify `extend_source` without `size`, the statement field ends at column 132.-

Description

This option specifies the size (column number) of the statement field of a source line in a fixed-form source file. This option is valid only for fixed-form files; it is ignored for free-form files.

When `size` is specified, it is the last column parsed as part of the statement field. Any columns after that are treated as comments.

If you do not specify `size`, it is the same as specifying `extend_source 132`.

Option	Description
<code>extend-source 72</code>	Specifies that the statement field ends at column 72.
<code>extend-source 80</code>	Specifies that the statement field ends at column 80.
<code>extend-source 132</code>	Specifies that the statement field ends at column 132.

Alternate Options

- `extend-source 72` Linux and Mac OS X: `-72`
Windows: `/4L72`
- `extend-source 80` Linux and Mac OS X: `-80`
Windows: `/4L80`
- `extend-source` Linux and Mac OS X: `-132`

extfor

Specifies file extensions to be processed by the compiler as Fortran files.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: None

Windows: /extfor:ext

Arguments

ext Are the file extensions to be processed as a Fortran file.

Default

OFF Only the file extensions recognized by the compiler are processed as Fortran files. For more information, see *Building Applications*.

Description

This option specifies file extensions (*ext*) to be processed by the compiler as Fortran files. It is useful if your source file has a nonstandard extension. You can specify one or more file extensions. A leading period before each extension is optional; for example, /extfor:myf95 and /extfor:.myf95 are equivalent.

Alternate Options

None

See Also

[source compiler option](#) Tells the compiler to compile the file as a Fortran source file.

extfpp

Specifies file extensions to be recognized as a file to be preprocessed by the Fortran preprocessor.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: None

Windows: `/extfpp:ext`

Arguments

ext Are the file extensions to be preprocessed by the Fortran preprocessor.

Default

OFF Only the file extensions recognized by the compiler are preprocessed by `fpp`. For more information, see *Building Applications*.

Description

This option specifies file extensions (*ext*) to be recognized as a file to be preprocessed by the Fortran preprocessor (`fpp`). It is useful if your source file has a nonstandard extension.

You can specify one or more file extensions. A leading period before each extension is optional; for example, `/extfpp:myfpp` and `/extfpp:.myfpp` are equivalent.

Alternate Options

None

extlnk

Specifies file extensions to be passed directly to the linker.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: None

Windows: `/extlnk:ext`

Arguments

ext Are the file extensions to be passed directly to the linker.

Default

OFF Only the file extensions recognized by the compiler are passed to the linker. For more information, see Building Applications.

Description

This option specifies file extensions (*ext*) to be passed directly to the linker. It is useful if your source file has a nonstandard extension.

You can specify one or more file extensions. A leading period before each extension is optional; for example, `/extlnk:myobj` and `/extlnk:.myobj` are equivalent.

Alternate Options

None

F (Windows*)

Specifies the stack reserve amount for the program.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: None

Windows: `/Fn`

Arguments

n Is the stack reserve amount. It can be specified as a decimal integer or by using a C-style convention for constants (for example, `/F0x1000`).

Default

OFF The stack size default is chosen by the operating system.

Description

This option specifies the stack reserve amount for the program. The amount (*n*) is passed to the linker.

Note that the linker property pages have their own option to do this.

Alternate Options

None

f66

Tells the compiler to apply FORTRAN 66 semantics.

IDE Equivalent

Windows: **Language > Enable FORTRAN 66 Semantics**

Linux: None

Mac OS X: **Language > Enable FORTRAN 66 Semantics**

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-f66`

Windows: `/f66`

Arguments

None

Default

OFF The compiler applies Fortran 95 semantics.

Description

This option tells the compiler to apply FORTRAN 66 semantics when interpreting language features. This causes the following to occur:

- DO loops are always executed at least once
- FORTRAN 66 EXTERNAL statement syntax and semantics are allowed

- If the OPEN statement STATUS specifier is omitted, the default changes to STATUS='NEW' instead of STATUS='UNKNOWN'
- If the OPEN statement BLANK specifier is omitted, the default changes to BLANK='ZERO' instead of BLANK='NULL'

Alternate Options

Linux and Mac OS X: -66

Windows: None

f77rtl

Tells the compiler to use the run-time behavior of FORTRAN 77.

IDE Equivalent

Windows: **Compatibility > Enable F77 Run-Time Compatibility**

Linux: None

Mac OS X: **Compatibility > Enable F77 Run-Time Compatibility**

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: -f77rtl

-nof77rtl

Windows: /f77rtl

/nof77rtl

Arguments

None

Default

`nof77rtl` The compiler uses the run-time behavior of Intel® Fortran.

Description

This option tells the compiler to use the run-time behavior of FORTRAN 77.

Specifying this option controls the following run-time behavior:

- When the unit is not connected to a file, some INQUIRE specifiers will return different values:
 - NUMBER= returns 0
 - ACCESS= returns 'UNKNOWN'
 - BLANK= returns 'UNKNOWN'
 - FORM= returns 'UNKNOWN'
 - PAD= defaults to 'NO' for formatted input.
 - NAMELIST and list-directed input of character strings must be delimited by apostrophes or quotes.
- When processing NAMELIST input:
 - Column 1 of each record is skipped.
 - The '\$' or '&' that appears prior to the group-name must appear in column 2 of the input record.

Alternate Options

None

Fa

See [asmfile](#)

FA

See [asmattr](#)

falias

Determines whether aliasing should be assumed in the program.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-falias`
`-fno-alias`

Windows: `None`

Arguments

None

Default

`-falias` Aliasing is assumed in the program.

Description

This option determines whether aliasing should be assumed in the program. If you do not want aliasing to be assumed in the program, specify `-fno-alias`.

Alternate Options

Linux and Mac OS X: `None`

Windows: `/Oa[-]`

See Also

[ffnalias](#) compiler option

falign-functions, Qfnalign

Tells the compiler to align functions on an optimal byte boundary.

IDE Equivalent

None

Architectures

IA-32, Intel® 64 architectures

Syntax

Linux and Mac OS X: `-falign-functions[=n]`

`-fno-align-functions`

Windows: `/Qfnalign[:n]`

`/Qfnalign-`

Arguments

n Is the byte boundary for function alignment.

Possible values are 2 or 16.

Default

`-fno-` The compiler aligns functions on 2-byte boundaries. This

`align-` is the same as specifying `-falign-functions=2`

`functions` (Linux and Mac OS X) or `/Qfnalign:2` (Windows).

or

`/Qfnalign-`

Description

This option tells the compiler to align functions on an optimal byte boundary. If you do not specify *n*, the compiler aligns the start of functions on 16-byte boundaries.

Alternate Options

None

falign-stack

Tells the compiler the stack alignment to use on entry to routines.

IDE Equivalent

None

Architectures

IA-32 architecture

Syntax

Linux and Mac OS X: `-falign-stack=mode`

Windows: None

Arguments

mode Is the method to use for stack alignment. Possible values are:

`default` Tells the compiler to use default heuristics for stack alignment. If alignment is required, the compiler dynamically aligns the stack.

`maintain-16-byte` Tells the compiler to not assume any specific stack alignment, but attempt to maintain alignment in case the stack is already aligned. If alignment is required, the compiler dynamically aligns the stack. This setting is compatible with GCC.

`assume-16-byte` Tells the compiler to assume the stack is aligned on 16-byte boundaries and continue to maintain 16-byte alignment.

This setting is compatible with
GCC.

Default

`-falign-` The compiler uses default heuristics for stack
`stack=default` alignment.

Description

This option tells the compiler the stack alignment to use on entry to routines.

Alternate Options

None

fast

Maximizes speed across the entire program.

IDE Equivalent

Windows: None

Linux: None

Mac OS X: None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-fast`

Windows: `/fast`

Arguments

None

Default

OFF The optimizations that maximize speed are not enabled.

Description

This option maximizes speed across the entire program.

It sets the following options:

- On systems using IA-64 architecture:
 Windows: `/O3` and `/Qipo`
 Linux: `-ipo`, `-O3`, and `-static`
- On systems using IA-32 architecture and Intel® 64 architecture:
 Mac OS X: `-ipo`, `-mdynamic-no-pic`, `-O3`, `-no-prec-div`, `-static`,
 and `-xHost`
 Windows: `/O3`, `/Qipo`, `/Qprec-div-`, and `/QxHost`
 Linux: `-ipo`, `-O3`, `-no-prec-div`, `-static`, and `-xHost`

When option `fast` is specified on systems using IA-32 architecture or Intel® 64 architecture, you can override the `-xHost` or `/QxHost` setting by specifying a different processor-specific `-x` or `/Qx` option on the command line. However, the last option specified on the command line takes precedence.

For example, if you specify `-fast -xSSE3` (Linux) or `/fast /QxSSE3` (Windows), option `-xSSE3` or `/QxSSE3` takes effect. However, if you specify `-xSSE3 -fast` (Linux) or `/QxSSE3 /fast` (Windows), option `-xHost` or `/QxHost` takes effect.



The options set by option `fast` may change from release to release.

Alternate Options

None

fast-transcendentals, Qfast-transcendentals

Enables the compiler to replace calls to transcendental functions with faster but less precise implementations.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-fast-transcendentals`
`-no-fast-transcendentals`

Windows: `/Qfast-transcendentals`
`/Qfast-transcendentals-`

Default

`-fast-transcendentals` The default depends on the setting of `-fp-model` (Linux and Mac OS X) or `/fp` (Windows).

or `/Qfast-transcendentals` The default is ON if default setting `-fp-model fast` or `/fp:fast` is in effect. However, if a value-safe option such as `-fp-model precise` or `/fp:precise` is specified, the default is OFF.

Description

This option enables the compiler to replace calls to transcendental functions with implementations that may be faster but less precise.

It tells the compiler to perform certain optimizations on transcendental functions, such as replacing individual calls to sine in a loop with a single call to a less precise vectorized sine library routine.

This option has an effect only when specified with one of the following options:

- Windows* OS: `/fp:except` or `/fp:precise`

- Linux* OS and Mac OS* X: `-fp-model except` or `-fp-model precise`
You cannot use this option with option `-fp-model strict` (Linux and Mac OS X) or `/fp:strict` (Windows).

Alternate Options

None

See Also

[fp-model, fp](#) compiler option

fcode-asm

Produces an assembly listing with machine code annotations.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-fcode-asm`

Windows: None

Arguments

None

Default

OFF No machine code annotations appear in the assembly listing file,
 if one is produced.

Description

This option produces an assembly listing file with machine code annotations.

The assembly listing file shows the hex machine instructions at the beginning of each line of assembly code. The file cannot be assembled; the filename is the name of the source file with an extension of `.cod`.

To use this option, you must also specify option `-S`, which causes an assembly listing to be generated.

Alternate Options

Linux and Mac OS X: None

Windows: `/asmattr:machine`

See Also

[S](#) compiler option

Fe

See [exe](#)

fexceptions

Enables exception handling table generation.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-fexceptions`
`-fno-exceptions`

Windows: `None`

Arguments

None

Default

`-fno-exceptions` Exception handling table generation is disabled.

Description

This option enables C++ exception handling table generation, preventing Fortran routines in mixed-language applications from interfering with exception handling between C++ routines. The `-fno-exceptions` option disables C++ exception handling table generation, resulting in smaller code. When this option is used, any use of C++ exception handling constructs (such as try blocks and throw statements) when a Fortran routine is in the call chain will produce an error.

Alternate Options

None

ffnalias

Specifies that aliasing should be assumed within functions.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-ffnalias`
`-fno-fnalias`

Windows: None

Arguments

None

Default

`-ffnalias` Aliasing is assumed within functions.

Description

This option specifies that aliasing should be assumed within functions.

The `-fno-fnalias` option specifies that aliasing should not be assumed within functions, but should be assumed across calls.

Alternate Options

Linux and Mac OS X: None

Windows: `/Ow[-]`

See Also

[fnalias](#) compiler option

FI

See [fixed](#).

inline

Tells the compiler to inline functions declared with `cDEC$ ATTRIBUTES FORCEINLINE`.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-inline`
`-fno-inline`

Windows: None

Arguments

None

Default

`-fno-inline` The compiler does not inline functions declared with
`cDEC$ ATTRIBUTES FORCEINLINE.`

Description

This option tells the compiler to inline functions declared with `cDEC$ ATTRIBUTES FORCEINLINE.`

Alternate Options

None

finline-functions

Enables function inlining for single file compilation.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-finline-functions`
`-fno-inline-functions`

Windows: None

Arguments

None

Default

`-finline-` Interprocedural optimizations occur. However, if you
`functions` specify `-O0`, the default is OFF.

Description

This option enables function inlining for single file compilation.

It enables the compiler to perform inline function expansion for calls to functions defined within the current source file.

The compiler applies a heuristic to perform the function expansion. To specify the size of the function to be expanded, use the `-finline-limit` option.

Alternate Options

Linux and Mac OS X: `-inline-level=2`

Windows: `/Ob2`

See Also

[ip, Qip](#) compiler option

[finline-limit](#) compiler option

finline-limit

Lets you specify the maximum size of a function to be inlined.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-finline-limit=n`

Windows: None

Arguments

n Must be an integer greater than or equal to zero. It is the maximum number of lines the function can have to be considered for inlining.

Default

OFF The compiler uses default heuristics when inlining functions.

Description

This option lets you specify the maximum size of a function to be inlined. The compiler inlines smaller functions, but this option lets you inline large functions. For example, to indicate a large function, you could specify 100 or 1000 for *n*. Note that parts of functions cannot be inlined, only whole functions. This option is a modification of the `-finline-functions` option, whose behavior occurs by default.

Alternate Options

None

See Also

[finline-functions](#) compiler option

finstrument-functions, Qinstrument-functions

Determines whether routine entry and exit points are instrumented.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-finstrument-functions`
`-fno-instrument-functions`

Windows: `/Qinstrument-functions`
`/Qinstrument-functions-`

Arguments

None

Default

`-fno-` Routine entry and exit points are not instrumented.
`instrument-`
`functions`
or `/Qinstrument-`
`functions-`

Description

This option determines whether routine entry and exit points are instrumented. It may increase execution time.

The following profiling functions are called with the address of the current routine and the address of where the routine was called (its "call site"):

- This function is called upon routine entry:

- On IA-32 architecture and Intel® 64 architecture:

```
void __cyg_profile_func_enter (void *this_fn,  
void *call_site);
```

- On IA-64 architecture:

```
void __cyg_profile_func_enter (void **this_fn,  
void *call_site);
```

- This function is called upon routine exit:

- On IA-32 architecture and Intel® 64 architecture:

```
void __cyg_profile_func_exit (void *this_fn,  
void *call_site);
```

- On IA-64 architecture:

```
void __cyg_profile_func_exit (void **this_fn,  
void *call_site);
```

On IA-64 architecture, the additional de-reference of the function pointer argument is required to obtain the routine entry point contained in the first word of the routine descriptor for indirect routine calls. The descriptor is documented in the Intel® Itanium® Software Conventions and Runtime Architecture Guide, section 8.4.2. You can find this design guide at web site <http://www.intel.com>. These functions can be used to gather more information, such as profiling information or timing information. Note that it is the user's responsibility to provide these profiling functions.

If you specify `-finstrument-functions` (Linux and Mac OS X) or `/Qinstrument-functions` (Windows), routine inlining is disabled. If you specify `-fno-instrument-functions` or `/Qinstrument-functions-`, inlining is not disabled.

This option is provided for compatibility with gcc.

Alternate Options

None

fixed

Specifies source files are in fixed format.

IDE Equivalent

Windows: **Language > Source File Format** (`/free`, `/fixed`)

Linux: None

Mac OS X: **Language > Source File Format** (`/free`, `/fixed`)

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-fixed`

`-nofixed`

Windows: /fixe
 /nofixe

Arguments

None

Default

OFF The source file format is determined from the file extension.

Description

This option specifies source files are in fixed format. If this option is not specified, format is determined as follows:

- Files with an extension of .f90, .F90, or .i90 are free-format source files.
- Files with an extension of .f, .for, .FOR, .ftn, .FTN, .fpp, .FPP, or .i are fixed-format files.

Note that on Linux and Mac OS X systems, file names and file extensions are case sensitive.

Alternate Options

Linux and Mac OS X: -FI

Windows: /nofree, /FI, /4Nf

fkeep-static-consts, Qkeep-static-consts

Tells the compiler to preserve allocation of variables that are not referenced in the source.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-fkeep-static-consts`
`-fno-keep-static-consts`

Windows: `/Qkeep-static-consts`
`/Qkeep-static-consts-`

Arguments

None

Default

`-fno-keep-static-consts` or `/Qkeep-static-consts-` If a variable is never referenced in a routine, the variable is discarded unless optimizations are disabled by option `-O0` (Linux and Mac OS X) or `/Od` (Windows).

Description

This option tells the compiler to preserve allocation of variables that are not referenced in the source.

The negated form can be useful when optimizations are enabled to reduce the memory usage of static data.

Alternate Options

None

fltconsistency

Enables improved floating-point consistency.

IDE Equivalent

Windows: **Floating-Point > Floating-Point Consistency**

Linux: None

Mac OS X: **Floating-Point > Improve Floating-Point Consistency**

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-fltconsistency`
`-nofltconsistency`

Windows: `/fltconsistency`
`/nofltconsistency`

Arguments

None

Default

`nofltconsistency` Improved floating-point consistency is not enabled.
This setting provides better accuracy and run-time performance at the expense of less consistent floating-point results.

Description

This option enables improved floating-point consistency and may slightly reduce execution speed. It limits floating-point optimizations and maintains declared precision. It also disables inlining of math library functions.

Floating-point operations are not reordered and the result of each floating-point operation is stored in the target variable rather than being kept in the floating-point processor for use in a subsequent calculation.

For example, the compiler can change floating-point division computations into multiplication by the reciprocal of the denominator. This change can alter the results of floating-point division computations slightly.

Floating-point intermediate results are kept in full 80 bits internal precision.

Additionally, all spills/reloads of the X87 floating point registers are done using

the internal formats; this prevents accidental loss of precision due to spill/reload behavior over which you have no control.

Specifying this option has the following effects on program compilation:

- On systems using IA-32 architecture or Intel® 64 architecture, floating-point user variables are not assigned to registers.
- On systems using IA-64 architecture, floating-point user variables may be assigned to registers. The expressions are evaluated using precision of source operands. The compiler will not use the Floating-point Multiply and Add (FMA) function to contract multiply and add/subtract operations in a single operation. The contractions can be enabled by using `-IPF_FMA` (Linux) or `/QIPF_fma` (Windows) option. The compiler will not speculate on floating-point operations that may affect the floating-point state of the machine.
- Floating-point arithmetic comparisons conform to IEEE 754.
- The exact operations specified in the code are performed. For example, division is never changed to multiplication by the reciprocal.
- The compiler performs floating-point operations in the order specified without reassociation.
- The compiler does not perform constant folding on floating-point values. Constant folding also eliminates any multiplication by 1, division by 1, and addition or subtraction of 0. For example, code that adds 0.0 to a number is executed exactly as written. Compile-time floating-point arithmetic is not performed to ensure that floating-point exceptions are also maintained.
- Whenever an expression is spilled, it is spilled as 80 bits (extended precision), not 64 bits (DOUBLE PRECISION). When assignments to type REAL and DOUBLE PRECISION are made, the precision is rounded from 80 bits down to 32 bits (REAL) or 64 bits (DOUBLE PRECISION). When you do not specify `/Op`, the extra bits of precision are not always rounded away before the variable is reused.
- Even if vectorization is enabled by the `-x` (Linux and Mac OS X) or `/Qx` (Windows) options, the compiler does not vectorize reduction loops (loops computing the dot product) and loops with mixed precision types. Similarly,

the compiler does not enable certain loop transformations. For example, the compiler does not transform reduction loops to perform partial summation or loop interchange.

This option causes performance degradation relative to using default floating-point optimization flags.

On Windows systems, an alternative is to use the `/Qprec` option, which should provide better than default floating-point precision while still delivering good floating-point performance.

The recommended method to control the semantics of floating-point calculations is to use option `-fp-model` (Linux and Mac OS X) or `/fp` (Windows).

Alternate Options

`fltconsistency` Linux and Mac OS X: `-mp` (this is a [deprecated](#) option), `-mieee-fp`

Windows: `/Op` (this is a [deprecated](#) option)

`nofltconsistency` Linux and Mac OS X: `-mno-ieee-fp`

Windows: None

See Also

[mp1, Qprec](#) compiler option

[fp-model, fp](#) compiler option

Building Applications: Using Compiler Optimizations

Fm

This option has been [deprecated](#).

See [map](#).

fma, Qfma

Enables the combining of floating-point multiplies and add/subtract operations.

IDE Equivalent

Windows: Floating Point > Contract Floating-Point Operations

Linux: None

Mac OS X: None

Architectures

IA-64 architecture

Syntax

Linux:	<code>-fma</code>
	<code>-no-fma</code>
Mac OS X:	None
Windows:	<code>/Qfma</code>
	<code>/Qfma-</code>

Arguments

None

Default

`-fma` or `/Qfma` Floating-point multiplies and add/subtract operations are combined.

However, if you specify `-mp` (Linux), `/Op` (Windows), `/fp:strict` (Windows), or `-fp-model strict` (Linux) but do not explicitly specify `-fma` or `/Qfma`, the default is `-no-fma` or `/Qfma-`.

Description

This option enables the combining of floating-point multiplies and add/subtract operations.

It also enables the contraction of floating-point multiply and add/subtract operations into a single operation. The compiler contracts these operations whenever possible.

Alternate Options

Linux: `-IPF-fma` (this is a [deprecated](#) option)

Windows: `/QIPF-fma` (this is a deprecated option)

See Also

[fp-model, fp](#) compiler option

Floating-point Operations: Floating-point Options Quick Reference

fmath-errno

Tells the compiler that `errno` can be reliably tested after calls to standard math library functions.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-fmath-errno`
`-fno-math-errno`

Windows: `None`

Arguments

None

Default

`-fno-math-errno` The compiler assumes that the program does not test `errno` after calls to standard math library functions.

Description

This option tells the compiler to assume that the program tests `errno` after calls to math library functions. This restricts optimization because it causes the compiler to treat most math functions as having side effects.

Option `-fno-math-errno` tells the compiler to assume that the program does not test `errno` after calls to math library functions. This frequently allows the compiler to generate faster code. Floating-point code that relies on IEEE exceptions instead of `errno` to detect errors can safely use this option to improve performance.

Alternate Options

None

fminshared

Specifies that a compilation unit is a component of a main program and should not be linked as part of a shareable object.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-fminshared`

Windows: None

Arguments

None

Default

OFF Source files are compiled together to form a single object file.

Description

This option specifies that a compilation unit is a component of a main program and should not be linked as part of a shareable object.

This option allows the compiler to optimize references to defined symbols without special visibility settings. To ensure that external and common symbol references are optimized, you need to specify visibility hidden or protected by using the `-fvisibility`, `-fvisibility-hidden`, or `-fvisibility-protected` option.

Also, the compiler does not need to generate position-independent code for the main program. It can use absolute addressing, which may reduce the size of the global offset table (GOT) and may reduce memory traffic.

Alternate Options

None

See Also

[fvisibility](#) compiler option

fnsplit, Qfnsplit

Enables function splitting.

IDE Equivalent

None

Architectures

`/Qfnsplit[-]`: IA-32 architecture, Intel® 64 architecture

`-[no-]fnsplit`: IA-64 architecture

Syntax

Linux: `-fnsplit`
 `-no-fnsplit`

Mac OS X: None
 Windows: /Qfnsplit
 /Qfnsplit-

Arguments

None

Default

`-no-fnsplit` Function splitting is not enabled unless `-prof-`
`or/Qfnsplit-` use (Linux) or `/Qprof-use` (Windows) is also
 specified.

Description

This option enables function splitting if `-prof-use` (Linux) or `/Qprof-use` (Windows) is also specified. Otherwise, this option has no effect.

It is enabled automatically if you specify `-prof-use` or `/Qprof-use`. If you do not specify one of those options, the default is `-no-fnsplit` (Linux) or `/Qfnsplit-` (Windows), which disables function splitting but leaves function grouping enabled.

To disable function splitting when you use `-prof-use` or `/Qprof-use`, specify `-no-fnsplit` or `/Qfnsplit-`.

Alternate Options

None

See Also

Optimizing Applications:
 Basic PGO Options
 Example of Profile-Guided Optimization

fomit-frame-pointer, Oy

Determines whether EBP is used as a general-purpose register in optimizations.

IDE Equivalent

Windows: **Optimization > Omit Frame Pointers**

Linux: None

Mac OS X: **Optimization > Provide Frame Pointer**

Architectures

`-f[no-]omit-frame-pointer`: IA-32 architecture, Intel® 64 architecture

`/Oy[-]`: IA-32 architecture

Syntax

Linux and Mac OS X: `-fomit-frame-pointer`

`-fno-omit-frame-pointer`

Windows: `/Oy`

`/Oy-`

Arguments

None

Default

`-fomit-frame-pointer` or `/Oy` EBP is used as a general-purpose register in optimizations. However, on Linux* and Mac OS X systems, the default is `-fno-omit-frame-pointer` if option `-O0` or `-g` is specified. On Windows* systems, the default is `/Oy-` if option `/Od` is specified.

Description

These options determine whether EBP is used as a general-purpose register in optimizations. Options `-fomit-frame-pointer` and `/Oy` allow this use.

Options `-fno-omit-frame-pointer` and `/Oy-` disallow it.

Some debuggers expect EBP to be used as a stack frame pointer, and cannot produce a stack backtrace unless this is so. The `-fno-omit-frame-pointer` and `/Oy-` options direct the compiler to generate code that maintains and uses EBP as a stack frame pointer for all functions so that a debugger can still produce a stack backtrace without doing the following:

- For `-fno-omit-frame-pointer`: turning off optimizations with `-O0`
- For `/Oy-`: turning off `/O1`, `/O2`, or `/O3` optimizations

The `-fno-omit-frame-pointer` option is set when you specify option `-O0` or the `-g` option. The `-fomit-frame-pointer` option is set when you specify option `-O1`, `-O2`, or `-O3`.

The `/Oy` option is set when you specify the `/O1`, `/O2`, or `/O3` option. Option `/Oy-` is set when you specify the `/Od` option.

Using the `-fno-omit-frame-pointer` or `/Oy-` option reduces the number of available general-purpose registers by 1, and can result in slightly less efficient code.

Alternate Options

Linux and Mac OS X: `-fp` (this is a [deprecated](#) option)

Windows: None

Fo

See [object](#)

fomit-frame-pointer, Oy

Determines whether EBP is used as a general-purpose register in optimizations.

IDE Equivalent

Windows: **Optimization > Omit Frame Pointers**

Linux: None

Mac OS X: **Optimization > Provide Frame Pointer**

Architectures

`-f[no-]omit-frame-pointer`: IA-32 architecture, Intel® 64 architecture

`/Oy[-]`: IA-32 architecture

Syntax

Linux and Mac OS X: `-fomit-frame-pointer`

`-fno-omit-frame-pointer`

Windows: `/Oy`

`/Oy-`

Arguments

None

Default

`-fomit-frame-pointer` or `/Oy` EBP is used as a general-purpose register in optimizations. However, on Linux* and Mac OS X systems, the default is `-fno-omit-frame-pointer` if option `-O0` or `-g` is specified. On Windows* systems, the default is `/Oy-` if option `/Od` is specified.

Description

These options determine whether EBP is used as a general-purpose register in optimizations. Options `-fomit-frame-pointer` and `/Oy` allow this use.

Options `-fno-omit-frame-pointer` and `/Oy-` disallow it.

Some debuggers expect EBP to be used as a stack frame pointer, and cannot produce a stack backtrace unless this is so. The `-fno-omit-frame-pointer` and `/Oy-` options direct the compiler to generate code that maintains and uses

EBP as a stack frame pointer for all functions so that a debugger can still produce a stack backtrace without doing the following:

- For `-fno-omit-frame-pointer`: turning off optimizations with `-O0`
- For `/Oy-`: turning off `/O1`, `/O2`, or `/O3` optimizations

The `-fno-omit-frame-pointer` option is set when you specify option `-O0` or the `-g` option. The `-fomit-frame-pointer` option is set when you specify option `-O1`, `-O2`, or `-O3`.

The `/Oy` option is set when you specify the `/O1`, `/O2`, or `/O3` option. Option `/Oy-` is set when you specify the `/Od` option.

Using the `-fno-omit-frame-pointer` or `/Oy-` option reduces the number of available general-purpose registers by 1, and can result in slightly less efficient code.

Alternate Options

Linux and Mac OS X: `-fp` (this is a [deprecated](#) option)

Windows: None

fp-model, fp

Controls the semantics of floating-point calculations.

IDE Equivalent

Windows: None

Linux: None

Mac OS X: **Floating Point > Floating Point Model** (`precise`, `fast`, `fast=2`, `strict`, `source`)

Floating Point > Reliable Floating Point Exceptions Model (`fp-model` `except`)

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-fp-modelkeyword`

Windows: `/fp:keyword`

Arguments

<i>keyword</i>	Specifies the semantics to be used. Possible values are:
<code>precise</code>	Enables value-safe optimizations on floating-point data and rounds intermediate results to source-defined precision.
<code>fast [=1 2]</code>	Enables more aggressive optimizations on floating-point data.
<code>strict</code>	Enables precise and except, disables contractions, and enables the property that allows modification of the floating-point environment.
<code>source</code>	Rounds intermediate results to source-defined precision and enables value-safe optimizations.
<code>[no-]except</code>	Determines whether floating-point exception semantics are used. (Linux and Mac OS X)
<code>or</code>	

```
except[-]
(Windows)
```

Default

`-fp-model` The compiler uses more aggressive optimizations on
`fast=1` floating-point calculations, except when `-O0` (Linux and
or Mac OS X) or `/Od` (Windows) is specified.
`/fp:fast=1`

Description

This option controls the semantics of floating-point calculations.

The *keywords* can be considered in groups:

- Group A: `precise`, `fast`, `strict`
- Group B: `source`
- Group C: `except` (or the negative form)

You can use more than one *keyword*. However, the following rules apply:

- You cannot specify `fast` and `except` together in the same compilation. You can specify any other combination of group A, group B, and group C. Since `fast` is the default, you must not specify `except` without a group A or group B *keyword*.
- You should specify only one *keyword* from group A and only one keyword from group B. If you try to specify more than one *keyword* from either group A or group B, the last (rightmost) one takes effect.
- If you specify `except` more than once, the last (rightmost) one takes effect.

Option	Description
<code>-fp-model precise</code> or <code>/fp:precise</code>	Tells the compiler to strictly adhere to value-safe optimizations when implementing floating-point calculations. It disables optimizations that can

Option	Description
	<p>change the result of floating-point calculations. These semantics ensure the accuracy of floating-point computations, but they may slow performance.</p> <p>The compiler assumes the default floating-point environment; you are not allowed to modify it.</p> <p>Floating-point exception semantics are disabled by default. To enable these semantics, you must also specify <code>-fp-model except</code> or <code>/fp:except</code>.</p> <p>For information on the semantics used to interpret floating-point calculations in the source code, see <code>precise</code> in <i>Floating-point Operations: Using the <code>-fp-model (/fp)</code> Option</i>.</p>
<pre>-fp-model fast[=1 2] or /fp:fast[=1 2]</pre>	<p>Tells the compiler to use more aggressive optimizations when implementing floating-point calculations. These optimizations increase speed, but may alter the accuracy of floating-point computations.</p> <p>Specifying <code>fast</code> is the same as specifying <code>fast=1</code>. <code>fast=2</code> may produce faster and less accurate results.</p> <p>Floating-point exception semantics are disabled by default and they cannot be enabled because you cannot specify</p>

Option	Description
	<p>fast and except together in the same compilation. To enable exception semantics, you must explicitly specify another keyword (see other keyword descriptions for details).</p> <p>For information on the semantics used to interpret floating-point calculations in the source code, see <code>fast</code> in <i>Floating-point Operations: Using the <code>-fp-model (/fp)</code> Option</i>.</p>
<code>-fp-model strict</code> or <code>/fp:strict</code>	<p>Tells the compiler to strictly adhere to value-safe optimizations when implementing floating-point calculations and enables floating-point exception semantics. This is the strictest floating-point model.</p> <p>The compiler does not assume the default floating-point environment; you are allowed to modify it.</p> <p>Floating-point exception semantics can be disabled by explicitly specifying <code>-fp-model no-exception</code> or <code>/fp:exception-</code>.</p> <p>For information on the semantics used to interpret floating-point calculations in the source code, see <code>strict</code> in <i>Floating-point Operations: Using the <code>-fp-model (/fp)</code> Option</i>.</p>
<code>-fp-model source</code> or <code>/fp:source</code>	<p>This option causes intermediate results</p>

Option	Description
	<p>to be rounded to the precision defined in the source code. It also implies keyword <code>precise</code> unless it is overridden by a keyword from Group A.</p> <p>The compiler assumes the default floating-point environment; you are not allowed to modify it.</p> <p>For information on the semantics used to interpret floating-point calculations in the source code, see <code>source</code> in <i>Floating-point Operations: Using the <code>-fp-model (/fp)</code> Option</i>.</p>



Note

This option cannot be used to change the default (source) precision for the calculation of intermediate results.



Note

On Windows and Linux operating systems on IA-32 architecture, the compiler, by default, implements floating-point (FP) arithmetic using SSE2 and SSE instructions. This can cause differences in floating-point results when compared to previous x87 implementations.

Alternate Options

None

Example

For examples of how to use this option, see *Floating-point Operations: Using the `-fp-model (/fp)` Option*.

See Also

[O](#) compiler option (specifically O0)

[Od](#) compiler option

[mp1, Qprec](#) compiler option

The MSDN article [Microsoft Visual C++ Floating-Point Optimization](#), which discusses concepts that apply to this option.

Floating-point Operations: Floating-Point Environment

fp-model, fp

Controls the semantics of floating-point calculations.

IDE Equivalent

Windows: None

Linux: None

Mac OS X: **Floating Point > Floating Point Model** (`precise`, `fast`, `fast=2`, `strict`, `source`)

Floating Point > Reliable Floating Point Exceptions Model (`fp-model` `except`)

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-fp-modelkeyword`

Windows: `/fp:keyword`

Arguments

keyword Specifies the semantics to be used. Possible values are:

`precise` Enables value-safe

	optimizations on floating-point data and rounds intermediate results to source-defined precision.
<code>fast [=1 2]</code>	Enables more aggressive optimizations on floating-point data.
<code>strict</code>	Enables precise and except, disables contractions, and enables the property that allows modification of the floating-point environment.
<code>source</code>	Rounds intermediate results to source-defined precision and enables value-safe optimizations.
<code>[no-]except</code> (Linux and Mac OS X) or <code>except[-]</code> (Windows)	Determines whether floating-point exception semantics are used.

Default

`-fp-model` The compiler uses more aggressive optimizations on
`fast=1` floating-point calculations, except when `-O0` (Linux and
or Mac OS X) or `/Od` (Windows) is specified.
`/fp:fast=1`

Description

This option controls the semantics of floating-point calculations.

The *keywords* can be considered in groups:

- Group A: `precise`, `fast`, `strict`
- Group B: `source`
- Group C: `except` (or the negative form)

You can use more than one *keyword*. However, the following rules apply:

- You cannot specify `fast` and `except` together in the same compilation. You can specify any other combination of group A, group B, and group C. Since `fast` is the default, you must not specify `except` without a group A or group B *keyword*.
- You should specify only one *keyword* from group A and only one keyword from group B. If you try to specify more than one *keyword* from either group A or group B, the last (rightmost) one takes effect.
- If you specify `except` more than once, the last (rightmost) one takes effect.

Option	Description
<code>-fp-model precise</code> or <code>/fp:precise</code>	<p>Tells the compiler to strictly adhere to value-safe optimizations when implementing floating-point calculations. It disables optimizations that can change the result of floating-point calculations. These semantics ensure the accuracy of floating-point computations, but they may slow performance.</p> <p>The compiler assumes the default floating-point environment; you are not allowed to modify it.</p> <p>Floating-point exception semantics are</p>

Option	Description
<code>-fp-model fast [=1 2] or /fp:fast [=1 2]</code>	<p>disabled by default. To enable these semantics, you must also specify <code>-fp-model except</code> or <code>/fp:except</code>.</p> <p>For information on the semantics used to interpret floating-point calculations in the source code, see <code>precise</code> in <i>Floating-point Operations: Using the <code>-fp-model (/fp)</code> Option</i>.</p> <p>Tells the compiler to use more aggressive optimizations when implementing floating-point calculations. These optimizations increase speed, but may alter the accuracy of floating-point computations.</p> <p>Specifying <code>fast</code> is the same as specifying <code>fast=1</code>. <code>fast=2</code> may produce faster and less accurate results.</p> <p>Floating-point exception semantics are disabled by default and they cannot be enabled because you cannot specify <code>fast</code> and <code>except</code> together in the same compilation. To enable exception semantics, you must explicitly specify another keyword (see other keyword descriptions for details).</p> <p>For information on the semantics used to interpret floating-point calculations in the source code, see <code>fast</code> in <i>Floating-point Operations: Using the <code>-fp-model</code></i></p>

Option	Description
<code>-fp-model strict</code> or <code>/fp:strict</code>	<p data-bbox="808 268 1003 310"><i>(/fp) Option.</i></p> <p data-bbox="808 342 1369 653">Tells the compiler to strictly adhere to value-safe optimizations when implementing floating-point calculations and enables floating-point exception semantics. This is the strictest floating-point model.</p> <p data-bbox="808 674 1354 821">The compiler does not assume the default floating-point environment; you are allowed to modify it.</p> <p data-bbox="808 842 1383 1272">Floating-point exception semantics can be disabled by explicitly specifying <code>-fp-model no-exception</code> or <code>/fp:exception-</code>. For information on the semantics used to interpret floating-point calculations in the source code, see <code>strict</code> in <i>Floating-point Operations: Using the -fp-model (/fp) Option</i>.</p>
<code>-fp-model source</code> or <code>/fp:source</code>	<p data-bbox="808 1308 1383 1566">This option causes intermediate results to be rounded to the precision defined in the source code. It also implies keyword <code>precise</code> unless it is overridden by a keyword from Group A.</p> <p data-bbox="808 1587 1359 1734">The compiler assumes the default floating-point environment; you are not allowed to modify it.</p> <p data-bbox="808 1755 1365 1837">For information on the semantics used to interpret floating-point calculations in</p>

Option	Description
	the source code, see <code>source</code> in <i>Floating-point Operations: Using the <code>-fp-model (/fp)</code> Option</i> .



Note

This option cannot be used to change the default (source) precision for the calculation of intermediate results.



Note

On Windows and Linux operating systems on IA-32 architecture, the compiler, by default, implements floating-point (FP) arithmetic using SSE2 and SSE instructions. This can cause differences in floating-point results when compared to previous x87 implementations.

Alternate Options

None

Example

For examples of how to use this option, see *Floating-point Operations: Using the `-fp-model (/fp)` Option*.

See Also

[O](#) compiler option (specifically O0)

[Od](#) compiler option

[mp1](#), [Qprec](#) compiler option

The MSDN article [Microsoft Visual C++ Floating-Point Optimization](#), which discusses concepts that apply to this option.

Floating-point Operations: Floating-Point Environment

fp-port, Qfp-port

Rounds floating-point results after floating-point operations.

IDE Equivalent

Windows: **Floating-Point > Round Floating-Point Results**

Linux: None

Mac OS X: **Floating-Point > Round Floating-Point Results**

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-fp-port`
`-no-fp-port`

Windows: `/Qfp-port`
`/Qfp-port-`

Arguments

None

Default

`-no-fp-port` or `/Qfp-port-` The default rounding behavior depends on the compiler's code generation decisions and the precision parameters of the operating system.

Description

This option rounds floating-point results after floating-point operations. Rounding to user-specified precision occurs at assignments and type conversions. This has some impact on speed.

The default is to keep results of floating-point operations in higher precision. This provides better performance but less consistent floating-point results.

Alternate Options

None

fp-relaxed, Qfp-relaxed

Enables use of faster but slightly less accurate code sequences for math functions.

IDE Equivalent

None

Architectures

IA-64 architecture

Syntax

Linux:	<code>-fp-relaxed</code> <code>-no-fp-relaxed</code>
Mac OS X:	None
Windows:	<code>/Qfp-relaxed</code> <code>/Qfp-relaxed-</code>

Arguments

None

Default

`-no-fp-relaxed` or `/Qfp-relaxed-` Default code sequences are used for math functions.

Description

This option enables use of faster but slightly less accurate code sequences for math functions, such as divide and sqrt. When compared to strict IEEE* precision,

this option slightly reduces the accuracy of floating-point calculations performed by these functions, usually limited to the least significant digit.

This option also enables the performance of more aggressive floating-point transformations, which may affect accuracy.

Alternate Options

Linux: `-IPF-fp-relaxed` (this is a [deprecated](#) option)

Windows: `/QIPF-fp-relaxed` (this is a deprecated option)

See Also

[fp-model](#), [fp](#) compiler option

fp-speculation, Qfp-speculation

Tells the compiler the mode in which to speculate on floating-point operations.

IDE Equivalent

Windows: **Floating Point > Floating-Point Speculation**

Linux: None

Mac OS X: **Floating Point > Floating-Point Speculation**

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-fp-speculation=mode`

Windows: `/Qfp-speculation:mode`

Arguments

mode Is the mode for floating-point operations. Possible values are:

`fast` Tells the compiler to speculate

	on floating-point operations.
<code>safe</code>	Tells the compiler to disable speculation if there is a possibility that the speculation may cause a floating-point exception.
<code>strict</code>	Tells the compiler to disable speculation on floating-point operations.
<code>off</code>	This is the same as specifying <code>strict</code> .

Default

`-fp-` The compiler speculates on floating-point
`speculation=fast` operations. This is also the behavior when
`/Qfp-` optimizations are enabled. However, if you specify
`speculation:fast` no optimizations (`-O0` on Linux; `/Od` on Windows),
the default is `-fp-speculation=safe` (Linux) or
`/Qfp-speculation:safe` (Windows).

Description

This option tells the compiler the mode in which to speculate on floating-point operations.

Alternate Options

None

See Also

Floating-point Operations: Floating-point Options Quick Reference

fp-stack-check, Qfp-stack-check

Tells the compiler to generate extra code after every function call to ensure that the floating-point stack is in the expected state.

IDE Equivalent

Windows: **Floating-Point > Check Floating-point Stack**

Linux: None

Mac OS X: **Floating-Point > Check Floating-point Stack**

Architectures

IA-32, Intel® 64 architectures

Syntax

Linux and Mac OS X: `-fp-stack-check`

Windows: `/Qfp-stack-check`

Arguments

None

Default

OFF There is no checking to ensure that the floating-point (FP) stack is in the expected state.

Description

This option tells the compiler to generate extra code after every function call to ensure that the floating-point (FP) stack is in the expected state.

By default, there is no checking. So when the FP stack overflows, a NaN value is put into FP calculations and the program's results differ. Unfortunately, the overflow point can be far away from the point of the actual bug. This option places code that causes an access violation exception immediately after an incorrect call occurs, thus making it easier to locate these issues.

Alternate Options

None

See Also

Floating-point Operations:

[Checking the Floating-point Stack State](#)

fpconstant

Tells the compiler that single-precision constants assigned to double-precision variables should be evaluated in double precision.

IDE Equivalent

Windows: **Floating-Point > Extend Precision of Single-Precision Constants**

Linux: None

Mac OS X: **Floating-Point > Extend Precision of Single-Precision Constants**

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-fpconstant`
`-nofpconstant`

Windows: `/fpconstant`
`/nofpconstant`

Arguments

None

Default

`nofpconstant` Single-precision constants assigned to double-precision variables are evaluated in single precision according to

Fortran 95/90 Standard rules.

Description

This option tells the compiler that single-precision constants assigned to double-precision variables should be evaluated in double precision.

This is extended precision. It does not comply with the Fortran 95/90 standard, which requires that single-precision constants assigned to double-precision variables be evaluated in single precision.

It allows compatibility with FORTRAN 77, where such extended precision was allowed. If this option is not used, certain programs originally created for FORTRAN 77 compilers may show different floating-point results because they rely on the extended precision for single-precision constants assigned to double-precision variables.

Alternate Options

None

Example

In the following example, if you specify `fpconstant`, identical values are assigned to D1 and D2. If you omit `fpconstant`, the compiler will obey the Fortran 95/90 Standard and assign a less precise value to D1:

```
REAL (KIND=8) D1, D2
DATA D1 /2.71828182846182/ ! REAL (KIND=4) value expanded to double
DATA D2 /2.71828182846182D0/ ! Double value assigned to double
```

fpe

Allows some control over floating-point exception handling for the main program at run-time.

IDE Equivalent

Windows: **Floating-Point > Floating-Point Exception Handling**

Linux: None

Mac OS X: **Floating-Point** > **Floating-Point Exception Handling**

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-fpem`

Windows: `/fpe:n`

Arguments

n Specifies the floating-point exception handling. Possible values are:

0 Floating-point invalid, divide-by-zero, and overflow exceptions are enabled. If any such exceptions occur, execution is aborted. This option sets the `-ftz` (Linux and Mac OS X) or `/Qftz` (Windows) option; therefore underflow results will be set to zero unless you explicitly specify `-no-ftz` (Linux and Mac OS X) or `/Qftz-` (Windows). On systems using IA-64 architecture, underflow behavior is equivalent to specifying option `-ftz` or `/Qftz`. On systems using IA-32 architecture or Intel® 64 architecture, underflow results from SSE instructions, as well as x87 instructions, will be set to zero. By contrast, option `-ftz` or `/Qftz` only sets SSE underflow results to zero.

To get more detailed location information about where the error occurred, use option `traceback`.

- 1 All floating-point exceptions are disabled. On systems using IA-64 architecture, underflow behavior is equivalent to specifying option `-ftz` or `/Qftz`. On systems using IA-32 architecture or Intel® 64 architecture, underflow results from SSE instructions, as well as x87 instructions, will be set to zero.
- 3 All floating-point exceptions are disabled. Floating-point underflow is gradual, unless you explicitly specify a compiler option that enables flush-to-zero, such as `-ftz` or `/Qftz, 03`, or `02` on systems using IA-32 architecture or Intel® 64 architecture. This setting provides full IEEE support.

Default

`-fpe3` All floating-point exceptions are disabled. Floating-point underflow is gradual, unless you explicitly specify a compiler `/fpe:3` option that enables flush-to-zero.

Description

This option allows some control over floating-point exception handling for the main program at run-time. This includes whether exceptional floating-point values are allowed and how precisely run-time exceptions are reported.

The `fpe` option affects how the following conditions are handled:

- When floating-point calculations result in a divide by zero, overflow, or invalid operation.
- When floating-point calculations result in an underflow.
- When a denormalized number or other exceptional number (positive infinity, negative infinity, or a NaN) is present in an arithmetic expression.

When enabled exceptions occur, execution is aborted and the cause of the abort reported to the user. If compiler option `traceback` is specified at compile time, detailed information about the location of the abort is also reported.

This option does not enable underflow exceptions, input denormal exceptions, or inexact exceptions.

Alternate Options

None

See Also

[ftz, Qftz](#) compiler option

[traceback](#) compiler option

[Using the -fpe or /fpe Compiler Option](#)

fpic

Determines whether the compiler generates position-independent code.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-fpic`
`-fno-pic`

Windows: None

Arguments

None

Default

`-fno-pic` or `-fpic` On systems using IA-32 or Intel® 64 architecture, the compiler does not generate position-independent code. On systems using IA-64 architecture, the compiler generates position-independent code.

Description

This option determines whether the compiler generates position-independent code.

Option `-fpic` specifies full symbol preemption. Global symbol definitions as well as global symbol references get default (that is, preemptable) visibility unless explicitly specified otherwise.

Option `-fno-pic` is only valid on systems using IA-32 or Intel® 64 architecture. On systems using IA-32 or Intel® 64 architecture, `-fpic` must be used when building shared objects.

This option can also be specified as `-fPIC`.

Alternate Options

None

fpie

Tells the compiler to generate position-independent code. The generated code can only be linked into executables.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux:	<code>-fpie</code>
Mac OS X:	None
Windows:	None

Arguments

None

Default

OFF	The compiler does not generate position-independent code for an executable-only object.
-----	---

Description

This option tells the compiler to generate position-independent code. It is similar to `-fpic`, but code generated by `-fpie` can only be linked into an executable. Because the object is linked into an executable, this option causes better optimization of some symbol references.

To ensure that run-time libraries are set up properly for the executable, you should also specify option `-pie` to the compiler driver on the link command line.

Option `-fpie` can also be specified as `-fPIE`.

Alternate Options

None

See Also

[fpic](#) compiler option

[pie](#) compiler option

fpp, Qfpp

Runs the Fortran preprocessor on source files before compilation.

IDE Equivalent

Windows: **Preprocessor > Preprocess Source File**

Linux: None

Mac OS X: Preprocessor > Preprocess Source File

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `fpp[n]`

`fpp["option"]`

`-nofpp`

Windows:

`/fpp[n]`

`/fpp["option"]`

`/nofpp`

`/Qfpp[n]`

`/Qfpp["option"]`

Arguments

n [Deprecated](#). Tells the compiler whether to run the preprocessor or not. Possible values are:

0 Tells the compiler not to run the preprocessor.

1, 2, or 3 Tells the compiler to run the preprocessor.

option Is a Fortran preprocessor (fpp) option; for example, "-macro=no", which disables macro expansion. The quotes are required. For a list of

fpp options, see *Fortran Preprocessor Options*.

Default

`nofpp` The Fortran preprocessor is not run on files before compilation.

Description

This option runs the Fortran preprocessor on source files before they are compiled.

If the option is specified with no argument, the compiler runs the preprocessor.

If 0 is specified for *n*, it is equivalent to `nofpp`. Note that argument *n* is [deprecated](#).

We recommend you use option `Qoption, fpp, "option"` to pass fpp options to the Fortran preprocessor.

Alternate Options

Linux and Mac OS X: `-cpp`

Windows: `/Qcpp`

See Also

[Fortran Preprocessor Options](#)

[Qoption](#) compiler option

fpscomp

Controls whether certain aspects of the run-time system and semantic language features within the compiler are compatible with Intel® Fortran or Microsoft* Fortran PowerStation.

IDE Equivalent

Windows: **Compatibility > Use Filenames from Command Line**

(`/fpscomp:[no]filesfromcmd`)

Compatibility > Use PowerStation I/O Format (`/fpscomp:[no]ioformat`)

Compatibility > Use PowerStation Portability Library (/fpscomp:[no]libs)**Compatibility > Use PowerStation List-Directed I/O Spacing**

(/fpscomp:[no]ldio_spacing)

Compatibility > Use PowerStation Logical Values

(/fpscomp:[no]logicals)

Compatibility > Use Other PowerStation Run-Time Behavior

(/fpscomp:[no]general)

Linux: None

Mac OS X: None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: -fpscomp [keyword]

-nofpscomp

Windows: /fpscomp[:keyword]

/nofpscomp

Arguments*keyword* Specifies the compatibility that the compiler should follow.

Possible values are:

none	Specifies that no options should be used for compatibility.
------	---

[no]filesfromcmd	Determines what compatibility is used when the OPEN statement FILE= specifier is blank.
------------------	---

[no]general	Determines what compatibility is used when semantics differences exist between Fortran
-------------	--

	PowerStation and Intel® Fortran.
<code>[no]ioformat</code>	Determines what compatibility is used for list-directed formatted and unformatted I/O.
<code>[no]libs</code>	Determines whether the portability library is passed to the linker.
<code>[no]ldio_spacing</code>	Determines whether a blank is inserted at run-time after a numeric value before a character value.
<code>[no]logicals</code>	Determines what compatibility is used for representation of LOGICAL values.
<code>all</code>	Specifies that all options should be used for compatibility.

Default

`fpscomp` The portability library is passed to the linker.

`libs`

Description

This option controls whether certain aspects of the run-time system and semantic language features within the compiler are compatible with Intel Fortran or Microsoft* Fortran PowerStation.

If you experience problems when porting applications from Fortran PowerStation, specify `fpscomp` (or `fpscomp all`). When porting applications from Intel Fortran, use `fpscomp none` or `fpscomp libs` (the default).

Option	Description
<code>fpscomp none</code>	<p>Specifies that no options should be used for compatibility with Fortran PowerStation. This is the same as specifying <code>nofpscomp</code>. Option <code>fpscomp none</code> enables full Intel® Fortran compatibility. If you omit <code>fpscomp</code>, the default is <code>fpscomp libs</code>. You cannot use the <code>fpscomp</code> and <code>vms</code> options in the same command.</p>
<code>fpscomp filesfromcmd</code>	<p>Specifies Fortran PowerStation behavior when the OPEN statement FILE= specifier is blank (FILE= ' '). It causes the following actions to be taken at run-time:</p> <ul style="list-style-type: none"> • The program reads a filename from the list of arguments (if any) in the command line that invoked the program. If any of the command-line arguments contain a null string (""), the program asks the user for the corresponding filename. Each additional OPEN statement with a blank FILE= specifier reads the next command-line argument. • If there are more nameless OPEN statements than command-line arguments, the program prompts for additional file names. • In a QuickWin application, a File Select dialog box appears to request file names. <p>To prevent the run-time system from using the filename specified on the command line when the OPEN statement FILE specifier is omitted, specify <code>fpscomp nofilesfromcmd</code>. This allows the application of Intel Fortran defaults, such as the FORTn environment variable and the FORT.n file name (where <i>n</i> is the unit number).</p> <p>The <code>fpscomp filesfromcmd</code> option affects the following Fortran features:</p> <ul style="list-style-type: none"> • The OPEN statement FILE specifier

Option	Description
	<p>For example, assume a program OPENTEST contains the following statements:</p> <pre>OPEN(UNIT = 2, FILE = ' ') OPEN(UNIT = 3, FILE = ' ') OPEN(UNIT = 4, FILE = ' ')</pre> <p>The following command line assigns the file TEST.DAT to unit 2, prompts the user for a filename to associate with unit 3, then prompts again for a filename to associate with unit 4:</p> <pre>opentest test.dat " "</pre> <ul style="list-style-type: none"> • Implicit file open statements such as the WRITE, READ, and ENDFILE statements Unopened files referred to in READ or WRITE statements are opened implicitly as if there had been an OPEN statement with a name specified as all blanks. The name is read from the command line.
<p><code>fpscomp</code> <code>general</code></p>	<p>Specifies that Fortran PowerStation semantics should be used when a difference exists between Intel Fortran and Fortran PowerStation. The <code>fpscomp general</code> option affects the following Fortran features:</p> <ul style="list-style-type: none"> • The BACKSPACE statement: • It allows files opened with <code>ACCESS='APPEND'</code> to be used with the BACKSPACE statement. • It allows files opened with <code>ACCESS='DIRECT'</code> to be used with the BACKSPACE statement. <p>Note: Allowing files that are not opened with sequential access (such as <code>ACCESS='DIRECT'</code>) to be used with the BACKSPACE statement violates the Fortran 95 standard and may be removed in the future.</p> <ul style="list-style-type: none"> • The READ statement: • It causes a READ from a formatted file opened for direct

Option	Description
	<p>access to read records that have the same record type format as Fortran PowerStation. This consists of accounting for the trailing Carriage Return/Line Feed pair (<CR><LF>) that is part of the record. It allows sequential reads from a formatted file opened for direct access.</p> <p>Note: Allowing files that are not opened with sequential access (such as ACCESS='DIRECT') to be used with the sequential READ statement violates the Fortran 95 standard and may be removed in the future.</p> <ul style="list-style-type: none"> • It allows the last record in a file opened with FORM='FORMATTED' and a record type of STREAM_LF or STREAM_CR that does not end with a proper record terminator (<line feed> or <carriage return>) to be read without producing an error. • It allows sequential reads from an unformatted file opened for direct access. • Note: Allowing files that are not opened with sequential access (such as ACCESS='DIRECT') to be read with the sequential READ statement violates the Fortran 95 standard and may be removed in the future. • The INQUIRE statement: • The CARRIAGECONTROL specifier returns the value "UNDEFINED" instead of "UNKNOWN" when the carriage control is not known. • The NAME specifier returns the file name "UNKNOWN" instead of filling the file name with spaces when the file name is not known. • The SEQUENTIAL specifier returns the value "YES" instead of "NO" for a direct access formatted file.

Option	Description
	<ul style="list-style-type: none"> • The UNFORMATTED specifier returns the value "NO" instead of "UNKNOWN" when it is not known whether unformatted I/O can be performed to the file. Note: Returning the value "NO" instead of "UNKNOWN" for this specifier violates the Fortran 95 standard and may be removed in the future. • The OPEN statement: • If a file is opened with an unspecified STATUS keyword value, and is not named (no FILE specifier), the file is opened as a scratch file. For example: OPEN (UNIT = 4) • In contrast, when fpscomp nogeneral is in effect with an unspecified STATUS value with no FILE specifier, the FORTn environment variable and the FORT.n file name are used (where n is the unit number). • If the STATUS value was not specified and if the name of the file is "USER", the file is marked for deletion when it is closed. • It allows a file to be opened with the APPEND and READONLY characteristics. • If the default for the CARRIAGECONTROL specifier is assumed, it gives "LIST" carriage control to direct access formatted files instead of "NONE". • If the default for the CARRIAGECONTROL specifier is assumed and the device type is a terminal file, the file is given the default carriage control value of "FORTRAN" instead of "LIST". • It gives an opened file the additional default of write sharing. • It gives the file a default block size of 1024 instead of 8192.

Option	Description
	<ul style="list-style-type: none"> • If the default for the MODE and ACTION specifier is assumed and there was an error opening the file, try opening the file as read only, then write only. • If a file that is being re-opened has a different file type than the current existing file, an error is returned. • It gives direct access formatted files the same record type as Fortran PowerStation. This means accounting for the trailing Carriage Return/Line Feed pair (<CR><LF>) that is part of the record. • The STOP statement: It writes the Fortran PowerStation output string and/or returns the same exit condition values. • The WRITE statement: • Writing to formatted direct files When writing to a formatted file opened for direct access, records are written in the same record type format as Fortran PowerStation. This consists of adding the trailing Carriage Return/Line Feed pair <CR><LF>) that is part of the record. It ignores the CARRIAGECONTROL specifier setting when writing to a formatted direct access file. • Interpreting Fortran carriage control characters When interpreting Fortran carriage control characters during formatted I/O, carriage control sequences are written that are the same as Fortran PowerStation. This is true for the "Space, 0, 1 and + " characters. • Performing non-advancing I/O to the terminal When performing non-advancing I/O to the terminal, output is written in the same format as Fortran PowerStation. • Interpreting the backslash (\) and dollar (\$) edit descriptors When interpreting backslash and dollar edit descriptors

Option	Description
<code>fpscomp</code> <code>ioformat</code>	<p>during formatted I/O, sequences are written the same as Fortran PowerStation.</p> <ul style="list-style-type: none"> Performing sequential writes <p>It allows sequential writes from an unformatted file opened for direct access.</p> <p>Note: Allowing files that are not opened with sequential access (such as <code>ACCESS='DIRECT'</code>) to be read with the sequential <code>WRITE</code> statement violates the Fortran 95 standard and may be removed in the future.</p> <p>Specifying <code>fpscomp general</code> sets <code>fpscomp ldio_spacing</code>.</p>
	<p>Specifies that Fortran PowerStation semantic conventions and record formats should be used for list-directed formatted and unformatted I/O. The <code>fpscomp ioformat</code> option affects the following Fortran features:</p> <ul style="list-style-type: none"> The <code>WRITE</code> statement: For formatted list-directed <code>WRITE</code> statements, formatted internal list-directed <code>WRITE</code> statements, and formatted namelist <code>WRITE</code> statements, the output line, field width values, and the list-directed data type semantics are determined according to the following sample for real constants (N below): <ul style="list-style-type: none"> For $1 \leq N < 10^{**7}$, use <code>F15.6</code> for single precision or <code>F24.15</code> for double. For $N < 1$ or $N \geq 10^{**7}$, use <code>E15.6E2</code> for single precision or <code>E24.15E3</code> for double. <p>See the Fortran PowerStation documentation for more detailed information about the other data types affected.</p> For unformatted <code>WRITE</code> statements, the unformatted file semantics are dictated according to the Fortran PowerStation

Option	Description
--------	-------------

documentation; these semantics are different from the Intel Fortran file format. See the Fortran PowerStation documentation for more detailed information.

The following table summarizes the default output formats for list-directed output with the intrinsic data types:

Data Type	Output Format with fpscomp noioformat	Output Format with fpscomp ioformat
BYTE	I5	I12
LOGICAL (all)	L2	L2
INTEGER(1)	I5	I12
INTEGER(2)	I7	I12
INTEGER(4)	I12	I12
INTEGER(8)	I22	I22
REAL(4)	1PG15.7E2	1PG16.6E2
REAL(8)	1PG24.15E3	1PG25.15E3
COMPLEX(4)	(' ',1PG14.7E2, ' ' ',1PG14.7E2, ') '	(' ',1PG16.6E2, ' ' ',1PG16.6E2, ') '
COMPLEX(8)	(' ',1PG23.15E3, ' ' ',1PG23.15E3, ') '	(' ',1PG25.15E3, ' ' ',1PG25.15E3, ') '
CHARACTER	Aw	Aw

- The READ statement:
- For formatted list-directed READ statements, formatted internal list-directed READ statements, and formatted namelist READ statements, the field width values and the list-directed semantics are dictated according to the following

Option	Description
	<p>sample for real constants (N below):</p> <p>For $1 \leq N < 10^{**7}$, use F15.6 for single precision or F24.15 for double.</p> <p>For $N < 1$ or $N \geq 10^{**7}$, use E15.6E2 for single precision or E24.15E3 for double.</p> <p>See the Fortran PowerStation documentation for more detailed information about the other data types affected.</p> <ul style="list-style-type: none"> For unformatted READ statements, the unformatted file semantics are dictated according to the Fortran PowerStation documentation; these semantics are different from the Intel Fortran file format. See the Fortran PowerStation documentation for more detailed information.
<p><code>fpscomp</code> <code>nolib</code></p>	<p>Prevents the portability library from being passed to the linker.</p>
<p><code>fpscomp</code> <code>ldio_spacing</code></p>	<p>Specifies that at run time a blank should not be inserted after a numeric value before a character value (undelimited character string). This representation is used by Intel Fortran releases before Version 8.0 and by Fortran PowerStation. If you specify <code>fpscomp general</code>, it sets <code>fpscomp ldio_spacing</code>.</p>
<p><code>fpscomp</code> <code>logicals</code></p>	<p>Specifies that integers with a non-zero value are treated as true, integers with a zero value are treated as false. The literal constant <code>.TRUE.</code> has an integer value of 1, and the literal constant <code>.FALSE.</code> has an integer value of 0. This representation is used by Intel Fortran releases before Version 8.0 and by Fortran PowerStation. The default is <code>fpscomp nologicals</code>, which specifies that odd integer values (low bit one) are treated as true and even integer values (low bit zero) are treated as false. The literal constant <code>.TRUE.</code> has an integer value of -1, and</p>

Option	Description
--------	-------------

the literal constant `.FALSE.` has an integer value of 0. This representation is used by Compaq* Visual Fortran. The internal representation of LOGICAL values is not specified by the Fortran standard. Programs which use integer values in LOGICAL contexts, or which pass LOGICAL values to procedures written in other languages, are non-portable and may not execute correctly. Intel recommends that you avoid coding practices that depend on the internal representation of LOGICAL values. The `fpscomp logical` option affects the results of all logical expressions and affects the return value for the following Fortran features:

- The INQUIRE statement specifiers OPENED, IOFOCUS, EXISTS, and NAMED
- The EOF intrinsic function
- The BTEST intrinsic function
- The lexical intrinsic functions LLT, LLE, LGT, and LGE

`fpscomp all` Specifies that all options should be used for compatibility with Fortran PowerStation. This is the same as specifying `fpscomp` with no keyword. Option `fpscomp all` enables full compatibility with Fortran PowerStation.

Alternate Options

None

See Also

Building Applications: Microsoft Fortran PowerStation Compatible Files

FR

See [free](#).

fr32

Disables the use of the high floating-point registers.

IDE Equivalent

None

Architectures

IA-64 architecture

Syntax

Linux: -`fr32`

Mac OS X: None

Windows: None

Arguments

None

Default

OFF The use of the high floating-point registers is enabled.

Description

This option disables the use of the high floating-point registers. Only the lower 32 floating-point registers are used.

Alternate Options

None

free

Specifies source files are in free format.

IDE Equivalent

Windows: **Language > Source File Format** (/free, /fixed)

Linux: None

Mac OS X: **Language > Source File Format** (/free, /fixed)

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: -free
 -nofree

Windows: /free
 /nofree

Arguments

None

Default

OFF The source file format is determined from the file extension.

Description

This option specifies source files are in free format. If this option is not specified, format is determined as follows:

- Files with an extension of .f90, .F90, or .i90 are free-format source files.
- Files with an extension of .f, .for, .FOR, .ftn, or .i are fixed-format files.

Alternate Options

Linux and Mac OS X: -FR

Windows: /nofixed, /FR, /4Yf

See Also

[fixed](#) compiler option

fsource-asm

Produces an assembly listing with source code annotations.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-fsource-asm`

Windows: `None`

Arguments

None

Default

OFF No source code annotations appear in the assembly listing file, if one is produced.

Description

This option produces an assembly listing file with source code annotations. The assembly listing file shows the source code as interspersed comments.

To use this option, you must also specify option `-s`, which causes an assembly listing to be generated.

Alternate Options

Linux and Mac OS X: None

Windows: `/asmattr:source, /FAs`

See Also

[s](#) compiler option

fstack-security-check, GS

Determines whether the compiler generates code that detects some buffer overruns.

IDE Equivalent

None

Architectures

IA-32, Intel® 64 architectures

Syntax

Linux and Mac OS X: `-fstack-security-check`
`-fno-stack-security-check`

Windows: `/GS`
`/GS-`

Arguments

None

Default

`-fno-stack-security-check` The compiler does not detect buffer overruns.
 or `/GS-`

Description

This option determines whether the compiler generates code that detects some buffer overruns that overwrite the return address. This is a common technique for exploiting code that does not enforce buffer size restrictions.

The `/GS` option is supported with Microsoft Visual Studio .NET 2003* and Microsoft Visual Studio 2005*.

Alternate Options

Linux and Mac OS X: `-f[no-]stack-protector`

Windows: None

fstack-security-check, GS

Determines whether the compiler generates code that detects some buffer overruns.

IDE Equivalent

None

Architectures

IA-32, Intel® 64 architectures

Syntax

Linux and Mac OS X: `-fstack-security-check`
`-fno-stack-security-check`

Windows: `/GS`
`/GS-`

Arguments

None

Default

`-fno-stack-security-check` The compiler does not detect buffer overruns.
or `/GS-`

Description

This option determines whether the compiler generates code that detects some buffer overruns that overwrite the return address. This is a common technique for exploiting code that does not enforce buffer size restrictions.

The `/GS` option is supported with Microsoft Visual Studio .NET 2003* and Microsoft Visual Studio 2005*.

Alternate Options

Linux and Mac OS X: `-f[no-]stack-protector`

Windows: None

fsyntax-only

See [syntax-only](#).

ftrapuv, Qtrapuv

Initializes stack local variables to an unusual value to aid error detection.

IDE Equivalent

Windows: **Data > Initialize stack variables to an unusual value**

Linux: None

Mac OS X: **Run-Time > Initialize Stack Variables to an Unusual Value**

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-ftrapuv`

Windows: `/Qtrapuv`

Arguments

None

Default

OFF The compiler does not initialize local variables.

Description

This option initializes stack local variables to an unusual value to aid error detection. Normally, these local variables should be initialized in the application. The option sets any uninitialized local variables that are allocated on the stack to a value that is typically interpreted as a very large integer or an invalid address. References to these variables are then likely to cause run-time errors that can help you detect coding errors.

This option sets option `-g` (Linux and Mac OS X) and `/zi` or `/z7` (Windows).

Alternate Options

None

See Also

[g, zi, z7](#) compiler options

ftz, Qftz

Flushes denormal results to zero.

IDE Equivalent

Windows: (IA-32 and IA-64 architectures): **Floating Point > Flush Denormal**

Results to Zero

(Intel® 64 architecture): None

Linux: None

Mac OS X: **Floating Point > Flush Denormal Results to Zero**

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-ftz`
 `-no-ftz`

Windows: `/Qftz`
 `/Qftz-`

Arguments

None

Default

Systems using IA-64 architecture: – On systems using IA-64
`no-ftz` or `/Qftz-` architecture, the compiler lets
 Systems using IA-32 architecture results gradually underflow. On
 and Intel® 64 architecture: `-ftz` or systems using IA-32 architecture
`/Qftz` and Intel® 64 architecture,
 denormal results are flushed to
 zero.

Description

This option flushes denormal results to zero when the application is in the gradual underflow mode. It may improve performance if the denormal values are not critical to your application's behavior.

This option sets or resets the FTZ and the DAZ hardware flags. If FTZ is ON, denormal results from floating-point calculations will be set to the value zero. If FTZ is OFF, denormal results remain as is. If DAZ is ON, denormal values used as input to floating-point instructions will be treated as zero. If DAZ is OFF, denormal instruction inputs remain as is. Systems using IA-64 architecture have FTZ but not DAZ. Systems using Intel® 64 architecture have both FTZ and DAZ. FTZ and DAZ are not supported on all IA-32 architectures.

When `-ftz` (Linux and Mac OS X) or `/Qftz` (Windows) is used in combination with an SSE-enabling option on systems using IA-32 architecture (for example, `xN` or `QxN`), the compiler will insert code in the main routine to set FTZ and DAZ.

When `-ftz` or `/Qftz` is used without such an option, the compiler will insert code to conditionally set FTZ/DAZ based on a run-time processor check. `-no-ftz` (Linux and Mac OS X) or `/Qftz-` (Windows) will prevent the compiler from inserting any code that might set FTZ or DAZ.

This option only has an effect when the main program is being compiled. It sets the FTZ/DAZ mode for the process. The initial thread and any threads subsequently created by that process will operate in FTZ/DAZ mode.

Options `-fpe0` and `-fpe1` (Linux and Mac OS X) set `-ftz`. Options `/fpe:0` and `/fpe:1` (Windows) set `/Qftz`.

On systems using IA-64 architecture, optimization option `O3` sets `-ftz` and `/Qftz`; optimization option `O2` sets `-no-ftz` (Linux) and `/Qftz-` (Windows).

On systems using IA-32 architecture and Intel® 64 architecture, every optimization option `O` level, except `O0`, sets `-ftz` and `/Qftz`.

If this option produces undesirable results of the numerical behavior of your program, you can turn the FTZ/DAZ mode off by using `-no-ftz` or `/Qftz-` in the command line while still benefiting from the `O3` optimizations.



Note

Options `-ftz` and `/Qftz` are performance options. Setting these options does not guarantee that all denormals in a program are flushed to zero. They only cause denormals generated at run time to be flushed to zero.

Alternate Options

None

See Also

[x, Qx](#) compiler option

Floating-point Operations: Using the `-fpe` or `/fpe` Compiler Option

Intrinsics Reference:

func-groups

This is a deprecated option. See [prof-func-groups](#).

funroll-loops

See [unroll](#), [Qunroll](#).

fverbose-asm

Produces an assembly listing with compiler comments, including options and version information.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-fverbose-asm`
`-fno-verbose-asm`

Windows: None

Arguments

None

Default

<code>-fno-verbose-asm</code>	No source code annotations appear in the assembly listing file, if one is produced.
-------------------------------	---

Description

This option produces an assembly listing file with compiler comments, including options and version information.

To use this option, you must also specify `-S`, which sets `-fverbose-asm`.

If you do not want this default when you specify `-S`, specify `-fno-verbose-asm`.

Alternate Options

None

See Also

[s](#) compiler option

fvisibility

Specifies the default visibility for global symbols or the visibility for symbols in a file.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-fvisibility=keyword`
`-fvisibility-keyword=file`

Windows: None

Arguments

keyword Specifies the visibility setting. Possible values are:

`default` Sets visibility to default.

`extern` Sets visibility to extern.

`hidden` Sets visibility to hidden.

`internal` Sets visibility to internal.

`protected` Sets visibility to protected.

file Is the pathname of a file containing the list of symbols whose visibility you want to set. The

symbols must be separated by whitespace
(spaces, tabs, or newlines).

Default

`-fvisibility=default` The compiler sets visibility of symbols to default.

Description

This option specifies the default visibility for global symbols (syntax `-fvisibility=keyword`) or the visibility for symbols in a file (syntax `-fvisibility-keyword=file`).

Visibility specified by `-fvisibility-keyword=file` overrides visibility specified by `-fvisibility=keyword` for symbols specified in a file.

Option	Description
<code>-fvisibility=default</code> <code>-fvisibility-default=file</code>	Sets visibility of symbols to default. This means other components can reference the symbol, and the symbol definition can be overridden (preempted) by a definition of the same name in another component.
<code>-fvisibility=extern</code> <code>-fvisibility-extern=file</code>	Sets visibility of symbols to extern. This means the symbol is treated as though it is defined in another component. It also means that the symbol can be overridden by a definition of the same name in another component.
<code>-fvisibility=hidden</code> <code>-fvisibility-hidden=file</code>	Sets visibility of symbols to hidden. This means that other components cannot directly reference the symbol. However, its address may be passed to other

Option	Description
	components indirectly.
<code>-fvisibility=internal</code> <code>-fvisibility-internal=<i>file</i></code>	Sets visibility of symbols to internal. This means the symbol cannot be referenced outside its defining component, either directly or indirectly.
<code>-fvisibility=protected</code> <code>-fvisibility-protected=<i>file</i></code>	CELL_TEXT

If an `-fvisibility` option is specified more than once on the command line, the last specification takes precedence over any others.

If a symbol appears in more than one visibility *file*, the setting with the least visibility takes precedence.

The following shows the precedence of the visibility settings (from greatest to least visibility):

- `extern`
- `default`
- `protected`
- `hidden`
- `internal`

Note that `extern` visibility only applies to functions. If a variable symbol is specified as `extern`, it is assumed to be `default`.

Alternate Options

None

Example

A file named `prot.txt` contains symbols a, b, c, d, and e. Consider the following:

```
-fvisibility-protected=prot.txt
```

This option sets `protected` visibility for all the symbols in the file. It has the same effect as specifying `fvisibility=protected` in the declaration for each of the symbols.

See Also

Optimizing Applications: Symbol Visibility Attribute Options (Linux* and Mac OS* X)

g, Zi, Z7

Tells the compiler to generate full debugging information in the object file.

IDE Equivalent

Windows: **General > Debug Information Format** (`/Z7`, `/Zd`, `/Zi`)

Linux: None

Mac OS X: **General > Generate Debug Information** (`-g`)

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-g`

Windows: `/Zi`

`/Z7`

Arguments

None

Default

OFF No debugging information is produced in the object file.

Description

This option tells the compiler to generate symbolic debugging information in the object file for use by debuggers.

The compiler does not support the generation of debugging information in assemblable files. If you specify this option, the resulting object file will contain debugging information, but the assemblable file will not.

This option turns off `o2` and makes `o0` (Linux and Mac OS X) or `od` (Windows) the default unless `o2` (or another `o` option) is explicitly specified in the same command line.

On Linux systems using Intel® 64 architecture and Linux and Mac OS X systems using IA-32 architecture, specifying the `-g` or `-O0` option sets the `-fno-omit-frame-pointer` option.

For more information on `zi` and `z7`, see keyword `full` in [debug \(Windows*\)](#).

Alternate Options

Linux and Mac OS X: None

Windows: `/debug:full` (or `/debug`)

See Also

[zd](#) compiler option

G2, G2-p9000

Optimizes application performance for systems using IA-64 architecture.

IDE Equivalent

Windows: **Optimization > Optimize For Intel® Processor**

Linux: None

Mac OS X: None

Architectures

IA-64 architecture

Syntax

Linux and Mac OS X: None

Windows: /G2
 /G2-p9000

Arguments

None

Default

/G2- Performance is optimized for Dual-Core Intel® Itanium® 2
 p9000 processor 9000 series.

Description

These options optimize application performance for a particular Intel® processor or family of processors. The compiler generates code that takes advantage of features of IA-64 architecture.

Option	Description
G2	Optimizes for Intel® Itanium® 2 processors.
G2-p9000	Optimizes for Dual-Core Intel® Itanium® 2 processor 9000 series. This option affects the order of the generated instructions, but the generated instructions are limited to Intel® Itanium® 2 processor instructions unless the program specifies and executes intrinsics specific to the Dual-Core Intel® Itanium® 2 processor 9000 series.

The resulting executable is backwards compatible and generated code is optimized for specific processors. For example, code generated with /G2-p9000

will run correctly on single-core Itanium® 2 processors, but it might not run as fast as if it had been generated using `/G2`.

Alternate Options

<code>/G2</code>	Linux: <code>-mtune=itanium2</code> Mac OS X: None Windows: None
<code>/G2-p9000</code>	Linux: <code>-mtune=itanium2-p9000, -mcpu=itanium2-p9000</code> (<code>-mcpu</code> is a deprecated option) Mac OS X: None Windows: None

Example

In the following example, the compiled binary of the source program `prog.f` is optimized for the Dual-Core Intel® Itanium® 2 processor 9000 series by default. The same binary will also run on single-core Itanium® 2 processors (unless the program specifies and executes intrinsics specific to the Dual-Core Intel® Itanium® 2 processor 9000 series). All lines in the code example are equivalent.

```
ifort prog.f  
ifort /G2-p9000 prog.f
```

In the following example, the compiled binary is optimized for single-core Itanium® 2 processors:

```
ifort /G2 prog.f
```

See Also

[mtune](#) compiler option

G5, G6, G7

Optimize application performance for systems using IA-32 architecture and Intel® 64 architecture. These are [deprecated](#) options.

IDE Equivalent

Windows: **Optimization > Optimize For Intel(R) Processor** (`/GB`, `/G5`, `/G6`, `/G7`)

Linux: None

Mac OS X: None

Architectures

IA-32, Intel® 64 architectures

Syntax

Linux and Mac OS X: None

Windows: /G5
 /G6
 /G7

Arguments

None

Default

/G7 On systems using IA-32 architecture and Intel® 64 architecture, performance is optimized for Intel® Pentium® 4 processors, Intel® Xeon® processors, Intel® Pentium® M processors, and Intel® Pentium® 4 processors with Streaming SIMD Extensions 3 (SSE3) instruction support.

Description

These options optimize application performance for a particular Intel® processor or family of processors. The compiler generates code that takes advantage of features of the specified processor.

Option	Description
G5	Optimizes for Intel® Pentium® and Pentium® with MMX™ technology processors.
G6	Optimizes for Intel® Pentium® Pro, Pentium® II and Pentium® III

Option	Description
	processors.
G7	Optimizes for Intel® Core™ Duo processors, Intel® Core™ Solo processors, Intel® Pentium® 4 processors, Intel® Xeon® processors based on the Intel® Core microarchitecture, Intel® Pentium® M processors, and Intel® Pentium® 4 processors with Streaming SIMD Extensions 3 (SSE3) instruction support.

- G7 Optimizes for Intel® Core™ Duo processors, Intel® Core™ Solo processors, Intel® Pentium® 4 processors, Intel® Xeon® processors based on the Intel® Core microarchitecture, Intel® Pentium® M processors, and Intel® Pentium® 4 processors with Streaming SIMD Extensions 3 (SSE3) instruction support.

On systems using Intel® 64 architecture, only option G7 is valid.

These options always generate code that is backwards compatible with Intel processors of the same architecture. For example, code generated with the G7 option runs correctly on Pentium III processors, although performance may be faster on Pentium III processors when compiled using or G6.

Alternate Options

Windows: /GB (an alternate for /G6; this option is also [deprecated](#))

Linux: None

Example

In the following example, the compiled binary of the source program prog.f is optimized, by default, for Intel® Pentium® 4 processors, Intel® Xeon® processors, Intel® Pentium® M processors, and Intel® Pentium® 4 processors with Streaming SIMD Extensions 3 (SSE3). The same binary will also run on Pentium, Pentium Pro, Pentium II, and Pentium III processors. All lines in the code example are equivalent.

```
ifort prog.f
ifort /G7 prog.f
```

In the following example, the compiled binary is optimized for Pentium processors and Pentium processors with MMX technology:

```
ifort /G5 prog.f
icl /G5 prog.c
```

See Also

[mtune](#) compiler option

gdwarf-2

Enables generation of debug information using the DWARF2 format.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-gdwarf-2`

Windows: None

Arguments

None

Default

OFF No debug information is generated. However, if compiler option `-g` is specified, debug information is generated in the latest DWARF format, which is currently DWARF2.

Description

This option enables generation of debug information using the DWARF2 format. This is currently the default when compiler option `-g` is specified.

Alternate Options

None

See Also

[g](#) compiler option

Ge

Enables stack-checking for all functions.

This option has been [deprecated](#).

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: None

Windows: /Ge

Arguments

None

Default

OFF Stack-checking for all functions is disabled.

Description

This option enables stack-checking for all functions.

Alternate Options

None

gen-interfaces

Tells the compiler to generate an interface block for each routine in a source file.

IDE Equivalent

Windows: **Diagnostics > Generate Interface Blocks**

Linux: None

Mac OS X: **Diagnostics > Generate Interface Blocks**

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-gen-interfaces` `[[no]source]`
`-nogen-interfaces`

Windows: `/gen-interfaces[:[no]source]`
`/nogen-interfaces`

Arguments

None

Default

`nogen-` The compiler does not generate interface blocks for `interfaces` routines in a source file.

Description

This option tells the compiler to generate an interface block for each routine (that is, for each SUBROUTINE and FUNCTION statement) defined in the source file. The compiler generates two files for each routine, a .mod file and a .f90 file, and places them in the current directory or in the directory specified by the include (`-I`) or `-module` option. The .f90 file is the text of the interface block; the .mod file is the interface block compiled into binary form.

If `source` is specified, the compiler creates the `*_mod.f90` as well as the `*_mod.mod` files. If `nosource` is specified, the compiler creates the `*_mod.mod` but not the `*_mod.f90` files. If neither is specified, it is the same as specifying `-gen-interfaces source` (Linux and Mac OS X) or `/gen-interfaces:source` (Windows).

Alternate Options

None

global-hoist, Qglobal-hoist

Enables certain optimizations that can move memory loads to a point earlier in the program execution than where they appear in the source.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-global-hoist`
`-no-global-hoist`
Windows: `/Qglobal-hoist`
`/Qglobal-hoist-`

Arguments

None

Default

`-global-hoist` Certain optimizations are enabled that can move memory loads.
or `/Qglobal-hoist`

Description

This option enables certain optimizations that can move memory loads to a point earlier in the program execution than where they appear in the source. In most cases, these optimizations are safe and can improve performance.

The `-no-global-hoist` (Linux and Mac OS X) or `/Qglobal-hoist-` (Windows) option is useful for some applications, such as those that use shared or dynamically mapped memory, which can fail if a load is moved too early in the execution stream (for example, before the memory is mapped).

Alternate Options

None

Gm

See keyword `cvf` in [iface](#).

Gs

Disables stack-checking for routines with more than a specified number of bytes of local variables and compiler temporaries.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: None

Windows: `/Gs [n]`

Arguments

n Is the number of bytes of local variables and compiler temporaries.

Default

4096 Stack checking is disabled for routines with more than 4KB of stack space allocated.

Description

This option disables stack-checking for routines with n or more bytes of local variables and compiler temporaries. If you do not specify n , you get the default of 4096.

Alternate Options

None

fstack-security-check, GS

Determines whether the compiler generates code that detects some buffer overruns.

IDE Equivalent

None

Architectures

IA-32, Intel® 64 architectures

Syntax

Linux and Mac OS X: `-fstack-security-check`
`-fno-stack-security-check`

Windows: `/GS`
`/GS-`

Arguments

None

Default

`-fno-stack-security-check` or `/GS-` The compiler does not detect buffer overruns.

Description

This option determines whether the compiler generates code that detects some buffer overruns that overwrite the return address. This is a common technique for exploiting code that does not enforce buffer size restrictions.

The `/GS` option is supported with Microsoft Visual Studio .NET 2003* and Microsoft Visual Studio 2005*.

Alternate Options

Linux and Mac OS X: `-f[no-]stack-protector`

Windows: None

Gz

See keyword `stdcall` in [iface](#)

heap-arrays

Puts automatic arrays and arrays created for temporary computations on the heap instead of the stack.

IDE Equivalent

Windows: **Optimization > Heap Arrays**

Linux: None

Mac OS X: **Optimization > Heap Arrays**

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-heap-arrays [size]`
`-no-heap-arrays`

Windows: `/heap-arrays[:size]`
`/heap-arrays-`

Arguments

size Is an integer value representing the size of the arrays in kilobytes. Any arrays known at compile-time to be larger than *size* are allocated on the heap instead of the stack.

Default

`-no-heap-arrays` The compiler puts automatic arrays and arrays created for temporary computations in temporary storage in the stack storage area.

or

`/heap-arrays-`

Description

This option puts automatic arrays and arrays created for temporary computations on the heap instead of the stack.

If `heap-arrays` is specified and *size* is omitted, all automatic and temporary arrays are put on the heap. If 10 is specified for *size*, all automatic and temporary arrays larger than 10 KB are put on the heap.

Alternate Options

None

Example

In Fortran, an automatic array gets its size from a run-time expression. For example:

```

RECURSIVE SUBROUTINE F( N )
INTEGER :: N
REAL :: X ( N )      ! an automatic array
REAL :: Y ( 1000 )  ! an explicit-shape local array on the stack

```

Array X in the example above is affected by the `heap-array` option. Array Y is not.

help

Displays all available compiler options or a category of compiler options.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-help[category]`

Windows: `/help[category]`

Arguments

category Is a category or class of options to display.

Possible values are:

<code>advanced</code>	Displays advanced optimization options that allow fine tuning of compilation or allow control over advanced features of the compiler.
<code>codegen</code>	Displays Code Generation options.
<code>compatibility</code>	Displays options affecting

	language compatibility.
<code>component</code>	Displays options for component control.
<code>data</code>	Displays options related to interpretation of data in programs or the storage of data.
<code>deprecated</code>	Displays options that have been deprecated.
<code>diagnostics</code>	Displays options that affect diagnostic messages displayed by the compiler.
<code>float</code>	Displays options that affect floating-point operations.
<code>help</code>	Displays all the available help categories.
<code>inline</code>	Displays options that affect inlining.
<code>ipo</code>	Displays Interprocedural Optimization (IPO) options
<code>language</code>	Displays options affecting the behavior of the compiler language features.
<code>link</code>	Displays linking or linker options.
<code>misc</code>	Displays miscellaneous

	options that do not fit within other categories.
<code>openmp</code>	Displays OpenMP and parallel processing options.
<code>opt</code>	Displays options that help you optimize code.
<code>output</code>	Displays options that provide control over compiler output.
<code>pgo</code>	Displays Profile Guided Optimization (PGO) options.
<code>preproc</code>	Displays options that affect preprocessing operations.
<code>reports</code>	Displays options for optimization reports.

Default

OFF No list is displayed unless this compiler option is specified.

Description

This option displays all available compiler options or a category of compiler options. If category is not specified, all available compiler options are displayed.

Alternate Options

Linux and Mac OS X: None

Windows: /?

homeparams

Tells the compiler to store parameters passed in registers to the stack.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: None

Windows: /homeparams

Arguments

None

Default

OFF Register parameters are not written to the stack.

Description

This option tells the compiler to store parameters passed in registers to the stack.

Alternate Options

None

I

Specifies an additional directory for the include path.

IDE Equivalent

Windows: **General > Additional Include Directories** (/include)

Preprocessor > Additional Include Directories (/include)

Linux: None

Mac OS X: Preprocessor > Additional Include Directories (/include)

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-Idir`

Windows: `/Idir`

Arguments

dir Is the directory to add to the include path.

Default

OFF The default include path is used.

Description

This option specifies an additional directory for the include path, which is searched for module files referenced in USE statements and include files referenced in INCLUDE statements. To specify multiple directories on the command line, repeat the option for each directory you want to add.

For all USE statements and for those INCLUDE statements whose file name does not begin with a device or directory name, the directories are searched in this order:

1. The directory containing the first source file.

Note that if `assume nosource_include` is specified, this directory will not be searched.

2. The current working directory where the compilation is taking place (if different from the above directory).

3. Any directory or directories specified using the I option. If multiple directories are specified, they are searched in the order specified on the command line, from left to right.
4. On Linux and Mac OS X systems, any directories indicated using environment variable FPATH. On Windows systems, any directories indicated using environment variable INCLUDE.

This option affects fpp preprocessor behavior and the USE statement.

Alternate Options

Linux and Mac OS X: None

Windows: `/include`

See Also

[x](#) compiler option

[assume](#) compiler option

i-dynamic

This is a deprecated option. See [shared-intel](#).

i-static

This is a deprecated option. See [static-intel](#).

i2, i4, i8

See [integer-size](#).

idirafter

Adds a directory to the second include file search path.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-idirafterdir`

Windows: None

Arguments

dir Is the name of the directory to add.

Default

OFF Include file search paths include certain default directories.

Description

This option adds a directory to the second include file search path (after `-I`).

Alternate Options

None

iface

Specifies the default calling convention and argument-passing convention for an application.

IDE Equivalent

Windows: **External Procedures > Calling Convention**

`(/iface:{cref|stdref|stdcall|cvf|default})`

External Procedures > String Length Argument Passing

`(/iface:[no]mixed_str_len_arg)`

Linux: None

Mac OS X: None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: None

Windows: `/iface:keyword`

Arguments

keyword Specifies the calling convention or the argument-passing convention. Possible values are:

<code>default</code>	Tells the compiler to use the default calling conventions.
<code>cref</code>	Tells the compiler to use calling conventions C, REFERENCE.
<code>cvf</code>	Tells the compiler to use calling conventions compatible with Compaq Visual Fortran*.
<code>[no]mixed_str_len_arg</code>	Determines the argument-passing convention for hidden-length character arguments.
<code>stdcall</code>	Tells the compiler to use calling convention STDCALL.
<code>stdref</code>	Tells the compiler to use calling conventions STDCALL, REFERENCE.

Default

`/iface:default` The default calling convention is used.

Description

This option specifies the default calling convention and argument-passing convention for an application.

The aspects of calling and argument passing controlled by this option are as follows:

- The calling mechanism (C or STDCALL): On IA-32 architecture, these mechanisms differ in how the stack register is adjusted when a procedure call returns. On Intel® 64 and IA-64 architectures, the only calling mechanism available is C; requests for the STDCALL mechanism are ignored.
- The argument passing mechanism (by value or by reference)
- Character-length argument passing (at the end of the argument list or after the argument address)
- The case of external names (uppercase or lowercase)
- The name decoration (prefix and suffix)

You can also use the ATTRIBUTES compiler directive to modify these conventions on an individual basis. Note that the effects of the ATTRIBUTES directive do not always match that of the `iface` option of the same name.

Option	Description
<code>/iface:default</code>	<p>Tells the compiler to use the default calling conventions. These conventions are as follows:</p> <ul style="list-style-type: none"> • The calling mechanism: C • The argument passing mechanism: by reference • Character-length argument passing: at end of argument list • The external name case: uppercase • The name decoration: Underscore prefix on

Option	Description
	IA-32 architecture, no prefix on Intel® 64 or IA-64 architecture; no suffix
<code>/iface:cref</code>	Tells the compiler to use the same conventions as <code>/iface:default</code> except that external names are lowercase.
<code>/iface:cvf</code>	<p>Tells the compiler to use calling conventions compatible with Compaq Visual Fortran* and Microsoft Fortran PowerStation. These conventions are as follows:</p> <ul style="list-style-type: none"> • The calling mechanism: STDCALL • The argument passing mechanism: by reference • Character-length argument passing: following the argument address • The external name case: uppercase • The name decoration: Underscore prefix on IA-32 architecture, no prefix on Intel® 64 or IA-64 architecture. On Windows* systems using IA-32 architecture, @n suffix where <i>n</i> is the number of bytes to be removed from the stack on exit from the procedure. No suffix on other systems.
<code>/iface:mixed_str_len_arg</code>	Specifies argument-passing conventions for hidden-length character arguments. This option tells the compiler that the hidden length passed for a character argument is to be placed immediately after its corresponding character argument in the argument list.

Option	Description
	<p>This is the method used by Compaq Visual Fortran*. When porting mixed-language programs that pass character arguments, either this option must be specified correctly or the order of hidden length arguments must be changed in the source code. This option can be used in addition to other <code>/iface</code> options.</p>
<code>/iface:stdcall</code>	<p>Tells the compiler to use the following conventions:</p> <ul style="list-style-type: none"> • The calling mechanism: STDCALL • The argument passing mechanism: by value • Character-length argument passing: at the end of the argument list • The external name case: uppercase • The name decoration: Underscore prefix on IA-32 architecture, no prefix on Intel® 64 or IA-64 architecture. On Windows* systems using IA-32 architecture, @n suffix where <i>n</i> is the number of bytes to be removed from the stack on exit from the procedure. No suffix on other systems.
<code>/iface:stdref</code>	<p>Tells the compiler to use the same conventions as <code>/iface:stdcall</code> except that argument passing is by reference.</p>



Caution

On Windows systems, if you specify option `/iface:ceref`, it overrides the default for external names and causes them to be lowercase. It is as if you specified "!dec\$ attributes c, reference" for the external name.

If you specify option `/iface:cref` and want external names to be uppercase, you must explicitly specify option `/names:uppercase`.



On systems using IA-32 architecture, there must be agreement between the calling program and the called procedure as to which calling mechanism (C or STDCALL) is used or unpredictable errors may occur. If you change the default mechanism to STDCALL, you must use the ATTRIBUTES DEFAULT directive to reset the calling conventions for routines specified with the USEROPEN keyword in an OPEN statement and for comparison routines passed to the QSORT library routine.

Alternate Options

<code>/iface:cvf</code>	Linux and Mac OS X: None Windows: <code>/Gm</code>
<code>/iface:mixed_str_len_arg</code>	Linux and Mac OS X: <code>-mixed-str-len-arg</code> Windows: None
<code>/iface:nomixed_str_len_arg</code>	Linux and Mac OS X: <code>-nomixed-str-len-arg</code> Windows: None
<code>/iface:stdcall</code>	Linux and Mac OS X: None Windows: <code>/Gz</code>

See Also

Building Applications: Programming with Mixed Languages Overview

Language Reference: ATTRIBUTES

implicitnone

See [warn](#).

include

See [1](#).

inline

Specifies the level of inline function expansion.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: None

Windows: `/inline[:keyword]`

Arguments

keyword Is the level of inline function expansion. Possible values are:

<code>none</code>	Disables inlining of user-defined functions. This is the same as specifying <code>manual</code> .
<code>manual</code>	Disables inlining of user-defined functions. Fortran statement functions are always inlined.
<code>size</code>	Enables inlining of any function. However, the compiler decides which functions are inlined. This option enables interprocedural optimizations and most speed optimizations.

<code>speed</code>	Enables inlining of any function. This is the same as specifying <code>all</code> .
<code>all</code>	Enables inlining of any function. However, the compiler decides which functions are inlined. This option enables interprocedural optimizations and all speed optimizations. This is the same as specifying <code>inline</code> with no <i>keyword</i> .

Default

OFF The compiler inlines certain functions by default.

Description

This option specifies the level of inline function expansion.

Alternate Options

<code>inline all</code> or <code>inline speed</code>	Linux and Mac OS X: None Windows: <code>/Ob2 /Ot</code>
<code>inline size</code>	Linux and Mac OS X: None Windows: <code>/Ob2 /Os</code>
<code>inline manual</code>	Linux and Mac OS X: None Windows: <code>/Ob0</code>
<code>inline none</code>	Linux and Mac OS X: None Windows: <code>/Ob0</code>

See Also

[finline-functions](#) compiler option

inline-debug-info, Qinline-debug-info

Produces enhanced source position information for inlined code. This is a deprecated option.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-inline-debug-info`

Windows: `/Qinline-debug-info`

Arguments

None

Default

OFF No enhanced source position information is produced for inlined code.

Description

This option produces enhanced source position information for inlined code. This leads to greater accuracy when reporting the source location of any instruction. It also provides enhanced debug information useful for function call traceback.

To use this option for debugging, you must also specify a debug enabling option, such as `-g` (Linux) or `/debug` (Windows).

Alternate Options

Linux and Mac OS X: `-debug inline-debug-info`

Windows: None

inline-factor, Qinline-factor

Specifies the percentage multiplier that should be applied to all inlining options that define upper limits.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-inline-factor=n`
`-no-inline-factor`

Windows: `/Qinline-factor=n`
`/Qinline-factor-`

Arguments

n Is a positive integer specifying the percentage value. The default value is 100 (a factor of 1).

Default

`-no-inline-factor` or `/Qinline-factor-` The compiler uses default heuristics for inline routine expansion.

Description

This option specifies the percentage multiplier that should be applied to all inlining options that define upper limits:

- `-inline-max-size` and `/Qinline-max-size`
- `-inline-max-total-size` and `/Qinline-max-total-size`

- `-inline-max-per-routine` and `/Qinline-max-per-routine`
- `-inline-max-per-compile` and `/Qinline-max-per-compile`

This option takes the default value for each of the above options and multiplies it by n divided by 100. For example, if 200 is specified, all inlining options that define upper limits are multiplied by a factor of 2. This option is useful if you do not want to individually increase each option limit.

If you specify `-no-inline-factor` (Linux and Mac OS X) or `/Qinline-factor-` (Windows), the following occurs:

- Every function is considered to be a small or medium function; there are no large functions.
- There is no limit to the size a routine may grow when inline expansion is performed.
- There is no limit to the number of times some routine may be inlined into a particular routine.
- There is no limit to the number of times inlining can be applied to a compilation unit.

To see compiler values for important inlining limits, specify compiler option `-opt-report` (Linux and Mac OS X) or `/Qopt-report` (Windows).



Caution

When you use this option to increase default limits, the compiler may do so much additional inlining that it runs out of memory and terminates with an "out of memory" message.

Alternate Options

None

See Also

[inline-max-size, Qinline-max-size](#) compiler option

[inline-max-total-size, Qinline-max-total-size](#) compiler option

[inline-max-per-routine, Qinline-max-per-routine](#) compiler option

[inline-max-per-compile, Qinline-max-per-compile](#) compiler option
[opt-report, Qopt-report](#) compiler option

Optimizing Applications:

Developer Directed Inline Expansion of User Functions

Compiler Directed Inline Expansion of User Functions

inline-forceinline, Qinline-forceinline

Specifies that an inline routine should be inlined whenever the compiler can do so.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-inline-forceinline`

Windows: `/Qinline-forceinline`

Default

OFF The compiler uses default heuristics for inline routine expansion.

Description

This option specifies that a inline routine should be inlined whenever the compiler can do so. This causes the routines marked with an inline keyword or directive to be treated as if they were "forceinline".



Note

Because C++ member functions whose definitions are included in the class declaration are considered inline functions by default, using this option will also make these member functions "forceinline" functions.

The "forceinline" condition can also be specified by using the directive `cDEC$ ATTRIBUTES FORCEINLINE`.

To see compiler values for important inlining limits, specify compiler option `-opt-report` (Linux and Mac OS) or `/Qopt-report` (Windows).



When you use this option to change the meaning of inline to "forceinline", the compiler may do so much additional inlining that it runs out of memory and terminates with an "out of memory" message.

Alternate Options

None

See Also

[opt-report](#), [Qopt-report](#) compiler option

Optimizing Applications:

Developer Directed Inline Expansion of User Functions

Compiler Directed Inline Expansion of User Functions

inline-level, Ob

Specifies the level of inline function expansion.

IDE Equivalent

Windows: **Optimization > Inline Function Expansion**

Linux: None

Mac OS X: **Optimization > Inline Function Expansion**

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-inline-level=n`

Windows: `/Obn`

Arguments

n Is the inline function expansion level. Possible values are 0, 1, and 2.

Default

`-inline-level=2` or `/Ob2` This is the default if option `O2` is specified or is in effect by default. On Windows systems, this is also the default if option `O3` is specified.

`-inline-level=0` or `/Ob0` This is the default if option `-O0` (Linux and Mac OS) or `/Od` (Windows) is specified.

Description

This option specifies the level of inline function expansion. Inlining procedures can greatly improve the run-time performance of certain programs.

Option	Description
<code>-inline-level=0</code> or <code>Ob0</code>	Disables inlining of user-defined functions. Note that statement functions are always inlined.
<code>-inline-level=1</code> or <code>Ob1</code>	Enables inlining when an inline keyword or an inline directive is specified.
<code>-inline-level=2</code> or <code>Ob2</code>	Enables inlining of any function at the compiler's discretion.

Alternate Options

Linux: `-Ob` (this is a [deprecated](#) option)

Mac OS X: None

Windows: None

See Also

[inline](#) compiler option

inline-max-per-compile, Qinline-max-per-compile

Specifies the maximum number of times inlining may be applied to an entire compilation unit.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-inline-max-per-compile=n`
`-no-inline-max-per-compile`

Windows: `/Qinline-max-per-compile=n`
`/Qinline-max-per-compile-`

Arguments

n Is a positive integer that specifies the number of times inlining may be applied.

Default

`-no-inline-max-per-compile` The compiler uses default heuristics for inline routine expansion.

or `/Qinline-max-per-compile-`

Description

This option the maximum number of times inlining may be applied to an entire compilation unit. It limits the number of times that inlining can be applied.

For compilations using Interprocedural Optimizations (IPO), the entire compilation is a compilation unit. For other compilations, a compilation unit is a file.

If you specify `-no-inline-max-per-compile` (Linux and Mac OS X) or `/Qinline-max-per-compile-` (Windows), there is no limit to the number of times inlining may be applied to a compilation unit.

To see compiler values for important inlining limits, specify compiler option `-opt-report` (Linux and Mac OS X) or `/Qopt-report` (Windows).



Caution

When you use this option to increase the default limit, the compiler may do so much additional inlining that it runs out of memory and terminates with an "out of memory" message.

Alternate Options

None

See Also

[inline-factor, Qinline-factor](#) compiler option

[opt-report, Qopt-report](#) compiler option

Optimizing Applications:

Developer Directed Inline Expansion of User Functions

Compiler Directed Inline Expansion of User Functions

inline-max-per-routine, Qinline-max-per-routine

Specifies the maximum number of times the inliner may inline into a particular routine.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-inline-max-per-routine=n`
`-no-inline-max-per-routine`

Windows: `/Qinline-max-per-routine=n`
`/Qinline-max-per-routine-`

Arguments

n Is a positive integer that specifies the maximum number of times the inliner may inline into a particular routine.

Default

`-no-inline-max-per-routine` The compiler uses default heuristics for inline routine expansion.

or `/Qinline-max-per-routine-`

Description

This option specifies the maximum number of times the inliner may inline into a particular routine. It limits the number of times that inlining can be applied to any routine.

If you specify `-no-inline-max-per-routine` (Linux and Mac OS X) or `/Qinline-max-per-routine-` (Windows), there is no limit to the number of times some routine may be inlined into a particular routine.

To see compiler values for important inlining limits, specify compiler option `-opt-report` (Linux and Mac OS X) or `/Qopt-report` (Windows).

To see compiler values for important inlining limits, specify compiler option `-opt-report` (Linux and Mac OS X) or `/Qopt-report` (Windows).



When you use this option to increase the default limit, the compiler may do so much additional inlining that it runs out of memory and terminates with an "out of memory" message.

Alternate Options

None

See Also

[inline-factor, Qinline-factor](#) compiler option

[opt-report, Qopt-report](#) compiler option

Optimizing Applications:

Developer Directed Inline Expansion of User Functions

Compiler Directed Inline Expansion of User Functions

inline-max-size, Qinline-max-size

Specifies the lower limit for the size of what the inliner considers to be a large routine.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-inline-max-size=n`
`-no-inline-max-size`

Windows: `/Qinline-max-size=n`

`/Qinline-max-size-`

Arguments

n Is a positive integer that specifies the minimum size of what the inliner considers to be a large routine.

Default

`-no-inline-max-size` The compiler uses default heuristics for
or `/Qinline-max-size-` inline routine expansion.

Description

This option specifies the lower limit for the size of what the inliner considers to be a large routine (a function or subroutine). The inliner classifies routines as small, medium, or large. This option specifies the boundary between what the inliner considers to be medium and large-size routines.

The inliner prefers to inline small routines. It has a preference against inlining large routines. So, any large routine is highly unlikely to be inlined.

If you specify `-no-inline-max-size` (Linux and Mac OS X) or `/Qinline-max-size-` (Windows), there are no large routines. Every routine is either a small or medium routine.

To see compiler values for important inlining limits, specify compiler option `-opt-report` (Linux and Mac OS X) or `/Qopt-report` (Windows).

To see compiler values for important inlining limits, specify compiler option `-opt-report` (Linux and Mac OS X) or `/Qopt-report` (Windows).



Caution

When you use this option to increase the default limit, the compiler may do so much additional inlining that it runs out of memory and terminates with an "out of memory" message.

Alternate Options

None

See Also

[inline-min-size, Qinline-min-size](#) compiler option

[inline-factor, Qinline-factor](#) compiler option

[opt-report, Qopt-report](#) compiler option

Optimizing Applications:

Developer Directed Inline Expansion of User Functions

Compiler Directed Inline Expansion of User Functions

inline-max-total-size, Qinline-max-total-size

Specifies how much larger a routine can normally grow when inline expansion is performed.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-inline-max-total-size=n`
`-no-inline-max-total-size`

Windows: `/Qinline-max-total-size=n`
`/Qinline-max-total-size-`

Arguments

n Is a positive integer that specifies the permitted increase in the routine's size when inline expansion is performed.

Default

`-no-inline-max-total-size` The compiler uses default heuristics for inline routine expansion.

or `/Qinline-max-total-size`

Description

This option specifies how much larger a routine can normally grow when inline expansion is performed. It limits the potential size of the routine. For example, if 2000 is specified for n , the size of any routine will normally not increase by more than 2000.

If you specify `-no-inline-max-total-size` (Linux and Mac OS X) or `/Qinline-max-total-size` (Windows), there is no limit to the size a routine may grow when inline expansion is performed.

To see compiler values for important inlining limits, specify compiler option `-opt-report` (Linux and Mac OS X) or `/Qopt-report` (Windows).

To see compiler values for important inlining limits, specify compiler option `-opt-report` (Linux and Mac OS X) or `/Qopt-report` (Windows).

**Caution**

When you use this option to increase the default limit, the compiler may do so much additional inlining that it runs out of memory and terminates with an "out of memory" message.

Alternate Options

None

See Also

[inline-factor](#), [Qinline-factor](#) compiler option

[opt-report](#), [Qopt-report](#) compiler option

Optimizing Applications:

inline-min-size, Qinline-min-size

Specifies the upper limit for the size of what the inliner considers to be a small routine.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-inline-min-size=n`
`-no-inline-min-size`

Windows: `/Qinline-min-size=n`
`/Qinline-min-size-`

Arguments

<i>n</i>	Is a positive integer that specifies the maximum size of what the inliner considers to be a small routine.
----------	--

Default

<code>-no-inline-min-size</code>	The compiler uses default heuristics for
<code>or/Qinline-min-size-</code>	inline routine expansion.

Description

This option specifies the upper limit for the size of what the inliner considers to be a small routine (a function or subroutine). The inliner classifies routines as small,

medium, or large. This option specifies the boundary between what the inliner considers to be small and medium-size routines.

The inliner has a preference to inline small routines. So, when a routine is smaller than or equal to the specified size, it is very likely to be inlined.

If you specify `-no-inline-min-size` (Linux and Mac OS X) or `/Qinline-min-size-` (Windows), there is no limit to the size of small routines. Every routine is a small routine; there are no medium or large routines.

To see compiler values for important inlining limits, specify compiler option `-opt-report` (Linux and Mac OS X) or `/Qopt-report` (Windows).

To see compiler values for important inlining limits, specify compiler option `-opt-report` (Linux and Mac OS X) or `/Qopt-report` (Windows).



Caution

When you use this option to increase the default limit, the compiler may do so much additional inlining that it runs out of memory and terminates with an "out of memory" message.

Alternate Options

None

See Also

[inline-min-size, Qinline-min-size](#) compiler option

[opt-report, Qopt-report](#) compiler option

Optimizing Applications:

Developer Directed Inline Expansion of User Functions

Compiler Directed Inline Expansion of User Functions

intconstant

Tells the compiler to use FORTRAN 77 semantics to determine the kind parameter for integer constants.

IDE Equivalent

Windows: **Compatibility > Use F77 Integer Constants**

Linux: None

Mac OS X: **Compatibility > Use F77 Integer Constants**

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-intconstant`
`-nointconstant`

Windows: `/intconstant`
`/nointconstant`

Arguments

None

Default

`nointconstant` The compiler uses the Fortran 95/90 default INTEGER type.

Description

This option tells the compiler to use FORTRAN 77 semantics to determine the kind parameter for integer constants.

With FORTRAN 77 semantics, the kind is determined by the value of the constant. All constants are kept internally by the compiler in the highest precision possible. For example, if you specify option `intconstant`, the compiler stores an integer constant of 14 internally as `INTEGER(KIND=8)` and converts the constant upon reference to the corresponding proper size. Fortran 95/90 specifies that integer constants with no explicit `KIND` are kept internally in the default `INTEGER` kind (`KIND=4` by default).

Note that the internal precision for floating-point constants is controlled by option [fpconstant](#).

Alternate Options

None

integer-size

Specifies the default KIND for integer and logical variables.

IDE Equivalent

Windows: **Data > Default Integer KIND**

Linux: None

Mac OS X: **Data > Default Integer KIND**

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-integer-size size`

Windows: `/integer-size:size`

Arguments

size Is the size for integer and logical variables. Possible values are:
16, 32, or 64.

Default

`integer-` Integer and logical variables are 4 bytes long
`size 32` (INTEGER(KIND=4) and LOGICAL(KIND=4)).

Description

This option specifies the default size (in bits) for integer and logical variables.

Option	Description
<code>integer-size 16</code>	Makes default integer and logical variables 2 bytes long. INTEGER and LOGICAL declarations are treated as (KIND=2).
<code>integer-size 32</code>	Makes default integer and logical variables 4 bytes long. INTEGER and LOGICAL declarations are treated as (KIND=4).
<code>integer-size 64</code>	Makes default integer and logical variables 8 bytes long. INTEGER and LOGICAL declarations are treated as (KIND=8).

Alternate Options

<code>integer-size 16</code>	Linux and Mac OS X: <code>-i2</code> Windows: <code>/4I2</code>
<code>integer-size 32</code>	Linux and Mac OS X: <code>-i4</code> Windows: <code>/4I4</code>
<code>integer-size 64</code>	Linux and Mac OS X: <code>-i8</code> Windows: <code>/4I8</code>

ip, Qip

Determines whether additional interprocedural optimizations for single-file compilation are enabled.

IDE Equivalent

Windows: **Optimization > Interprocedural Optimization**

Linux: None

Mac OS X: None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-ip`
`-no-ip`

Windows: `/Qip`
`/Qip-`

Arguments

None

Default

OFF Some limited interprocedural optimizations occur, including inline function expansion for calls to functions defined within the current source file. These optimizations are a subset of full intra-file interprocedural optimizations. Note that this setting is not the same as `-no-ip` (Linux and Mac OS X) or `/Qip-` (Windows).

Description

This option determines whether additional interprocedural optimizations for single-file compilation are enabled.

Options `-ip` (Linux and Mac OS X) and `/Qip` (Windows) enable additional interprocedural optimizations for single-file compilation.

Options `-no-ip` (Linux and Mac OS X) and `/Qip-` (Windows) may not disable inlining. To ensure that inlining of user-defined functions is disabled, specify `-inline-level=0` or `-fno-inline` (Linux and Mac OS X), or specify `/Ob0` (Windows).

Alternate Options

None

See Also

[finline-functions](#) compiler option

ip-no-inlining, Qip-no-inlining

Disables full and partial inlining enabled by interprocedural optimization options.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-ip-no-inlining`

Windows: `/Qip-no-inlining`

Arguments

None

Default

OFF Inlining enabled by interprocedural optimization options is performed.

Description

This option disables full and partial inlining enabled by the following interprocedural optimization options:

- On Linux and Mac OS X systems: `-ip` or `-ipo`
- On Windows systems: `/Qip`, `/Qipo`, or `/Ob2`

It has no effect on other interprocedural optimizations.

On Windows systems, this option also has no effect on user-directed inlining specified by option `/Ob1`.

Alternate Options

None

`ip-no-pinlining`, `Qip-no-pinlining`

Disables partial inlining enabled by interprocedural optimization options.

IDE Equivalent

None

Architectures

IA-32, Intel® 64 architectures

Syntax

Linux and Mac OS X: `-ip-no-pinlining`

Windows: `/Qip-no-pinlining`

Arguments

None

Default

OFF Inlining enabled by interprocedural optimization options is performed.

Description

This option disables partial inlining enabled by the following interprocedural optimization options:

- On Linux and Mac OS X systems: `-ip` or `-ipo`
- On Windows systems: `/Qip` or `/Qipo`

It has no effect on other interprocedural optimizations.

Alternate Options

None

IPF-flt-eval-method0, QIPF-flt-eval-method0

Tells the compiler to evaluate the expressions involving floating-point operands in the precision indicated by the variable types declared in the program. This is a deprecated option.

IDE Equivalent

None

Architectures

IA-64 architecture

Syntax

Linux: `-IPF-flt-eval-method0`

Mac OS X: `None`

Windows: `/QIPF-flt-eval-method0`

Arguments

None

Default

OFF Expressions involving floating-point operands are evaluated by default rules.

Description

This option tells the compiler to evaluate the expressions involving floating-point operands in the precision indicated by the variable types declared in the program. By default, intermediate floating-point expressions are maintained in higher precision.

The recommended method to control the semantics of floating-point calculations is to use option `-fp-model` (Linux) or `/fp` (Windows).

Instead of using `-IPF-flt-eval-method0` (Linux) or `/QIPF-flt-eval-method0` (Windows), you can use `-fp-model source` (Linux) or `/fp:source` (Windows).

Alternate Options

None

See Also

[fp-model, fp](#) compiler option

IPF-fltacc, QIPF-fltacc

Disables optimizations that affect floating-point accuracy. This is a deprecated option.

IDE Equivalent

Windows: **Floating Point > Floating-Point Accuracy**

Linux: None

Mac OS X: None

Architectures

IA-64 architecture

Syntax

Linux: `-IPF-fltacc`
 `-no-IPF-fltacc`

Mac OS X: `None`

Windows: `/QIPF-fltacc`
 `/QIPF-fltacc-`

Arguments

None

Default

`-no-IPF-fltacc` Optimizations are enabled that affect floating-point or `/QIPF-fltacc-` accuracy.

Description

This option disables optimizations that affect floating-point accuracy.

If the default setting is used, the compiler may apply optimizations that reduce floating-point accuracy.

You can use this option to improve floating-point accuracy, but at the cost of disabling some optimizations.

The recommended method to control the semantics of floating-point calculations is to use option `-fp-model` (Linux) or `/fp` (Windows).

Instead of using `-IPF-fltacc` (Linux) or `/QIPF-fltacc` (Windows), you can use `-fp-model precise` (Linux) or `/fp:precise` (Windows).

Instead of using `-no-IPF-fltacc` (Linux) or `/QIPF-fltacc-` (Windows), you can use `-fp-model fast` (Linux) or `/fp:fast` (Windows).

Alternate Options

None

See Also

[fp-model, fp](#) compiler option

IPF-fma, QIPF-fma

See [fma, Qfma](#).

IPF-fp-relaxed, QIPF-fp-relaxed

See [fp-relaxed, Qfp-relaxed](#).

ipo, Qipo

Enables interprocedural optimization between files.

IDE Equivalent

Windows: **Optimization > Interprocedural Optimization**

General > Whole Program Optimization

Linux: None

Mac OS X: None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-ipo[n]`

Windows: `/Qipo[n]`

Arguments

n Is an optional integer that specifies the number of object files the compiler should create. The integer must be greater than or equal to 0.

Default

OFF Multifile interprocedural optimization is not enabled.

Description

This option enables interprocedural optimization between files. This is also called multifile interprocedural optimization (multifile IPO) or Whole Program Optimization (WPO).

When you specify this option, the compiler performs inline function expansion for calls to functions defined in separate files.

You cannot specify the names for the files that are created.

If *n* is 0, the compiler decides whether to create one or more object files based on an estimate of the size of the application. It generates one object file for small applications, and two or more object files for large applications.

If n is greater than 0, the compiler generates n object files, unless n exceeds the number of source files (m), in which case the compiler generates only m object files.

If you do not specify n , the default is 0.

Alternate Options

None

See Also

Optimizing Applications:

Interprocedural Optimization (IPO) Quick Reference

Interprocedural Optimization (IPO) Overview

Using IPO

ipo-c, Qipo-c

Tells the compiler to optimize across multiple files and generate a single object file.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-ipo-c`

Windows: `/Qipo-c`

Arguments

None

Default

OFF The compiler does not generate a multifile object file.

Description

This option tells the compiler to optimize across multiple files and generate a single object file (named ipo_out.o on Linux and Mac OS X systems; ipo_out.obj on Windows systems).

It performs the same optimizations as `-ipo` (Linux and Mac OS X) or `/Qipo` (Windows), but compilation stops before the final link stage, leaving an optimized object file that can be used in further link steps.

Alternate Options

None

See Also

[ipo, Qipo](#) compiler option

ipo-jobs, Qipo-jobs

Specifies the number of commands (jobs) to be executed simultaneously during the link phase of Interprocedural Optimization (IPO).

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-ipo-jobsn`

Windows: `/Qipo-jobs:n`

Arguments

n Is the number of commands (jobs) to run simultaneously. The number must be greater than or equal to 1.

Default

`-ipo-jobs1` One command (job) is executed in an
or `/Qipo-jobs:1` interprocedural optimization parallel build.

Description

This option specifies the number of commands (jobs) to be executed simultaneously during the link phase of Interprocedural Optimization (IPO). It should only be used if the link-time compilation is generating more than one object. In this case, each object is generated by a separate compilation, which can be done in parallel.

This option can be affected by the following compiler options:

- `-ipo` (Linux and Mac OS X) or `/Qipo` (Windows) when applications are large enough that the compiler decides to generate multiple object files.
- `-ipon` (Linux and Mac OS X) or `/Qipon` (Windows) when *n* is greater than 1.
- `-ipo-separate` (Linux) or `/Qipo-separate` (Windows)



Caution

Be careful when using this option. On a multi-processor system with lots of memory, it can speed application build time. However, if *n* is greater than the number of processors, or if there is not enough memory to avoid thrashing, this option can increase application build time.

Alternate Options

None

See Also

[ipo, Qipo](#) compiler option

[ipo-separate, Qipo-separate](#) compiler option

ipo-S, Qipo-S

Tells the compiler to optimize across multiple files and generate a single assembly file.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-ipo-S`

Windows: `/Qipo-S`

Arguments

None

Default

OFF The compiler does not generate a multifile assembly file.

Description

This option tells the compiler to optimize across multiple files and generate a single assembly file (named `ipo_out.s` on Linux and Mac OS X systems; `ipo_out.asm` on Windows systems).

It performs the same optimizations as `-ipo` (Linux and Mac OS X) or `/Qipo` (Windows), but compilation stops before the final link stage, leaving an optimized assembly file that can be used in further link steps.

Alternate Options

None

See Also

[ipo, Qipo](#) compiler option

ipo-separate, Qipo-separate

Tells the compiler to generate one object file for every source file.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux:	<code>-ipo-separate</code>
Mac OS X:	None
Windows:	<code>/Qipo-separate</code>

Arguments

None

Default

OFF The compiler decides whether to create one or more object files.

Description

This option tells the compiler to generate one object file for every source file. It overrides any `-ipo` (Linux) or `/Qipo` (Windows) specification.

Alternate Options

None

See Also

[ipo, Qipo](#) compiler option

isystem

Specifies a directory to add to the start of the system include path.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-isystemdir`

Windows: None

Arguments

dir Is the directory to add to the system include path.

Default

OFF The default system include path is used.

Description

This option specifies a directory to add to the system include path. The compiler searches the specified directory for include files after it searches all directories specified by the `-I` compiler option but before it searches the standard system directories. This option is provided for compatibility with gcc.

Alternate Options

None

ivdep-parallel, Qivdep-parallel

Tells the compiler that there is no loop-carried memory dependency in the loop following an IVDEP directive.

IDE Equivalent

Windows: **Optimization > IVDEP Directive Memory Dependency**

Linux: None

Mac OS X: None

Architectures

IA-64 architecture

Syntax

Linux: `-ivdep-parallel`

Mac OS X: `None`

Windows: `/Qivdep-parallel`

Arguments

None

Default

OFF There may be loop-carried memory dependency in a loop that follows an IVDEP directive.

Description

This option tells the compiler that there is no loop-carried memory dependency in the loop following an IVDEP. There may be loop-carried memory dependency in a loop that follows an IVDEP directive.

This has the same effect as specifying the IVDEP:LOOP directive.

Alternate Options

None

See Also

Optimizing Applications: Absence of Loop-carried Memory Dependency with IVDEP Directive

I

Tells the linker to search for a specified library when linking.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-lstring`

Windows: None

Arguments

string Specifies the library (*libstring*) that the linker should search.

Default

OFF The linker searches for standard libraries in standard directories.

Description

This option tells the linker to search for a specified library when linking.

When resolving references, the linker normally searches for libraries in several standard directories, in directories specified by the `L` option, then in the library specified by the `l` option.

The linker searches and processes libraries and object files in the order they are specified. So, you should specify this option following the last object file it applies to.

Alternate Options

None

See Also

[L](#) compiler option

L

Tells the linker to search for libraries in a specified directory before searching the standard directories.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-Ldir`

Windows: None

Arguments

dir Is the name of the directory to search for libraries.

Default

OFF The linker searches the standard directories for libraries.

Description

This option tells the linker to search for libraries in a specified directory before searching for them in the standard directories.

Alternate Options

None

See Also

[1](#) compiler option

LD

See [dll](#).

libdir

Controls whether linker options for search libraries are included in object files generated by the compiler.

IDE Equivalent

Windows: **Libraries > Disable Default Library Search Rules**

(/libdir:[no]automatic)

Libraries > Disable OBJCOMMENT Library Name in Object

(/libdir:[no]user)

Linux: None

Mac OS X: None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: None

Windows: /libdir[:*keyword*]
 /nolibdir

Arguments

keyword Specifies the linker search options. Possible values are:

none	Prevents any linker search options from being included into the object
------	--

file. This is the same as specifying
`/nolibdir`.

<code>[no]automatic</code>	Determines whether linker search options for libraries automatically determined by the <code>ifort</code> command driver (default libraries) are included in the object file.
<code>[no]user</code>	Determines whether linker search options for libraries specified by the <code>OBJCOMMENT</code> source directives are included in the object file.
<code>all</code>	Causes linker search options for the following libraries: <ul style="list-style-type: none">• Libraries automatically determined by the <code>ifort</code> command driver (default libraries)• Libraries specified by the <code>OBJCOMMENT</code> directive to be included in the object file This is the same as specifying <code>/libdir</code> .

Default

`/libdir:all` Linker search options for libraries automatically determined by the `ifort` command driver (default libraries) and libraries specified by the `OBJCOMMENT` directive are included in the object file.

Description

This option controls whether linker options for search libraries (`/DEFAULTTLIB:library`) are included in object files generated by the compiler.

The linker option `/DEFAULTLIB:library` adds one library to the list of libraries that the linker searches when resolving references. A library specified with `/DEFAULTLIB:library` is searched after libraries specified on the command line and before default libraries named in `.obj` files.

Alternate Options

`/libdir:none` Linux and Mac OS X: None

Windows: `/Z1`

libs

Tells the compiler which type of run-time library to link to.

IDE Equivalent

Windows: **Libraries > Runtime Library** (`/libs:{static|dll|qwin|qwins}`,
`/threads, /dbglibs`)

Linux: None

Mac OS X: None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: None

Windows: `/libs[:keyword]`

Arguments

keyword Specifies the type of run-time library to link to. Possible values are:

<code>static</code>	Specifies a single-threaded, static library (same as specifying <code>/libs</code>).
<code>dll</code>	Specifies a single-threaded, dynamic-link (DLL) library.
<code>qwin</code>	Specifies the Fortran QuickWin library.
<code>qwins</code>	Specifies the Fortran Standard Graphics library.

Default

`/libs:static` The compiler links to a single-threaded, static run-time or `/libs` library.

Description

This option tells the compiler which type of run-time library to link to.

The library can be statically or dynamically loaded, multithreaded (`/threads`) or single-threaded, or debug (`/dbglibs`) or nondebug.

If you use the `/libs:dll` option and an unresolved reference is found in the DLL, it gets resolved when the program is executed, during program loading, reducing executable program size.

If you use the `/libs:qwin` or `/libs:qwins` option with the `/dll` option, the compiler issues a warning.

You cannot use the `/libs:qwin` option and options `/libs:dll /threads`.

The following table shows which options to specify for different run-time libraries:

Type of Library	Options Required	Alternate Option
Single-threaded, static	<code>/libs:static</code> /ML or <code>/libs</code> or	

Type of Library	Options Required	Alternate Option
	/static	
Multithreaded	/libs:static /MT /threads	
Debug single-threaded	/libs:static /MLd /dbglibs	
Debug multithreaded	/libs:static /MTd /threads /dbglibs	
Single-threaded, dynamic-link libraries (DLLs)	/libs:dll	/MDs
Debug single-threaded, dynamic-link libraries (DLLs)	/libs:dll /dbglibs	/MDsd
Multithreaded DLLs	/libs:dll /threads	/MD
Multithreaded debug DLLs	/libs:dll /threads /dbglibs	/MDd
Fortran QuickWin multi-doc applications	/libs:qwin	/MW
Fortran standard graphics (QuickWin single-doc) applications	/libs:qwins	/MWS
Debug Fortran QuickWin multi-doc applications	/libs:qwin /dbglibs	None
Debug Fortran standard graphics (QuickWin single-doc) applications	/libs:qwins /dbglibs	None

Alternate Options

Intel® Fortran Compiler User and Reference Guides

<code>/libs:dll</code>	Linux and Mac OS X:None Windows: /MDs
<code>/libs:static</code>	Linux and Mac OS X: None Windows: /ML
<code>/libs:qwin</code>	Linux and Mac OS X: None Windows: /MW
<code>/libs:qwins</code>	Linux and Mac OS X: None Windows: /MWS

[threads](#) compiler option

[dbglibs](#) compiler option

Building Applications: Specifying Consistent Library Types; Programming with Mixed Languages Overview

link

Passes user-specified options directly to the linker at compile time.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: None

Windows: `/link`

Arguments

None

Default

OFF No user-specified options are passed directly to the linker.

Description

This option passes user-specified options directly to the linker at compile time. All options that appear following `/link` are passed directly to the linker.

Alternate Options

None

See Also

[xlinker](#) compiler option

logo

Displays the compiler version information.

IDE Equivalent

Windows: **General > Suppress Startup Banner** (`/nologo`)

Linux: None

Mac OS X: **General > Show Startup Banner** (`-v`)

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-logo`
`-nologo`

Windows: `/logo`
`/nologo`

Arguments

None

Default

Linux and Mac OS X: The compiler version information is not displayed.

`nologo`

Windows: `logo` The compiler version information is displayed.

Description

This option displays the startup banner, which contains the following compiler version information:

- ID: unique identification number for the compiler
- x.y.z: version of the compiler
- years: years for which the software is copyrighted

This option can be placed anywhere on the command line.

Alternate Options

Linux and Mac OS X: `-v`

Windows: None

lowercase, Qlowercase

See [names](#).

m

Tells the compiler to generate optimized code specialized for the processor that executes your program.

IDE Equivalent

None

Architectures

IA-32, Intel® 64 architectures

Syntax

Linux and Mac OS X: `-m[processor]`

Windows: None

Arguments

<i>processor</i>	Indicates the processor for which code is generated. Possible values are:
<code>ia32</code>	Generates code that will run on any Pentium or later processor. Disables any default extended instruction settings, and any previously set extended instruction settings. This value is only available on Linux systems using IA-32 architecture.
<code>sse</code>	This is the same as specifying <code>ia32</code> .
<code>sse2</code>	Generates code for Intel® Streaming SIMD Extensions 2 (Intel® SSE2). This value is only available on Linux systems.
<code>sse3</code>	Generates code for Intel® Streaming SIMD Extensions 3 (Intel® SSE3).
<code>ssse3</code>	Generates code for Intel® Supplemental Streaming SIMD Extensions 3 (Intel® SSSE3).
<code>sse4.1</code>	Generates code for Intel®

Streaming SIMD Extensions 4
Vectorizing Compiler and
Media Accelerators.

Default

Linux For more information on the default values, see Arguments
systems: – above.

`msse2`

Mac OS X

systems

using IA-32

architecture:

`-msse3`

Mac OS X

systems

using Intel®

64

architecture:

`-mssse3`

Description

This option tells the compiler to generate optimized code specialized for the processor that executes your program.

Code generated with the values `ia32`, `sse`, `sse2` or `sse3` should execute on any compatible non-Intel processor with support for the corresponding instruction set.

Alternate Options

Linux and Mac OS X: None

Windows: `/arch`

See Also

[x, Qx](#) compiler option

[ax, Qax](#) compiler option

[arch](#) compiler option

m32, m64

Tells the compiler to generate code for a specific architecture.

IDE Equivalent

None

Architectures

IA-32, Intel® 64 architectures

Syntax

Linux and Mac OS X: -m32

-m64

Windows: None

Arguments

None

Default

OFF The compiler's behavior depends on the host system.

Description

These options tell the compiler to generate code for a specific architecture.

Option	Description
-m32	Tells the compiler to generate code for IA-32 architecture.
-m64	Tells the compiler to generate code for Intel® 64 architecture.

The `-m32` and `-m64` options are the same as Mac OS* X options `-arch i386` and `-arch x86_64`, respectively. Note that these options are provided for compatibility with `gcc`. They are not related to the Intel® Fortran compiler option `arch`.

Alternate Options

None

map

Tells the linker to generate a link map file.

IDE Equivalent

Windows: **Linker > Debug > Generate Map File**

Linux: None

Mac OS X: None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: None

Windows: `/map[:file]`
 `/nomap`

Arguments

file Is the name for the link map file. It can be a file name or a directory name.

Default

`/nomap` No link map is generated.

Description

This option tells the linker to generate a link map file.

Alternate Options

None

map-opts, Qmap-opts

Maps one or more compiler options to their equivalent on a different operating system.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux: `-map-opts`

Mac OS X: `None`

Windows: `/Qmap-opts`

Arguments

None

Default

OFF No platform mappings are performed.

Description

This option maps one or more compiler options to their equivalent on a different operating system. The result is output to `stdout`.

On Windows systems, the options you provide are presumed to be Windows options, so the options that are output to `stdout` will be Linux equivalents.

On Linux systems, the options you provide are presumed to be Linux options, so the options that are output to `stdout` will be Windows equivalents.

The tool can be invoked from the compiler command line or it can be used directly.

No compilation is performed when the option mapping tool is used.

This option is useful if you have both compilers and want to convert scripts or makefiles.



Compiler options are mapped to their equivalent on the architecture you are using.

For example, if you are using a processor with IA-32 architecture, you will only see equivalent options that are available on processors with IA-32 architecture.

Alternate Options

None

Example

The following command line invokes the option mapping tool, which maps the Linux options to Windows-based options, and then outputs the results to

`stdout`:

```
ifort -map-opts -xP -O2
```

The following command line invokes the option mapping tool, which maps the Windows options to Linux-based options, and then outputs the results to

`stdout`:

```
ifort /Qmap-opts /QxP /O2
```

See Also

Building Applications: Using the Option Mapping Tool

march

Tells the compiler to generate code for a specified processor.

IDE Equivalent

None

Architectures

IA-32, Intel® 64 architectures

Syntax

Linux: `-march=processor`

Mac OS X: None

Windows: None

Arguments

processor Is the processor for which the compiler should generate code. Possible values are:

`pentium3` Generates code for Intel® Pentium® III processors.

`pentium4` Generates code for Intel® Pentium® 4 processors.

`core2` Generates code for the Intel® Core 2™ processor family.

Default

OFF or
-
`march=pentium4` On IA-32 architecture, the compiler does not generate processor-specific code unless it is told to do so. On systems using Intel® 64 architecture, the compiler generates code for Intel Pentium 4 processors.

Description

This option tells the compiler to generate code for a specified processor.

Specifying `-march=pentium4` sets `-mtune=pentium4`.

For compatibility, a number of historical *processor* values are also supported, but the generated code will not differ from the default.

Alternate Options

None

mcmmodel

Tells the compiler to use a specific memory model to generate code and store data.

IDE Equivalent

None

Architectures

Intel® 64 architecture

Syntax

Linux: `-mcmmodel=mem_model`

Mac OS X: None

Windows: None

Arguments

mem_model Is the memory model to use. Possible values are:

<code>small</code>	Tells the compiler to restrict code and data to the first 2GB of address space. All accesses of code and data can be done with Instruction Pointer (IP)-
--------------------	--

	relative addressing.
medium	Tells the compiler to restrict code to the first 2GB; it places no memory restriction on data. Accesses of code can be done with IP-relative addressing, but accesses of data must be done with absolute addressing.
large	Places no memory restriction on code or data. All accesses of code and data must be done with absolute addressing.

Default

- On systems using Intel® 64 architecture, the compiler `mcmmodel=small` restricts code and data to the first 2GB of address space. Instruction Pointer (IP)-relative addressing can be used to access code and data.

Description

This option tells the compiler to use a specific memory model to generate code and store data. It can affect code size and performance. If your program has COMMON blocks and local data with a total size smaller than 2GB, `-mcmmodel=small` is sufficient. COMMONs larger than 2GB require `-mcmmodel=medium` or `-mcmmodel=large`. Allocation of memory larger than 2GB can be done with any setting of `-mcmmodel`.

IP-relative addressing requires only 32 bits, whereas absolute addressing requires 64-bits. IP-relative addressing is somewhat faster. So, the `small` memory model has the least impact on performance.



When you specify `-mmodel=medium` or `-mmodel=large`, you must also specify compiler option `-shared-intel` to ensure that the correct dynamic versions of the Intel run-time libraries are used.

Alternate Options

None

Example

The following example shows how to compile using `-mmodel`:

```
ifort -shared-intel -mmodel=medium -o prog prog.f
```

See Also

[shared-intel](#) compiler option

[fpic](#) compiler option

mcpu

This is a deprecated option. See [mtune](#).

MD

Tells the linker to search for unresolved references in a multithreaded, debug, dynamic-link run-time library.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: None

Windows: /MD

/MDd

Arguments

None

Default

OFF The linker searches for unresolved references in a single-threaded, static run-time library.

Description

This option tells the linker to search for unresolved references in a multithreaded, debug, dynamic-link (DLL) run-time library. This is the same as specifying options `/libs:dll /threads /dbglibs`. You can also specify `/MDd`, where `d` indicates a debug version.

Alternate Options

None

See Also

[libs](#) compiler option

MDs

Tells the linker to search for unresolved references in a single-threaded, dynamic-link run-time library.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: None

Windows: /MDs

/MDsd

Arguments

None

Default

OFF The linker searches for unresolved references in a single-threaded, static run-time library.

Description

This option tells the linker to search for unresolved references in a single-threaded, dynamic-link (DLL) run-time library.

You can also specify /MDsd, where *d* indicates a debug version.

Alternate Options

/MDs Linux and Mac OS X: None

Windows: /libs:dll

See Also

[libs](#) compiler option

mdynamic-no-pic

Generates code that is not position-independent but has position-independent external references.

IDE Equivalent

None

Architectures

IA-32 architecture

Syntax

Linux: None
Mac OS X: -mdynamic-no-pic
Windows: None

Arguments

None

Default

OFF All references are generated as position independent.

Description

This option generates code that is not position-independent but has position-independent external references.

The generated code is suitable for building executables, but it is not suitable for building shared libraries.

This option may reduce code size and produce more efficient code. It overrides the `-fpic` compiler option.

Alternate Options

None

See Also

[fpic](#) compiler option

MG

See [winapp](#).

mieee-fp

See [fltconsistency](#).

minstruction, Qinstruction

Determines whether MOVBE instructions are generated for Intel processors.

IDE Equivalent

None

Architectures

IA-32, Intel® 64 architectures

Syntax

Linux and Mac OS X: `-minstruction=[no]movbe`

Windows: `/Qinstruction:[no]movbe`

Arguments

None

Default

<code>-minstruction=movbe</code>	The compiler generates MOVBE instructions for Intel® Atom™ processors.
<code>or /Qinstruction:movbe</code>	

Description

This option determines whether MOVBE instructions are generated for Intel processors. To use this option, you must also specify `-xSSE3_ATOM` (Linux and Mac OS X) or `/QxSSE3_ATOM` (Windows).

If `-minstruction=movbe` or `/Qinstruction:movbe` is specified, the following occurs:

- MOVBE instructions are generated that are specific to the Intel® Atom™ processor.

- The options are ON by default when `-xSSE3_ATOM` or `/QxSSE3_ATOM` is specified.
- Generated executables can only be run on Intel® Atom™ processors or processors that support Intel® Streaming SIMD Extensions 3 (Intel® SSE3) and MOVBE.

If `-minstruction=nomovbe` or `/Qinstruction:nomovbe` is specified, the following occurs:

- The compiler optimizes code for the Intel® Atom™ processor, but it does not generate MOVBE instructions.
- Generated executables can be run on non-Intel® Atom™ processors that support Intel® SSE3.

Alternate Options

None

See Also

[x, Qx](#) compiler option

mixed-str-len-arg

See [iface](#).

ML

Tells the linker to search for unresolved references in a single-threaded, static run-time library.

This option has been [deprecated](#).

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: None

Windows: /ML
 /MLd

Arguments

None

Default

Systems using Intel® On systems using Intel® 64 architecture, the linker searches for unresolved references in a multithreaded, static run-time library. On systems using IA-32 architecture and IA-64 architecture, the linker searches for unresolved references in a single-threaded, static run-time library.

Description

This option tells the linker to search for unresolved references in a single-threaded, static run-time library. You can also specify `/MLd`, where `d` indicates a debug version.

Alternate Options

Linux: None

Mac OS X: None

Windows: `/libs:static`

See Also

[libs](#) compiler option

module

Specifies the directory where module files should be placed when created and where they should be searched for.

IDE Equivalent

Windows: **Output > Module Path**

Linux: None

Mac OS X: **Output Files > Module Path**

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-module path`

Windows: `/module:path`

Arguments

path Is the directory for module files.

Default

OFF The compiler places module files in the current directory.

Description

This option specifies the directory (path) where module (.mod) files should be placed when created and where they should be searched for (USE statement).

Alternate Options

None

mp

See [fltconsistency](#)

multiple-processes, MP

Creates multiple processes that can be used to compile large numbers of source files at the same time.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-multiple-processes[=n]`

Windows: `/MP[:n]`

Arguments

n Is the maximum number of processes that the compiler should create.

Default

OFF A single process is used to compile source files.

Description

This option creates multiple processes that can be used to compile large numbers of source files at the same time. It can improve performance by reducing the time it takes to compile source files on the command line.

This option causes the compiler to create one or more copies of itself, each in a separate process. These copies simultaneously compile the source files.

If *n* is not specified for this option, the default value is as follows:

- On Windows OS, the value is based on the setting of the `NUMBER_OF_PROCESSORS` environment variable.
- On Linux OS and Mac OS X, the value is 2.

This option applies to compilations, but not to linking or link-time code generation.

Alternate Options

None

mp1, Qprec

Improves floating-point precision and consistency.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-mp1`

Windows: `/Qprec`

Arguments

None

Default

OFF The compiler provides good accuracy and run-time performance at the expense of less consistent floating-point results.

Description

This option improves floating-point consistency. It ensures the out-of-range check of operands of transcendental functions and improves the accuracy of floating-point compares.

This option prevents the compiler from performing optimizations that change NaN comparison semantics and causes all values to be truncated to declared precision before they are used in comparisons. It also causes the compiler to use library routines that give better precision results compared to the X87 transcendental instructions.

This option disables fewer optimizations and has less impact on performance than option `fltconsistency` or `mp`.

Alternate Options

None

See Also

[fltconsistency](#) compiler option

[mp](#) compiler option

mrelax

Tells the compiler to pass linker option `-relax` to the linker.

IDE Equivalent

None

Architectures

IA-64 architecture

Syntax

Linux and Mac OS X: `-mrelax`
`-mno-relax`

Mac OS X: None

Windows: None

Arguments

None

Default

`-mno-relax` The compiler does not pass `-relax` to the linker.

Description

This option tells the compiler to pass linker option `-relax` to the linker.

Alternate Options

None

MT

Tells the linker to search for unresolved references in a multithreaded, static run-time library.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: None

Windows: /MT
 /MTd

Arguments

None

Default

Systems using Intel® On systems using Intel® 64 architecture, the linker
64 architecture: /MT searches for unresolved references in a
/noreentrancy multithreaded, static run-time library. On systems
IA-32 architecture using IA-32 architecture and IA-64 architecture,
and IA-64 the linker searches for unresolved references in a
architecture: OFF single-threaded, static run-time library. However,
on systems using IA-32 architecture, if option
QVC8 is in effect, the linker searches for
unresolved references in threaded libraries.

Description

This option tells the linker to search for unresolved references in a multithreaded, static run-time library. This is the same as specifying options `/libs:static` `/threads` `/noreentrancy`. You can also specify `/MTd`, where `d` indicates a debug version.

Alternate Options

None

See Also

[Qvc](#) compiler option

[libs](#) compiler option

[threads](#) compiler option

[reentrancy](#) compiler option

mtune

Performs optimizations for specific processors.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-mtune=processor`

Windows: None

Arguments

processor Is the processor for which the compiler should perform optimizations. Possible values are:

`generic` Generates code for the compiler's default behavior.

Intel(R) Fortran Compiler User and Reference Guides

<code>core2</code>	Optimizes for the Intel® Core 2 processor family, including support for MMX™, Intel® SSE, SSE2, SSE3 and SSSE3 instruction sets.
<code>pentium</code>	Optimizes for Intel® Pentium® processors.
<code>pentium- mmx</code>	Optimizes for Intel® Pentium® with MMX technology.
<code>pentiumpro</code>	Optimizes for Intel® Pentium® Pro, Intel Pentium II, and Intel Pentium III processors.
<code>pentium4</code>	Optimizes for Intel® Pentium® 4 processors.
<code>pentium4m</code>	Optimizes for Intel® Pentium® 4 processors with MMX technology.
<code>itanium2</code>	Optimizes for Intel® Itanium® 2 processors.
<code>itanium2- p9000</code>	Optimizes for the Dual-Core Intel® Itanium® 2 processor 9000 series. This option affects the order of the generated instructions, but the generated instructions are limited to Intel® Itanium® 2 processor instructions unless the program uses (executes)

intrinsic specific to the Dual-Core Intel® Itanium® 2 processor 9000 series.

Default

- `generic` On systems using IA-32 and Intel® 64 architectures, code is generated for the compiler's default behavior.
- `itanium2-p9000` On systems using IA-64 architecture, the compiler optimizes for the Dual-Core Intel® Itanium® 2 processor 9000 series.

Description

This option performs optimizations for specific processors.

The resulting executable is backwards compatible and generated code is optimized for specific processors. For example, code generated with `-mtune=itanium2-p9000` will run correctly on single-core Itanium® 2 processors, but it might not run as fast as if it had been generated using `-mtune=itanium2`.

The following table shows on which architecture you can use each value.

<i>processor</i> Value	Architecture		
	IA-32 architecture	Intel® 64 architecture	IA-64 architecture
<code>generic</code>	X	X	X
<code>core2</code>	X	X	
<code>pentium</code>	X		
<code>pentium-mmx</code>	X		
<code>pentiumpro</code>	X		
<code>pentium4</code>	X		

<i>processor</i> Value	Architecture		
	IA-32 architecture	Intel® 64 architecture	IA-64 architecture
pentium4m	X		
itanium2			X
itanium2- p9000			X

Alternate Options

-mtune	Linux: -mcpu (this is a deprecated option) Mac OS X: None Windows: None
-mtune=itanium2	Linux: -mcpu=itanium2 (-mcpu is a deprecated option) Mac OS X: None Windows: /G2
- mtune=itanium2- p9000	Linux: -mcpu=itanium2-p9000 (-mcpu is a deprecated option) Mac OS X: None Windows: /G2-p9000

multiple-processes, MP

Creates multiple processes that can be used to compile large numbers of source files at the same time.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-multiple-processes[=n]`

Windows: `/MP[:n]`

Arguments

n Is the maximum number of processes that the compiler should create.

Default

OFF A single process is used to compile source files.

Description

This option creates multiple processes that can be used to compile large numbers of source files at the same time. It can improve performance by reducing the time it takes to compile source files on the command line.

This option causes the compiler to create one or more copies of itself, each in a separate process. These copies simultaneously compile the source files.

If *n* is not specified for this option, the default value is as follows:

- On Windows OS, the value is based on the setting of the `NUMBER_OF_PROCESSORS` environment variable.
- On Linux OS and Mac OS X, the value is 2.

This option applies to compilations, but not to linking or link-time code generation.

Alternate Options

None

MW

See [libs](#).

MWs

See [libs](#).

names

Specifies how source code identifiers and external names are interpreted.

IDE Equivalent

Windows: **External Procedures > Name Case Interpretation**

Linux: None

Mac OS X: **External Procedures > Name Case Interpretation**

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-names keyword`

Windows: `/names:keyword`

Arguments

keyword Specifies how to interpret the identifiers and external names in source code. Possible values are:

<code>lowercase</code>	Causes the compiler to ignore case differences in identifiers and to convert external names to lowercase.
<code>uppercase</code>	Causes the compiler to ignore case differences in identifiers and to convert external names to uppercase.
<code>as_is</code>	Causes the compiler to distinguish case differences in identifiers and to preserve the case of external names.

Default

<code>lowercase</code>	This is the default on Linux and Mac OS X systems. The compiler ignores case differences in identifiers and converts external names to lowercase.
<code>uppercase</code>	This is the default on Windows systems. The compiler ignores case differences in identifiers and converts external names to uppercase.

Description

This option specifies how source code identifiers and external names are interpreted. It can be useful in mixed-language programming.

This naming convention applies whether names are being defined or referenced. You can use the ALIAS directive to specify an alternate external name to be used when referring to external subprograms.



Caution

On Windows systems, if you specify option `/iface:cref`, it overrides the default for external names and causes them to be lowercase. It is as if you specified `!dec$ attributes c, reference` for the external name. If you specify option `/iface:cref` and want external names to be uppercase, you must explicitly specify option `/names:uppercase`.

Alternate Options

<code>names lowercase</code>	Linux and Mac OS X: <code>-lowercase</code> Windows: <code>/Qlowercase</code>
<code>names uppercase</code>	Linux and Mac OS X: <code>-uppercase</code> Windows: <code>/Quppercase</code>

See Also

[iface](#) compiler option

[ALIAS Directive](#)

nbs

See [assume](#).

no-bss-init, Qnobss-init

Tells the compiler to place in the DATA section any variables explicitly initialized with zeros.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-no-bss-init`

Windows: `/Qnobss-init`

Arguments

None

Default

OFF Variables explicitly initialized with zeros are placed in the BSS section.

Description

This option tells the compiler to place in the DATA section any variables explicitly initialized with zeros.

Alternate Options

Linux and Mac OS X: `-nobss-init` (this is a [deprecated](#) option)

Windows: None

nodefaultlibs

Prevents the compiler from using standard libraries when linking.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-nodefaultlibs`

Windows: None

Arguments

None

Default

OFF The standard libraries are linked.

Description

This option prevents the compiler from using standard libraries when linking.

Alternate Options

None

See Also

[nostdlib](#) compiler option

nodefine

See [D](#).

nofor-main

Specifies that the main program is not written in Fortran.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-nofor-main`

Windows: None

Arguments

None

Default

OFF The compiler assumes the main program is written in Fortran.

Description

This option specifies that the main program is not written in Fortran. It is a link-time option that prevents the compiler from linking `for_main.o` into applications. For example, if the main program is written in C and calls a Fortran subprogram, specify `-nofor-main` when compiling the program with the `ifort` command. If you omit this option, the main program must be a Fortran program.

Alternate Options

None

noinclude

See [X](#).

nolib-inline

Disables inline expansion of standard library or intrinsic functions.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-nolib-inline`

Windows: None

Arguments

None

Default

OFF The compiler inlines many standard library and intrinsic functions.

Description

This option disables inline expansion of standard library or intrinsic functions. It prevents the unexpected results that can arise from inline expansion of these functions.

Alternate Options

None

nostartfiles

Prevents the compiler from using standard startup files when linking.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-nostartfiles`

Windows: None

Arguments

None

Default

OFF The compiler uses standard startup files when linking.

Description

This option prevents the compiler from using standard startup files when linking.

Alternate Options

None

See Also

[nostdlib](#) compiler option

nostdinc

See [X](#).

nostdlib

Prevents the compiler from using standard libraries and startup files when linking.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-nostdlib`

Windows: None

Arguments

None

Default

OFF The compiler uses standard startup files and standard libraries when linking.

Description

This option prevents the compiler from using standard libraries and startup files when linking.

Alternate Options

None

See Also

[nodefaultlibs](#) compiler option

[nostartfiles](#) compiler option

nus

See [assume](#).

o

Specifies the name for an output file.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-ofile`

Windows: None

Arguments

file Is the name for the output file. The space before *file* is optional.

Default

OFF The compiler uses the default file name for an output file.

Description

This option specifies the name for an output file as follows:

- If `-c` is specified, it specifies the name of the generated object file.
- If `-S` is specified, it specifies the name of the generated assembly listing file.
- If `-preprocess-only` or `-P` is specified, it specifies the name of the generated preprocessor file.

Otherwise, it specifies the name of the executable file.

**Note**

If you misspell a compiler option beginning with "o", such as `-openmp`, `-opt-report`, etc., the compiler interprets the misspelled option as an `-ofile` option. For example, say you misspell "`-opt-report`" as "`-opt-reprt`"; in this case, the compiler interprets the misspelled option as "`-opt-reprt`", where `pt-reprt` is the output file name.

Alternate Options

Linux and Mac OS X: None

Windows: `/Fe`, `/exe`

See Also

[Fe](#) compiler option

[object](#) compiler option

O

Specifies the code optimization for applications.

IDE Equivalent

Windows: **General > Optimization** (`/Od`, `/O1`, `/O2`, `/O3`, `/fast`)

Optimization > Optimization (`/Od`, `/O1`, `/O2`, `/O3`, `/fast`)

Linux: None

Mac OS X: **General > Optimization Level** (`-O`)

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-O[n]`

Windows: `/O[n]`

Arguments

n Is the optimization level. Possible values are 1, 2, or 3. On Linux and Mac OS X systems, you can also specify 0.

Default

`O2` Optimizes for code speed. This default may change depending on which other compiler options are specified. For details, see

below.

Description

This option specifies the code optimization for applications.

Option	Description
o (Linux and Mac OS X)	This is the same as specifying O2.
o0 (Linux and Mac OS X)	<p>Disables all optimizations. On systems using IA-32 architecture and Intel® 64 architecture, this option sets option <code>-fno-omit-frame-pointer</code> and option <code>-fmath-errno</code>.</p> <p>Option <code>-o0</code> also implies option <code>-mp</code>. So, intermediate floating-point results are evaluated at extended precision. On IA-32 or Intel® 64 architecture, this may cause the compiler to use x87 instructions instead of SSE instructions. You can use option <code>-fp-model</code> to independently control the evaluation precision for intermediate results.</p> <p>Option <code>-mp</code> is deprecated. Therefore, the default behavior of <code>-o0</code> may change in a future compiler release.</p> <p>This option causes certain <code>warn</code> options to be ignored. This is the default if you specify option <code>-debug</code> (with no keyword).</p>
o1	<p>Enables optimizations for speed and disables some optimizations that increase code size and affect speed. To limit code size, this option:</p> <ul style="list-style-type: none"> • Enables global optimization; this includes data-flow analysis, code motion, strength reduction and test

Option	Description
	<p>replacement, split-lifetime analysis, and instruction scheduling.</p> <ul style="list-style-type: none"> On systems using IA-64 architecture, it disables software pipelining, loop unrolling, and global code scheduling. <p>On systems using IA-64 architecture, this option also enables optimizations for server applications (straight-line and branch-like code with a flat profile). The <code>O1</code> option sets the following options:</p> <ul style="list-style-type: none"> On Linux and Mac OS X systems: <ul style="list-style-type: none"> <code>-funroll-loops0</code>, <code>-nofltnconsistency</code> (same as <code>-mno-ieee-fp</code>), <code>-fomit-frame-pointer</code>, <code>-ftz</code> On Windows systems using IA-32 architecture: <ul style="list-style-type: none"> <code>/Qunroll0</code>, <code>/nofltnconsistency</code> (same as <code>/Op-</code>), <code>/Oy</code>, <code>/Os</code>, <code>/Ob2</code>, <code>/Qftz</code> On Windows systems using Intel® 64 architecture and IA-64 architecture: <ul style="list-style-type: none"> <code>/Qunroll0</code>, <code>/nofltnconsistency</code> (same as <code>/Op-</code>), <code>/Os</code>, <code>/Ob2</code>, <code>/Qftz</code> <p>The <code>O1</code> option may improve performance for applications with very large code size, many branches, and execution time not dominated by code within loops.</p>
O2	<p>Enables optimizations for speed. This is the generally recommended optimization level.</p> <p>Vectorization is enabled at O2 and higher levels.</p> <p>On systems using IA-64 architecture, this option enables optimizations for speed, including global code scheduling, software pipelining, predication, and</p>

Option	Description
	<p>speculation.</p> <p>This option also enables:</p> <ul style="list-style-type: none"> • Inlining of intrinsics • Intra-file interprocedural optimization, which includes: <ul style="list-style-type: none"> ○ inlining ○ constant propagation ○ forward substitution ○ routine attribute propagation ○ variable address-taken analysis ○ dead static function elimination ○ removal of unreferenced variables • The following capabilities for performance gain: <ul style="list-style-type: none"> ○ constant propagation ○ copy propagation ○ dead-code elimination ○ global register allocation ○ global instruction scheduling and control speculation ○ loop unrolling ○ optimized code selection ○ partial redundancy elimination ○ strength reduction/induction variable simplification ○ variable renaming ○ exception handling optimizations ○ tail recursions ○ peephole optimizations ○ structure assignment lowering and optimizations ○ dead store elimination <p>On Windows systems, this option is the same as the <code>Ox</code> option.</p>

Option	Description
	<p>The <code>o2</code> option sets the following options:</p> <ul style="list-style-type: none"> On Windows systems using IA-32 architecture: <code>/Og, /Ot, /Oy, /Ob2, /Gs, and /Qftz</code> On Windows systems using Intel® 64 architecture: <code>/Og, /Ot, /Ob2, /Gs, and /Qftz</code> <p>On Linux and Mac OS X systems, if <code>-g</code> is specified, <code>o2</code> is turned off and <code>o0</code> is the default unless <code>o2</code> (or <code>o1</code> or <code>o3</code>) is explicitly specified in the command line together with <code>-g</code>. This option sets other options that optimize for code speed. The options set are determined by the compiler depending on which architecture and operating system you are using.</p>
o3	<p>Enables <code>o2</code> optimizations plus more aggressive optimizations, such as prefetching, scalar replacement, and loop and memory access transformations. Enables optimizations for maximum speed, such as:</p> <ul style="list-style-type: none"> Loop unrolling, including instruction scheduling Code replication to eliminate branches Padding the size of certain power-of-two arrays to allow more efficient cache use. <p>On Windows systems, the <code>o3</code> option sets the <code>/Ob2</code> option.</p> <p>On Linux and Mac OS X systems, the <code>o3</code> option sets option <code>-fomit-frame-pointer</code>.</p> <p>On systems using IA-32 architecture or Intel® 64 architecture, when <code>o3</code> is used with options <code>-ax</code> or <code>-x</code> (Linux) or with options <code>/Qax</code> or <code>/Qx</code> (Windows), the compiler performs more aggressive data dependency</p>

Option	Description
	<p>analysis than for O2, which may result in longer compilation times.</p> <p>On systems using IA-64 architecture, the O3 option enables optimizations for technical computing applications (loop-intensive code): loop optimizations and data prefetch.</p> <p>The O3 optimizations may not cause higher performance unless loop and memory access transformations take place. The optimizations may slow down code in some cases compared to O2 optimizations.</p> <p>The O3 option is recommended for applications that have loops that heavily use floating-point calculations and process large data sets.</p>

The last O option specified on the command line takes precedence over any others.



Note

The options set by the O option may change from release to release.

Alternate Options

O1	Linux and Mac OS X: None Windows: /Od, /optimize:0, /nooptimize
O2	Linux and Mac OS X: None Windows: /optimize:1, /optimize:2
O3	Linux and Mac OS X: None Windows: /Ox, /optimize:3, /optimize:4
O4	Linux and Mac OS X: None

Windows: `/optimize:5`

See Also

[od](#) compiler option

[Op](#) compiler option

[fp-model, fp](#) compiler option

[fast](#) compiler option

See Also

Optimizing Applications:

Compiler Optimizations Overview

Optimization Options Summary

Efficient Compilation

inline-level, Ob

Specifies the level of inline function expansion.

IDE Equivalent

Windows: **Optimization > Inline Function Expansion**

Linux: None

Mac OS X: **Optimization > Inline Function Expansion**

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-inline-level=n`

Windows: `/Obn`

Arguments

n Is the inline function expansion level. Possible

values are 0, 1, and 2.

Default

`-inline-` This is the default if option `O2` is specified or is in effect by
`level=2` or `/Ob2` default. On Windows systems, this is also the default if
`/Od` option `O3` is specified.

`-inline-` This is the default if option `-O0` (Linux and Mac OS) or
`level=0` or `/Od` (Windows) is specified.
`/Ob0`

Description

This option specifies the level of inline function expansion. Inlining procedures can greatly improve the run-time performance of certain programs.

Option	Description
<code>-inline-level=0</code> or <code>Ob0</code>	Disables inlining of user-defined functions. Note that statement functions are always inlined.
<code>-inline-level=1</code> or <code>Ob1</code>	Enables inlining when an inline keyword or an inline directive is specified.
<code>-inline-level=2</code> or <code>Ob2</code>	Enables inlining of any function at the compiler's discretion.

Alternate Options

Linux: `-Ob` (this is a [deprecated](#) option)

Mac OS X: None

Windows: None

See Also

[inline](#) compiler option

object

Specifies the name for an object file.

IDE Equivalent

Windows: **Output Files > Object File Name**

Linux: None

Mac OS X: None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: None

Windows: `/object:file`

Arguments

file Is the name for the object file. It can be a file or directory name.

Default

OFF An object file has the same name as the name of the first source file and a file extension of .obj.

Description

This option specifies the name for an object file.

If you specify this option and you omit `/c` or `/compile-only`, the `/object` option gives the object file its name.

On Linux and Mac OS X systems, this option is equivalent to specifying option `-ofile -c`.

Alternate Options

Linux and Mac OS X: None

Windows: `/Fo`

See Also

[o](#) compiler option

Od

Disables all optimizations.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `None`

Windows: `/Od`

Arguments

None

Default

OFF The compiler performs default optimizations.

Description

This option disables all optimizations. It can be used for selective optimizations, such as a combination of `/Od` and `/Og` (disables all optimizations except global optimizations), or `/Od` and `/Ob1` (disables all optimizations, but enables inlining).

This option also causes certain `/warn` options to be ignored.

On IA-32 architecture, this option sets the `/Oy-` option.

Option `/Od` also implies option `/Op`. So, intermediate floating-point results are evaluated at extended precision. On IA-32 and Intel® 64 architecture, this may cause the compiler to use x87 instructions instead of SSE instructions. You can

use option `/fp` to independently control the evaluation precision for intermediate results.

Option `/Op` is deprecated. Therefore, the default behavior of `/Od` may change in a future compiler release.

Alternate Options

Linux and Mac OS X: `-O0`

Windows: `/optimize:0`

See Also

[O](#) compiler option (see `O0`)

[fp-model](#), [fp](#) compiler option

Og

Enables global optimizations.

IDE Equivalent

Windows: **Optimization > Global Optimizations**

Linux: None

Mac OS X: None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: None

Windows: `/Og`

`/Og-`

Arguments

None

Default

`/Og` Global optimizations are enabled unless `/Od` is specified.

Description

This option enables global optimizations.

Alternate Options

None

onetrip, Qonetrip

Tells the compiler to follow the FORTRAN 66 Standard and execute DO loops at least once.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-onetrip`

Windows: `/Qonetrip`

Arguments

None

Default

OFF The compiler applies the current Fortran Standard semantics, which allows zero-trip DO loops.

Description

This option tells the compiler to follow the FORTRAN 66 Standard and execute DO loops at least once.

Alternate Options

Linux and Mac OS X: -1

Windows: /1

Op

This is a deprecated option. See [fltconsistency](#).

openmp, Qopenmp

Enables the parallelizer to generate multi-threaded code based on the OpenMP* directives.

IDE Equivalent

Windows: **Language > Process OpenMP Directives**

Linux: None

Mac OS X: None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: -openmp

Windows: /Qopenmp

Arguments

None

Default

OFF No OpenMP multi-threaded code is generated by the compiler.

Description

This option enables the parallelizer to generate multi-threaded code based on the OpenMP* directives. The code can be executed in parallel on both uniprocessor and multiprocessor systems.

If you use this option, multithreaded libraries are used, but option `fpp` is not automatically invoked.

This option sets option `automatic`.

This option works with any optimization level. Specifying no optimization (`-O0` on Linux or `/Od` on Windows) helps to debug OpenMP applications.

**Note**

On Mac OS X systems, when you enable OpenMP*, you must also set the `DYLD_LIBRARY_PATH` environment variable within Xcode or an error will be displayed.

Alternate Options

None

See Also

[openmp-stubs, Qopenmp-stubs](#) compiler option

openmp-lib, Qopenmp-lib

Lets you specify an OpenMP* run-time library to use for linking.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux: `-openmp-lib type`
Mac OS X: `None`
Windows: `/Qopenmp-lib:type`

Arguments

type Specifies the type of library to use; it implies compatibility levels. Possible values are:

`legacy` Tells the compiler to use the legacy OpenMP* run-time library (libguide). This setting does not provide compatibility with object files created using other compilers. This is a [deprecated](#) option.

`compat` Tells the compiler to use the compatibility OpenMP* run-time library (libiomp). This setting provides compatibility with object files created using Microsoft* and GNU* compilers.

Default

`-openmp-lib compat` The compiler uses the compatibility
`or/Qopenmp-lib:compat` OpenMP* run-time library (libiomp).

Description

This option lets you specify an OpenMP* run-time library to use for linking. The legacy OpenMP run-time library is not compatible with object files created using OpenMP run-time libraries supported in other compilers.

The compatibility OpenMP run-time library is compatible with object files created using the Microsoft* OpenMP run-time library (vcomp) and GNU OpenMP run-time library (libgomp).

To use the compatibility OpenMP run-time library, compile and link your application using the `-openmp-lib compat` (Linux) or `/Qopenmp-lib:compat` (Windows) option. To use this option, you must also specify one of the following compiler options:

- Linux OS: `-openmp`, `-openmp-profile`, or `-openmp-stubs`
- Windows OS: `/Qopenmp`, `/Qopenmp-profile`, or `/Qopenmp-stubs`

On Windows* systems, the compatibility OpenMP* run-time library lets you combine OpenMP* object files compiled with the Microsoft* C/C++ compiler with OpenMP* object files compiled with the Intel C/C++ or Fortran compilers. The linking phase results in a single, coherent copy of the run-time library.

On Linux* systems, the compatibility Intel OpenMP* run-time library lets you combine OpenMP* object files compiled with the GNU* gcc or gfortran compilers with similar OpenMP* object files compiled with the Intel C/C++ or Fortran compilers. The linking phase results in a single, coherent copy of the run-time library.

You cannot link object files generated by the Intel® Fortran compiler to object files compiled by the GNU Fortran compiler, regardless of the presence or absence of the `-openmp` (Linux) or `/Qopenmp` (Windows) compiler option. This is because the Fortran run-time libraries are incompatible.



Note

The compatibility OpenMP run-time library is not compatible with object files created using versions of the Intel compiler earlier than 10.0.

Alternate Options

None

See Also

[openmp, Qopenmp](#) compiler option

[openmp-stubs, Qopenmp-stubs](#) compiler option

[openmp-profile, Qopenmp-profile](#) compiler option

openmp-link, Qopenmp-link

Controls whether the compiler links to static or dynamic OpenMP run-time libraries.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-openmp-link library`

Windows: `/Qopenmp-link:library`

Arguments

<i>library</i>	Specifies the OpenMP library to use. Possible values are:
<code>static</code>	Tells the compiler to link to static OpenMP run-time libraries.
<code>dynamic</code>	Tells the compiler to link to dynamic OpenMP run-time libraries.

Default

`-openmp-link dynamic` or `/Qopenmp-` The compiler links to dynamic OpenMP run-time libraries. However, if option

`link:dynamic` `static` is specified, the compiler links to static OpenMP run-time libraries.

Description

This option controls whether the compiler links to static or dynamic OpenMP run-time libraries.

To link to the static OpenMP run-time library (RTL) and create a purely static executable, you must specify `-openmp-link static` (Linux and Mac OS X) or `/Qopenmp-link:static` (Windows). However, we strongly recommend you use the default setting, `-openmp-link dynamic` (Linux and Mac OS X) or `/Qopenmp-link:dynamic` (Windows).



Note

Compiler options `-static-intel` and `-shared-intel` (Linux and Mac OS X) have no effect on which OpenMP run-time library is linked.

Alternate Options

None

openmp-profile, Qopenmp-profile

Enables analysis of OpenMP* applications if Intel® Thread Profiler is installed.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux: `-openmp-profile`

Mac OS X: `None`

Windows: `/Qopenmp-profile`

Arguments

None

Default

OFF OpenMP applications are not analyzed.

Description

This option enables analysis of OpenMP* applications. To use this option, you must have previously installed Intel® Thread Profiler, which is one of the Intel® Threading Analysis Tools.

This option can adversely affect performance because of the additional profiling and error checking invoked to enable compatibility with the threading tools. Do not use this option unless you plan to use the Intel® Thread Profiler.

For more information about Intel® Thread Profiler (including an evaluation copy) open the page associated with threading tools at [Intel® Software Development Products](#).

Alternate Options

None

openmp-report, Qopenmp-report

Controls the OpenMP* parallelizer's level of diagnostic messages.

IDE Equivalent

Windows: **Compilation Diagnostics > OpenMP Diagnostic Level**

Linux: None

Mac OS X: **Compiler Diagnostics > OpenMP Report**

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-openmp-report [n]`

Windows: `/Qopenmp-report [n]`

Arguments

<i>n</i>	Is the level of diagnostic messages to display. Possible values are:
0	No diagnostic messages are displayed.
1	Diagnostic messages are displayed indicating loops, regions, and sections successfully parallelized.
2	The same diagnostic messages are displayed as specified by <code>openmp_report1</code> plus diagnostic messages indicating successful handling of MASTER constructs, SINGLE constructs, CRITICAL constructs, ORDERED constructs, ATOMIC directives, and so forth.

Default

<code>-openmp-report1</code>	This is the default if you do not specify <i>n</i> .
<code>or/Qopenmp-report1</code>	The compiler displays diagnostic messages

indicating loops, regions, and sections successfully parallelized. If you do not specify the option on the command line, the default is to display no messages.

Description

This option controls the OpenMP* parallelizer's level of diagnostic messages. To use this option, you must also specify `-openmp` (Linux and Mac OS X) or `/Qopenmp` (Windows).

If this option is specified on the command line, the report is sent to `stdout`. On Windows systems, if this option is specified from within the IDE, the report is included in the build log if the Generate Build Logs option is selected.

Alternate Options

None

See Also

[openmp, Qopenmp](#) compiler option

Optimizing Applications:

Using Parallelism

OpenMP* Report

openmp-stubs, Qopenmp-stubs

Enables compilation of OpenMP programs in sequential mode.

IDE Equivalent

Windows: **Language > Process OpenMP Directives**

Linux: None

Mac OS X: **Language > Process OpenMP Directives**

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-openmp-stubs`

Windows: `/Qopenmp-stubs`

Arguments

None

Default

OFF The library of OpenMP function stubs is not linked.

Description

This option enables compilation of OpenMP programs in sequential mode. The OpenMP directives are ignored and a stub OpenMP library is linked.

Alternate Options

None

See Also

[openmp, Qopenmp](#) compiler option

openmp-threadprivate, Qopenmp-threadprivate

Lets you specify an OpenMP* threadprivate implementation.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux: `-openmp-threadprivate type`
Mac OS X: None
Windows: `/Qopenmp-threadprivate:type`

Arguments

type Specifies the type of threadprivate implementation.

Possible values are:

`legacy` Tells the compiler to use the legacy OpenMP* threadprivate implementation used in the previous releases of the Intel® compiler. This setting does not provide compatibility with the implementation used by other compilers.

`compat` Tells the compiler to use the compatibility OpenMP* threadprivate implementation based on applying the thread-local attribute to each threadprivate variable. This setting provides compatibility with the implementation provided by the Microsoft* and GNU* compilers.

Default

`-openmp-threadprivate legacy` or `/Qopenmp-` The compiler uses the legacy OpenMP* threadprivate implementation used in the previous releases of the Intel® compiler.

`threadprivate:legacy`

Description

This option lets you specify an OpenMP* threadprivate implementation.

The legacy OpenMP run-time library is not compatible with object files created using OpenMP run-time libraries supported in other compilers.

To use this option, you must also specify one of the following compiler options:

- Linux OS: `-openmp`, `-openmp-profile`, or `-openmp-stubs`
- Windows OS: `/Qopenmp`, `/Qopenmp-profile`, or `/Qopenmp-stubs`

The value specified for this option is independent of the value used for option `-openmp-lib` (Linux) or `/Qopenmp-lib` (Windows).

Alternate Options

None

opt-block-factor, Qopt-block-factor

Lets you specify a loop blocking factor.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-opt-block-factor=n`

Windows: `/Qopt-block-factor:n`

Arguments

n Is the blocking factor. It must be an integer. The compiler may ignore the blocking factor if the value

is 0 or 1.

Default

OFF The compiler uses default heuristics for loop blocking.

Description

This option lets you specify a loop blocking factor.

Alternate Options

None

opt-jump-tables, Qopt-jump-tables

Enables or disables generation of jump tables for switch statements.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-opt-jump-tables=keyword`
`-no-opt-jump-tables`

Windows: `/Qopt-jump-tables:keyword`
`/Qopt-jump-tables-`

Arguments

keyword Is the instruction for generating jump tables.
Possible values are:

<code>never</code>	Tells the compiler to never generate jump tables. All
--------------------	---

switch statements are implemented as chains of if-then-elses. This is the same as specifying `-no-opt-jump-tables` (Linux and Mac OS) or `/Qopt-jump-tables-` (Windows).

<code>default</code>	The compiler uses default heuristics to determine when to generate jump tables.
<code>large</code>	Tells the compiler to generate jump tables up to a certain pre-defined size (64K entries).
<code>n</code>	Must be an integer. Tells the compiler to generate jump tables up to n entries in size.

Default

`-opt-jump-tables=default` or `/Qopt-jump-tables:default` The compiler uses default heuristics to determine when to generate jump tables for switch statements.

Description

This option enables or disables generation of jump tables for switch statements. When the option is enabled, it may improve performance for programs with large switch statements.

Alternate Options

None

opt-loadpair, Qopt-loadpair

Enables or disables loadpair optimization.

IDE Equivalent

None

Architectures

IA-64 architecture

Syntax

Linux:	<code>-opt-loadpair</code> <code>-no-opt-loadpair</code>
Mac OS X:	None
Windows:	<code>/Qopt-loadpair</code> <code>/Qopt-loadpair-</code>

Arguments

None

Default

<code>-no-opt-loadpair</code>	Loadpair optimization is disabled
<code>or/Qopt-loadpair-</code>	unless option <code>O3</code> is specified.

Description

This option enables or disables loadpair optimization.

When `-O3` is specified on IA-64 architecture, loadpair optimization is enabled by default. To disable loadpair generation, specify `-no-opt-loadpair` (Linux) or `/Qopt-loadpair-` (Windows).

Alternate Options

None

opt-malloc-options

Lets you specify an alternate algorithm for malloc().

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-opt-malloc-options=n`

Windows: None

Arguments

- n* Specifies the algorithm to use for malloc().
Possible values are:
- 0 Tells the compiler to use the default algorithm for malloc(). This is the default.
 - 1 Causes the following adjustments to the malloc() algorithm:
 M_MMAP_MAX=2 and
 M_TRIM_THRESHOLD=0x10000000.
 - 2 Causes the following adjustments to the malloc() algorithm:
 M_MMAP_MAX=2 and
 M_TRIM_THRESHOLD=0x40000000.
 - 3 Causes the following adjustments to the malloc() algorithm:
 M_MMAP_MAX=0 and

M_TRIM_THRESHOLD=-1.

- 4 Causes the following adjustments to the malloc() algorithm: M_MMAP_MAX=0, M_TRIM_THRESHOLD=-1, M_TOP_PAD=4096.

Default

`-opt-malloc-options=0` The compiler uses the default algorithm when malloc() is called. No call is made to mallopt().

Description

This option lets you specify an alternate algorithm for malloc(). If you specify a non-zero value for *n*, it causes alternate configuration parameters to be set for how malloc() allocates and frees memory. It tells the compiler to insert calls to mallopt() to adjust these parameters to malloc() for dynamic memory allocation. This may improve speed.

Alternate Options

None

See Also

malloc(3) man page
mallopt function (defined in malloc.h)

opt-mem-bandwidth, Qopt-mem-bandwidth

Enables performance tuning and heuristics that control memory bandwidth use among processors.

IDE Equivalent

None

Architectures

IA-64 architecture

Syntax

Linux: `-opt-mem-bandwidthn`
 Mac OS X: `None`
 Windows: `/Qopt-mem-bandwidthn`

Arguments

n Is the level of optimizing for memory bandwidth usage. Possible values are:

- 0 Enables a set of performance tuning and heuristics in compiler optimizations that is optimal for serial code.
- 1 Enables a set of performance tuning and heuristics in compiler optimizations for multithreaded code generated by the compiler.
- 2 Enables a set of performance tuning and heuristics in compiler optimizations for parallel code such as Windows Threads, pthreads, and MPI code, besides multithreaded code generated by the compiler.

Default

<code>-opt-mem-bandwidth0</code> or <code>/Qopt-mem-bandwidth0</code>	For serial (non-parallel) compilation, a set of performance tuning and heuristics in compiler optimizations is enabled that is optimal for serial code.
<code>-opt-mem-bandwidth1</code> or <code>/Qopt-mem-bandwidth1</code>	If you specify compiler option <code>-parallel</code> (Linux) or <code>/Qparallel</code> (Windows), <code>-openmp</code> (Linux) or <code>/Qopenmp</code> (Windows), or Cluster OpenMP option <code>-cluster-openmp</code> (Linux), a set of performance tuning and heuristics in compiler optimizations for multithreaded code generated by the compiler is enabled.

Description

This option enables performance tuning and heuristics that control memory bandwidth use among processors. It allows the compiler to be less aggressive with optimizations that might consume more bandwidth, so that the bandwidth can be well-shared among multiple processors for a parallel program.

For values of n greater than 0, the option tells the compiler to enable a set of performance tuning and heuristics in compiler optimizations such as prefetching, privatization, aggressive code motion, and so forth, for reducing memory bandwidth pressure and balancing memory bandwidth traffic among threads.

This option can improve performance for threaded or parallel applications on multiprocessors or multicore processors, especially when the applications are bounded by memory bandwidth.

Alternate Options

None

See Also

[parallel, Qparallel](#) compiler option

[openmp, Qopenmp](#) compiler option

opt-mod-versioning, Qopt-mod-versioning

Enables or disables versioning of modulo operations for certain types of operands.

IDE Equivalent

None

Architectures

IA-64 architecture

Syntax

Linux: `-opt-mod-versioning`
 `-no-opt-mod-versioning`

Mac OS X: None

Windows: `/Qopt-mod-versioning`
 `/Qopt-mod-versioning-`

Arguments

None

Default

`-no-opt-mod-versioning` Versioning of modulo operations is
`or/Qopt-mod-versioning-` disabled.

Description

This option enables or disables versioning of modulo operations for certain types of operands. It is used for optimization tuning.

Versioning of modulo operations may improve performance for $x \bmod y$ when modulus y is a power of 2.

Alternate Options

None

opt-multi-version-aggressive, Qopt-multi-version-aggressive

Tells the compiler to use aggressive multi-versioning to check for pointer aliasing and scalar replacement.

IDE Equivalent

None

Architectures

IA-32, Intel® 64 architectures

Syntax

Linux and Mac OS X: `-opt-multi-version-aggressive`
`-no-opt-multi-version-aggressive`

Windows: `/Qopt-multi-version-aggressive`
`/Qopt-multi-version-aggressive-`

Arguments

None

Default

<code>-no-opt-multi-version-aggressive</code>	The compiler uses default heuristics when checking for pointer aliasing and scalar replacement.
<code>or/Qopt-multi-version-aggressive-</code>	

Description

This option tells the compiler to use aggressive multi-versioning to check for pointer aliasing and scalar replacement. This option may improve performance.

Alternate Options

None

opt-prefetch, Qopt-prefetch

Enables or disables prefetch insertion optimization.

IDE Equivalent

Windows: **Optimization > Prefetch Insertion**

Linux: None

Mac OS X: **Optimization > Enable Prefetch Insertion**

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-opt-prefetch[=n]`
`-no-opt-prefetch`

Windows: `/Qopt-prefetch[:n]`
`/Qopt-prefetch-`

Arguments

<i>n</i>	Is the level of detail in the report. Possible values are:
0	Disables software prefetching. This is the same as specifying <code>-no-opt-prefetch</code> (Linux and Mac OS X) or <code>/Qopt-prefetch-</code> (Windows).
1 to 4	Enables different levels of software prefetching. If you do not specify a value for <i>n</i> , the

default is 2 on IA-32 and Intel® 64 architecture; the default is 3 on IA-64 architecture. Use lower values to reduce the amount of prefetching.

Default

IA-64 architecture: `-opt-prefetch`
or `/Qopt-prefetch`

On IA-64 architecture, prefetch insertion optimization is enabled.

IA-32 architecture and Intel® 64 architecture:
`-no-opt-prefetch`
or `/Qopt-prefetch-`

On IA-32 architecture and Intel® 64 architecture, prefetch insertion optimization is disabled.

Description

This option enables or disables prefetch insertion optimization. The goal of prefetching is to reduce cache misses by providing hints to the processor about when data should be loaded into the cache.

On IA-64 architecture, this option is enabled by default if you specify option `01` or higher. To disable prefetching at these optimization levels, specify `-no-opt-prefetch` (Linux and Mac OS X) or `/Qopt-prefetch-` (Windows).

On IA-32 architecture and Intel® 64 architecture, this option enables prefetching when higher optimization levels are specified.

Alternate Options

Linux and Mac OS X: `-prefetch` (this is a [deprecated](#) option)

Windows: `/Qprefetch` (this is a [deprecated](#) option)

opt-prefetch-initial-values, Qopt-prefetch-initial-values

Enables or disables prefetches that are issued before a loop is entered.

IDE Equivalent

None

Architectures

IA-64 architecture

Syntax

Linux: `-opt-prefetch-initial-values`
 `-no-opt-prefetch-initial-values`

Mac OS X: None

Windows: `/Qopt-prefetch-initial-values`
 `/Qopt-prefetch-initial-values-`

Arguments

None

Default

`-opt-prefetch-initial-values` Prefetches are issued before a loop is entered.

or `/Qopt-prefetch-initial-values`
`values`

Description

This option enables or disables prefetches that are issued before a loop is entered. These prefetches target the initial iterations of the loop.

When `-O1` or higher is specified on IA-64 architecture, prefetches are issued before a loop is entered. To disable these prefetches, specify `-no-opt-prefetch-initial-values` (Linux) or `/Qopt-prefetch-initial-values-` (Windows).

Alternate Options

None

opt-prefetch-issue-excl-hint, Qopt-prefetch-issue-excl-hint

Determines whether the compiler issues prefetches for stores with exclusive hint.

IDE Equivalent

None

Architectures

IA-64 architecture

Syntax

Linux:	<code>-opt-prefetch-issue-excl-hint</code> <code>-no-opt-prefetch-issue-excl-hint</code>
Mac OS X:	None
Windows:	<code>/Qopt-prefetch-issue-excl-hint</code> <code>/Qopt-prefetch-issue-excl-hint-</code>

Arguments

None

Default

<code>-no-opt-prefetch-issue-excl-hint</code>	The compiler does not issue prefetches for stores with exclusive hint.
<code>or/Qopt-prefetch-issue-excl-hint-</code>	

Description

This option determines whether the compiler issues prefetches for stores with exclusive hint. If option `-opt-prefetch-issue-excl-hint` (Linux) or

`/Qopt-prefetch-issue-excl-hint` (Windows) is specified, the prefetches will be issued if the compiler determines it is beneficial to do so.

When prefetches are issued for stores with exclusive-hint, the cache-line is in "exclusive-mode". This saves on cache-coherence traffic when other processors try to access the same cache-line. This feature can improve performance tuning.

Alternate Options

None

opt-prefetch-next-iteration, Qopt-prefetch-next-iteration

Enables or disables prefetches for a memory access in the next iteration of a loop.

IDE Equivalent

None

Architectures

IA-64 architecture

Syntax

Linux: `-opt-prefetch-next-iteration`
 `-no-opt-prefetch-next-iteration`

Mac OS X: None

Windows: `/Qopt-prefetch-next-iteration`
 `/Qopt-prefetch-next-iteration-`

Arguments

None

Default

`-opt-prefetch-next-` Prefetches are issued for a memory

`iteration` access in the next iteration of a
`or/Qopt-prefetch-next-` loop.
`iteration`

Description

This option enables or disables prefetches for a memory access in the next iteration of a loop. It is typically used in a pointer-chasing loop.

When `-O1` or higher is specified on IA-64 architecture, prefetches are issued for a memory access in the next iteration of a loop. To disable these prefetches, specify `-no-opt-prefetch-next-iteration` (Linux) or `/Qopt-prefetch-next-iteration-` (Windows).

Alternate Options

None

opt-ra-region-strategy, Qopt-ra-region-strategy

Selects the method that the register allocator uses to partition each routine into regions.

IDE Equivalent

None

Architectures

IA-32, Intel® 64 architectures

Syntax

Linux and Mac OS X: `-opt-ra-region-strategy[=keyword]`

Windows: `/Qopt-ra-region-strategy[:keyword]`

Arguments

keyword Is the method used for partitioning. Possible

values are:

<code>routine</code>	Creates a single region for each routine.
<code>block</code>	Partitions each routine into one region per basic block.
<code>trace</code>	Partitions each routine into one region per trace.
<code>region</code>	Partitions each routine into one region per loop.
<code>default</code>	The compiler determines which method is used for partitioning.

Default

<code>-opt-ra-region-strategy=default</code> or <code>/Qopt-ra-region-strategy:default</code>	The compiler determines which method is used for partitioning. This is also the default if <code>keyword</code> is not specified.
--	---

Description

This option selects the method that the register allocator uses to partition each routine into regions.

When setting `default` is in effect, the compiler attempts to optimize the tradeoff between compile-time performance and generated code performance.

This option is only relevant when optimizations are enabled (O1 or higher).

Alternate Options

None

See Also

[o](#) compiler option

opt-report, Qopt-report

Tells the compiler to generate an optimization report to stderr.

IDE Equivalent

Windows: Diagnostics > Optimization Diagnostics Level

Linux: None

Mac OS X: None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-opt-report [n]`

Windows: `/Qopt-report[:n]`

Arguments

<i>n</i>	Is the level of detail in the report. Possible values are:
0	Tells the compiler to generate no optimization report.
1	Tells the compiler to generate a report with the minimum level of detail.
2	Tells the compiler to generate a report with the medium level of detail.
3	Tells the compiler to generate a report with the maximum

level of detail.

Default

`-opt-report 2` or `/Qopt-report:2` If you do not specify *n*, the compiler generates a report with medium detail. If you do not specify the option on the command line, the compiler does not generate an optimization report.

Description

This option tells the compiler to generate an optimization report to `stderr`.

Alternate Options

None

See Also

[opt-report-file, Qopt-report-file](#) compiler option
Optimizing Applications: Optimizer Report Generation

opt-report-file, Qopt-report-file

Specifies the name for an optimization report.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-opt-report-file=file`

Windows: `/Qopt-report-file:file`

Arguments

file Is the name for the optimization report.

Default

OFF No optimization report is generated.

Description

This option specifies the name for an optimization report. If you use this option, you do not have to specify `-opt-report` (Linux and Mac OS X) or `/Qopt-report` (Windows).

Alternate Options

None

See Also

[opt-report, Qopt-report](#) compiler option

Optimizing Applications: Optimizer Report Generation

opt-report-help, Qopt-report-help

Displays the optimizer phases available for report generation.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-opt-report-help`

Windows: `/Qopt-report-help`

Arguments

None

Default

OFF No optimization reports are generated.

Description

This option displays the optimizer phases available for report generation using `-opt-report-phase` (Linux and Mac OS X) or `/Qopt-report-phase` (Windows). No compilation is performed.

Alternate Options

None

See Also

[opt-report, Qopt-report](#) compiler option

[opt-report-phase, Qopt-report-phase](#) compiler option

opt-report-phase, Qopt-report-phase

Specifies an optimizer phase to use when optimization reports are generated.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-opt-report-phase=phase`

Windows: `/Qopt-report-phase:phase`

Arguments

phase Is the phase to generate reports for. Some of the possible values are:

<code>ipo</code>	The Interprocedural Optimizer phase
<code>hlo</code>	The High Level Optimizer phase
<code>hpo</code>	The High Performance Optimizer phase
<code>ilo</code>	The Intermediate Language Scalar Optimizer phase
<code>ecg</code>	The Code Generator phase (Windows and Linux systems using IA-64 architecture only)
<code>ecg_swp</code>	The software pipelining component of the Code Generator phase (Windows and Linux systems using IA-64 architecture only)
<code>pgo</code>	The Profile Guided Optimization phase
<code>all</code>	All optimizer phases

Default

OFF No optimization reports are generated.

Description

This option specifies an optimizer phase to use when optimization reports are generated. To use this option, you must also specify `-opt-report` (Linux and Mac OS X) or `/Qopt-report` (Windows).

This option can be used multiple times on the same command line to generate reports for multiple optimizer phases.

When one of the logical names for optimizer phases is specified for phase, all reports from that optimizer phase are generated.

To find all phase possibilities, use option `-opt-report-help` (Linux and Mac OS X) or `/Qopt-report-help` (Windows).

Alternate Options

None

See Also

[opt-report](#), [Qopt-report](#) compiler option

opt-report-routine, Qopt-report-routine

Tells the compiler to generate reports on the routines containing specified text.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-opt-report-routine=string`

Windows: `/Qopt-report-routine:string`

Arguments

string Is the text (string) to look for.

Default

OFF No optimization reports are generated.

Description

This option tells the compiler to generate reports on the routines containing specified text as part of their name.

Alternate Options

None

See Also

[opt-report](#), [Qopt-report](#) compiler option

opt-streaming-stores, Qopt-streaming-stores

Enables generation of streaming stores for optimization.

IDE Equivalent

None

Architectures

IA-32, Intel® 64 architectures

Syntax

Linux and Mac OS X: `-opt-streaming-stores keyword`

Windows: `/Qopt-streaming-stores:keyword`

Arguments

keyword Specifies whether streaming stores are generated.

Possible values are:

<code>always</code>	Enables generation of streaming stores for optimization. The compiler optimizes under the assumption that the application
---------------------	---

	is memory bound.
never	Disables generation of streaming stores for optimization. Normal stores are performed.
auto	Lets the compiler decide which instructions to use.

Default

```
-opt-          The compiler decides whether to use streaming stores or
streaming-    normal stores.
stores auto
or/Qopt-
streaming-
stores:auto
```

Description

This option enables generation of streaming stores for optimization. This method stores data with instructions that use a non-temporal buffer, which minimizes memory hierarchy pollution.

For this option to be effective, the compiler must be able to generate SSE2 (or higher) instructions. For more information, see compiler option `x` or `ax`.

This option may be useful for applications that can benefit from streaming stores.

Alternate Options

None

See Also

[ax, Qax](#) compiler option

[x, Qx](#) compiler option

[opt-mem-bandwidth](#), [Qopt-mem-bandwidth](#), [Qx](#) compiler option

Optimizing Applications: Vectorization Support

opt-subscript-in-range, Qopt-subscript-in-range

Determines whether the compiler assumes no overflows in the intermediate computation of subscript expressions in loops.

IDE Equivalent

None

Architectures

IA-32, Intel® 64 architectures

Syntax

Linux and Mac OS X: `-opt-subscript-in-range`
`-no-opt-subscript-in-range`

Windows: `/Qopt-subscript-in-range`
`/Qopt-subscript-in-range-`

Arguments

None

Default

`-no-opt-subscript-in-range` The compiler assumes overflows in the intermediate computation of subscript expressions in loops.

or `/Qopt-subscript-in-range-`

Description

This option determines whether the compiler assumes no overflows in the intermediate computation of subscript expressions in loops.

If you specify `-opt-subscript-in-range` (Linux and Mac OS X) or `/Qopt-subscript-in-range` (Windows), the compiler ignores any data type conversions used and it assumes no overflows in the intermediate computation of subscript expressions. This feature can enable more loop transformations.

Alternate Options

None

Example

The following shows an example where these options can be useful. `m` is declared as type `integer(kind=8)` (64-bits) and all other variables inside the subscript are declared as type `integer(kind=4)` (32-bits):

```
A[ i + j + ( n + k ) * m ]
```

optimize

See [O](#).

Os

Enables optimizations that do not increase code size and produces smaller code size than `O2`.

IDE Equivalent

Windows: **Optimization > Favor Size or Speed**

Linux: None

Mac OS X: None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-Os`

Windows: `/Os`

Arguments

None

Default

OFF Optimizations are made for code speed. However, if `O1` is specified, `Os` is the default.

Description

This option enables optimizations that do not increase code size and produces smaller code size than `O2`. It disables some optimizations that increase code size for a small speed benefit.

This option tells the compiler to favor transformations that reduce code size over transformations that produce maximum performance.

Alternate Options

None

See Also

[O](#) compiler option

[Ot](#) compiler option

Ot

Enables all speed optimizations.

IDE Equivalent

Windows: **Optimization > Favor Size or Speed** (`/Ot`, `/Os`)

Linux: None

Mac OS X: None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: None

Windows: /O_t

Arguments

None

Default

/O_t Optimizations are made for code speed.

If O_d is specified, all optimizations are disabled. If O₁ is specified, O_s is the default.

Description

This option enables all speed optimizations.

Alternate Options

None

See Also

[O](#) compiler option

[Os](#) compiler option

Ox

See [O](#).

fomit-frame-pointer, Oy

Determines whether EBP is used as a general-purpose register in optimizations.

IDE Equivalent

Windows: **Optimization > Omit Frame Pointers**

Linux: None

Mac OS X: **Optimization > Provide Frame Pointer**

Architectures

`-f[no-]omit-frame-pointer`: IA-32 architecture, Intel® 64 architecture

`/Oy[-]`: IA-32 architecture

Syntax

Linux and Mac OS X: `-fomit-frame-pointer`

`-fno-omit-frame-pointer`

Windows: `/Oy`

`/Oy-`

Arguments

None

Default

`-fomit-frame-pointer` or `/Oy` EBP is used as a general-purpose register in optimizations. However, on Linux* and Mac OS X systems, the default is `-fno-omit-frame-pointer` if option `-O0` or `-g` is specified. On Windows* systems, the default is `/Oy-` if option `/Od` is specified.

Description

These options determine whether EBP is used as a general-purpose register in optimizations. Options `-fomit-frame-pointer` and `/Oy` allow this use.

Options `-fno-omit-frame-pointer` and `/Oy-` disallow it.

Some debuggers expect EBP to be used as a stack frame pointer, and cannot produce a stack backtrace unless this is so. The `-fno-omit-frame-pointer`

and `/Oy-` options direct the compiler to generate code that maintains and uses `EBP` as a stack frame pointer for all functions so that a debugger can still produce a stack backtrace without doing the following:

- For `-fno-omit-frame-pointer`: turning off optimizations with `-O0`
- For `/Oy-`: turning off `/O1`, `/O2`, or `/O3` optimizations

The `-fno-omit-frame-pointer` option is set when you specify option `-O0` or the `-g` option. The `-fomit-frame-pointer` option is set when you specify option `-O1`, `-O2`, or `-O3`.

The `/Oy` option is set when you specify the `/O1`, `/O2`, or `/O3` option. Option `/Oy-` is set when you specify the `/Od` option.

Using the `-fno-omit-frame-pointer` or `/Oy-` option reduces the number of available general-purpose registers by 1, and can result in slightly less efficient code.

Alternate Options

Linux and Mac OS X: `-fp` (this is a [deprecated](#) option)

Windows: None

p

Compiles and links for function profiling with `gprof(1)`.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-p`

Windows: None

Arguments

None

Default

OFF Files are compiled and linked without profiling.

Description

This option compiles and links for function profiling with gprof(1).

Alternate Options

Linux and Mac OS X: `-pg` (only available on systems using IA-32 architecture or Intel® 64 architecture), `-qp` (this is a [deprecated](#) option)

Windows: None

P

See [preprocess-only](#).

pad, Qpad

Enables the changing of the variable and array memory layout.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-pad`
`-nopad`

Windows: `/Qpad`
`/Qpad-`

Arguments

None

Default

`-nopad` Variable and array memory layout is performed by default
or `pad` methods.
`/Qpad-`

Description

This option enables the changing of the variable and array memory layout. This option is effectively not different from the `align` option when applied to structures and derived types. However, the scope of `pad` is greater because it applies also to common blocks, derived types, sequence types, and structures.

Alternate Options

None

See Also

[align](#) compiler option

pad-source, Qpad-source

Specifies padding for fixed-form source records.

IDE Equivalent

Windows: **Language > Pad Fixed Form Source Lines**

Linux: None

Mac OS X: **Language > Pad Fixed Form Source Lines**

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-pad-source`
`-nopad-source`

Windows: `/pad-source`
`/nopad-source`
`/Qpad-source`
`/Qpad-source-`

Arguments

None

Default

`-nopad-` Fixed-form source records are not padded.
source
or
`/Qpad-`
source-

Description

This option specifies padding for fixed-form source records. It tells the compiler that fixed-form source lines shorter than the statement field width are to be padded with spaces to the end of the statement field. This affects the interpretation of character and Hollerith literals that are continued across source records.

The default value setting causes a warning message to be displayed if a character or Hollerith literal that ends before the statement field ends is continued onto the next source record. To suppress this warning message, specify option `-warn nousage` (Linux and Mac OS X) or `/warn:nousage` (Windows).

Specifying `pad-source` or `/Qpad-source` can prevent warning messages associated with option `-warn usage` (Linux and Mac OS X) or `/warn:usage` (Windows).

Alternate Options

None

See Also

[warn](#) compiler option

par-report, Qpar-report

Controls the diagnostic information reported by the auto-parallelizer.

IDE Equivalent

Windows: **Compilation Diagnostics > Auto-Parallelizer Diagnostic Level**

Linux: None

Mac OS X: **Diagnostics > Auto-Parallelizer Report**

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-par-report[n]`

Windows: `/Qpar-report[n]`

Arguments

n Is a value denoting which diagnostic messages to report. Possible values are:

0 Tells the auto-parallelizer to report no diagnostic information.

1 Tells the auto-parallelizer to report diagnostic messages for loops successfully auto-

- parallelized. The compiler also issues a "LOOP AUTO-PARALLELIZED" message for parallel loops.
- 2 Tells the auto-parallelizer to report diagnostic messages for loops successfully and unsuccessfully auto-parallelized.
- 3 Tells the auto-parallelizer to report the same diagnostic messages specified by 2 plus additional information about any proven or assumed dependencies inhibiting auto-parallelization (reasons for not parallelizing).

Default

`-par-report1` or `/Qpar-report1` If you do not specify n , the compiler displays diagnostic messages for loops successfully auto-parallelized. If you do not specify the option on the command line, the default is to display no parallel diagnostic messages.

Description

This option controls the diagnostic information reported by the auto-parallelizer (parallel optimizer). To use this option, you must also specify `-parallel` (Linux and Mac OS X) or `/Qparallel` (Windows).

If this option is specified on the command line, the report is sent to `stdout`.

On Windows systems, if this option is specified from within the IDE, the report is included in the build log if the Generate Build Logs option is selected.

Alternate Options

None

par-runtime-control, Qpar-runtime-control

Generates code to perform run-time checks for loops that have symbolic loop bounds.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-par-runtime-control`
`-no-par-runtime-control`

Windows: `/Qpar-runtime-control`
`/Qpar-runtime-control-`

Arguments

None

Default

`-no-par-runtime-control` or `/Qpar-runtime-control-` The compiler uses default heuristics when checking loops.

Description

This option generates code to perform run-time checks for loops that have symbolic loop bounds.

If the granularity of a loop is greater than the parallelization threshold, the loop will be executed in parallel.

If you do not specify this option, the compiler may not parallelize loops with symbolic loop bounds if the compile-time granularity estimation of a loop can not ensure it is beneficial to parallelize the loop.

Alternate Options

None

par-schedule, Qpar-schedule

Lets you specify a scheduling algorithm or a tuning method for loop iterations.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-par-schedule-keyword[=n]`

Windows: `/Qpar-schedule-keyword[[:]n]`

Arguments

keyword Specifies the scheduling algorithm or tuning method. Possible values are:

`auto` Lets the compiler or run-time system determine the scheduling algorithm.

`static` Divides iterations into

contiguous pieces.

`static-balanced` Divides iterations into even-sized chunks.

`static-steal` Divides iterations into even-sized chunks, but allows threads to steal parts of chunks from neighboring threads.

`dynamic` Gets a set of iterations dynamically.

`guided` Specifies a minimum number of iterations.

`guided-analytical` Divides iterations by using exponential distribution or dynamic distribution.

`runtime` Defers the scheduling decision until run time.

n Is the size of the chunk or the number of iterations for each chunk. This setting can only be specified for `static`, `dynamic`, and `guided`. For more information, see the descriptions of each keyword below.

Default

`static-balanced` Iterations are divided into even-sized chunks and the chunks are assigned to the threads in the team in a round-robin fashion in the order of the thread number.

Description

This option lets you specify a scheduling algorithm or a tuning method for loop iterations. It specifies how iterations are to be divided among the threads of the team.

This option affects performance tuning and can provide better performance during auto-parallelization.

Option	Description
<code>-par-schedule-auto</code> or <code>/Qpar-schedule-auto</code>	Lets the compiler or run-time system determine the scheduling algorithm. Any possible mapping may occur for iterations to threads in the team.
<code>-par-schedule-static</code> or <code>/Qpar-schedule-static</code>	Divides iterations into contiguous pieces (chunks) of size n . The chunks are assigned to threads in the team in a round-robin fashion in the order of the thread number. Note that the last chunk to be assigned may have a smaller number of iterations. If no n is specified, the iteration space is divided into chunks that are approximately equal in size, and each thread is assigned at most one chunk.
<code>-par-schedule-static-balanced</code> or <code>/Qpar-schedule-static-balanced</code>	Divides iterations into even-sized chunks. The chunks are assigned to the threads in the team in a round-robin fashion in the order of the thread number.
<code>-par-schedule-static-steal</code> or <code>/Qpar-schedule-static-steal</code>	Divides iterations into even-sized chunks, but when a thread completes its chunk, it can steal parts of chunks

Option	Description
	<p>assigned to neighboring threads.</p> <p>Each thread keeps track of L and U, which represent the lower and upper bounds of its chunks respectively.</p> <p>Iterations are executed starting from the lower bound, and simultaneously, L is updated to represent the new lower bound.</p>
<p><code>-par-schedule-dynamic</code> or <code>/Qpar-schedule-dynamic</code></p>	<p>Can be used to get a set of iterations dynamically. Assigns iterations to threads in chunks as the threads request them. The thread executes the chunk of iterations, then requests another chunk, until no chunks remain to be assigned.</p> <p>As each thread finishes a piece of the iteration space, it dynamically gets the next set of iterations. Each chunk contains n iterations, except for the last chunk to be assigned, which may have fewer iterations. If no n is specified, the default is 1.</p>
<p><code>-par-schedule-guided</code> or <code>/Qpar-schedule-guided</code></p>	<p>Can be used to specify a minimum number of iterations. Assigns iterations to threads in chunks as the threads request them. The thread executes the chunk of iterations, then requests another chunk, until no chunks remain to be assigned.</p>

Option	Description
	<p>For a chunk of size 1, the size of each chunk is proportional to the number of unassigned iterations divided by the number of threads, decreasing to 1. For an n with value k (greater than 1), the size of each chunk is determined in the same way with the restriction that the chunks do not contain fewer than k iterations (except for the last chunk to be assigned, which may have fewer than k iterations). If no n is specified, the default is 1.</p>
<code>-par-schedule-guided-analytical</code> or <code>/Qpar-schedule-guided-analytical</code>	<p>Divides iterations by using exponential distribution or dynamic distribution. The method depends on run-time implementation. Loop bounds are calculated with faster synchronization and chunks are dynamically dispatched at run time by threads in the team.</p>
<code>-par-schedule-runtime</code> or <code>/Qpar-schedule-runtime</code>	<p>Defers the scheduling decision until run time. The scheduling algorithm and chunk size are then taken from the setting of environment variable <code>OMP_SCHEDULE</code>.</p>

Alternate Options

None

par-threshold, Qpar-threshold

Sets a threshold for the auto-parallelization of loops.

IDE Equivalent

Windows: **Optimization > Threshold For Auto-Parallelization**

Linux: None

Mac OS X: **Optimization > Threshold For Auto-Parallelization**

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-par-threshold[n]`

Windows: `/Qpar-threshold[[:]n]`

Arguments

n Is an integer whose value is the threshold for the auto-parallelization of loops. Possible values are 0 through 100.

If *n* is 0, loops get auto-parallelized always, regardless of computation work volume.

If *n* is 100, loops get auto-parallelized when performance gains are predicted based on the compiler analysis data. Loops get auto-parallelized only if profitable parallel execution is almost certain.

The intermediate 1 to 99 values represent the percentage probability for profitable speed-up. For example, *n*=50 directs the compiler to parallelize only if there is a 50% probability of the code speeding up if executed in parallel.

Default

<code>-par-</code>	Loops get auto-parallelized only if profitable
<code>threshold100</code>	parallel execution is almost certain. This is also the
<code>or/Qpar-</code>	default if you do not specify <i>n</i> .
<code>threshold100</code>	

Description

This option sets a threshold for the auto-parallelization of loops based on the probability of profitable execution of the loop in parallel. To use this option, you must also specify `-parallel` (Linux and Mac OS X) or `/Qparallel` (Windows). This option is useful for loops whose computation work volume cannot be determined at compile-time. The threshold is usually relevant when the loop trip count is unknown at compile-time.

The compiler applies a heuristic that tries to balance the overhead of creating multiple threads versus the amount of work available to be shared amongst the threads.

Alternate Options

None

parallel, Qparallel

Tells the auto-parallelizer to generate multithreaded code for loops that can be safely executed in parallel.

IDE Equivalent

Windows: **Optimization > Parallelization**

Linux: None

Mac OS X: **Optimization > Parallelization**

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-parallel`

Windows: `/Qparallel`

Arguments

None

Default

OFF Multithreaded code is not generated for loops that can be safely executed in parallel.

Description

This option tells the auto-parallelizer to generate multithreaded code for loops that can be safely executed in parallel.

To use this option, you must also specify option `O2` or `O3`.



Note

On Mac OS X systems, when you enable automatic parallelization, you must also set the `DYLD_LIBRARY_PATH` environment variable within Xcode or an error will be displayed.

Alternate Options

None

See Also

[Q compiler option](#)

pc, Qpc

Enables control of floating-point significand precision.

IDE Equivalent

None

Architectures

IA-32, Intel® 64 architectures

Syntax

Linux and Mac OS X: `-pcn`

Windows: `/Qpcn`

Arguments

n Is the floating-point significand precision. Possible values are:

32 Rounds the significand to 24 bits (single precision).

64 Rounds the significand to 53 bits (double precision).

80 Rounds the significand to 64 bits (extended precision).

Default

`-pc80` On Linux* and Mac OS* X systems, the floating-point or `/Qpc64` significand is rounded to 64 bits. On Windows* systems, the floating-point significand is rounded to 53 bits.

Description

This option enables control of floating-point significand precision. Some floating-point algorithms are sensitive to the accuracy of the significand, or fractional part of the floating-point value. For example, iterative operations like division and finding the square root can run faster if you lower the precision with the this option.

Note that a change of the default precision control or rounding mode, for example, by using the `-pc32` (Linux and Mac OS X) or `/Qpc32` (Windows) option or by user intervention, may affect the results returned by some of the mathematical functions.

Alternate Options

None

See Also

Floating-point Operations: Floating-point Options Quick Reference

pdbfile

Specifies that any debug information generated by the compiler should be saved to a program database file.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: None

Windows: `/pdbfile[:file]`
 `/nopdbfile`

Arguments

file Is the name of the program database file.

Default

`/nopdbfile` Debug information generated by the compiler is not saved

to a program database file.

Description

This option specifies that any debug information generated by the compiler should be saved to a program database file. To use this option, you must also specify `/debug:full` (or the equivalent).

If *file* is not specified, the default file name used is the name of your file with an extension of `.pdb`.

The compiler places debug information in the object file if you specify `/nopdbfile` or omit both `/pdbfile` and `/debug:full` (or the equivalent).

Alternate Options

None

See Also

[debug \(Windows*\)](#) compiler option

pg

See [p](#).

pie

Produces a position-independent executable on processors that support it.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux: `-pie`

Mac OS X: None

Windows: None

Arguments

None

Default

OFF The driver does not set up special run-time libraries and the linker does not perform the optimizations on executables.

Description

This option produces a position-independent executable on processors that support it. It is both a compiler option and a linker option. When used as a compiler option, this option ensures the linker sets up run-time libraries correctly. Normally the object linked has been compiled with option `-fpie`. When you specify `-pie`, it is recommended that you specify the same options that were used during compilation of the object.

Alternate Options

None

See Also

[fpie](#) compiler option

prec-div, Qprec-div

Improves precision of floating-point divides.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-prec-div`
`-no-prec-div`
Windows: `/Qprec-div`
`/Qprec-div-`

Arguments

None

Default

`-prec-div` The compiler uses this method for floating-point divides.
or `/Qprec-div`
`div`

Description

This option improves precision of floating-point divides. It has a slight impact on speed.

With some optimizations, such as `-xSSE2` (Linux) or `/QxSSE2` (Windows), the compiler may change floating-point division computations into multiplication by the reciprocal of the denominator. For example, A/B is computed as $A * (1/B)$ to improve the speed of the computation.

However, sometimes the value produced by this transformation is not as accurate as full IEEE division. When it is important to have fully precise IEEE division, use this option to disable the floating-point division-to-multiplication optimization. The result is more accurate, with some loss of performance.

If you specify `-no-prec-div` (Linux and Mac OS X) or `/Qprec-div-` (Windows), it enables optimizations that give slightly less precise results than full IEEE division.

Alternate Options

None

See Also

Floating-point Operations: Floating-point Options Quick Reference

prec-sqrt, Qprec-sqrt

Improves precision of square root implementations.

IDE Equivalent

None

Architectures

IA-32, Intel® 64 architectures

Syntax

Linux and Mac OS X: `-prec-sqrt`
`-no-prec-sqrt`

Windows: `/Qprec-sqrt`
`/Qprec-sqrt-`

Arguments

None

Default

`-no-prec-sqrt` The compiler uses a faster but less precise implementation of square root.

or `/Qprec-sqrt-` Note that the default is `-prec-sqrt` or `/Qprec-sqrt` if any of the following options are specified: `/Od`, `/Op`, or `/Qprec` on Windows systems; `-O0`, `-mp` (or `-fltconsistency`), or `-mp1` on Linux and Mac OS X

systems.

Description

This option improves precision of square root implementations. It has a slight impact on speed.

This option inhibits any optimizations that can adversely affect the precision of a square root computation. The result is fully precise square root implementations, with some loss of performance.

Alternate Options

None

preprocess-only

Causes the Fortran preprocessor to send output to a file.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-preprocess-only`

Windows: `/preprocess-only`

Arguments

None

Default

OFF Preprocessed source files are output to the compiler.

Description

This option causes the Fortran preprocessor to send output to a file.

The source file is preprocessed by the Fortran preprocessor, and the result for each source file is output to a corresponding .i or .i90 file.

Note that the source file is not compiled.

Alternate Options

Linux and Mac OS X: `-P`

Windows: `/P`

print-multi-lib

Prints information about where system libraries should be found.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-print-multi-lib`

Windows: `None`

Arguments

None

Default

OFF No information is printed unless the option is specified.

Description

This option prints information about where system libraries should be found, but no compilation occurs. It is provided for compatibility with gcc.

Alternate Options

None

prof-data-order, Qprof-data-order

Enables or disables data ordering if profiling information is enabled.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux:	-prof-data-order -no-prof-data-order
Mac OS X:	None
Windows:	/Qprof-data-order /Qprof-data-order-

Arguments

None

Default

-no-prof- Data ordering is disabled.
data-order
or/Qprof-
data-
order-

Description

This option enables or disables data ordering if profiling information is enabled. It controls the use of profiling information to order static program data items.

For this option to be effective, you must do the following:

- For instrumentation compilation, you must specify `-prof-gen=globdata` (Linux) or `/Qprof-gen:globdata` (Windows).
- For feedback compilation, you must specify `-prof-use` (Linux) or `/Qprof-use` (Windows). You must not use multi-file optimization by specifying options such as option `-ipo` (Linux) or `/Qipo` (Windows), or option `-ipo-c` (Linux) or `/Qipo-c` (Windows).

Alternate Options

None

See Also

[prof-gen, Qprof-gen](#) compiler option

[prof-use, Qprof-use](#) compiler option

[prof-func-order, Qprof-func-order](#) compiler option

prof-dir, Qprof-dir

Specifies a directory for profiling information output files.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-prof-dir dir`

Windows: `/Qprof-dir dir`

Arguments

dir Is the name of the directory.

Default

OFF Profiling output files are placed in the directory where the program is compiled.

Description

This option specifies a directory for profiling information output files (*.dyn and *.dpi). The specified directory must already exist.

You should specify this option using the same directory name for both instrumentation and feedback compilations. If you move the .dyn files, you need to specify the new path.

Alternate Options

None

See Also

Floating-point Operations:
Profile-guided Optimization (PGO) Quick Reference
Coding Guidelines for Intel(R) Architectures

prof-file, Qprof-file

Specifies an alternate file name for the profiling summary files.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-prof-file file`

Windows: `/Qprof-file file`

Arguments

file Is the name of the profiling summary file.

Default

OFF The profiling summary files have the file name pgopti.*

Description

This option specifies an alternate file name for the profiling summary files. The *file* is used as the base name for files created by different profiling passes.

If you add this option to profmerge, the .dpi file will be named *file.dpi* instead of pgopti.dpi.

If you specify `-prof-genx` (Linux and Mac OS X) or `/Qprof-genx` (Windows) with this option, the .spi and .spl files will be named *file.spi* and *file.spl* instead of pgopti.spi and pgopti.spl.

If you specify `-prof-use` (Linux and Mac OS X) or `/Qprof-use` (Windows) with this option, the .dpi file will be named *file.dpi* instead of pgopti.dpi.

Alternate Options

None

See Also

[prof-gen, Qprof-gen](#) compiler option

[prof-use, Qprof-use](#) compiler option

Optimizing Applications:

Profile-guided Optimizations Overview

Coding Guidelines for Intel(R) Architectures

Profile an Application

prof-func-groups

Enables or disables function grouping if profiling information is enabled.

IDE Equivalent

None

Architectures

IA-32, Intel® 64 architectures

Syntax

Linux: `-prof-func-groups`
 `-no-prof-func-groups`

Mac OS X: None

Windows: None

Arguments

None

Default

`-no-prof-func-` Function grouping is disabled.
`groups`

Description

This option enables or disables function grouping if profiling information is enabled.

A "function grouping" is a profiling optimization in which entire routines are placed either in the cold code section or the hot code section.

If profiling information is enabled by option `-prof-use`, option `-prof-func-groups` is set and function grouping is enabled. However, if you explicitly enable `-prof-func-order` (Linux) or `/Qprof-func-order` (Windows), function ordering is performed instead of function grouping.

If you want to disable function grouping when profiling information is enabled, specify `-no-prof-func-groups`.

To set the hotness threshold for function grouping, use option `-prof-hotness-threshold` (Linux) or `/Qprof-hotness-threshold` (Windows).

Alternate Options

`-func-groups` (this is a [deprecated](#) option)

See Also

[prof-use](#), [Qprof-use](#) compiler option

[prof-func-order](#), [Qprof-func-order](#) compiler option

[prof-hotness-threshold](#), [Qprof-hotness-threshold](#) compiler option

prof-func-order, Qprof-func-order

Enables or disables function ordering if profiling information is enabled.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux: `-prof-func-order`
 `-no-prof-func-order`

Mac OS X: None

Windows: `/Qprof-func-order`
 `/Qprof-func-order-`

Arguments

None

Default

`-no-prof-` Function ordering is disabled.

`func-order`

or `/Qprof-`

`func-`

`order-`

Description

This option enables or disables function ordering if profiling information is enabled.

For this option to be effective, you must do the following:

- For instrumentation compilation, you must specify `-prof-gen=srcpos` (Linux) or `/Qprof-gen:srcpos` (Windows).
- For feedback compilation, you must specify `-prof-use` (Linux) or `/Qprof-use` (Windows). You must not use multi-file optimization by specifying options such as option `-ipo` (Linux) or `/Qipo` (Windows), or option `-ipo-c` (Linux) or `/Qipo-c` (Windows).

If you enable profiling information by specifying option `-prof-use` (Linux) or `/Qprof-use` (Windows), `-prof-func-groups` (Linux) and `/Qprof-func-groups` (Windows) are set and function grouping is enabled. However, if you explicitly enable `-prof-func-order` (Linux) or `/Qprof-func-order` (Windows), function ordering is performed instead of function grouping.

On Linux* systems, this option is only available for Linux linker 2.15.94.0.1, or later.

To set the hotness threshold for function grouping and function ordering, use option `-prof-hotness-threshold` (Linux) or `/Qprof-hotness-threshold` (Windows).

Alternate Options

None

The following example shows how to use this option on a Windows system:

```
ifort /Qprof-gen:globdata file1.f90 file2.f90 /exe:instrumented.exe
```

```
./instrumented.exe
ifort /Qprof-use /Qprof-func-order file1.f90 file2.f90
/exe:feedback.exe
```

The following example shows how to use this option on a Linux system:

```
ifort -prof-gen:globdata file1.f90 file2.f90 -o instrumented
./instrumented.exe
ifort -prof-use -prof-func-order file1.f90 file2.f90 -o feedback
```

See Also

[prof-hotness-threshold](#), [Qprof-hotness-threshold](#) compiler option

[prof-gen](#), [Qprof-gen](#) compiler option

[prof-use](#), [Qprof-use](#) compiler option

[prof-data-order](#), [Qprof-data-order](#) compiler option

[prof-func-groups](#) compiler option

prof-gen, Qprof-gen

Produces an instrumented object file that can be used in profile-guided optimization.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-prof-gen[=keyword]`

`-no-prof-gen`

Windows: `/Qprof-gen[:keyword]`

`/Qprof-gen-`

Arguments

keyword Specifies details for the instrumented file. Possible values are:

`default` Produces an instrumented

	object file. This is the same as specifying <code>-prof-gen</code> (Linux* and Mac OS* X) or <code>/Qprof-gen</code> (Windows*) with no keyword.
<code>srcpos</code>	Produces an instrumented object file that includes extra source position information. This option is the same as option <code>-prof-genx</code> (Linux* and Mac OS* X) or <code>/Qprof-genx</code> (Windows*), which are deprecated .
<code>globdata</code>	Produces an instrumented object file that includes information for global data layout.

Default

`-no-prof-gen` or `/Qprof-gen-` Profile generation is disabled.

Description

This option produces an instrumented object file that can be used in profile-guided optimization. It gets the execution count of each basic block.

If you specify keyword `srcpos` or `globdata`, a static profile information file (.spi) is created. These settings may increase the time needed to do a parallel build using `-prof-gen`, because of contention writing the .spi file.

These options are used in phase 1 of the Profile Guided Optimizer (PGO) to instruct the compiler to produce instrumented code in your object files in preparation for instrumented execution.

Alternate Options

None

See Also

Optimizing Applications:

Basic PGO Options

Example of Profile-Guided Optimization

prof-genx, Qprof-genx

This is a deprecated option. See [prof-gen](#) keyword `srcpos`.

prof-hotness-threshold, Qprof-hotness-threshold

Lets you set the hotness threshold for function grouping and function ordering.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux: `-prof-hotness-threshold=n`

Mac OS X: `None`

Windows: `/Qprof-hotness-threshold:n`

Arguments

n Is the hotness threshold. *n* is a percentage having a value between 0 and 100 inclusive. If you specify 0, there will be no hotness threshold setting in effect for function grouping and function ordering.

Default

OFF The compiler's default hotness threshold setting of 10 percent is in effect for function grouping and function ordering.

Description

This option lets you set the hotness threshold for function grouping and function ordering.

The "hotness threshold" is the percentage of functions in the application that should be placed in the application's hot region. The hot region is the most frequently executed part of the application. By grouping these functions together into one hot region, they have a greater probability of remaining resident in the instruction cache. This can enhance the application's performance.

For this option to take effect, you must specify option `-prof-use` (Linux) or `/Qprof-use` (Windows) and one of the following:

- On Linux systems: `-prof-func-groups` or `-prof-func-order`
- On Windows systems: `/Qprof-func-order`

Alternate Options

None

See Also

[prof-use, Qprof-use](#) compiler option

[prof-func-groups](#) compiler option

[prof-func-order, Qprof-func-order](#) compiler option

prof-src-dir, Qprof-src-dir

Determines whether directory information of the source file under compilation is considered when looking up profile data records.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-prof-src-dir`
`-no-prof-src-dir`

Windows: `/Qprof-src-dir`
`/Qprof-src-dir-`

Arguments

None

Default

`-prof-src-dir` Directory information is used when looking up profile data records in the .dpi file.
 or `/Qprof-src-dir`

Description

This option determines whether directory information of the source file under compilation is considered when looking up profile data records in the .dpi file. To use this option, you must also specify option `-prof-use` (Linux and Mac OS X) or `/Qprof-use` (Windows).

If the option is enabled, directory information is considered when looking up the profile data records within the .dpi file. You can specify directory information by using one of the following options:

- Linux and Mac OS X: `-prof-src-root` or `-prof-src-root-cwd`
- Windows: `/Qprof-src-root` or `/Qprof-src-root-cwd`

If the option is disabled, directory information is ignored and only the name of the file is used to find the profile data record.

Note that options `-prof-src-dir` (Linux and Mac OS X) and `/Qprof-src-dir` (Windows) control how the names of the user's source files get represented within the .dyn or .dpi files. Options `-prof-dir` (Linux and Mac OS X) and `/Qprof-dir` (Windows) specify the location of the .dyn or the .dpi files.

Alternate Options

None

See Also

[prof-use, Qprof-use](#) compiler option

[prof-src-root, Qprof-src-root](#) compiler option

[prof-src-root-cwd, Qprof-src-root-cwd](#) compiler option

prof-src-root, Qprof-src-root

Lets you use relative directory paths when looking up profile data and specifies a directory as the base.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-prof-src-root=dir`

Windows: `/Qprof-src-root:dir`

Arguments

dir Is the base for the relative paths.

Default

OFF The setting of relevant options determines the path used when looking up profile data records.

Description

This option lets you use relative directory paths when looking up profile data in .dpi files. It lets you specify a directory as the base. The paths are relative to a base directory specified during the `-prof-gen` (Linux and Mac OS X) or `/Qprof-gen` (Windows) compilation phase.

This option is available during the following phases of compilation:

- Linux and Mac OS X: `-prof-gen` and `-prof-use` phases
- Windows: `/Qprof-gen` and `/Qprof-use` phases

When this option is specified during the `-prof-gen` or `/Qprof-gen` phase, it stores information into the .dyn or .dpi file. Then, when .dyn files are merged together or the .dpi file is loaded, only the directory information below the root directory is used for forming the lookup key.

When this option is specified during the `-prof-use` or `/Qprof-use` phase, it specifies a root directory that replaces the root directory specified at the `-prof-gen` or `/Qprof-gen` phase for forming the lookup keys.

To be effective, this option or option `-prof-src-root-cwd` (Linux and Mac OS X) or `/Qprof-src-root-cwd` (Windows) must be specified during the `-prof-gen` or `/Qprof-gen` phase. In addition, if one of these options is not specified, absolute paths are used in the .dpi file.

Alternate Options

None

Consider the initial `-prof-gen` compilation of the source file

`c:\user1\feature_foo\myproject\common\glob.f90`:

```
ifort -prof-gen -prof-src-root=c:\user1\feature_foo\myproject -c
common\glob.f90
```

For the `-prof-use` phase, the file `glob.f90` could be moved into the directory `c:\user2\feature_bar\myproject\common\glob.f90` and profile information would be found from the `.dpi` when using the following:

```
ifort -prof-use -prof-src-root=c:\user2\feature_bar\myproject -c
common\glob.f90
```

If you do not use option `-prof-src-root` during the `-prof-gen` phase, by default, the `-prof-use` compilation can only find the profile data if the file is compiled in the `c:\user1\feature_foo\my_project\common` directory.

See Also

[prof-gen, Qprof-gen](#) compiler option

[prof-use, Qprof-use](#) compiler option

[prof-src-dir, Qprof-src-dir](#) compiler option

[prof-src-root-cwd, Qprof-src-root-cwd](#) compiler option

prof-src-root-cwd, Qprof-src-root-cwd

Lets you use relative directory paths when looking up profile data and specifies the current working directory as the base.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-prof-src-root-cwd`

Windows: `/Qprof-src-root-cwd`

Arguments

None

Default

OFF The setting of relevant options determines the path used when looking up profile data records.

Description

This option lets you use relative directory paths when looking up profile data in .dpi files. It specifies the current working directory as the base. To use this option, you must also specify option `-prof-use` (Linux and Mac OS) or `/Qprof-use` (Windows).

This option is available during the following phases of compilation:

- Linux and Mac OS X: `-prof-gen` and `-prof-use` phases
- Windows: `/Qprof-gen` and `/Qprof-use` phases

When this option is specified during the `-prof-gen` or `/Qprof-gen` phase, it stores information into the .dyn or .dpi file. Then, when .dyn files are merged together or the .dpi file is loaded, only the directory information below the root directory is used for forming the lookup key.

When this option is specified during the `-prof-use` or `/Qprof-use` phase, it specifies a root directory that replaces the root directory specified at the `-prof-gen` or `/Qprof-gen` phase for forming the lookup keys.

To be effective, this option or option `-prof-src-root` (Linux and Mac OS X) or `/Qprof-src-root` (Windows) must be specified during the `-prof-gen` or `/Qprof-gen` phase. In addition, if one of these options is not specified, absolute paths are used in the .dpi file.

Alternate Options

None

See Also

[prof-gen, Qprof-gen](#) compiler option

[prof-use, Qprof-use](#) compiler option

[prof-src-dir, Qprof-src-dir](#) compiler option

[prof-src-root, Qprof-src-root](#) compiler option

prof-use, Qprof-use

Enables the use of profiling information during optimization.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-prof-use[=arg]`

`-no-prof-use`

Windows: `/Qprof-use[:arg]`

`/Qprof-use-`

Arguments

arg Specifies additional instructions. Possible values are:

`weighted` Tells the profmerge utility to apply a weighting to the .dyn file values when creating the .dpi file to normalize the data counts when the training runs have different execution durations. This argument only has an effect when the compiler invokes the profmerge

utility to create the .dpi file.

This argument does not have an effect if the .dpi file was previously created without weighting.

[no]merge Enables or disables automatic invocation of the profmerge utility. The default is merge. Note that you cannot specify both weighted and nomerge. If you try to specify both values, a warning will be displayed and nomerge takes precedence.

default Enables the use of profiling information during optimization. The profmerge utility is invoked by default. This value is the same as specifying -prof-use (Linux and Mac OS X) or /Qprof-use (Windows) with no argument.

Default

-no- Profiling information is not used during optimization.
 prof-
 use or
 /Qprof-
 use-

Description

This option enables the use of profiling information (including function splitting and function grouping) during optimization. It enables option `-fnsplit` (Linux) or `/Qfnsplit` (Windows).

This option instructs the compiler to produce a profile-optimized executable and it merges available profiling output files into a `pgopti.dpi` file.

Note that there is no way to turn off function grouping if you enable it using this option.

To set the hotness threshold for function grouping and function ordering, use option `-prof-hotness-threshold` (Linux) or `/Qprof-hotness-threshold` (Windows).

Alternate Options

None

See Also

[prof-hotness-threshold, Qprof-hotness-threshold](#) compiler option

Optimizing Applications:

Basic PGO Options

Example of Profile-Guided Optimization

ansi-alias, Qansi-alias

Tells the compiler to assume that the program adheres to Fortran Standard type aliasability rules.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-ansi-alias`
`-no-ansi-alias`

Windows: `/Qansi-alias`
`/Qansi-alias-`

Arguments

None

Default

`-ansi-alias` Programs adhere to Fortran Standard type aliasability rules.

or

`/Qansi-alias`

Description

This option tells the compiler to assume that the program adheres to type aliasability rules defined in the Fortran Standard.

For example, an object of type real cannot be accessed as an integer. For information on the rules for data types and data type constants, see "Data Types, Constants, and Variables" in the Language Reference.

This option directs the compiler to assume the following:

- Arrays are not accessed out of arrays' bounds.
- Pointers are not cast to non-pointer types and vice-versa.
- References to objects of two different scalar types cannot alias. For example, an object of type integer cannot alias with an object of type real or an object of type real cannot alias with an object of type double precision.

If your program adheres to the Fortran Standard type aliasability rules, this option enables the compiler to optimize more aggressively. If it doesn't adhere to these rules, then you should disable the option with `-no-ansi-alias` (Linux and Mac

OS X) or `/Qansi-alias-` (Windows) so the compiler does not generate incorrect code.

Alternate Options

None

auto, Qauto

See [automatic](#).

auto-scalar, Qauto-scalar

Causes scalar variables of intrinsic types INTEGER, REAL, COMPLEX, and LOGICAL that do not have the SAVE attribute to be allocated to the run-time stack.

IDE Equivalent

Windows: **Data > Local Variable Storage** (`/Qsave`, `/Qauto`, `/Qauto_scalar`)

Linux: None

Mac OS X: None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-auto-scalar`

Windows: `/Qauto-scalar`

Arguments

None

Default

`-auto-` Scalar variables of intrinsic types INTEGER, REAL,

`scalar` COMPLEX, and LOGICAL that do not have the SAVE attribute or are allocated to the run-time stack. Note that if option `/Qauto-recursive`, `-openmp` (Linux and Mac OS X), or `/Qopenmp` `scalar` (Windows) is specified, the default is `automatic`.

Description

This option causes allocation of scalar variables of intrinsic types INTEGER, REAL, COMPLEX, and LOGICAL to the run-time stack. It is as if they were declared with the AUTOMATIC attribute.

It does not affect variables that have the SAVE attribute (which include initialized locals) or that appear in an EQUIVALENCE statement or in a common block.

This option may provide a performance gain for your program, but if your program depends on variables having the same value as the last time the routine was invoked, your program may not function properly. Variables that need to retain their values across subroutine calls should appear in a SAVE statement.

You cannot specify option `save`, `auto`, or `automatic` with this option.



Note

On Windows NT* systems, there is a performance penalty for addressing a stack frame that is too large. This penalty may be incurred with `/automatic`, `/auto`, or `/Qauto` because arrays are allocated on the stack along with scalars. However, with `/Qauto-scalar`, you would have to have more than 32K bytes of local scalar variables before you incurred the performance penalty. `/Qauto-scalar` enables the compiler to make better choices about which variables should be kept in registers during program execution.

Alternate Options

None

See Also

[auto](#) compiler option

[save](#) compiler option

autodouble, Qautodouble

See [real-size](#).

ax, Qax

Tells the compiler to generate multiple, processor-specific auto-dispatch code paths for Intel processors if there is a performance benefit.

IDE Equivalent

Windows: **Code Generation > Add Processor-Optimized Code Path**

Optimization > Generate Alternate Code Paths

Linux: None

Mac OS X: **Code Generation > Add Processor-Optimized Code Path**

Architectures

IA-32, Intel® 64 architectures

Syntax

Linux and Mac OS X: *-axprocessor*

Windows: */Qaxprocessor*

Arguments

processor Indicates the processor for which code is generated. The following descriptions refer to Intel® Streaming SIMD Extensions (Intel® SSE) and Supplemental Streaming SIMD Extensions (Intel® SSSE). Possible values are:

SSE4.2	Can generate Intel® SSE4 Efficient Accelerated String and
--------	---

Text Processing instructions supported by Intel® Core™ i7 processors. Can generate Intel® SSE4 Vectorizing Compiler and Media Accelerator, Intel® SSSE3, SSE3, SSE2, and SSE instructions and it can optimize for the Intel® Core™ processor family.

SSE4 . 1 Can generate Intel® SSE4 Vectorizing Compiler and Media Accelerator instructions for Intel processors. Can generate Intel® SSSE3, SSE3, SSE2, and SSE instructions and it can optimize for Intel® 45nm Hi-k next generation Intel® Core™ microarchitecture. This replaces value S, which is [deprecated](#).

SSSE3 Can generate Intel® SSSE3, SSE3, SSE2, and SSE instructions for Intel processors and it can optimize for the Intel® Core™2 Duo processor family. This replaces value T, which is [deprecated](#).

SSE3	Can generate Intel® SSE3, SSE2, and SSE instructions for Intel processors and it can optimize for processors based on Intel® Core™ microarchitecture and Intel NetBurst® microarchitecture. This replaces value P, which is deprecated .
SSE2	Can generate Intel® SSE2 and SSE instructions for Intel processors, and it can optimize for Intel® Pentium® 4 processors, Intel® Pentium® M processors, and Intel® Xeon® processors with Intel® SSE2. This value is not available on Mac OS X systems. This replaces value N, which is deprecated .

Default

OFF No auto-dispatch code is generated. Processor-specific code is generated and is controlled by the setting of compiler option `-m` (Linux), compiler option `/arch` (Windows), or compiler option `-x` (Mac OS* X).

Description

This option tells the compiler to generate multiple, processor-specific auto-dispatch code paths for Intel processors if there is a performance benefit. It also

generates a baseline code path. The baseline code is usually slower than the specialized code.

The baseline code path is determined by the architecture specified by the `-x` (Linux and Mac OS X) or `/Qx` (Windows) option. While there are defaults for the `-x` or `/Qx` option that depend on the operating system being used, you can specify an architecture for the baseline code that is higher or lower than the default. The specified architecture becomes the effective minimum architecture for the baseline code path.

If you specify both the `-ax` and `-x` options (Linux and Mac OS X) or the `/Qax` and `/Qx` options (Windows), the baseline code will only execute on processors compatible with the processor type specified by the `-x` or `/Qx` option.

This option tells the compiler to find opportunities to generate separate versions of functions that take advantage of features of the specified Intel® processor. If the compiler finds such an opportunity, it first checks whether generating a processor-specific version of a function is likely to result in a performance gain. If this is the case, the compiler generates both a processor-specific version of a function and a baseline version of the function. At run time, one of the versions is chosen to execute, depending on the Intel processor in use. In this way, the program can benefit from performance gains on more advanced Intel processors, while still working properly on older processors.

You can use more than one of the processor values by combining them. For example, you can specify `-axSSE4.1,SSSE3` (Linux and Mac OS X) or `/QaxSSE4.1,SSSE3` (Windows). You cannot combine the old style, deprecated options and the new options. For example, you cannot specify `-axSSE4.1,T` (Linux and Mac OS X) or `/QaxSSE4.1,T` (Windows).

Previous values `W` and `K` are deprecated. The details on replacements are as follows:

- Mac OS X systems: On these systems, there is no exact replacement for `W` or `K`. You can upgrade to the default option `-msse3` (IA-32 architecture) or option `-mssse3` (Intel® 64 architecture).

Intel® Fortran Compiler User and Reference Guides

- Windows and Linux systems: The replacement for W is `-msse2` (Linux) or `/arch:SSE2` (Windows). There is no exact replacement for K. However, on Windows systems, `/QaxK` is interpreted as `/arch:IA32`; on Linux systems, `-axK` is interpreted as `-mia32`. You can also do one of the following:
- Upgrade to option `-msse2` (Linux) or option `/arch:SSE2` (Windows). This will produce one code path that is specialized for Intel® SSE2. It will not run on earlier processors
- Specify the two option combination `-mia32 -axSSE2` (Linux) or `/arch:IA32 /QaxSSE2` (Windows). This combination will produce an executable that runs on any processor with IA-32 architecture but with an additional specialized Intel® SSE2 code path.

The `-ax` and `/Qax` options enable additional optimizations not enabled with option `-m` or option `/arch`.

Alternate Options

None

See Also

[x, Qx](#) compiler option

[m](#) compiler option

[arch](#) compiler option

Qchkstk

Enables stack probing when the stack is dynamically expanded at run-time.

IDE Equivalent

Windows: **Run-time > Enable Stack Check Upon Expansion**

Linux: None

Mac OS X: None

Architectures

IA-64 architecture

Syntax

Linux and Mac OS X: None

Windows: /Qchkstk
 /Qchkstk-

Arguments

None

Default

/Qchkstk Stack probing is enabled when the stack is dynamically expanded at run-time.

Description

This option enables stack probing when the stack is dynamically expanded at run-time.

It instructs the compiler to generate a call to `_chkstk`. The call will probe the requested memory and detect possible stack overflow.

To cancel the call to `_chkstk`, specify `/Qchkstk-`.

Alternate Options

None

common-args, Qcommon-args

See [assume](#).

complex-limited-range, Qcomplex-limited-range

Determines whether the use of basic algebraic expansions of some arithmetic operations involving data of type COMPLEX is enabled.

IDE Equivalent

Windows: **Floating point > Limit COMPLEX Range**

Linux: None

Mac OS X: **Floating point > Limit COMPLEX Range**

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-complex-limited-range`
`-no-complex-limited-range`

Windows: `/Qcomplex-limited-range`
`/Qcomplex-limited-range-`

Arguments

None

Default

`-no-complex-limited-range` Basic algebraic expansions of some arithmetic operations involving data of type `COMPLEX` are disabled.
`or/Qcomplex-limited-range-`

Description

This option determines whether the use of basic algebraic expansions of some arithmetic operations involving data of type `COMPLEX` is enabled.

When the option is enabled, this can cause performance improvements in programs that use a lot of `COMPLEX` arithmetic. However, values at the extremes of the exponent range may not compute correctly.

Alternate Options

None

cpp, Qcpp

See [fpp](#), [Qfpp](#).

d-lines, Qd-lines

Compiles debug statements.

IDE Equivalent

Windows: **Language > Compile Lines With D in Column 1**

Linux: None

Mac OS X: **Language > Compile Lines With D in Column 1**

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-d-lines`
`-nod-lines`

Windows: `/d-lines`
`/nod-lines`
`/Qd-lines`

Arguments

None

Default

`nod-` Debug lines are treated as comment lines.
`lines`

Description

This option compiles debug statements. It specifies that lines in fixed-format files that contain a D in column 1 (debug statements) should be treated as source code.

Alternate Options

Linux and Mac OS X: `-DD`

Windows: None

diag, Qdiag

Controls the display of diagnostic information.

IDE Equivalent

Windows: **Diagnostics > Disable Specific Diagnostics** (/Qdiag-disable id)

Diagnostics > Level of Static Analysis (/Qdiag-enable[:sv1,sv2, sv3])

Linux: None

Mac OS X: **Diagnostics > Disable Specific Diagnostics** (-diag-disable id)

Diagnostics > Level of Static Analysis (-diag-enable [sv1,sv2, sv3])

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-diag-type diag-list`

Windows: `/Qdiag-type:diag-list`

Arguments

type Is an action to perform on diagnostics. Possible values are:

`enable` Enables a diagnostic message or a group of messages.

`disable` Disables a diagnostic message or a group of messages.

<code>error</code>	Tells the compiler to change diagnostics to errors.
<code>warning</code>	Tells the compiler to change diagnostics to warnings.
<code>remark</code>	Tells the compiler to change diagnostics to remarks (comments).

diag-list

Is a diagnostic group or ID value. Possible values are:

<code>driver</code>	Specifies diagnostic messages issued by the compiler driver.
<code>vec</code>	Specifies diagnostic messages issued by the vectorizer.
<code>par</code>	Specifies diagnostic messages issued by the auto-parallelizer (parallel optimizer).
<code>sv[n]</code>	Specifies diagnostic messages issued by the Static Verifier. <i>n</i> can be any of the following: 1, 2, 3. For more details on these values, see below.
<code>warn</code>	Specifies diagnostic messages that have a "warning" severity level.

<code>error</code>	Specifies diagnostic messages that have an "error" severity level.
<code>remark</code>	Specifies diagnostic messages that are remarks or comments.
<code>cpu-dispatch</code>	Specifies the CPU dispatch remarks for diagnostic messages. These remarks are enabled by default. This diagnostic group is only available on IA-32 architecture and Intel® 64 architecture.
<code>id[,id, ...]</code>	Specifies the ID number of one or more messages. If you specify more than one message number, they must be separated by commas. There can be no intervening white space between each id.
<code>tag[,tag, ...]</code>	Specifies the mnemonic name of one or more messages. If you specify more than one mnemonic name, they must be separated by commas.

There can be no
intervening white space
between each tag.

Default

OFF The compiler issues certain diagnostic messages
by default.

Description

This option controls the display of diagnostic information. Diagnostic messages are output to stderr unless compiler option `-diag-file` (Linux and Mac OS X) or `/Qdiag-file` (Windows) is specified.

When *diag-list* value "warn" is used with the Static Verifier (sv) diagnostics, the following behavior occurs:

- Option `-diag-enable warn` (Linux and Mac OS X) and `/Qdiag-enable:warn` (Windows) enable all Static Verifier diagnostics except those that have an "error" severity level. They enable all Static Verifier warnings, cautions, and remarks.
- Option `-diag-disable warn` (Linux and Mac OS X) and `/Qdiag-disable:warn` (Windows) disable all Static Verifier diagnostics except those that have an "error" severity level. They suppress all Static Verifier warnings, cautions, and remarks.

The following table shows more information on values you can specify for *diag-list* item sv.

<i>diag-list</i> Item	Description
sv[n]	The value of <i>n</i> for Static Verifier messages can be any of the following:
1	Produces the diagnostics with severity level set to all critical errors.

<i>diag-list</i> Item	Description
2	Produces the diagnostics with severity level set to all errors. This is the default if <i>n</i> is not specified.
3	Produces the diagnostics with severity level set to all errors and warnings.

To control the diagnostic information reported by the vectorizer, use the `-vec-report` (Linux and Mac OS X) or `/Qvec-report` (Windows) option.

To control the diagnostic information reported by the auto-parallelizer, use the `-par-report` (Linux and Mac OS X) or `/Qpar-report` (Windows) option.

Alternate Options

enable vec	Linux and Mac OS X: <code>-vec-report</code> Windows: <code>/Qvec-report</code>
disable vec	Linux and Mac OS X: <code>-vec-report0</code> Windows: <code>/Qvec-report0</code>
enable par	Linux and Mac OS X: <code>-par-report</code> Windows: <code>/Qpar-report</code>
disable par	Linux and Mac OS X: <code>-par-report0</code> Windows: <code>/Qpar-report0</code>

Example

The following example shows how to enable diagnostic IDs 117, 230 and 450:

```
-diag-enable 117,230,450      ! Linux and Mac OS X systems
/Qdiag-enable:117,230,450    ! Windows systems
```

The following example shows how to change vectorizer diagnostic messages to warnings:

```
-diag-enable vec -diag-warning vec      ! Linux and Mac OS X systems
/Qdiag-enable:vec /Qdiag-warning:vec    ! Windows systems
```

Note that you need to enable the vectorizer diagnostics before you can change them to warnings.

The following example shows how to disable all auto-parallelizer diagnostic messages:

```
-diag-disable par      ! Linux and Mac OS X systems
/Qdiag-disable:par     ! Windows systems
```

The following example shows how to produce Static Verifier diagnostic messages for all critical errors:

```
-diag-enable sv1      ! Linux and Mac OS X systems
/Qdiag-enable:sv1    ! Windows system
```

The following example shows how to cause Static Verifier diagnostics (and default diagnostics) to be sent to a file:

```
-diag-enable sv -diag-file=stat_ver_msg ! Linux and Mac OS X systems
/Qdiag-enable:sv /Qdiag-file:stat_ver_msg ! Windows systems
```

Note that you need to enable the Static Verifier diagnostics before you can send them to a file. In this case, the diagnostics are sent to file `stat_ver_msg.diag`. If a file name is not specified, the diagnostics are sent to `name-of-the-first-source-file.diag`.

The following example shows how to change all diagnostic warnings and remarks to errors:

```
-diag-error warn,remark ! Linux and Mac OS X systems
/Qdiag-error:warn,remark ! Windows systems
```

See Also

[diag-dump, Qdiag-dump](#) compiler option

[diag-id-numbers, Qdiag-id-numbers](#) compiler option

[diag-file, Qdiag-file](#) compiler option

[par-report, Qpar-report](#) compiler option

[vec-report, Qvec-report](#) compiler option

diag-dump, Qdiag-dump

Tells the compiler to print all enabled diagnostic messages and stop compilation.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-diag-dump`

Windows: /Qdiag-dump

Arguments

None

Default

OFF The compiler issues certain diagnostic messages by default.

Description

This option tells the compiler to print all enabled diagnostic messages and stop compilation. The diagnostic messages are output to `stdout`.

This option prints the enabled diagnostics from all possible diagnostics that the compiler can issue, including any default diagnostics.

If `-diag-enable diag-list` (Linux and Mac OS X) or `/Qdiag-enable diag-list` (Windows) is specified, the print out will include the *diag-list* diagnostics.

Alternate Options

None

Example

The following example adds vectorizer diagnostic messages to the printout of default diagnostics:

```
-diag-enable vec -diag-dump                ! Linux and Mac OS X systems  
/Qdiag-enable:vec /Qdiag-dump            ! Windows systems
```

See Also

[diag, Qdiag](#) compiler option

diag-enable sv-include, Qdiag-enable sv-include

Tells the Static Verifier to analyze include files and source files when issuing diagnostic messages.

IDE Equivalent

Windows: **Diagnostics > Analyze Include Files**

Linux: None

Mac OS X: **Diagnostics > Analyze Include Files**

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-diag-enable sv-include`

Windows: `/Qdiag-enable sv-include`

Arguments

None

Default

OFF The compiler issues certain diagnostic messages by default. If the Static Verifier is enabled, include files are not analyzed by default.

Description

This option tells the Static Verifier to analyze include files and source files when issuing diagnostic messages. Normally, when Static Verifier diagnostics are enabled, only source files are analyzed.

To use this option, you must also specify `-diag-enable sv` (Linux and Mac OS X) or `/Qdiag-enable:sv` (Windows) to enable the Static Verifier diagnostics.

Alternate Options

None

Example

The following example shows how to cause include files to be analyzed as well as source files:

```
-diag-enable sv -diag-enable sv-include      ! Linux and Mac OS systems  
/Qdiag-enable:sv /Qdiag-enable:sv-include  ! Windows systems
```

In the above example, the first compiler option enables Static Verifier messages.

The second compiler option causes include files referred to by the source file to be analyzed also.

See Also

[diag, Qdiag](#) compiler option

diag-error-limit, Qdiag-error-limit

Specifies the maximum number of errors allowed before compilation stops.

IDE Equivalent

Windows: **Compilation Diagnostics > Error Limit**

Linux: None

Mac OS X: **Compiler Diagnostics > Error Limit**

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-diag-error-limit`*n*
`-no-diag-error-limit`

Windows: `/Qdiag-error-limit:`*n*
`/Qdiag-error-limit-`

Arguments

n Is the maximum number of error-level or fatal-level compiler errors allowed.

Default

30 A maximum of 30 error-level and fatal-level messages are allowed.

Description

This option specifies the maximum number of errors allowed before compilation stops. It indicates the maximum number of error-level or fatal-level compiler errors allowed for a file specified on the command line.

If you specify `-no-diag-error-limit` (Linux and Mac OS X) or `/Qdiag-error-limit-` (Windows) on the command line, there is no limit on the number of errors that are allowed.

If the maximum number of errors is reached, a warning message is issued and the next file (if any) on the command line is compiled.

Alternate Options

Linux and Mac OS X: `-error-limit` and `-noerror-limit`

Windows: `/error-limit` and `/noerror-limit`

diag-file, Qdiag-file

Causes the results of diagnostic analysis to be output to a file.

IDE Equivalent

Windows: **Diagnostics > Diagnostics File**

Linux: None

Mac OS X: **Diagnostics > Diagnostics File**

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-diag-file[=file]`

Windows: `/Qdiag-file[:file]`

Arguments

file Is the name of the file for output.

Default

OFF Diagnostic messages are output to stderr.

Description

This option causes the results of diagnostic analysis to be output to a file. The file is placed in the current working directory.

If *file* is specified, the name of the file is *file.diag*. The file can include a file extension; for example, if *file.ext* is specified, the name of the file is *file.ext*.

If *file* is not specified, the name of the file is *name-of-the-first-source-file.diag*. This is also the name of the file if the name specified for file conflicts with a source file name provided in the command line.



If you specify `-diag-file` (Linux and Mac OS X) or `/Qdiag-file` (Windows) and you also specify `-diag-file-append` (Linux and Mac OS X) or `/Qdiag-file-append` (Windows), the last option specified on the command line takes precedence.

Alternate Options

None

Example

The following example shows how to cause diagnostic analysis to be output to a file named *my_diagnostics.diag*:

```
-diag-file=my_diagnostics      ! Linux and Mac OS X systems  
/Qdiag-file:my_diagnostics    ! Windows systems
```

See Also

[diag, Qdiag](#) compiler option

[diag-file-append, Qdiag-file-append](#) compiler option

diag-file-append, Qdiag-file-append

Causes the results of diagnostic analysis to be appended to a file.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-diag-file-append[=file]`

Windows: `/Qdiag-file-append[:file]`

Arguments

file Is the name of the file to be appended to. It can include a path.

Default

OFF Diagnostic messages are output to `stderr`.

Description

This option causes the results of diagnostic analysis to be appended to a file. If you do not specify a path, the driver will look for *file* in the current working directory.

If *file* is not found, then a new file with that name is created in the current working directory. If the name specified for file conflicts with a source file name provided in the command line, the name of the file is *name-of-the-first-source-file.diag*.



Note

If you specify `-diag-file-append` (Linux and Mac OS X) or `/Qdiag-file-append` (Windows) and you also specify `-diag-file` (Linux and Mac OS X) or `/Qdiag-file` (Windows), the last option specified on the command line takes precedence.

Alternate Options

None

Example

The following example shows how to cause diagnostic analysis to be appended to a file named `my_diagnostics.txt`:

```
-diag-file-append=my_diagnostics.txt      ! Linux and Mac OS X systems  
/Qdiag-file-append:my_diagnostics.txt    ! Windows systems
```

See Also

[diag, Qdiag](#) compiler option

[diag-file, Qdiag-file](#) compiler option

diag-id-numbers, Qdiag-id-numbers

Determines whether the compiler displays diagnostic messages by using their ID number values.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-diag-id-numbers`
`-no-diag-id-numbers`

Windows: `/Qdiag-id-numbers`
`/Qdiag-id-numbers-`

Arguments

None

Default

`-diag-id-numbers` The compiler displays diagnostic messages by using their ID number values.
`or/Qdiag-id-numbers`

Description

This option determines whether the compiler displays diagnostic messages by using their ID number values. If you specify `-no-diag-id-numbers` (Linux and Mac OS X) or `/Qdiag-id-numbers-` (Windows), mnemonic names are output for driver diagnostics only.

Alternate Options

None

See Also

[diag, Qdiag](#) compiler option

diag-once, Qdiag-once

Tells the compiler to issue one or more diagnostic messages only once.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-diag-onceid[,id, ...]`

Windows: `/Qdiag-once:id[,id, ...]`

Arguments

id Is the ID number of the diagnostic message. If you specify more than one message number, they must be separated by commas. There can be no intervening white space between each *id*.

Default

OFF The compiler issues certain diagnostic messages by default.

Description

This option tells the compiler to issue one or more diagnostic messages only once.

Alternate Options

None

dps, Qdps

See [altparam](#).

dyncom, Qdyncom

Enables dynamic allocation of common blocks at run time.

IDE Equivalent

Windows: **Data > Dynamic Common Blocks**

Linux: None

Mac OS X: **Data > Dynamic Common Blocks**

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-dyncom "common1,common2,..."`

Windows: `/Qdyncom "common1,common2,..."`

Arguments

`common1,common2,...` Are the names of the common blocks to be dynamically allocated. The list of names must be within quotes.

Default

OFF Common blocks are not dynamically allocated at run time.

Description

This option enables dynamic allocation of the specified common blocks at run time. For example, to enable dynamic allocation of common blocks a, b, and c at run time, use this syntax:

```
/Qdyncom "a,b,c"      ! on Windows systems
-dyncom "a,b,c"      ! on Linux and Mac OS X systems
```

The following are some limitations that you should be aware of when using this option:

- An entity in a dynamic common cannot be initialized in a DATA statement.
- Only named common blocks can be designated as dynamic COMMON.
- An entity in a dynamic common block must not be used in an EQUIVALENCE expression with an entity in a static common block or a DATA-initialized variable.

Alternate Options

None

See Also

Building Applications: Allocating Common Blocks

Qextend-source

See [extend-source](#).

fast-transcendentals, Qfast-transcendentals

Enables the compiler to replace calls to transcendental functions with faster but less precise implementations.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-fast-transcendentals`
`-no-fast-transcendentals`

Windows: `/Qfast-transcendentals`
`/Qfast-transcendentals-`

Default

`-fast-transcendentals` The default depends on the setting of `-fp-model` (Linux and Mac OS X) or `/fp` (Windows).

`/Qfast-transcendentals` The default is ON if default setting `-fp-model fast` or `/fp:fast` is in effect. However, if a value-safe option such as `-fp-model precise` or `/fp:precise` is specified, the default is OFF.

Description

This option enables the compiler to replace calls to transcendental functions with implementations that may be faster but less precise.

It tells the compiler to perform certain optimizations on transcendental functions, such as replacing individual calls to sine in a loop with a single call to a less precise vectorized sine library routine.

This option has an effect only when specified with one of the following options:

- Windows* OS: `/fp:except` or `/fp:precise`
- Linux* OS and Mac OS* X: `-fp-model except` or `-fp-model precise`

You cannot use this option with option `-fp-model strict` (Linux and Mac OS X) or `/fp:strict` (Windows).

Alternate Options

None

See Also

[fp-model, fp](#) compiler option

fma, Qfma

Enables the combining of floating-point multiplies and add/subtract operations.

IDE Equivalent

Windows: **Floating Point > Contract Floating-Point Operations**

Linux: None

Mac OS X: None

Architectures

IA-64 architecture

Syntax

Linux: `-fma`
 `-no-fma`

Mac OS X: `None`

Windows: `/Qfma`

/Qfma-

Arguments

None

Default

-fma Floating-point multiplies and add/subtract operations are combined.
or /Qfma
However, if you specify `-mp` (Linux), `/Op` (Windows), `/fp:strict` (Windows), or `-fp-model strict` (Linux) but do not explicitly specify `-fma` or `/Qfma`, the default is `-no-fma` or `/Qfma-`.

Description

This option enables the combining of floating-point multiplies and add/subtract operations.

It also enables the contraction of floating-point multiply and add/subtract operations into a single operation. The compiler contracts these operations whenever possible.

Alternate Options

Linux: `-IPF-fma` (this is a [deprecated](#) option)

Windows: `/QIPF-fma` (this is a deprecated option)

See Also

[fp-model](#), [fp](#) compiler option

Floating-point Operations: Floating-point Options Quick Reference

falign-functions, Qfalign

Tells the compiler to align functions on an optimal byte boundary.

IDE Equivalent

None

Architectures

IA-32, Intel® 64 architectures

SyntaxLinux and Mac OS X: `-falign-functions[=n]``-fno-align-functions`Windows: `/Qfnalign[:n]``/Qfnalign-`**Arguments**

n Is the byte boundary for function alignment.
Possible values are 2 or 16.

Default

`-fno-align-functions` The compiler aligns functions on 2-byte boundaries. This is the same as specifying `-falign-functions=2` (Linux and Mac OS X) or `/Qfnalign:2` (Windows).
or
`/Qfnalign-`

Description

This option tells the compiler to align functions on an optimal byte boundary. If you do not specify *n*, the compiler aligns the start of functions on 16-byte boundaries.

Alternate Options

None

fnsplit, Qfnsplit

Enables function splitting.

IDE Equivalent

None

Architectures

`/Qfnsplit[-]`: IA-32 architecture, Intel® 64 architecture

`-[no-]fnsplit`: IA-64 architecture

Syntax

Linux: `-fnsplit`
 `-no-fnsplit`

Mac OS X: None

Windows: `/Qfnsplit`
 `/Qfnsplit-`

Arguments

None

Default

`-no-fnsplit` Function splitting is not enabled unless `-prof-`
`or/Qfnsplit-` use (Linux) or `/Qprof-use` (Windows) is also
 specified.

Description

This option enables function splitting if `-prof-use` (Linux) or `/Qprof-use` (Windows) is also specified. Otherwise, this option has no effect.

It is enabled automatically if you specify `-prof-use` or `/Qprof-use`. If you do not specify one of those options, the default is `-no-fnsplit` (Linux) or `/Qfnsplit-` (Windows), which disables function splitting but leaves function grouping enabled.

To disable function splitting when you use `-prof-use` or `/Qprof-use`, specify `-no-fnsplit` or `/Qfnsplit-`.

Alternate Options

None

See Also

Optimizing Applications:

Basic PGO Options

Example of Profile-Guided Optimization

fp-port, Qfp-port

Rounds floating-point results after floating-point operations.

IDE Equivalent

Windows: **Floating-Point > Round Floating-Point Results**

Linux: None

Mac OS X: **Floating-Point > Round Floating-Point Results**

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-fp-port`
`-no-fp-port`

Windows: `/Qfp-port`
`/Qfp-port-`

Arguments

None

Default

`-no-fp-
port
or/Qfp-
port-` The default rounding behavior depends on the compiler's code generation decisions and the precision parameters of the operating system.

Description

This option rounds floating-point results after floating-point operations. Rounding to user-specified precision occurs at assignments and type conversions. This has some impact on speed.

The default is to keep results of floating-point operations in higher precision. This provides better performance but less consistent floating-point results.

Alternate Options

None

fp-relaxed, Qfp-relaxed

Enables use of faster but slightly less accurate code sequences for math functions.

IDE Equivalent

None

Architectures

IA-64 architecture

Syntax

Linux:	<code>-fp-relaxed</code> <code>-no-fp-relaxed</code>
Mac OS X:	None
Windows:	<code>/Qfp-relaxed</code> <code>/Qfp-relaxed-</code>

Arguments

None

Default

`-no-fp-relaxed` Default code sequences are used for math
or `/Qfp-relaxed-` functions.

Description

This option enables use of faster but slightly less accurate code sequences for math functions, such as divide and sqrt. When compared to strict IEEE* precision, this option slightly reduces the accuracy of floating-point calculations performed by these functions, usually limited to the least significant digit.

This option also enables the performance of more aggressive floating-point transformations, which may affect accuracy.

Alternate Options

Linux: `-IPF-fp-relaxed` (this is a [deprecated](#) option)

Windows: `/QIPF-fp-relaxed` (this is a deprecated option)

See Also

[fp-model](#), [fp](#) compiler option

fp-speculation, Qfp-speculation

Tells the compiler the mode in which to speculate on floating-point operations.

IDE Equivalent

Windows: **Floating Point > Floating-Point Speculation**

Linux: None

Mac OS X: **Floating Point > Floating-Point Speculation**

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-fp-speculation=mode`

Windows: `/Qfp-speculation:mode`

Arguments

mode Is the mode for floating-point operations. Possible values are:

<code>fast</code>	Tells the compiler to speculate on floating-point operations.
<code>safe</code>	Tells the compiler to disable speculation if there is a possibility that the speculation may cause a floating-point exception.
<code>strict</code>	Tells the compiler to disable speculation on floating-point operations.
<code>off</code>	This is the same as specifying <code>strict</code> .

Default

`-fp-speculation=fast` or `/Qfp-speculation:fast` The compiler speculates on floating-point operations. This is also the behavior when optimizations are enabled. However, if you specify no optimizations (`-O0` on Linux; `/Od` on Windows), the default is `-fp-speculation=safe` (Linux) or `/Qfp-speculation:safe` (Windows).

Description

This option tells the compiler the mode in which to speculate on floating-point operations.

Alternate Options

None

See Also

Floating-point Operations: Floating-point Options Quick Reference

fp-stack-check, Qfp-stack-check

Tells the compiler to generate extra code after every function call to ensure that the floating-point stack is in the expected state.

IDE Equivalent

Windows: **Floating-Point > Check Floating-point Stack**

Linux: None

Mac OS X: **Floating-Point > Check Floating-point Stack**

Architectures

IA-32, Intel® 64 architectures

Syntax

Linux and Mac OS X: `-fp-stack-check`

Windows: `/Qfp-stack-check`

Arguments

None

Default

OFF There is no checking to ensure that the floating-point (FP) stack

is in the expected state.

Description

This option tells the compiler to generate extra code after every function call to ensure that the floating-point (FP) stack is in the expected state.

By default, there is no checking. So when the FP stack overflows, a NaN value is put into FP calculations and the program's results differ. Unfortunately, the overflow point can be far away from the point of the actual bug. This option places code that causes an access violation exception immediately after an incorrect call occurs, thus making it easier to locate these issues.

Alternate Options

None

See Also

Floating-point Operations:

[Checking the Floating-point Stack State](#)

fpp, Qfpp

Runs the Fortran preprocessor on source files before compilation.

IDE Equivalent

Windows: **Preprocessor > Preprocess Source File**

Linux: None

Mac OS X: Preprocessor > Preprocess Source File

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `fpp[n]`

`fpp[= "option"]`

```

-nofpp
Windows: /fpp[n]
          /fpp[ : "option" ]
          /nofpp
          /Qfpp[n]
          /Qfpp[ : "option" ]

```

Arguments

n [Deprecated](#). Tells the compiler whether to run the preprocessor or not. Possible values are:

0 Tells the compiler not to run the preprocessor.

1, 2, or 3 Tells the compiler to run the preprocessor.

option Is a Fortran preprocessor (fpp) option; for example, "-macro=no", which disables macro expansion. The quotes are required. For a list of fpp options, see *Fortran Preprocessor Options*.

Default

`nofpp` The Fortran preprocessor is not run on files before compilation.

Description

This option runs the Fortran preprocessor on source files before they are compiled.

If the option is specified with no argument, the compiler runs the preprocessor.

If 0 is specified for *n*, it is equivalent to `nofpp`. Note that argument *n* is [deprecated](#).

We recommend you use option `Qoption, fpp, "option"` to pass fpp options to the Fortran preprocessor.

Alternate Options

Linux and Mac OS X: `-cpp`

Windows: `/Qcpp`

See Also

[Fortran Preprocessor Options](#)

[Qoption](#) compiler option

ftz, Qftz

Flushes denormal results to zero.

IDE Equivalent

Windows: (IA-32 and IA-64 architectures): **Floating Point > Flush Denormal Results to Zero**

(Intel® 64 architecture): None

Linux: None

Mac OS X: **Floating Point > Flush Denormal Results to Zero**

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-ftz`

`-no-ftz`

Windows: `/Qftz`

`/Qftz-`

Arguments

None

Default

Systems using IA-64 architecture: `-ftz` or `/Qftz-` On systems using IA-64 architecture, the compiler lets results gradually underflow. On systems using IA-32 architecture and Intel® 64 architecture: `-ftz` or `/Qftz` and Intel® 64 architecture, denormal results are flushed to zero.

Description

This option flushes denormal results to zero when the application is in the gradual underflow mode. It may improve performance if the denormal values are not critical to your application's behavior.

This option sets or resets the FTZ and the DAZ hardware flags. If FTZ is ON, denormal results from floating-point calculations will be set to the value zero. If FTZ is OFF, denormal results remain as is. If DAZ is ON, denormal values used as input to floating-point instructions will be treated as zero. If DAZ is OFF, denormal instruction inputs remain as is. Systems using IA-64 architecture have FTZ but not DAZ. Systems using Intel® 64 architecture have both FTZ and DAZ. FTZ and DAZ are not supported on all IA-32 architectures.

When `-ftz` (Linux and Mac OS X) or `/Qftz` (Windows) is used in combination with an SSE-enabling option on systems using IA-32 architecture (for example, `xN` or `QxN`), the compiler will insert code in the main routine to set FTZ and DAZ.

When `-ftz` or `/Qftz` is used without such an option, the compiler will insert code to conditionally set FTZ/DAZ based on a run-time processor check. `-no-ftz` (Linux and Mac OS X) or `/Qftz-` (Windows) will prevent the compiler from inserting any code that might set FTZ or DAZ.

This option only has an effect when the main program is being compiled. It sets the FTZ/DAZ mode for the process. The initial thread and any threads subsequently created by that process will operate in FTZ/DAZ mode.

Options `-fpe0` and `-fpe1` (Linux and Mac OS X) set `-ftz`. Options `/fpe:0` and `/fpe:1` (Windows) set `/Qftz`.

On systems using IA-64 architecture, optimization option O3 sets `-ftz` and `/Qftz`; optimization option O2 sets `-no-ftz` (Linux) and `/Qftz-` (Windows).

On systems using IA-32 architecture and Intel® 64 architecture, every optimization option O level, except O0, sets `-ftz` and `/Qftz`.

If this option produces undesirable results of the numerical behavior of your program, you can turn the FTZ/DAZ mode off by using `-no-ftz` or `/Qftz-` in the command line while still benefiting from the O3 optimizations.



Note

Options `-ftz` and `/Qftz` are performance options. Setting these options does not guarantee that all denormals in a program are flushed to zero. They only cause denormals generated at run time to be flushed to zero.

Alternate Options

None

See Also

[x, Qx](#) compiler option

Floating-point Operations: Using the `-fpe` or `/fpe` Compiler Option

Intrinsics Reference:

global-hoist, Qglobal-hoist

Enables certain optimizations that can move memory loads to a point earlier in the program execution than where they appear in the source.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-global-hoist`
`-no-global-hoist`
Windows: `/Qglobal-hoist`
`/Qglobal-hoist-`

Arguments

None

Default

`-global-hoist` Certain optimizations are enabled that can move memory loads.
or `/Qglobal-hoist`
`-no-global-hoist`

Description

This option enables certain optimizations that can move memory loads to a point earlier in the program execution than where they appear in the source. In most cases, these optimizations are safe and can improve performance.

The `-no-global-hoist` (Linux and Mac OS X) or `/Qglobal-hoist-` (Windows) option is useful for some applications, such as those that use shared or dynamically mapped memory, which can fail if a load is moved too early in the execution stream (for example, before the memory is mapped).

Alternate Options

None

QIA64-fr32

Disables use of high floating-point registers.

IDE Equivalent

Windows: **Floating Point > Disable Use of High Floating-Point Registers**

Linux: None

Mac OS X: None

Architectures

IA-64 architecture

Syntax

Linux and Mac OS X: None

Windows: `/QIA64-fr32`

Arguments

None

Default

OFF Use of high floating-point registers is enabled.

Description

This option disables use of high floating-point registers.

Alternate Options

None

Qlfist

See [rcd](#), [Qrcd](#).

inline-debug-info, Qinline-debug-info

Produces enhanced source position information for inlined code. This is a deprecated option.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-inline-debug-info`

Windows: `/Qinline-debug-info`

Arguments

None

Default

OFF No enhanced source position information is produced for inlined code.

Description

This option produces enhanced source position information for inlined code. This leads to greater accuracy when reporting the source location of any instruction. It also provides enhanced debug information useful for function call traceback.

To use this option for debugging, you must also specify a debug enabling option, such as `-g` (Linux) or `/debug` (Windows).

Alternate Options

Linux and Mac OS X: `-debug inline-debug-info`

Windows: None

Qinline-dllimport

Determines whether dllimport functions are inlined.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: None

Windows: /Qinline-dllimport
 /Qinline-dllimport-

Arguments

None

Default

/Qinline- The dllimport functions are inlined.
dllimport

Description

This option determines whether dllimport functions are inlined. To disable dllimport functions from being inlined, specify /Qinline-dllimport-.

Alternate Options

None

inline-factor, Qinline-factor

Specifies the percentage multiplier that should be applied to all inlining options that define upper limits.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-inline-factor=n`
`-no-inline-factor`

Windows: `/Qinline-factor=n`
`/Qinline-factor-`

Arguments

n Is a positive integer specifying the percentage value. The default value is 100 (a factor of 1).

Default

`-no-inline-factor` or `/Qinline-factor-` The compiler uses default heuristics for inline routine expansion.

Description

This option specifies the percentage multiplier that should be applied to all inlining options that define upper limits:

- `-inline-max-size` and `/Qinline-max-size`
- `-inline-max-total-size` and `/Qinline-max-total-size`
- `-inline-max-per-routine` and `/Qinline-max-per-routine`
- `-inline-max-per-compile` and `/Qinline-max-per-compile`

This option takes the default value for each of the above options and multiplies it by *n* divided by 100. For example, if 200 is specified, all inlining options that define upper limits are multiplied by a factor of 2. This option is useful if you do not want to individually increase each option limit.

If you specify `-no-inline-factor` (Linux and Mac OS X) or `/Qinline-factor-` (Windows), the following occurs:

- Every function is considered to be a small or medium function; there are no large functions.

- There is no limit to the size a routine may grow when inline expansion is performed.
- There is no limit to the number of times some routine may be inlined into a particular routine.
- There is no limit to the number of times inlining can be applied to a compilation unit.

To see compiler values for important inlining limits, specify compiler option `-opt-report` (Linux and Mac OS X) or `/Qopt-report` (Windows).



Caution

When you use this option to increase default limits, the compiler may do so much additional inlining that it runs out of memory and terminates with an "out of memory" message.

Alternate Options

None

See Also

[inline-max-size, Qinline-max-size](#) compiler option

[inline-max-total-size, Qinline-max-total-size](#) compiler option

[inline-max-per-routine, Qinline-max-per-routine](#) compiler option

[inline-max-per-compile, Qinline-max-per-compile](#) compiler option

[opt-report, Qopt-report](#) compiler option

Optimizing Applications:

Developer Directed Inline Expansion of User Functions

Compiler Directed Inline Expansion of User Functions

inline-forceinline, Qinline-forceinline

Specifies that an inline routine should be inlined whenever the compiler can do so.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-inline-forceinline`

Windows: `/Qinline-forceinline`

Default

OFF The compiler uses default heuristics for inline routine expansion.

Description

This option specifies that a inline routine should be inlined whenever the compiler can do so. This causes the routines marked with an inline keyword or directive to be treated as if they were "forceinline".



Note

Because C++ member functions whose definitions are included in the class declaration are considered inline functions by default, using this option will also make these member functions "forceinline" functions.

The "forceinline" condition can also be specified by using the directive `cDEC$ ATTRIBUTES FORCEINLINE`.

To see compiler values for important inlining limits, specify compiler option `-opt-report` (Linux and Mac OS) or `/Qopt-report` (Windows).



Caution

When you use this option to change the meaning of inline to "forceinline", the compiler may do so much additional inlining that it runs out of memory and terminates with an "out of memory" message.

Alternate Options

None

See Also

[opt-report](#), [Qopt-report](#) compiler option

Optimizing Applications:

Developer Directed Inline Expansion of User Functions

Compiler Directed Inline Expansion of User Functions

inline-max-per-compile, Qinline-max-per-compile

Specifies the maximum number of times inlining may be applied to an entire compilation unit.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-inline-max-per-compile=n`
`-no-inline-max-per-compile`

Windows: `/Qinline-max-per-compile=n`
`/Qinline-max-per-compile-`

Arguments

n Is a positive integer that specifies the number of times inlining may be applied.

Default

`-no-inline-max-per-compile` The compiler uses default heuristics for inline routine expansion.

or `/Qinline-max-per-compile-`

Description

This option the maximum number of times inlining may be applied to an entire compilation unit. It limits the number of times that inlining can be applied. For compilations using Interprocedural Optimizations (IPO), the entire compilation is a compilation unit. For other compilations, a compilation unit is a file.

If you specify `-no-inline-max-per-compile` (Linux and Mac OS X) or `/Qinline-max-per-compile-` (Windows), there is no limit to the number of times inlining may be applied to a compilation unit.

To see compiler values for important inlining limits, specify compiler option `-opt-report` (Linux and Mac OS X) or `/Qopt-report` (Windows).



Caution

When you use this option to increase the default limit, the compiler may do so much additional inlining that it runs out of memory and terminates with an "out of memory" message.

Alternate Options

None

See Also

[inline-factor, Qinline-factor](#) compiler option

[opt-report, Qopt-report](#) compiler option

Optimizing Applications:

Developer Directed Inline Expansion of User Functions

Compiler Directed Inline Expansion of User Functions

inline-max-per-routine, Qinline-max-per-routine

Specifies the maximum number of times the inliner may inline into a particular routine.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-inline-max-per-routine=n`
`-no-inline-max-per-routine`

Windows: `/Qinline-max-per-routine=n`
`/Qinline-max-per-routine-`

Arguments

n Is a positive integer that specifies the maximum number of times the inliner may inline into a particular routine.

Default

`-no-inline-max-per-routine` The compiler uses default heuristics for inline routine expansion.
`or/Qinline-max-per-routine-`

Description

This option specifies the maximum number of times the inliner may inline into a particular routine. It limits the number of times that inlining can be applied to any routine.

If you specify `-no-inline-max-per-routine` (Linux and Mac OS X) or `/Qinline-max-per-routine-` (Windows), there is no limit to the number of times some routine may be inlined into a particular routine.

To see compiler values for important inlining limits, specify compiler option `-opt-report` (Linux and Mac OS X) or `/Qopt-report` (Windows).

To see compiler values for important inlining limits, specify compiler option `-opt-report` (Linux and Mac OS X) or `/Qopt-report` (Windows).



When you use this option to increase the default limit, the compiler may do so much additional inlining that it runs out of memory and terminates with an "out of memory" message.

Alternate Options

None

See Also

[inline-factor, Qinline-factor](#) compiler option

[opt-report, Qopt-report](#) compiler option

Optimizing Applications:

Developer Directed Inline Expansion of User Functions

Compiler Directed Inline Expansion of User Functions

inline-max-size, Qinline-max-size

Specifies the lower limit for the size of what the inliner considers to be a large routine.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-inline-max-size=n`
`-no-inline-max-size`

Windows: `/Qinline-max-size=n`
`/Qinline-max-size-`

Arguments

n Is a positive integer that specifies the minimum size of what the inliner considers to be a large routine.

Default

`-no-inline-max-size` The compiler uses default heuristics for
`or/Qinline-max-size-` inline routine expansion.

Description

This option specifies the lower limit for the size of what the inliner considers to be a large routine (a function or subroutine). The inliner classifies routines as small, medium, or large. This option specifies the boundary between what the inliner considers to be medium and large-size routines.

The inliner prefers to inline small routines. It has a preference against inlining large routines. So, any large routine is highly unlikely to be inlined.

If you specify `-no-inline-max-size` (Linux and Mac OS X) or `/Qinline-max-size-` (Windows), there are no large routines. Every routine is either a small or medium routine.

To see compiler values for important inlining limits, specify compiler option `-opt-report` (Linux and Mac OS X) or `/Qopt-report` (Windows).

To see compiler values for important inlining limits, specify compiler option `-opt-report` (Linux and Mac OS X) or `/Qopt-report` (Windows).



Caution

When you use this option to increase the default limit, the compiler may do so much additional inlining that it runs out of memory and terminates with an "out of memory" message.

Alternate Options

None

See Also

[inline-min-size, Qinline-min-size](#) compiler option

[inline-factor, Qinline-factor](#) compiler option

[opt-report, Qopt-report](#) compiler option

Optimizing Applications:

Developer Directed Inline Expansion of User Functions

Compiler Directed Inline Expansion of User Functions

inline-max-total-size, Qinline-max-total-size

Specifies how much larger a routine can normally grow when inline expansion is performed.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-inline-max-total-size=n`
`-no-inline-max-total-size`

Windows: `/Qinline-max-total-size=n`
`/Qinline-max-total-size-`

Arguments

n Is a positive integer that specifies the permitted increase in the routine's size when inline expansion is performed.

Default

`-no-inline-max-total-size` The compiler uses default heuristics for inline routine expansion.
or `/Qinline-max-total-size`

Description

This option specifies how much larger a routine can normally grow when inline expansion is performed. It limits the potential size of the routine. For example, if 2000 is specified for *n*, the size of any routine will normally not increase by more than 2000.

If you specify `-no-inline-max-total-size` (Linux and Mac OS X) or `/Qinline-max-total-size` (Windows), there is no limit to the size a routine may grow when inline expansion is performed.

To see compiler values for important inlining limits, specify compiler option `-opt-report` (Linux and Mac OS X) or `/Qopt-report` (Windows).

To see compiler values for important inlining limits, specify compiler option `-opt-report` (Linux and Mac OS X) or `/Qopt-report` (Windows).



Caution

When you use this option to increase the default limit, the compiler may do so much additional inlining that it runs out of memory and terminates with an "out of memory" message.

Alternate Options

None

See Also

[inline-factor, Qinline-factor](#) compiler option

[opt-report, Qopt-report](#) compiler option

Optimizing Applications:

Developer Directed Inline Expansion of User Functions

Compiler Directed Inline Expansion of User Functions

inline-min-size, Qinline-min-size

Specifies the upper limit for the size of what the inliner considers to be a small routine.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-inline-min-size=n`
`-no-inline-min-size`

Windows: `/Qinline-min-size=n`
`/Qinline-min-size-`

Arguments

n Is a positive integer that specifies the maximum size of what the inliner considers to be a small routine.

Default

`-no-inline-min-size` The compiler uses default heuristics for
or `/Qinline-min-size-` inline routine expansion.

Description

This option specifies the upper limit for the size of what the inliner considers to be a small routine (a function or subroutine). The inliner classifies routines as small, medium, or large. This option specifies the boundary between what the inliner considers to be small and medium-size routines.

The inliner has a preference to inline small routines. So, when a routine is smaller than or equal to the specified size, it is very likely to be inlined.

If you specify `-no-inline-min-size` (Linux and Mac OS X) or `/Qinline-min-size-` (Windows), there is no limit to the size of small routines. Every routine is a small routine; there are no medium or large routines.

To see compiler values for important inlining limits, specify compiler option `-opt-report` (Linux and Mac OS X) or `/Qopt-report` (Windows).

To see compiler values for important inlining limits, specify compiler option `-opt-report` (Linux and Mac OS X) or `/Qopt-report` (Windows).



Caution

When you use this option to increase the default limit, the compiler may do so much additional inlining that it runs out of memory and terminates with an "out of memory" message.

Alternate Options

None

See Also

[inline-min-size, Qinline-min-size](#) compiler option

[opt-report, Qopt-report](#) compiler option

Optimizing Applications:

Developer Directed Inline Expansion of User Functions

Compiler Directed Inline Expansion of User Functions

Qinstall

Specifies the root directory where the compiler installation was performed.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-Qinstalldir`

Windows: None

Arguments

<i>dir</i>	Is the root directory where the installation was performed.
------------	---

Default

OFF The default root directory for compiler installation is searched for the compiler.

Description

This option specifies the root directory where the compiler installation was performed. It is useful if you want to use a different compiler or if you did not use the ifortvars shell script to set your environment variables.

Alternate Options

None

minstruction, Qinstruction

Determines whether MOVBE instructions are generated for Intel processors.

IDE Equivalent

None

Architectures

IA-32, Intel® 64 architectures

Syntax

Linux and Mac OS X: `-minstruction=[no]movbe`

Windows: `/Qinstruction:[no]movbe`

Arguments

None

Default

<code>-minstruction=movbe</code>	The compiler generates MOVBE instructions for Intel® Atom™ processors.
<code>or /Qinstruction:movbe</code>	

Description

This option determines whether MOVBE instructions are generated for Intel processors. To use this option, you must also specify `-xSSE3_ATOM` (Linux and Mac OS X) or `/QxSSE3_ATOM` (Windows).

If `-minstruction=movbe` or `/Qinstruction:movbe` is specified, the following occurs:

- MOVBE instructions are generated that are specific to the Intel® Atom™ processor.
- The options are ON by default when `-xSSE3_ATOM` or `/QxSSE3_ATOM` is specified.

- Generated executables can only be run on Intel® Atom™ processors or processors that support Intel® Streaming SIMD Extensions 3 (Intel® SSE3) and MOVBE.

If `-minstruction=nomovbe` or `/Qinstruction:nomovbe` is specified, the following occurs:

- The compiler optimizes code for the Intel® Atom™ processor, but it does not generate MOVBE instructions.
- Generated executables can be run on non-Intel® Atom™ processors that support Intel® SSE3.

Alternate Options

None

See Also

[x, Qx](#) compiler option

finstrument-functions, Qinstrument-functions

Determines whether routine entry and exit points are instrumented.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-finstrument-functions`
`-fno-instrument-functions`

Windows: `/Qinstrument-functions`
`/Qinstrument-functions-`

Arguments

None

Default

`-fno-` Routine entry and exit points are not instrumented.
`instrument-`
`functions`
`or/Qinstrument-`
`functions-`

Description

This option determines whether routine entry and exit points are instrumented. It may increase execution time.

The following profiling functions are called with the address of the current routine and the address of where the routine was called (its "call site"):

- This function is called upon routine entry:

- On IA-32 architecture and Intel® 64 architecture:

```
void __cyg_profile_func_enter (void *this_fn,  
void *call_site);
```

- On IA-64 architecture:

```
void __cyg_profile_func_enter (void **this_fn,  
void *call_site);
```

- This function is called upon routine exit:

- On IA-32 architecture and Intel® 64 architecture:

```
void __cyg_profile_func_exit (void *this_fn,  
void *call_site);
```

- On IA-64 architecture:

```
void __cyg_profile_func_exit (void **this_fn,  
void *call_site);
```

On IA-64 architecture, the additional de-reference of the function pointer argument is required to obtain the routine entry point contained in the first word of the routine descriptor for indirect routine calls. The descriptor is documented in the Intel® Itanium® Software Conventions and Runtime Architecture Guide, section 8.4.2. You can find this design guide at web site <http://www.intel.com>

These functions can be used to gather more information, such as profiling information or timing information. Note that it is the user's responsibility to provide these profiling functions.

If you specify `-finstrument-functions` (Linux and Mac OS X) or `/Qinstrument-functions` (Windows), routine inlining is disabled. If you specify `-fno-instrument-functions` or `/Qinstrument-functions-`, inlining is not disabled.

This option is provided for compatibility with gcc.

Alternate Options

None

ip, Qip

Determines whether additional interprocedural optimizations for single-file compilation are enabled.

IDE Equivalent

Windows: **Optimization > Interprocedural Optimization**

Linux: None

Mac OS X: None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-ip`
`-no-ip`

Windows: `/Qip`
`/Qip-`

Arguments

None

Default

OFF Some limited interprocedural optimizations occur, including inline function expansion for calls to functions defined within the current source file. These optimizations are a subset of full intra-file interprocedural optimizations. Note that this setting is not the same as `-no-ip` (Linux and Mac OS X) or `/Qip-` (Windows).

Description

This option determines whether additional interprocedural optimizations for single-file compilation are enabled.

Options `-ip` (Linux and Mac OS X) and `/Qip` (Windows) enable additional interprocedural optimizations for single-file compilation.

Options `-no-ip` (Linux and Mac OS X) and `/Qip-` (Windows) may not disable inlining. To ensure that inlining of user-defined functions is disabled, specify `-inline-level=0` or `-fno-inline` (Linux and Mac OS X), or specify `/Ob0` (Windows).

Alternate Options

None

See Also

[finline-functions](#) compiler option

ip-no-inlining, Qip-no-inlining

Disables full and partial inlining enabled by interprocedural optimization options.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-ip-no-inlining`

Windows: `/Qip-no-inlining`

Arguments

None

Default

OFF Inlining enabled by interprocedural optimization options is performed.

Description

This option disables full and partial inlining enabled by the following interprocedural optimization options:

- On Linux and Mac OS X systems: `-ip` or `-ipo`
- On Windows systems: `/Qip`, `/Qipo`, or `/Ob2`

It has no effect on other interprocedural optimizations.

On Windows systems, this option also has no effect on user-directed inlining specified by option `/Ob1`.

Alternate Options

None

`ip-no-pinlining`, `Qip-no-pinlining`

Disables partial inlining enabled by interprocedural optimization options.

IDE Equivalent

None

Architectures

IA-32, Intel® 64 architectures

Syntax

Linux and Mac OS X: `-ip-no-pinlining`

Windows: `/Qip-no-pinlining`

Arguments

None

Default

OFF Inlining enabled by interprocedural optimization options is performed.

Description

This option disables partial inlining enabled by the following interprocedural optimization options:

- On Linux and Mac OS X systems: `-ip` or `-ipo`
- On Windows systems: `/Qip` or `/Qipo`

It has no effect on other interprocedural optimizations.

Alternate Options

None

IPF-flt-eval-method0, QIPF-flt-eval-method0

Tells the compiler to evaluate the expressions involving floating-point operands in the precision indicated by the variable types declared in the program. This is a deprecated option.

IDE Equivalent

None

Architectures

IA-64 architecture

SyntaxLinux: `-IPF-flt-eval-method0`Mac OS X: `None`Windows: `/QIPF-flt-eval-method0`**Arguments**

None

Default

OFF Expressions involving floating-point operands are evaluated by default rules.

Description

This option tells the compiler to evaluate the expressions involving floating-point operands in the precision indicated by the variable types declared in the program. By default, intermediate floating-point expressions are maintained in higher precision.

The recommended method to control the semantics of floating-point calculations is to use option `-fp-model` (Linux) or `/fp` (Windows).

Instead of using `-IPF-flt-eval-method0` (Linux) or `/QIPF-flt-eval-method0` (Windows), you can use `-fp-model source` (Linux) or `/fp:source` (Windows).

Alternate Options

None

See Also

[fp-model](#), [fp](#) compiler option

IPF-fltacc, QIPF-fltacc

Disables optimizations that affect floating-point accuracy. This is a deprecated option.

IDE Equivalent

Windows: **Floating Point > Floating-Point Accuracy**

Linux: None

Mac OS X: None

Architectures

IA-64 architecture

Syntax

Linux: `-IPF-fltacc`
 `-no-IPF-fltacc`

Mac OS X: None

Windows: `/QIPF-fltacc`
 `/QIPF-fltacc-`

Arguments

None

Default

`-no-IPF-fltacc` Optimizations are enabled that affect floating-point
or `/QIPF-fltacc-` accuracy.

Description

This option disables optimizations that affect floating-point accuracy.

If the default setting is used, the compiler may apply optimizations that reduce floating-point accuracy.

You can use this option to improve floating-point accuracy, but at the cost of disabling some optimizations.

The recommended method to control the semantics of floating-point calculations is to use option `-fp-model` (Linux) or `/fp` (Windows).

Instead of using `-IPF-fltacc` (Linux) or `/QIPF-fltacc` (Windows), you can use `-fp-model precise` (Linux) or `/fp:precise` (Windows).

Instead of using `-no-IPF-fltacc` (Linux) or `/QIPF-fltacc-` (Windows), you can use `-fp-model fast` (Linux) or `/fp:fast` (Windows).

Alternate Options

None

See Also

[fp-model](#), [fp](#) compiler option

IPF-fma, QIPF-fma

See [fma](#), [Qfma](#).

IPF-fp-relaxed, QIPF-fp-relaxed

See [fp-relaxed](#), [Qfp-relaxed](#).

ipo, Qipo

Enables interprocedural optimization between files.

IDE Equivalent

Windows: **Optimization > Interprocedural Optimization**

General > Whole Program Optimization

Linux: None

Mac OS X: None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-ipo[n]`

Windows: `/Qipo[n]`

Arguments

n Is an optional integer that specifies the number of object files the compiler should create. The integer must be greater than or equal to 0.

Default

OFF Multifile interprocedural optimization is not enabled.

Description

This option enables interprocedural optimization between files. This is also called multifile interprocedural optimization (multifile IPO) or Whole Program Optimization (WPO).

When you specify this option, the compiler performs inline function expansion for calls to functions defined in separate files.

You cannot specify the names for the files that are created.

If *n* is 0, the compiler decides whether to create one or more object files based on an estimate of the size of the application. It generates one object file for small applications, and two or more object files for large applications.

If *n* is greater than 0, the compiler generates *n* object files, unless *n* exceeds the number of source files (*m*), in which case the compiler generates only *m* object files.

If you do not specify *n*, the default is 0.

Alternate Options

None

See Also

Optimizing Applications:

Interprocedural Optimization (IPO) Quick Reference

Interprocedural Optimization (IPO) Overview

Using IPO

ipo-c, Qipo-c

Tells the compiler to optimize across multiple files and generate a single object file.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-ipo-c`

Windows: `/Qipo-c`

Arguments

None

Default

OFF The compiler does not generate a multifile object file.

Description

This option tells the compiler to optimize across multiple files and generate a single object file (named `ipo_out.o` on Linux and Mac OS X systems; `ipo_out.obj` on Windows systems).

It performs the same optimizations as `-ipo` (Linux and Mac OS X) or `/Qipo` (Windows), but compilation stops before the final link stage, leaving an optimized object file that can be used in further link steps.

Alternate Options

None

See Also

[ipo, Qipo](#) compiler option

ipo-jobs, Qipo-jobs

Specifies the number of commands (jobs) to be executed simultaneously during the link phase of Interprocedural Optimization (IPO).

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-ipo-jobs n`

Windows: `/Qipo-jobs: n`

Arguments

n Is the number of commands (jobs) to run simultaneously. The number must be greater than or equal to 1.

Default

`-ipo-jobs1` One command (job) is executed in an
or `/Qipo-jobs:1` interprocedural optimization parallel build.

Description

This option specifies the number of commands (jobs) to be executed simultaneously during the link phase of Interprocedural Optimization (IPO). It

should only be used if the link-time compilation is generating more than one object. In this case, each object is generated by a separate compilation, which can be done in parallel.

This option can be affected by the following compiler options:

- `-ipo` (Linux and Mac OS X) or `/Qipo` (Windows) when applications are large enough that the compiler decides to generate multiple object files.
- `-ipon` (Linux and Mac OS X) or `/Qipon` (Windows) when n is greater than 1.
- `-ipo-separate` (Linux) or `/Qipo-separate` (Windows)



Caution

Be careful when using this option. On a multi-processor system with lots of memory, it can speed application build time. However, if n is greater than the number of processors, or if there is not enough memory to avoid thrashing, this option can increase application build time.

Alternate Options

None

See Also

[ipo, Qipo](#) compiler option

[ipo-separate, Qipo-separate](#) compiler option

ipo-S, Qipo-S

Tells the compiler to optimize across multiple files and generate a single assembly file.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-ipo-S`

Windows: `/Qipo-S`

Arguments

None

Default

OFF The compiler does not generate a multifile assembly file.

Description

This option tells the compiler to optimize across multiple files and generate a single assembly file (named `ipo_out.s` on Linux and Mac OS X systems; `ipo_out.asm` on Windows systems).

It performs the same optimizations as `-ipo` (Linux and Mac OS X) or `/Qipo` (Windows), but compilation stops before the final link stage, leaving an optimized assembly file that can be used in further link steps.

Alternate Options

None

See Also

[ipo, Qipo](#) compiler option

ipo-separate, Qipo-separate

Tells the compiler to generate one object file for every source file.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux: `-ipo-separate`
Mac OS X: `None`
Windows: `/Qipo-separate`

Arguments

None

Default

OFF The compiler decides whether to create one or more object files.

Description

This option tells the compiler to generate one object file for every source file. It overrides any `-ipo` (Linux) or `/Qipo` (Windows) specification.

Alternate Options

None

See Also

[ipo, Qipo](#) compiler option

ivdep-parallel, Qivdep-parallel

Tells the compiler that there is no loop-carried memory dependency in the loop following an IVDEP directive.

IDE Equivalent

Windows: **Optimization > IVDEP Directive Memory Dependency**

Linux: None

Mac OS X: None

Architectures

IA-64 architecture

Syntax

Linux: `-ivdep-parallel`
Mac OS X: `None`
Windows: `/Qivdep-parallel`

Arguments

None

Default

OFF There may be loop-carried memory dependency in a loop that follows an IVDEP directive.

Description

This option tells the compiler that there is no loop-carried memory dependency in the loop following an IVDEP. There may be loop-carried memory dependency in a loop that follows an IVDEP directive.

This has the same effect as specifying the IVDEP:LOOP directive.

Alternate Options

None

See Also

Optimizing Applications: Absence of Loop-carried Memory Dependency with IVDEP Directive

fkeep-static-consts, Qkeep-static-consts

Tells the compiler to preserve allocation of variables that are not referenced in the source.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-fkeep-static-consts`
`-fno-keep-static-consts`

Windows: `/Qkeep-static-consts`
`/Qkeep-static-consts-`

Arguments

None

Default

<code>-fno-keep-</code>	If a variable is never referenced in a routine, the
<code>static-consts</code> or	
<code>/Qkeep-static-</code>	variable is discarded unless optimizations are
<code>consts-</code>	
	disabled by option <code>-O0</code> (Linux and Mac OS X) or
	<code>/Od</code> (Windows).

Description

This option tells the compiler to preserve allocation of variables that are not referenced in the source.

The negated form can be useful when optimizations are enabled to reduce the memory usage of static data.

Alternate Options

None

Qlocation

Specifies the directory for supporting tools.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-Qlocation,string,dir`

Windows: `/Qlocation,string,dir`

Arguments

string Is the name of the tool.

dir Is the directory (path) where the tool is located.

Default

OFF The compiler looks for tools in a default area.

Description

This option specifies the directory for supporting tools.

string can be any of the following:

- `f` - Indicates the Intel Fortran compiler.
- `fpp` (or `cpp`) - Indicates the Intel Fortran preprocessor.
- `asm` - Indicates the assembler.
- `link` - Indicates the linker.
- `prof` - Indicates the profiler.
- On Windows systems, the following is also available:
 - `masm` - Indicates the Microsoft assembler.
- On Linux and Mac OS X systems, the following are also available:
 - `as` - Indicates the assembler.
 - `gas` - Indicates the GNU assembler.
 - `ld` - Indicates the loader.

- o `gld` - Indicates the GNU loader.
- o `lib` - Indicates an additional library.
- o `crt` - Indicates the `crt%.o` files linked into executables to contain the place to start execution.

Alternate Options

None

Example

The following command provides the path for the `fpp` tool:

```
ifort -Qlocation,fpp,/usr/preproc myprog.f
```

See Also

[Qoption](#) compiler option

lowercase, Qlowercase

See [names](#).

map-opts, Qmap-opts

Maps one or more compiler options to their equivalent on a different operating system.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux: `-map-opts`

Mac OS X: `None`

Windows: `/Qmap-opts`

Arguments

None

Default

OFF No platform mappings are performed.

Description

This option maps one or more compiler options to their equivalent on a different operating system. The result is output to `stdout`.

On Windows systems, the options you provide are presumed to be Windows options, so the options that are output to `stdout` will be Linux equivalents.

On Linux systems, the options you provide are presumed to be Linux options, so the options that are output to `stdout` will be Windows equivalents.

The tool can be invoked from the compiler command line or it can be used directly.

No compilation is performed when the option mapping tool is used.

This option is useful if you have both compilers and want to convert scripts or makefiles.



Note

Compiler options are mapped to their equivalent on the architecture you are using.

For example, if you are using a processor with IA-32 architecture, you will only see equivalent options that are available on processors with IA-32 architecture.

Alternate Options

None

Example

The following command line invokes the option mapping tool, which maps the Linux options to Windows-based options, and then outputs the results to stdout:

```
ifort -map-opts -xP -O2
```

The following command line invokes the option mapping tool, which maps the Windows options to Linux-based options, and then outputs the results to stdout:

```
ifort /Qmap-opts /QxP /O2
```

See Also

Building Applications: Using the Option Mapping Tool

no-bss-init, Qnobss-init

Tells the compiler to place in the DATA section any variables explicitly initialized with zeros.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-no-bss-init`

Windows: `/Qnobss-init`

Arguments

None

Default

OFF Variables explicitly initialized with zeros are placed in the BSS section.

Description

This option tells the compiler to place in the DATA section any variables explicitly initialized with zeros.

Alternate Options

Linux and Mac OS X: `-nobss-init` (this is a [deprecated](#) option)

Windows: None

onetrip, Qonetrip

Tells the compiler to follow the FORTRAN 66 Standard and execute DO loops at least once.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-onetrip`

Windows: `/Qonetrip`

Arguments

None

Default

OFF The compiler applies the current Fortran Standard semantics, which allows zero-trip DO loops.

Description

This option tells the compiler to follow the FORTRAN 66 Standard and execute DO loops at least once.

Alternate Options

Linux and Mac OS X: -1

Windows: /1

openmp, Qopenmp

Enables the parallelizer to generate multi-threaded code based on the OpenMP* directives.

IDE Equivalent

Windows: **Language > Process OpenMP Directives**

Linux: None

Mac OS X: None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: -openmp

Windows: /Qopenmp

Arguments

None

Default

OFF No OpenMP multi-threaded code is generated by the compiler.

Description

This option enables the parallelizer to generate multi-threaded code based on the OpenMP* directives. The code can be executed in parallel on both uniprocessor and multiprocessor systems.

If you use this option, multithreaded libraries are used, but option `fpp` is not automatically invoked.

This option sets option `automatic`.

This option works with any optimization level. Specifying no optimization (`-O0` on Linux or `/Od` on Windows) helps to debug OpenMP applications.



Note

On Mac OS X systems, when you enable OpenMP*, you must also set the `DYLD_LIBRARY_PATH` environment variable within Xcode or an error will be displayed.

Alternate Options

None

See Also

[openmp-stubs, Qopenmp-stubs](#) compiler option

openmp-lib, Qopenmp-lib

Lets you specify an OpenMP* run-time library to use for linking.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux: `-openmp-lib type`

Mac OS X: `None`

Windows: `/Qopenmp-lib:type`

Arguments

type

Specifies the type of library to use; it implies compatibility levels. Possible values are:

`legacy` Tells the compiler to use the legacy OpenMP* run-time library (libguide). This setting does not provide compatibility with object files created using other compilers. This is a [deprecated](#) option.

`compat` Tells the compiler to use the compatibility OpenMP* run-time library (libiomp). This setting provides compatibility with object files created using Microsoft* and GNU* compilers.

Default

`-openmp-lib compat` The compiler uses the compatibility
`or/Qopenmp-lib:compat` OpenMP* run-time library (libiomp).

Description

This option lets you specify an OpenMP* run-time library to use for linking. The legacy OpenMP run-time library is not compatible with object files created using OpenMP run-time libraries supported in other compilers.

The compatibility OpenMP run-time library is compatible with object files created using the Microsoft* OpenMP run-time library (vcomp) and GNU OpenMP run-time library (libgomp).

To use the compatibility OpenMP run-time library, compile and link your application using the `-openmp-lib compat` (Linux) or `/Qopenmp-lib:compat` (Windows) option. To use this option, you must also specify one of the following compiler options:

- Linux OS: `-openmp`, `-openmp-profile`, or `-openmp-stubs`
- Windows OS: `/Qopenmp`, `/Qopenmp-profile`, or `/Qopenmp-stubs`

On Windows* systems, the compatibility OpenMP* run-time library lets you combine OpenMP* object files compiled with the Microsoft* C/C++ compiler with OpenMP* object files compiled with the Intel C/C++ or Fortran compilers. The linking phase results in a single, coherent copy of the run-time library.

On Linux* systems, the compatibility Intel OpenMP* run-time library lets you combine OpenMP* object files compiled with the GNU* gcc or gfortran compilers with similar OpenMP* object files compiled with the Intel C/C++ or Fortran compilers. The linking phase results in a single, coherent copy of the run-time library.

You cannot link object files generated by the Intel® Fortran compiler to object files compiled by the GNU Fortran compiler, regardless of the presence or absence of the `-openmp` (Linux) or `/Qopenmp` (Windows) compiler option. This is because the Fortran run-time libraries are incompatible.



Note

The compatibility OpenMP run-time library is not compatible with object files created using versions of the Intel compiler earlier than 10.0.

Alternate Options

None

See Also

[openmp, Qopenmp](#) compiler option

[openmp-stubs, Qopenmp-stubs](#) compiler option

[openmp-profile, Qopenmp-profile](#) compiler option

openmp-link, Qopenmp-link

Controls whether the compiler links to static or dynamic OpenMP run-time libraries.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-openmp-link library`

Windows: `/Qopenmp-link:library`

Arguments

<i>library</i>	Specifies the OpenMP library to use. Possible values are:
<code>static</code>	Tells the compiler to link to static OpenMP run-time libraries.
<code>dynamic</code>	Tells the compiler to link to dynamic OpenMP run-time libraries.

Default

`-openmp-link dynamic` or `/Qopenmp-` The compiler links to dynamic OpenMP run-time libraries. However, if option

`link:dynamic` `static` is specified, the compiler links to static OpenMP run-time libraries.

Description

This option controls whether the compiler links to static or dynamic OpenMP run-time libraries.

To link to the static OpenMP run-time library (RTL) and create a purely static executable, you must specify `-openmp-link static` (Linux and Mac OS X) or `/Qopenmp-link:static` (Windows). However, we strongly recommend you use the default setting, `-openmp-link dynamic` (Linux and Mac OS X) or `/Qopenmp-link:dynamic` (Windows).



Note

Compiler options `-static-intel` and `-shared-intel` (Linux and Mac OS X) have no effect on which OpenMP run-time library is linked.

Alternate Options

None

openmp-profile, Qopenmp-profile

Enables analysis of OpenMP* applications if Intel® Thread Profiler is installed.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux: `-openmp-profile`

Mac OS X: `None`

Windows: `/Qopenmp-profile`

Arguments

None

Default

OFF OpenMP applications are not analyzed.

Description

This option enables analysis of OpenMP* applications. To use this option, you must have previously installed Intel® Thread Profiler, which is one of the Intel® Threading Analysis Tools.

This option can adversely affect performance because of the additional profiling and error checking invoked to enable compatibility with the threading tools. Do not use this option unless you plan to use the Intel® Thread Profiler.

For more information about Intel® Thread Profiler (including an evaluation copy) open the page associated with threading tools at [Intel® Software Development Products](#).

Alternate Options

None

openmp-report, Qopenmp-report

Controls the OpenMP* parallelizer's level of diagnostic messages.

IDE Equivalent

Windows: **Compilation Diagnostics > OpenMP Diagnostic Level**

Linux: None

Mac OS X: **Compiler Diagnostics > OpenMP Report**

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-openmp-report [n]`

Windows: `/Qopenmp-report [n]`

Arguments

<i>n</i>	Is the level of diagnostic messages to display. Possible values are:
0	No diagnostic messages are displayed.
1	Diagnostic messages are displayed indicating loops, regions, and sections successfully parallelized.
2	The same diagnostic messages are displayed as specified by <code>openmp_report1</code> plus diagnostic messages indicating successful handling of MASTER constructs, SINGLE constructs, CRITICAL constructs, ORDERED constructs, ATOMIC directives, and so forth.

Default

<code>-openmp-report1</code>	This is the default if you do not specify <i>n</i> .
<code>or/Qopenmp-report1</code>	The compiler displays diagnostic messages

indicating loops, regions, and sections successfully parallelized. If you do not specify the option on the command line, the default is to display no messages.

Description

This option controls the OpenMP* parallelizer's level of diagnostic messages. To use this option, you must also specify `-openmp` (Linux and Mac OS X) or `/Qopenmp` (Windows).

If this option is specified on the command line, the report is sent to `stdout`. On Windows systems, if this option is specified from within the IDE, the report is included in the build log if the Generate Build Logs option is selected.

Alternate Options

None

See Also

[openmp, Qopenmp](#) compiler option

Optimizing Applications:

Using Parallelism

OpenMP* Report

openmp-stubs, Qopenmp-stubs

Enables compilation of OpenMP programs in sequential mode.

IDE Equivalent

Windows: **Language > Process OpenMP Directives**

Linux: None

Mac OS X: **Language > Process OpenMP Directives**

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-openmp-stubs`

Windows: `/Qopenmp-stubs`

Arguments

None

Default

OFF The library of OpenMP function stubs is not linked.

Description

This option enables compilation of OpenMP programs in sequential mode. The OpenMP directives are ignored and a stub OpenMP library is linked.

Alternate Options

None

See Also

[openmp, Qopenmp](#) compiler option

openmp-threadprivate, Qopenmp-threadprivate

Lets you specify an OpenMP* threadprivate implementation.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux: `-openmp-threadprivate type`
 Mac OS X: None
 Windows: `/Qopenmp-threadprivate:type`

Arguments

type Specifies the type of threadprivate implementation.

Possible values are:

`legacy` Tells the compiler to use the legacy OpenMP* threadprivate implementation used in the previous releases of the Intel® compiler. This setting does not provide compatibility with the implementation used by other compilers.

`compat` Tells the compiler to use the compatibility OpenMP* threadprivate implementation based on applying the thread-local attribute to each threadprivate variable. This setting provides compatibility with the implementation provided by the Microsoft* and GNU* compilers.

Default

`-openmp-threadprivate legacy` or `/Qopenmp-` The compiler uses the legacy OpenMP* threadprivate implementation used in the previous releases of the Intel® compiler.

`threadprivate:legacy`

Description

This option lets you specify an OpenMP* threadprivate implementation.

The legacy OpenMP run-time library is not compatible with object files created using OpenMP run-time libraries supported in other compilers.

To use this option, you must also specify one of the following compiler options:

- Linux OS: `-openmp`, `-openmp-profile`, or `-openmp-stubs`
- Windows OS: `/Qopenmp`, `/Qopenmp-profile`, or `/Qopenmp-stubs`

The value specified for this option is independent of the value used for option `-openmp-lib` (Linux) or `/Qopenmp-lib` (Windows).

Alternate Options

None

opt-block-factor, Qopt-block-factor

Lets you specify a loop blocking factor.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-opt-block-factor=n`

Windows: `/Qopt-block-factor:n`

Arguments

n Is the blocking factor. It must be an integer. The compiler may ignore the blocking factor if the value

is 0 or 1.

Default

OFF The compiler uses default heuristics for loop blocking.

Description

This option lets you specify a loop blocking factor.

Alternate Options

None

opt-jump-tables, Qopt-jump-tables

Enables or disables generation of jump tables for switch statements.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-opt-jump-tables=keyword`
`-no-opt-jump-tables`

Windows: `/Qopt-jump-tables:keyword`
`/Qopt-jump-tables-`

Arguments

keyword Is the instruction for generating jump tables.
 Possible values are:

<code>never</code>	Tells the compiler to never generate jump tables. All
--------------------	---

switch statements are implemented as chains of if-then-elses. This is the same as specifying `-no-opt-jump-tables` (Linux and Mac OS) or `/Qopt-jump-tables-` (Windows).

<code>default</code>	The compiler uses default heuristics to determine when to generate jump tables.
<code>large</code>	Tells the compiler to generate jump tables up to a certain pre-defined size (64K entries).
<code>n</code>	Must be an integer. Tells the compiler to generate jump tables up to n entries in size.

Default

`-opt-jump-tables=default` or `/Qopt-jump-tables:default` The compiler uses default heuristics to determine when to generate jump tables for switch statements.

Description

This option enables or disables generation of jump tables for switch statements. When the option is enabled, it may improve performance for programs with large switch statements.

Alternate Options

None

opt-loadpair, Qopt-loadpair

Enables or disables loadpair optimization.

IDE Equivalent

None

Architectures

IA-64 architecture

Syntax

Linux:	<code>-opt-loadpair</code> <code>-no-opt-loadpair</code>
Mac OS X:	None
Windows:	<code>/Qopt-loadpair</code> <code>/Qopt-loadpair-</code>

Arguments

None

Default

<code>-no-opt-loadpair</code>	Loadpair optimization is disabled
<code>or/Qopt-loadpair-</code>	unless option <code>O3</code> is specified.

Description

This option enables or disables loadpair optimization.

When `-O3` is specified on IA-64 architecture, loadpair optimization is enabled by default. To disable loadpair generation, specify `-no-opt-loadpair` (Linux) or `/Qopt-loadpair-` (Windows).

Alternate Options

None

opt-mem-bandwidth, Qopt-mem-bandwidth

Enables performance tuning and heuristics that control memory bandwidth use among processors.

IDE Equivalent

None

Architectures

IA-64 architecture

Syntax

Linux: `-opt-mem-bandwidthn`

Mac OS X: `None`

Windows: `/Qopt-mem-bandwidthn`

Arguments

n Is the level of optimizing for memory bandwidth usage. Possible values are:

- | | |
|---|--|
| 0 | Enables a set of performance tuning and heuristics in compiler optimizations that is optimal for serial code. |
| 1 | Enables a set of performance tuning and heuristics in compiler optimizations for multithreaded code generated by the compiler. |
| 2 | Enables a set of performance tuning and heuristics in compiler optimizations for |

parallel code such as Windows Threads, pthreads, and MPI code, besides multithreaded code generated by the compiler.

Default

<code>-opt-mem-bandwidth0</code> or <code>/Qopt-mem-bandwidth0</code>	For serial (non-parallel) compilation, a set of performance tuning and heuristics in compiler optimizations is enabled that is optimal for serial code.
<code>-opt-mem-bandwidth1</code> or <code>/Qopt-mem-bandwidth1</code>	If you specify compiler option <code>-parallel</code> (Linux) or <code>/Qparallel</code> (Windows), <code>-openmp</code> (Linux) or <code>/Qopenmp</code> (Windows), or Cluster OpenMP option <code>-cluster-openmp</code> (Linux), a set of performance tuning and heuristics in compiler optimizations for multithreaded code generated by the compiler is enabled.

Description

This option enables performance tuning and heuristics that control memory bandwidth use among processors. It allows the compiler to be less aggressive with optimizations that might consume more bandwidth, so that the bandwidth can be well-shared among multiple processors for a parallel program. For values of n greater than 0, the option tells the compiler to enable a set of performance tuning and heuristics in compiler optimizations such as prefetching, privatization, aggressive code motion, and so forth, for reducing memory bandwidth pressure and balancing memory bandwidth traffic among threads.

This option can improve performance for threaded or parallel applications on multiprocessors or multicore processors, especially when the applications are bounded by memory bandwidth.

Alternate Options

None

See Also

[parallel, Qparallel](#) compiler option

[openmp, Qopenmp](#) compiler option

opt-mod-versioning, Qopt-mod-versioning

Enables or disables versioning of modulo operations for certain types of operands.

IDE Equivalent

None

Architectures

IA-64 architecture

Syntax

Linux: `-opt-mod-versioning`
 `-no-opt-mod-versioning`

Mac OS X: None

Windows: `/Qopt-mod-versioning`
 `/Qopt-mod-versioning-`

Arguments

None

Default

`-no-opt-mod-versioning` Versioning of modulo operations is
or `/Qopt-mod-versioning-` disabled.

Description

This option enables or disables versioning of modulo operations for certain types of operands. It is used for optimization tuning.

Versioning of modulo operations may improve performance for $x \bmod y$ when modulus y is a power of 2.

Alternate Options

None

opt-multi-version-aggressive, Qopt-multi-version-aggressive

Tells the compiler to use aggressive multi-versioning to check for pointer aliasing and scalar replacement.

IDE Equivalent

None

Architectures

IA-32, Intel® 64 architectures

Syntax

Linux and Mac OS X: `-opt-multi-version-aggressive`
 `-no-opt-multi-version-aggressive`
Windows: `/Qopt-multi-version-aggressive`
 `/Qopt-multi-version-aggressive-`

Arguments

None

Default

`-no-opt-multi-version-aggressive` or `/Qopt-multi-version-aggressive-` The compiler uses default heuristics when checking for pointer aliasing and scalar replacement.

Description

This option tells the compiler to use aggressive multi-versioning to check for pointer aliasing and scalar replacement. This option may improve performance.

Alternate Options

None

opt-prefetch, Qopt-prefetch

Enables or disables prefetch insertion optimization.

IDE Equivalent

Windows: **Optimization > Prefetch Insertion**

Linux: None

Mac OS X: **Optimization > Enable Prefetch Insertion**

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-opt-prefetch[=n]`
`-no-opt-prefetch`

Windows: `/Qopt-prefetch[:n]`
`/Qopt-prefetch-`

Arguments

n Is the level of detail in the report. Possible values

are:

- | | |
|--------|--|
| 0 | Disables software prefetching. This is the same as specifying <code>-no-opt-prefetch</code> (Linux and Mac OS X) or <code>/Qopt-prefetch-</code> (Windows). |
| 1 to 4 | Enables different levels of software prefetching. If you do not specify a value for <i>n</i> , the default is 2 on IA-32 and Intel® 64 architecture; the default is 3 on IA-64 architecture. Use lower values to reduce the amount of prefetching. |

Default

IA-64 architecture: <code>-opt-prefetch</code> or <code>/Qopt-prefetch</code>	On IA-64 architecture, prefetch insertion optimization is enabled.
--	--

IA-32 architecture and Intel® 64 architecture: <code>-no-opt-prefetch</code> or <code>/Qopt-prefetch-</code>	On IA-32 architecture and Intel® 64 architecture, prefetch insertion optimization is disabled.
--	--

Description

This option enables or disables prefetch insertion optimization. The goal of prefetching is to reduce cache misses by providing hints to the processor about when data should be loaded into the cache.

On IA-64 architecture, this option is enabled by default if you specify option `O1` or higher. To disable prefetching at these optimization levels, specify `-no-opt-prefetch` (Linux and Mac OS X) or `/Qopt-prefetch-` (Windows).

On IA-32 architecture and Intel® 64 architecture, this option enables prefetching when higher optimization levels are specified.

Alternate Options

Linux and Mac OS X: `-prefetch` (this is a [deprecated](#) option)

Windows: `/Qprefetch` (this is a [deprecated](#) option)

`opt-prefetch-initial-values`, `Qopt-prefetch-initial-values`

Enables or disables prefetches that are issued before a loop is entered.

IDE Equivalent

None

Architectures

IA-64 architecture

Syntax

Linux: `-opt-prefetch-initial-values`
`-no-opt-prefetch-initial-values`

Mac OS X: `None`

Windows: `/Qopt-prefetch-initial-values`
`/Qopt-prefetch-initial-values-`

Arguments

None

Default

`-opt-prefetch-initial-` Prefetches are issued before a loop

values is entered.

or /Qopt-prefetch-initial-values

Description

This option enables or disables prefetches that are issued before a loop is entered. These prefetches target the initial iterations of the loop.

When `-O1` or higher is specified on IA-64 architecture, prefetches are issued before a loop is entered. To disable these prefetches, specify `-no-opt-prefetch-initial-values` (Linux) or `/Qopt-prefetch-initial-values-` (Windows).

Alternate Options

None

opt-prefetch-issue-excl-hint, Qopt-prefetch-issue-excl-hint

Determines whether the compiler issues prefetches for stores with exclusive hint.

IDE Equivalent

None

Architectures

IA-64 architecture

Syntax

Linux: `-opt-prefetch-issue-excl-hint`
`-no-opt-prefetch-issue-excl-hint`

Mac OS X: None

Windows: `/Qopt-prefetch-issue-excl-hint`
`/Qopt-prefetch-issue-excl-hint-`

Arguments

None

Default

<code>-no-opt-prefetch-issue-excl-hint</code>	The compiler does not issue
<code>/Qopt-prefetch-issue-excl-hint-</code>	prefetches for stores with exclusive
	hint.

Description

This option determines whether the compiler issues prefetches for stores with exclusive hint. If option `-opt-prefetch-issue-excl-hint` (Linux) or `/Qopt-prefetch-issue-excl-hint` (Windows) is specified, the prefetches will be issued if the compiler determines it is beneficial to do so.

When prefetches are issued for stores with `exclusive-hint`, the cache-line is in "exclusive-mode". This saves on cache-coherence traffic when other processors try to access the same cache-line. This feature can improve performance tuning.

Alternate Options

None

opt-prefetch-next-iteration, Qopt-prefetch-next-iteration

Enables or disables prefetches for a memory access in the next iteration of a loop.

IDE Equivalent

None

Architectures

IA-64 architecture

Syntax

Linux: `-opt-prefetch-next-iteration`
`-no-opt-prefetch-next-iteration`

Mac OS X: None

Windows: `/Qopt-prefetch-next-iteration`
`/Qopt-prefetch-next-iteration-`

Arguments

None

Default

`-opt-prefetch-next-iteration` Prefetches are issued for a memory access in the next iteration of a loop.

or `/Qopt-prefetch-next-iteration`

Description

This option enables or disables prefetches for a memory access in the next iteration of a loop. It is typically used in a pointer-chasing loop.

When `-O1` or higher is specified on IA-64 architecture, prefetches are issued for a memory access in the next iteration of a loop. To disable these prefetches, specify `-no-opt-prefetch-next-iteration` (Linux) or `/Qopt-prefetch-next-iteration-` (Windows).

Alternate Options

None

opt-ra-region-strategy, Qopt-ra-region-strategy

Selects the method that the register allocator uses to partition each routine into regions.

IDE Equivalent

None

Architectures

IA-32, Intel® 64 architectures

Syntax

Linux and Mac OS X: `-opt-ra-region-strategy[=keyword]`

Windows: `/Qopt-ra-region-strategy[:keyword]`

Arguments

<i>keyword</i>	Is the method used for partitioning. Possible values are:
<code>routine</code>	Creates a single region for each routine.
<code>block</code>	Partitions each routine into one region per basic block.
<code>trace</code>	Partitions each routine into one region per trace.
<code>region</code>	Partitions each routine into one region per loop.
<code>default</code>	The compiler determines which method is used for partitioning.

Default

<code>-opt-ra-region-strategy=default</code> <code>or/Qopt-ra-region-strategy:default</code>	The compiler determines which method is used for partitioning. This is also the default if <i>keyword</i> is not specified.
---	---

Description

This option selects the method that the register allocator uses to partition each routine into regions.

When setting `default` is in effect, the compiler attempts to optimize the tradeoff between compile-time performance and generated code performance.

This option is only relevant when optimizations are enabled (`O1` or higher).

Alternate Options

None

See Also

[o](#) compiler option

opt-report, Qopt-report

Tells the compiler to generate an optimization report to `stderr`.

IDE Equivalent

Windows: Diagnostics > Optimization Diagnostics Level

Linux: None

Mac OS X: None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-opt-report[n]`

Windows: `/Qopt-report[:n]`

Arguments

n Is the level of detail in the report. Possible values are:

0 Tells the compiler to generate

- no optimization report.
- | | |
|---|---|
| 1 | Tells the compiler to generate a report with the minimum level of detail. |
| 2 | Tells the compiler to generate a report with the medium level of detail. |
| 3 | Tells the compiler to generate a report with the maximum level of detail. |

Default

`-opt-` If you do not specify n , the compiler generates a report
`report 2` or with medium detail. If you do not specify the option on the
`/Qopt-` command line, the compiler does not generate an
`report:2` optimization report.

Description

This option tells the compiler to generate an optimization report to `stderr`.

Alternate Options

None

See Also

[opt-report-file, Qopt-report-file](#) compiler option
Optimizing Applications: Optimizer Report Generation

opt-report-file, Qopt-report-file

Specifies the name for an optimization report.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-opt-report-file=file`

Windows: `/Qopt-report-file:file`

Arguments

file Is the name for the optimization report.

Default

OFF No optimization report is generated.

Description

This option specifies the name for an optimization report. If you use this option, you do not have to specify `-opt-report` (Linux and Mac OS X) or `/Qopt-report` (Windows).

Alternate Options

None

See Also

[opt-report](#), [Qopt-report](#) compiler option

Optimizing Applications: Optimizer Report Generation

opt-report-help, Qopt-report-help

Displays the optimizer phases available for report generation.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-opt-report-help`

Windows: `/Qopt-report-help`

Arguments

None

Default

OFF No optimization reports are generated.

Description

This option displays the optimizer phases available for report generation using `-opt-report-phase` (Linux and Mac OS X) or `/Qopt-report-phase` (Windows). No compilation is performed.

Alternate Options

None

See Also

[opt-report, Qopt-report](#) compiler option

[opt-report-phase, Qopt-report-phase](#) compiler option

opt-report-phase, Qopt-report-phase

Specifies an optimizer phase to use when optimization reports are generated.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-opt-report-phase=phase`

Windows: `/Qopt-report-phase:phase`

Arguments

phase Is the phase to generate reports for. Some of the possible values are:

<code>ipo</code>	The Interprocedural Optimizer phase
<code>hlo</code>	The High Level Optimizer phase
<code>hpo</code>	The High Performance Optimizer phase
<code>ilo</code>	The Intermediate Language Scalar Optimizer phase
<code>ecg</code>	The Code Generator phase (Windows and Linux systems using IA-64 architecture only)
<code>ecg_swp</code>	The software pipelining component of the Code Generator phase (Windows and Linux systems using IA-64 architecture only)
<code>pgo</code>	The Profile Guided Optimization phase

`all` All optimizer phases

Default

OFF No optimization reports are generated.

Description

This option specifies an optimizer phase to use when optimization reports are generated. To use this option, you must also specify `-opt-report` (Linux and Mac OS X) or `/Qopt-report` (Windows).

This option can be used multiple times on the same command line to generate reports for multiple optimizer phases.

When one of the logical names for optimizer phases is specified for phase, all reports from that optimizer phase are generated.

To find all phase possibilities, use option `-opt-report-help` (Linux and Mac OS X) or `/Qopt-report-help` (Windows).

Alternate Options

None

See Also

[opt-report, Qopt-report](#) compiler option

opt-report-routine, Qopt-report-routine

Tells the compiler to generate reports on the routines containing specified text.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-opt-report-routine=string`

Windows: `/Qopt-report-routine:string`

Arguments

string Is the text (string) to look for.

Default

OFF No optimization reports are generated.

Description

This option tells the compiler to generate reports on the routines containing specified text as part of their name.

Alternate Options

None

See Also

[opt-report, Qopt-report](#) compiler option

opt-streaming-stores, Qopt-streaming-stores

Enables generation of streaming stores for optimization.

IDE Equivalent

None

Architectures

IA-32, Intel® 64 architectures

Syntax

Linux and Mac OS X: `-opt-streaming-stores keyword`

Windows: `/Qopt-streaming-stores:keyword`

Arguments

<i>keyword</i>	Specifies whether streaming stores are generated. Possible values are:
<code>always</code>	Enables generation of streaming stores for optimization. The compiler optimizes under the assumption that the application is memory bound.
<code>never</code>	Disables generation of streaming stores for optimization. Normal stores are performed.
<code>auto</code>	Lets the compiler decide which instructions to use.

Default

`-opt-streaming-stores auto`
The compiler decides whether to use streaming stores or normal stores.

`or/Qopt-streaming-stores:auto`

Description

This option enables generation of streaming stores for optimization. This method stores data with instructions that use a non-temporal buffer, which minimizes memory hierarchy pollution.

For this option to be effective, the compiler must be able to generate SSE2 (or higher) instructions. For more information, see compiler option `x` or `ax`.

This option may be useful for applications that can benefit from streaming stores.

Alternate Options

None

See Also

[ax, Qax](#) compiler option

[x, Qx](#) compiler option

[opt-mem-bandwidth, Qopt-mem-bandwidth, Qx](#) compiler option

Optimizing Applications: Vectorization Support

opt-subscript-in-range, Qopt-subscript-in-range

Determines whether the compiler assumes no overflows in the intermediate computation of subscript expressions in loops.

IDE Equivalent

None

Architectures

IA-32, Intel® 64 architectures

Syntax

Linux and Mac OS X: `-opt-subscript-in-range`
`-no-opt-subscript-in-range`

Windows: `/Qopt-subscript-in-range`
`/Qopt-subscript-in-range-`

Arguments

None

Default

`-no-opt-` The compiler assumes overflows in the

`subscript-in-range` intermediate computation of subscript expressions in loops.

`/Qopt-`

`subscript-in-range-`

Description

This option determines whether the compiler assumes no overflows in the intermediate computation of subscript expressions in loops.

If you specify `-opt-subscript-in-range` (Linux and Mac OS X) or `/Qopt-subscript-in-range` (Windows), the compiler ignores any data type conversions used and it assumes no overflows in the intermediate computation of subscript expressions. This feature can enable more loop transformations.

Alternate Options

None

Example

The following shows an example where these options can be useful. `m` is declared as type `integer(kind=8)` (64-bits) and all other variables inside the subscript are declared as type `integer(kind=4)` (32-bits):

```
A[ i + j + ( n + k ) * m ]
```

Option

Passes options to a specified tool.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-Qoption,string,options`

Windows: `/Qoption,string,options`

Arguments

string Is the name of the tool.

options Are one or more comma-separated, valid options for the designated tool.

Default

OFF No options are passed to tools.

Description

This option passes options to a specified tool.

If an argument contains a space or tab character, you must enclose the entire argument in quotation marks (" "). You must separate multiple arguments with commas.

string can be any of the following:

- fpp (or cpp) - Indicates the Intel Fortran preprocessor.
- asm - Indicates the assembler.
- link - Indicates the linker.
- prof - Indicates the profiler.
- On Windows systems, the following is also available:
 - masm - Indicates the Microsoft assembler.
- On Linux and Mac OS X systems, the following are also available:
 - as - Indicates the assembler.
 - gas - Indicates the GNU assembler.
 - ld - Indicates the loader.
 - gld - Indicates the GNU loader.
 - lib - Indicates an additional library.

- o crt - Indicates the crt%.o files linked into executables to contain the place to start execution.

Alternate Options

None

Example

On Linux and Mac OS X systems:

The following example directs the linker to link with an alternative library:

```
ifort -Qoption,link,-L.,-Lmylib prog1.f
```

The following example passes a compiler option to the assembler to generate a listing file:

```
ifort -Qoption,as,"-as=myprogram.lst" -use-asm myprogram.f90
```

On Windows systems:

The following example directs the linker to create a memory map when the compiler produces the executable file from the source being compiled:

```
ifort /Qoption,link,/map:prog1.map prog1.f
```

The following example passes a compiler option to the assembler:

```
ifort /Quse_asm /Qoption,masm,"/WX" myprogram.f90
```

See Also

[Qlocation](#) compiler option

qp

See [p](#).

pad, Qpad

Enables the changing of the variable and array memory layout.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-pad`
 `-nopad`

Windows: `/Qpad`
 `/Qpad-`

Arguments

None

Default

`-nopad` Variable and array memory layout is performed by default
 or methods.
`/Qpad-`

Description

This option enables the changing of the variable and array memory layout.
 This option is effectively not different from the `align` option when applied to
 structures and derived types. However, the scope of `pad` is greater because it
 applies also to common blocks, derived types, sequence types, and structures.

Alternate Options

None

See Also

[align](#) compiler option

pad-source, Qpad-source

Specifies padding for fixed-form source records.

IDE Equivalent

Windows: **Language > Pad Fixed Form Source Lines**

Linux: None

Mac OS X: **Language > Pad Fixed Form Source Lines**

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-pad-source`
`-nopad-source`

Windows: `/pad-source`
`/nopad-source`
`/Qpad-source`
`/Qpad-source-`

Arguments

None

Default

`-nopad-` Fixed-form source records are not padded.

`source`

or

`/Qpad-`

`source-`

Description

This option specifies padding for fixed-form source records. It tells the compiler that fixed-form source lines shorter than the statement field width are to be padded with spaces to the end of the statement field. This affects the interpretation of character and Hollerith literals that are continued across source records.

The default value setting causes a warning message to be displayed if a character or Hollerith literal that ends before the statement field ends is continued onto the next source record. To suppress this warning message, specify option `-warn nousage` (Linux and Mac OS X) or `/warn:nousage` (Windows).

Specifying `pad-source` or `/Qpad-source` can prevent warning messages associated with option `-warn usage` (Linux and Mac OS X) or `/warn:usage` (Windows).

Alternate Options

None

See Also

[warn](#) compiler option

Qpar-adjust-stack

Tells the compiler to generate code to adjust the stack size for a fiber-based main thread.

IDE Equivalent

None

Architectures

IA-32, Intel® 64 architectures

Syntax

Linux and Mac OS X: None

Windows: `/Qpar-adjust-stack:n`

Arguments

n Is the stack size (in bytes) for the fiber-based main thread. It must be a number equal to or greater

than zero.

Default

`/Qpar-adjust-stack:0` No adjustment is made to the main thread stack size.

Description

This option tells the compiler to generate code to adjust the stack size for a fiber-based main thread. This can reduce the stack size of threads.

For this option to be effective, you must also specify option `/Qparallel`.

Alternate Options

None

See Also

[parallel, Qparallel](#) compiler option

par-report, Qpar-report

Controls the diagnostic information reported by the auto-parallelizer.

IDE Equivalent

Windows: **Compilation Diagnostics > Auto-Parallelizer Diagnostic Level**

Linux: None

Mac OS X: **Diagnostics > Auto-Parallelizer Report**

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-par-report[n]`

Windows: `/Qpar-report[n]`

Arguments

n Is a value denoting which diagnostic messages to report. Possible values are:

- | | |
|---|---|
| 0 | Tells the auto-parallelizer to report no diagnostic information. |
| 1 | Tells the auto-parallelizer to report diagnostic messages for loops successfully auto-parallelized. The compiler also issues a "LOOP AUTO-PARALLELIZED" message for parallel loops. |
| 2 | Tells the auto-parallelizer to report diagnostic messages for loops successfully and unsuccessfully auto-parallelized. |
| 3 | Tells the auto-parallelizer to report the same diagnostic messages specified by 2 plus additional information about any proven or assumed dependencies inhibiting auto-parallelization (reasons for not parallelizing). |

Default

`-par-` If you do not specify *n*, the compiler displays diagnostic

`report1` messages for loops successfully auto-parallelized. If you
`or/Qpar-` do not specify the option on the command line, the default
`report1` is to display no parallel diagnostic messages.

Description

This option controls the diagnostic information reported by the auto-parallelizer (parallel optimizer). To use this option, you must also specify `-parallel` (Linux and Mac OS X) or `/Qparallel` (Windows).

If this option is specified on the command line, the report is sent to `stdout`.

On Windows systems, if this option is specified from within the IDE, the report is included in the build log if the Generate Build Logs option is selected.

Alternate Options

None

par-runtime-control, Qpar-runtime-control

Generates code to perform run-time checks for loops that have symbolic loop bounds.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-par-runtime-control`
`-no-par-runtime-control`

Windows: `/Qpar-runtime-control`
`/Qpar-runtime-control-`

Arguments

None

Default

`-no-par-` The compiler uses default heuristics when
`runtime-control` checking loops.
`or/Qpar-runtime-`
`control-`

Description

This option generates code to perform run-time checks for loops that have symbolic loop bounds.

If the granularity of a loop is greater than the parallelization threshold, the loop will be executed in parallel.

If you do not specify this option, the compiler may not parallelize loops with symbolic loop bounds if the compile-time granularity estimation of a loop can not ensure it is beneficial to parallelize the loop.

Alternate Options

None

par-schedule, Qpar-schedule

Lets you specify a scheduling algorithm or a tuning method for loop iterations.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-par-schedule-keyword[=n]`

Windows: `/Qpar-schedule-keyword[[:]n]`

Arguments

<i>keyword</i>	Specifies the scheduling algorithm or tuning method. Possible values are:
<code>auto</code>	Lets the compiler or run-time system determine the scheduling algorithm.
<code>static</code>	Divides iterations into contiguous pieces.
<code>static-balanced</code>	Divides iterations into even-sized chunks.
<code>static-steal</code>	Divides iterations into even-sized chunks, but allows threads to steal parts of chunks from neighboring threads.
<code>dynamic</code>	Gets a set of iterations dynamically.
<code>guided</code>	Specifies a minimum number of iterations.
<code>guided-analytical</code>	Divides iterations by using exponential distribution or dynamic distribution.
<code>runtime</code>	Defers the scheduling decision until run time.
<i>n</i>	Is the size of the chunk or the number of iterations for each chunk. This setting can only be specified

for static, dynamic, and guided. For more information, see the descriptions of each keyword below.

Default

`static-` Iterations are divided into even-sized chunks and the chunks balanced are assigned to the threads in the team in a round-robin fashion in the order of the thread number.

Description

This option lets you specify a scheduling algorithm or a tuning method for loop iterations. It specifies how iterations are to be divided among the threads of the team.

This option affects performance tuning and can provide better performance during auto-parallelization.

Option	Description
<code>-par-schedule-auto</code> or <code>/Qpar-schedule-auto</code>	Lets the compiler or run-time system determine the scheduling algorithm. Any possible mapping may occur for iterations to threads in the team.
<code>-par-schedule-static</code> or <code>/Qpar-schedule-static</code>	Divides iterations into contiguous pieces (chunks) of size n . The chunks are assigned to threads in the team in a round-robin fashion in the order of the thread number. Note that the last chunk to be assigned may have a smaller number of iterations. If no n is specified, the iteration space is divided into chunks that are

Option	Description
	approximately equal in size, and each thread is assigned at most one chunk.
<code>-par-schedule-static-balanced</code> or <code>/Qpar-schedule-static-balanced</code>	Divides iterations into even-sized chunks. The chunks are assigned to the threads in the team in a round-robin fashion in the order of the thread number.
<code>-par-schedule-static-steal</code> or <code>/Qpar-schedule-static-steal</code>	Divides iterations into even-sized chunks, but when a thread completes its chunk, it can steal parts of chunks assigned to neighboring threads. Each thread keeps track of L and U, which represent the lower and upper bounds of its chunks respectively. Iterations are executed starting from the lower bound, and simultaneously, L is updated to represent the new lower bound.
<code>-par-schedule-dynamic</code> or <code>/Qpar-schedule-dynamic</code>	Can be used to get a set of iterations dynamically. Assigns iterations to threads in chunks as the threads request them. The thread executes the chunk of iterations, then requests another chunk, until no chunks remain to be assigned. As each thread finishes a piece of the iteration space, it dynamically gets the next set of iterations. Each chunk

Option	Description
	contains n iterations, except for the last chunk to be assigned, which may have fewer iterations. If no n is specified, the default is 1.
<code>-par-schedule-guided</code> or <code>/Qpar-schedule-guided</code>	<p>Can be used to specify a minimum number of iterations. Assigns iterations to threads in chunks as the threads request them. The thread executes the chunk of iterations, then requests another chunk, until no chunks remain to be assigned.</p> <p>For a chunk of size 1, the size of each chunk is proportional to the number of unassigned iterations divided by the number of threads, decreasing to 1.</p> <p>For an n with value k (greater than 1), the size of each chunk is determined in the same way with the restriction that the chunks do not contain fewer than k iterations (except for the last chunk to be assigned, which may have fewer than k iterations). If no n is specified, the default is 1.</p>
<code>-par-schedule-guided-analytical</code> or <code>/Qpar-schedule-guided-analytical</code>	Divides iterations by using exponential distribution or dynamic distribution. The method depends on run-time implementation. Loop bounds are calculated with faster synchronization and chunks are dynamically dispatched

Option	Description
-par-schedule-runtime or /Qpar-schedule-runtime	<p>at run time by threads in the team.</p> <p>Defers the scheduling decision until run time. The scheduling algorithm and chunk size are then taken from the setting of environment variable OMP_SCHEDULE.</p>

Alternate Options

None

par-threshold, Qpar-threshold

Sets a threshold for the auto-parallelization of loops.

IDE EquivalentWindows: **Optimization > Threshold For Auto-Parallelization**

Linux: None

Mac OS X: **Optimization > Threshold For Auto-Parallelization****Architectures**

IA-32, Intel® 64, IA-64 architectures

SyntaxLinux and Mac OS X: `-par-threshold[n]`Windows: `/Qpar-threshold[[:]n]`**Arguments**

n Is an integer whose value is the threshold for the auto-parallelization of loops. Possible values are 0 through 100.

If n is 0, loops get auto-parallelized always, regardless of computation work volume.

If n is 100, loops get auto-parallelized when performance gains are predicted based on the compiler analysis data. Loops get auto-parallelized only if profitable parallel execution is almost certain.

The intermediate 1 to 99 values represent the percentage probability for profitable speed-up. For example, $n=50$ directs the compiler to parallelize only if there is a 50% probability of the code speeding up if executed in parallel.

Default

```
-par-
threshold100
or/Qpar-
threshold100
```

Loops get auto-parallelized only if profitable parallel execution is almost certain. This is also the default if you do not specify n .

Description

This option sets a threshold for the auto-parallelization of loops based on the probability of profitable execution of the loop in parallel. To use this option, you must also specify `-parallel` (Linux and Mac OS X) or `/Qparallel` (Windows). This option is useful for loops whose computation work volume cannot be determined at compile-time. The threshold is usually relevant when the loop trip count is unknown at compile-time.

The compiler applies a heuristic that tries to balance the overhead of creating multiple threads versus the amount of work available to be shared amongst the threads.

Alternate Options

None

parallel, Qparallel

Tells the auto-parallelizer to generate multithreaded code for loops that can be safely executed in parallel.

IDE Equivalent

Windows: **Optimization > Parallelization**

Linux: None

Mac OS X: **Optimization > Parallelization**

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-parallel`

Windows: `/Qparallel`

Arguments

None

Default

OFF Multithreaded code is not generated for loops that can be safely executed in parallel.

Description

This option tells the auto-parallelizer to generate multithreaded code for loops that can be safely executed in parallel.

To use this option, you must also specify option `O2` or `O3`.



On Mac OS X systems, when you enable automatic parallelization, you must also set the `DYLD_LIBRARY_PATH` environment variable within Xcode or an error will be displayed.

Alternate Options

None

See Also

[o](#) compiler option

pc, Qpc

Enables control of floating-point significand precision.

IDE Equivalent

None

Architectures

IA-32, Intel® 64 architectures

Syntax

Linux and Mac OS X: `-pcn`

Windows: `/Qpcn`

Arguments

n Is the floating-point significand precision. Possible values are:

32 Rounds the significand to 24 bits (single precision).

64 Rounds the significand to 53 bits (double precision).

80 Rounds the significand to 64

bits (extended precision).

Default

`-pc80` On Linux* and Mac OS* X systems, the floating-point or `/Qpc64` significand is rounded to 64 bits. On Windows* systems, the floating-point significand is rounded to 53 bits.

Description

This option enables control of floating-point significand precision.

Some floating-point algorithms are sensitive to the accuracy of the significand, or fractional part of the floating-point value. For example, iterative operations like division and finding the square root can run faster if you lower the precision with the this option.

Note that a change of the default precision control or rounding mode, for example, by using the `-pc32` (Linux and Mac OS X) or `/Qpc32` (Windows) option or by user intervention, may affect the results returned by some of the mathematical functions.

Alternate Options

None

See Also

Floating-point Operations: Floating-point Options Quick Reference

mp1, Qprec

Improves floating-point precision and consistency.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-mp1`

Windows: `/Qprec`

Arguments

None

Default

OFF The compiler provides good accuracy and run-time performance at the expense of less consistent floating-point results.

Description

This option improves floating-point consistency. It ensures the out-of-range check of operands of transcendental functions and improves the accuracy of floating-point compares.

This option prevents the compiler from performing optimizations that change NaN comparison semantics and causes all values to be truncated to declared precision before they are used in comparisons. It also causes the compiler to use library routines that give better precision results compared to the X87 transcendental instructions.

This option disables fewer optimizations and has less impact on performance than option `fltconsistency` or `mp`.

Alternate Options

None

See Also

[fltconsistency](#) compiler option

[mp](#) compiler option

prec-div, Qprec-div

Improves precision of floating-point divides.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-prec-div`
`-no-prec-div`

Windows: `/Qprec-div`
`/Qprec-div-`

Arguments

None

Default

`-prec-div` The compiler uses this method for floating-point divides.
or `/Qprec-div`

Description

This option improves precision of floating-point divides. It has a slight impact on speed.

With some optimizations, such as `-xSSE2` (Linux) or `/QxSSE2` (Windows), the compiler may change floating-point division computations into multiplication by the reciprocal of the denominator. For example, A/B is computed as $A * (1/B)$ to improve the speed of the computation.

However, sometimes the value produced by this transformation is not as accurate as full IEEE division. When it is important to have fully precise IEEE division, use this option to disable the floating-point division-to-multiplication optimization. The result is more accurate, with some loss of performance. If you specify `-no-prec-div` (Linux and Mac OS X) or `/Qprec-div-` (Windows), it enables optimizations that give slightly less precise results than full IEEE division.

Alternate Options

None

See Also

Floating-point Operations: Floating-point Options Quick Reference

prec-sqrt, Qprec-sqrt

Improves precision of square root implementations.

IDE Equivalent

None

Architectures

IA-32, Intel® 64 architectures

Syntax

Linux and Mac OS X: `-prec-sqrt`
`-no-prec-sqrt`

Windows: `/Qprec-sqrt`
`/Qprec-sqrt-`

Arguments

None

Default

`-no-prec-sqrt` The compiler uses a faster but less precise implementation of square root.

or `/Qprec-sqrt` Note that the default is `-prec-sqrt` or `/Qprec-sqrt` if any of the following options are specified: `/Od`, `/Op`, or `/Qprec` on Windows systems; `-O0`, `-mp` (or `-fltconsistency`), or `-mp1` on Linux and Mac OS X systems.

Description

This option improves precision of square root implementations. It has a slight impact on speed.

This option inhibits any optimizations that can adversely affect the precision of a square root computation. The result is fully precise square root implementations, with some loss of performance.

Alternate Options

None

prof-data-order, Qprof-data-order

Enables or disables data ordering if profiling information is enabled.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux: `-prof-data-order`
`-no-prof-data-order`

Mac OS X: None
 Windows: /Qprof-data-order
 /Qprof-data-order-

Arguments

None

Default

-no-prof- Data ordering is disabled.
 data-order
 or/Qprof-
 data-
 order-

Description

This option enables or disables data ordering if profiling information is enabled. It controls the use of profiling information to order static program data items.

For this option to be effective, you must do the following:

- For instrumentation compilation, you must specify `-prof-gen=globdata` (Linux) or `/Qprof-gen:globdata` (Windows).
- For feedback compilation, you must specify `-prof-use` (Linux) or `/Qprof-use` (Windows). You must not use multi-file optimization by specifying options such as option `-ipo` (Linux) or `/Qipo` (Windows), or option `-ipo-c` (Linux) or `/Qipo-c` (Windows).

Alternate Options

None

See Also

[prof-gen, Qprof-gen](#) compiler option
[prof-use, Qprof-use](#) compiler option

[prof-func-order, Qprof-func-order](#) compiler option

prof-dir, Qprof-dir

Specifies a directory for profiling information output files.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-prof-dir dir`

Windows: `/Qprof-dir dir`

Arguments

dir Is the name of the directory.

Default

OFF Profiling output files are placed in the directory where the program is compiled.

Description

This option specifies a directory for profiling information output files (*.dyn and *.dpi). The specified directory must already exist.

You should specify this option using the same directory name for both instrumentation and feedback compilations. If you move the .dyn files, you need to specify the new path.

Alternate Options

None

See Also

Floating-point Operations:

Profile-guided Optimization (PGO) Quick Reference

Coding Guidelines for Intel(R) Architectures

prof-file, Qprof-file

Specifies an alternate file name for the profiling summary files.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-prof-file file`

Windows: `/Qprof-file file`

Arguments

file Is the name of the profiling summary file.

Default

OFF The profiling summary files have the file name pgopti.*

Description

This option specifies an alternate file name for the profiling summary files. The *file* is used as the base name for files created by different profiling passes.

If you add this option to profmerge, the .dpi file will be named *file.dpi* instead of pgopti.dpi.

If you specify `-prof-genx` (Linux and Mac OS X) or `/Qprof-genx` (Windows) with this option, the `.spi` and `.spl` files will be named *file.spi* and *file.spl* instead of `pgopti.spi` and `pgopti.spl`.

If you specify `-prof-use` (Linux and Mac OS X) or `/Qprof-use` (Windows) with this option, the `.dpi` file will be named *file.dpi* instead of `pgopti.dpi`.

Alternate Options

None

See Also

[prof-gen, Qprof-gen](#) compiler option

[prof-use, Qprof-use](#) compiler option

Optimizing Applications:

Profile-guided Optimizations Overview

Coding Guidelines for Intel(R) Architectures

Profile an Application

prof-func-order, Qprof-func-order

Enables or disables function ordering if profiling information is enabled.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux: `-prof-func-order`
 `-no-prof-func-order`

Mac OS X: `None`

Windows: `/Qprof-func-order`
 `/Qprof-func-order-`

Arguments

None

Default`-no-prof-` Function ordering is disabled.`func-order`or `/Qprof-``func-``order-`**Description**

This option enables or disables function ordering if profiling information is enabled.

For this option to be effective, you must do the following:

- For instrumentation compilation, you must specify `-prof-gen=srcpos` (Linux) or `/Qprof-gen:srcpos` (Windows).
- For feedback compilation, you must specify `-prof-use` (Linux) or `/Qprof-use` (Windows). You must not use multi-file optimization by specifying options such as option `-ipo` (Linux) or `/Qipo` (Windows), or option `-ipo-c` (Linux) or `/Qipo-c` (Windows).

If you enable profiling information by specifying option `-prof-use` (Linux) or `/Qprof-use` (Windows), `-prof-func-groups` (Linux) and `/Qprof-func-groups` (Windows) are set and function grouping is enabled. However, if you explicitly enable `-prof-func-order` (Linux) or `/Qprof-func-order` (Windows), function ordering is performed instead of function grouping.

On Linux* systems, this option is only available for Linux linker 2.15.94.0.1, or later.

To set the hotness threshold for function grouping and function ordering, use option `-prof-hotness-threshold` (Linux) or `/Qprof-hotness-threshold` (Windows).

Alternate Options

None

The following example shows how to use this option on a Windows system:

```
ifort /Qprof-gen:globdata file1.f90 file2.f90 /exe:instrumented.exe  
      ./instrumented.exe  
ifort /Qprof-use /Qprof-func-order file1.f90 file2.f90  
/exe:feedback.exe
```

The following example shows how to use this option on a Linux system:

```
ifort -prof-gen:globdata file1.f90 file2.f90 -o instrumented  
      ./instrumented.exe  
ifort -prof-use -prof-func-order file1.f90 file2.f90 -o feedback
```

See Also

[prof-hotness-threshold, Qprof-hotness-threshold](#) compiler option

[prof-gen, Qprof-gen](#) compiler option

[prof-use, Qprof-use](#) compiler option

[prof-data-order, Qprof-data-order](#) compiler option

[prof-func-groups](#) compiler option

prof-gen, Qprof-gen

Produces an instrumented object file that can be used in profile-guided optimization.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-prof-gen[=keyword]`

`-no-prof-gen`

Windows: `/Qprof-gen[:keyword]`

`/Qprof-gen-`

Arguments

keyword

Specifies details for the instrumented file. Possible values are:

default	Produces an instrumented object file. This is the same as specifying <code>-prof-gen</code> (Linux* and Mac OS* X) or <code>/Qprof-gen</code> (Windows*) with no keyword.
srcpos	Produces an instrumented object file that includes extra source position information. This option is the same as option <code>-prof-genx</code> (Linux* and Mac OS* X) or <code>/Qprof-genx</code> (Windows*), which are deprecated .
globdata	Produces an instrumented object file that includes information for global data layout.

Default

`-no-` Profile generation is disabled.
`prof-`
`gen` or
`/Qprof-`
`gen-`

Description

This option produces an instrumented object file that can be used in profile-guided optimization. It gets the execution count of each basic block.

If you specify keyword `srcpos` or `globdata`, a static profile information file (.spi) is created. These settings may increase the time needed to do a parallel build using `-prof-gen`, because of contention writing the .spi file.

These options are used in phase 1 of the Profile Guided Optimizer (PGO) to instruct the compiler to produce instrumented code in your object files in preparation for instrumented execution.

Alternate Options

None

See Also

Optimizing Applications:

Basic PGO Options

Example of Profile-Guided Optimization

prof-genx, Qprof-genx

This is a deprecated option. See [prof-gen](#) keyword `srcpos`.

prof-hotness-threshold, Qprof-hotness-threshold

Lets you set the hotness threshold for function grouping and function ordering.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux: `-prof-hotness-threshold=n`

Mac OS X: None
 Windows: /Qprof-hotness-threshold:*n*

Arguments

n Is the hotness threshold. *n* is a percentage having a value between 0 and 100 inclusive. If you specify 0, there will be no hotness threshold setting in effect for function grouping and function ordering.

Default

OFF The compiler's default hotness threshold setting of 10 percent is in effect for function grouping and function ordering.

Description

This option lets you set the hotness threshold for function grouping and function ordering.

The "hotness threshold" is the percentage of functions in the application that should be placed in the application's hot region. The hot region is the most frequently executed part of the application. By grouping these functions together into one hot region, they have a greater probability of remaining resident in the instruction cache. This can enhance the application's performance.

For this option to take effect, you must specify option `-prof-use` (Linux) or `/Qprof-use` (Windows) and one of the following:

- On Linux systems: `-prof-func-groups` or `-prof-func-order`
- On Windows systems: `/Qprof-func-order`

Alternate Options

None

See Also

[prof-use, Qprof-use](#) compiler option

[prof-func-groups](#) compiler option

[prof-func-order, Qprof-func-order](#) compiler option

prof-src-dir, Qprof-src-dir

Determines whether directory information of the source file under compilation is considered when looking up profile data records.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-prof-src-dir`
`-no-prof-src-dir`

Windows: `/Qprof-src-dir`
`/Qprof-src-dir-`

Arguments

None

Default

`-prof-src-dir` Directory information is used when looking up profile data records in the .dpi file.
or `/Qprof-src-dir`

Description

This option determines whether directory information of the source file under compilation is considered when looking up profile data records in the .dpi file. To

use this option, you must also specify option `-prof-use` (Linux and Mac OS X) or `/Qprof-use` (Windows).

If the option is enabled, directory information is considered when looking up the profile data records within the `.dpi` file. You can specify directory information by using one of the following options:

- Linux and Mac OS X: `-prof-src-root` or `-prof-src-root-cwd`
- Windows: `/Qprof-src-root` or `/Qprof-src-root-cwd`

If the option is disabled, directory information is ignored and only the name of the file is used to find the profile data record.

Note that options `-prof-src-dir` (Linux and Mac OS X) and `/Qprof-src-dir` (Windows) control how the names of the user's source files get represented within the `.dyn` or `.dpi` files. Options `-prof-dir` (Linux and Mac OS X) and `/Qprof-dir` (Windows) specify the location of the `.dyn` or the `.dpi` files.

Alternate Options

None

See Also

[prof-use](#), [Qprof-use](#) compiler option

[prof-src-root](#), [Qprof-src-root](#) compiler option

[prof-src-root-cwd](#), [Qprof-src-root-cwd](#) compiler option

prof-src-root, Qprof-src-root

Lets you use relative directory paths when looking up profile data and specifies a directory as the base.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-prof-src-root=dir`

Windows: `/Qprof-src-root:dir`

Arguments

dir Is the base for the relative paths.

Default

OFF The setting of relevant options determines the path used when looking up profile data records.

Description

This option lets you use relative directory paths when looking up profile data in .dpi files. It lets you specify a directory as the base. The paths are relative to a base directory specified during the `-prof-gen` (Linux and Mac OS X) or `/Qprof-gen` (Windows) compilation phase.

This option is available during the following phases of compilation:

- Linux and Mac OS X: `-prof-gen` and `-prof-use` phases
- Windows: `/Qprof-gen` and `/Qprof-use` phases

When this option is specified during the `-prof-gen` or `/Qprof-gen` phase, it stores information into the .dyn or .dpi file. Then, when .dyn files are merged together or the .dpi file is loaded, only the directory information below the root directory is used for forming the lookup key.

When this option is specified during the `-prof-use` or `/Qprof-use` phase, it specifies a root directory that replaces the root directory specified at the `-prof-gen` or `/Qprof-gen` phase for forming the lookup keys.

To be effective, this option or option `-prof-src-root-cwd` (Linux and Mac OS X) or `/Qprof-src-root-cwd` (Windows) must be specified during the `-prof-gen` or `/Qprof-gen` phase. In addition, if one of these options is not specified, absolute paths are used in the .dpi file.

Alternate Options

None

Consider the initial `-prof-gen` compilation of the source file

`c:\user1\feature_foo\myproject\common\glob.f90`:

```
ifort -prof-gen -prof-src-root=c:\user1\feature_foo\myproject -c
common\glob.f90
```

For the `-prof-use` phase, the file `glob.f90` could be moved into the directory

`c:\user2\feature_bar\myproject\common\glob.f90` and profile information would be

found from the `.dpi` when using the following:

```
ifort -prof-use -prof-src-root=c:\user2\feature_bar\myproject -c
common\glob.f90
```

If you do not use option `-prof-src-root` during the `-prof-gen` phase, by default, the `-prof-use` compilation can only find the profile data if the file is compiled in the `c:\user1\feature_foo\my_project\common` directory.

See Also

[prof-gen, Qprof-gen](#) compiler option

[prof-use, Qprof-use](#) compiler option

[prof-src-dir, Qprof-src-dir](#) compiler option

[prof-src-root-cwd, Qprof-src-root-cwd](#) compiler option

prof-src-root-cwd, Qprof-src-root-cwd

Lets you use relative directory paths when looking up profile data and specifies the current working directory as the base.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-prof-src-root-cwd`

Windows: `/Qprof-src-root-cwd`

Arguments

None

Default

OFF The setting of relevant options determines the path used when looking up profile data records.

Description

This option lets you use relative directory paths when looking up profile data in .dpi files. It specifies the current working directory as the base. To use this option, you must also specify option `-prof-use` (Linux and Mac OS) or `/Qprof-use` (Windows).

This option is available during the following phases of compilation:

- Linux and Mac OS X: `-prof-gen` and `-prof-use` phases
- Windows: `/Qprof-gen` and `/Qprof-use` phases

When this option is specified during the `-prof-gen` or `/Qprof-gen` phase, it stores information into the .dyn or .dpi file. Then, when .dyn files are merged together or the .dpi file is loaded, only the directory information below the root directory is used for forming the lookup key.

When this option is specified during the `-prof-use` or `/Qprof-use` phase, it specifies a root directory that replaces the root directory specified at the `-prof-gen` or `/Qprof-gen` phase for forming the lookup keys.

To be effective, this option or option `-prof-src-root` (Linux and Mac OS X) or `/Qprof-src-root` (Windows) must be specified during the `-prof-gen` or `/Qprof-gen` phase. In addition, if one of these options is not specified, absolute paths are used in the .dpi file.

Alternate Options

None

See Also

[prof-gen](#), [Qprof-gen](#) compiler option

[prof-use](#), [Qprof-use](#) compiler option

[prof-src-dir](#), [Qprof-src-dir](#) compiler option

[prof-src-root](#), [Qprof-src-root](#) compiler option

prof-use, Qprof-use

Enables the use of profiling information during optimization.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-prof-use[=arg]`
`-no-prof-use`

Windows: `/Qprof-use[:arg]`
`/Qprof-use-`

Arguments

arg Specifies additional instructions. Possible values are:

`weighted` Tells the profmerge utility to apply a weighting to the .dyn file values when creating the .dpi file to normalize the data counts when the training runs have different execution durations. This argument only

has an effect when the compiler invokes the `profmerge` utility to create the `.dpi` file. This argument does not have an effect if the `.dpi` file was previously created without weighting.

`[no]merge` Enables or disables automatic invocation of the `profmerge` utility. The default is `merge`. Note that you cannot specify both `weighted` and `nomerge`. If you try to specify both values, a warning will be displayed and `nomerge` takes precedence.

`default` Enables the use of profiling information during optimization. The `profmerge` utility is invoked by default. This value is the same as specifying `-prof-use` (Linux and Mac OS X) or `/Qprof-use` (Windows) with no argument.

Default

`-no-prof-use` or `Profiling information is not used during optimization.`

`/Qprof-`
`use-`

Description

This option enables the use of profiling information (including function splitting and function grouping) during optimization. It enables option `-fnsplit` (Linux) or `/Qfnsplit` (Windows).

This option instructs the compiler to produce a profile-optimized executable and it merges available profiling output files into a `pgopti.dpi` file.

Note that there is no way to turn off function grouping if you enable it using this option.

To set the hotness threshold for function grouping and function ordering, use option `-prof-hotness-threshold` (Linux) or `/Qprof-hotness-threshold` (Windows).

Alternate Options

None

See Also

[prof-hotness-threshold, Qprof-hotness-threshold](#) compiler option

Optimizing Applications:

Basic PGO Options

Example of Profile-Guided Optimization

rcd, Qrcd

Enables fast float-to-integer conversions.

IDE Equivalent

None

Architectures

IA-32, Intel® 64 architectures

Syntax

Linux and Mac OS X: `-rct`

Windows: `/Qrct`

Arguments

None

Default

OFF Floating-point values are truncated when a conversion to an integer is involved. On Windows, this is the same as specifying `/QIfist-`.

Description

This option enables fast float-to-integer conversions. It can improve the performance of code that requires floating-point-to-integer conversions. The system default floating-point rounding mode is round-to-nearest. However, the Fortran language requires floating-point values to be truncated when a conversion to an integer is involved. To do this, the compiler must change the rounding mode to truncation before each floating-point-to-integer conversion and change it back afterwards.

This option disables the change to truncation of the rounding mode for all floating-point calculations, including floating point-to-integer conversions. This option can improve performance, but floating-point conversions to integer will not conform to Fortran semantics.

Alternate Options

Linux and Mac OS X: None

Windows: `/QIfist`

rct, Qrct

Sets the internal FPU rounding control to Truncate.

IDE Equivalent

None

Architectures

IA-32, Intel® 64 architectures

Syntax

Linux and Mac OS X: `-rct`

Windows: `/Qrct`

Arguments

None

Default

OFF The compiler uses the default setting for the FPU rounding control.

Description

This option sets the internal FPU rounding control to Truncate.

Alternate Options

Linux and Mac OS X: None

Windows: `/rounding-mode:chopped`

safe-cray-ptr, Qsafe-cray-ptr

Tells the compiler that Cray* pointers do not alias other variables.

IDE Equivalent

Windows: **Data > Assume Cray Pointers Do Not Share Memory Locations**

Linux: None

Mac OS X: **Data > Assume Cray Pointers Do Not Share Memory Locations**

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-safe-cray-ptr`

Windows: `/Qsafe-cray-ptr`

Arguments

None

Default

OFF The compiler assumes that Cray pointers alias other variables.

Description

This option tells the compiler that Cray pointers do not alias (that is, do not specify sharing memory with) other variables.

Alternate Options

None

Example

Consider the following:

```
pointer (pb, b)
pb = getstorage()
do i = 1, n
b(i) = a(i) + 1
enddo
```

By default, the compiler assumes that `b` and `a` are aliased. To prevent such an assumption, specify the `-safe-cray-ptr` (Linux and Mac OS X) or `/Qsafe-cray-ptr` (Windows) option, and the compiler will treat `b(i)` and `a(i)` as independent of each other.

However, if the variables are intended to be aliased with Cray pointers, using the option produces incorrect results. In the following example, you should not use the option:

```
pointer (pb, b)
pb = loc(a(2))
do i=1, n
b(i) = a(i) +1
enddo
```

save, Qsave

Causes variables to be placed in static memory.

IDE Equivalent

Windows: **Data > Local Variable Storage**

Linux: None

Mac OS X: **Data > Local Variable Storage**

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-save`

Windows: `/Qsave`

Arguments

None

Default

`-auto-` Scalar variables of intrinsic types INTEGER, REAL,
`scalar` COMPLEX, and LOGICAL are allocated to the run-time stack.
 or Note that if option `recursive`, `-openmp` (Linux and Mac OS
`/Qauto-` X), or `/Qopenmp` (Windows) is specified, the default is -
`scalar` `automatic` (Linux) or `/Qauto` (Windows).

Description

This option saves all variables in static allocation except local variables within a recursive routine and variables declared as AUTOMATIC.

If you want all local, non-SAVEd variables to be allocated to the run-time stack, specify option `automatic`.

Alternate Options

Linux and Mac OS X: `-noautomatic`, `-noauto`

Windows: `/noautomatic`, `/noauto`, `/4Na`

See Also

[automatic](#) compiler option

[auto_scalar](#) compiler option

save-temps, Qsave-temps

Tells the compiler to save intermediate files created during compilation.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-save-temps`

`-no-save-temps`

Windows: `/Qsave-temps`

`/Qsave-temps-`

Arguments

None

Default

Linux and Mac OS X: `-no-save-temps` On Linux and Mac OS X systems, the compiler deletes intermediate files after compilation is completed.

Windows: .obj files are saved On Windows systems, the compiler saves only intermediate object files after compilation is completed.

Description

This option tells the compiler to save intermediate files created during compilation. The names of the files saved are based on the name of the source file; the files are saved in the current working directory.

If `-save-temps` or `/Qsave-temps` is specified, the following occurs:

- The object .o file (Linux and Mac OS X) or .obj file (Windows) is saved.
- The assembler .s file (Linux and Mac OS X) or .asm file (Windows) is saved if you specified `-use-asm` (Linux or Mac OS X) or `/Quse-asm` (Windows).
- The .i or .i90 file is saved if the fpp preprocessor is invoked.

If `-no-save-temps` is specified on Linux or Mac OS X systems, the following occurs:

- The .o file is put into `/tmp` and deleted after calling `ld`.
- The preprocessed file is not saved after it has been used by the compiler.

If `/Qsave-temps-` is specified on Windows systems, the following occurs:

- The .obj file is not saved after the linker step.
- The preprocessed file is not saved after it has been used by the compiler.

**Note**

This option only saves intermediate files that are normally created during compilation.

Alternate Options

None

Example

If you compile program *my_foo.F* on a Linux or Mac OS X system and you specify option `-save-temps` and option `-use-asm`, the compilation will produce files *my_foo.o*, *my_foo.s*, and *my_foo.i*.

If you compile program *my_foo.fpp* on a Windows system and you specify option `/Qsave-temps` and option `/Quse-asm`, the compilation will produce files *my_foo.obj*, *my_foo.asm*, and *my_foo.i*.

scalar-rep, Qscalar-rep

Enables scalar replacement performed during loop transformation.

IDE Equivalent

None

Architectures

IA-32 architecture

Syntax

Linux and Mac OS X: `-scalar-rep`
`-no-scalar-rep`

Windows: `/Qscalar-rep`
`/Qscalar-rep-`

Arguments

None

Default

`-no-scalar-rep` Scalar replacement is not performed during loop transformation.

`or/Qscalar-rep-`

Description

This option enables scalar replacement performed during loop transformation. To use this option, you must also specify `Q3`.

Alternate Options

None

See Also

[Q compiler option](#)

Qsalign

Specifies stack alignment for functions.

IDE Equivalent

None

Architectures

IA-32 architecture

Syntax

Linux and Mac OS X: None

Windows: `/Qsalign[n]`

Arguments

n Is the byte size of aligned variables. Possible values are:

8	Specifies that alignment should occur for functions with 8-byte aligned variables. At this setting the compiler aligns the stack to 16 bytes if there is any
---	--

	16-byte or 8-byte data on the stack. For 8-byte data, the compiler only aligns the stack if the alignment will produce a performance advantage.
16	Specifies that alignment should occur for functions with 16-byte aligned variables. At this setting, the compiler only aligns the stack for 16-byte data. No attempt is made to align for 8-byte data.

Default

`/Qsfalign8` Alignment occurs for functions with 8-byte aligned variables.

Description

This option specifies stack alignment for functions. It lets you disable the normal optimization that aligns a stack for 8-byte data.

If you do not specify *n*, stack alignment occurs for all functions. If you specify `/Qsfalign-`, no stack alignment occurs for any function.

Alternate Options

None

sox, Qsox

Tells the compiler to save the compiler options and version number in the executable.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-sox`

`-no-sox`

Windows: `/Qsox`

`/Qsox-`

Arguments

None

Default

`-no-sox` The compiler does not save the compiler options and
or `/Qsox-` version number in the executable.

Description

This option tells the compiler to save the compiler options and version number in the executable. The size of the executable on disk is increased slightly by the inclusion of these information strings.

This option forces the compiler to embed in each object file or assembly output a string that contains information about the compiler version and compilation options for each source file that has been compiled. When you link the object files into an executable file, the linker places each of the information strings into the header of the executable. It is then possible to use a tool, such as a strings utility, to determine what options were used to build the executable file.

If `-no-sox` or `/Qsox-` is specified, this extra information is not put into the object or assembly output generated by the compiler.

Alternate Options

None

tcheck, Qtcheck

Enables analysis of threaded applications.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux: `-tcheck`

Mac OS X: `None`

Windows: `/Qtcheck`

Arguments

None

Default

OFF Threaded applications are not instrumented by the compiler for analysis by Intel® Thread Checker.

Description

This option enables analysis of threaded applications.

To use this option, you must have Intel® Thread Checker installed, which is one of the Intel® Threading Analysis Tools. If you do not have this tool installed, the compilation will fail. Remove the `-tcheck` (Linux) or `/Qtcheck` (Windows) option from the command line and recompile.

For more information about Intel® Thread Checker (including an evaluation copy), open the page associated with threading tools at [Intel® Software Development Products](#).

Alternate Options

None

tcollect, Qtcollect

Inserts instrumentation probes calling the Intel® Trace Collector API.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux: `-tcollect[lib]`

Mac OS X: `None`

Windows: `/Qtcollect[:lib]`

Arguments

lib Is one of the Intel® Trace Collector libraries; for example, VT, VTcs, VTmc, or VTfs. If you do not specify *lib*, the default library is VT.

Default

OFF Instrumentation probes are not inserted into compiled applications.

Description

This option inserts instrumentation probes calling the Intel® Trace Collector API. To use this option, you must have the Intel® Trace Collector installed and set up through one of its set-up scripts. This tool is a component of the Intel® Trace Analyzer and Collector.

This option provides a flexible and convenient way of instrumenting functions of a compiled application. For every function, the entry and exit points are instrumented at compile time to let the Intel® Trace Collector record functions beyond the default MPI calls. For non-MPI applications (for example, threaded or serial), you must ensure that the Intel® Trace Collector is properly initialized (VT_initialize/VT_init).



Be careful with full instrumentation because this feature can produce very large trace files.

For more details, see the Intel® Trace Collector User Guide.

Alternate Options

None

See Also

[tcollect-filter, Qtcollect-filter](#) compiler option

tcollect-filter, Qtcollect-filter

Lets you enable or disable the instrumentation of specified functions.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux: `-tcollect-filter file`

Mac OS X: `None`

Windows: `/Qtcollect-filter:file`

Arguments

file Is a configuration file that lists filters, one per line. Each filter consists of a regular expression string and a switch. Strings with leading or trailing white spaces must be quoted. Other strings do not have to be quoted. The switch value can be ON, on, OFF, or off.

Default

OFF Functions are not instrumented. However, if option `-tcollect` (Linux) or `/Qtcollect` (Windows) is specified, the filter setting is `.* ON` and all functions get instrumented.

Description

This option lets you enable or disable the instrumentation of specified functions. During instrumentation, the regular expressions in the file are matched against the function names. The switch specifies whether matching functions are to be instrumented or not. Multiple filters are evaluated from top to bottom with increasing precedence.

The names of the functions to match against are formatted as follows:

- The source file name is followed by a colon-separated function name. Source file names should contain the full path, if available. For example:

```
/home/joe/src/file.f:FOO_bar
```

- Classes and function names are separated by double colons. For example:

```
/home/joe/src/file.fpp:app::foo::bar
```

You can use option `-opt-report` (Linux) or `/Qopt-report` (Windows) to get a full list of file and function names that the compiler recognizes from the compilation unit. This list can be used as the basis for filtering in the configuration file.

To use this option, you must have the Intel® Trace Collector installed and set up through one of its set-up scripts. This tool is a component of the Intel® Trace Analyzer and Collector.

For more details, see the Intel® Trace Collector User Guide.

Alternate Options

None

Consider the following filters in a configuration file:

```
'.*' OFF '.*vector.*' ON
```

The above will cause instrumentation of only those functions having the string 'vector' in their names. No other function will be instrumented. Note that reversing the order of the two lines will prevent instrumentation of all functions.

To get a list of the file or routine strings that can be matched by the regular expression filters, generate an optimization report with tcollect information. For example:

```
Windows OS: ifort /Qtcollect /Qopt-report /Qopt-report-phase tcollect  
Linux OS: ifort -tcollect -opt-report -opt-report-phase tcollect
```

See Also

[tcollect, Qtcollect](#) compiler option

tprofile, Qtprofile

Generates instrumentation to analyze multi-threading performance.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux: -tprofile

Mac OS X: None

Windows: `/Qtprofile`

Arguments

None

Default

OFF Instrumentation is not generated by the compiler for analysis by Intel® Thread Profiler.

Description

This option generates instrumentation to analyze multi-threading performance. To use this option, you must have Intel® Thread Profiler installed, which is one of the Intel® Threading Analysis Tools. If you do not have this tool installed, the compilation will fail. Remove the `-tprofile` (Linux) or `/Qtprofile` (Windows) option from the command line and recompile.

For more information about Intel® Thread Profiler (including an evaluation copy), open the page associated with threading tools at [Intel® Software Development Products](#).

Alternate Options

None

ftrapuv, Qtrapuv

Initializes stack local variables to an unusual value to aid error detection.

IDE Equivalent

Windows: **Data > Initialize stack variables to an unusual value**

Linux: None

Mac OS X: **Run-Time > Initialize Stack Variables to an Unusual Value**

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-ftrapuv`

Windows: `/Qtrapuv`

Arguments

None

Default

OFF The compiler does not initialize local variables.

Description

This option initializes stack local variables to an unusual value to aid error detection. Normally, these local variables should be initialized in the application. The option sets any uninitialized local variables that are allocated on the stack to a value that is typically interpreted as a very large integer or an invalid address. References to these variables are then likely to cause run-time errors that can help you detect coding errors.

This option sets option `-g` (Linux and Mac OS X) and `/zi` or `/z7` (Windows).

Alternate Options

None

See Also

[g, zi, z7](#) compiler options

unroll, Qunroll

Tells the compiler the maximum number of times to unroll loops.

IDE Equivalent

Windows: **Optimization > Loop Unroll Count**

Linux: None

Mac OS X: **Optimization > Loop Unroll Count**

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-unroll[=n]`

Windows: `/Qunroll[:n]`

Arguments

n Is the maximum number of times a loop can be unrolled. To disable loop unrolling, specify 0. On systems using IA-64 architecture, you can only specify a value of 0.

Default

`-unroll` The compiler uses default heuristics when unrolling loops.
or `/Qunroll`

Description

This option tells the compiler the maximum number of times to unroll loops. If you do not specify *n*, the optimizer determines how many times loops can be unrolled.

Alternate Options

Linux and Mac OS X: `-funroll-loops`

Windows: `/unroll`

See Also

Optimizing Applications: Loop Unrolling

unroll-aggressive, Qunroll-aggressive

Determines whether the compiler uses more aggressive unrolling for certain loops.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-unroll-aggressive`
`-no-unroll-aggressive`

Windows: `/Qunroll-aggressive`
`/Qunroll-aggressive-`

Arguments

None

Default

`-no-unroll-aggressive` The compiler uses default heuristics when
or `/Qunroll-` unrolling loops.
`aggressive-`

Description

This option determines whether the compiler uses more aggressive unrolling for certain loops. The positive form of the option may improve performance.

On IA-32 architecture and Intel® 64 architecture, this option enables aggressive, complete unrolling for loops with small constant trip counts.

On IA-64 architecture, this option enables additional complete unrolling for loops that have multiple exits or outer loops that have a small constant trip count.

Alternate Options

None

uppercase, Quppercase

See [names](#).

use-asm, Quse-asm

Tells the compiler to produce objects through the assembler.

IDE Equivalent

None

Architectures

`-use-asm`: IA-32 architecture, Intel® 64 architecture, IA-64 architecture

`/Quse-asm`: IA-64 architecture

Syntax

Linux and Mac OS X: `-use-asm`

`-no-use-asm`

Windows: `/Quse-asm`

`/Quse-asm-`

Arguments

None

Default

`-no-use-asm` The compiler produces objects directly.

`asm`

or `/Quse-`

`asm-`

Description

This option tells the compiler to produce objects through the assembler.

Alternate Options

None

Quse-msasm-symbols

Tells the compiler to use a dollar sign ("\$\$") when producing symbol names.

IDE Equivalent

None

Architectures

IA-32, Intel® 64 architectures

Syntax

Linux and Mac OS X: None

Windows: `/Quse-msasm-symbols`

Arguments

None

Default

OFF The compiler uses a period (".") when producing symbol names

Description

This option tells the compiler to use a dollar sign ("\$\$") when producing symbol names.

Use this option if you require symbols in your .asm files to contain characters that are accepted by the MS assembler.

Alternate Options

None

Quse-vcdebug

Tells the compiler to issue debug information compatible with the Visual C++ debugger.

IDE Equivalent

None

Architectures

IA-32 architecture

Syntax

Linux and Mac OS X: None

Windows: /Quse-vcdebug

Arguments

None

Default

OFF Debug information is issued that is compatible with Fortran debuggers.

Description

This option tells the compiler to issue debug information compatible with the Visual C++ debugger. It prevents the compiler from issuing the extended information used by Fortran debuggers.

Alternate Options

None

Qvc

Specifies compatibility with Microsoft* Visual C++ or Microsoft* Visual Studio.

IDE Equivalent

None

Architectures

IA-32, Intel® 64 architectures

Syntax

Linux and Mac OS X: None

Windows: /*Qvc7.1*
 /*Qvc8*
 /*Qvc9*

Arguments

None

Default

varies When the compiler is installed, it detects which version of Visual Studio is on your system. *Qvc* defaults to the form of the option that is compatible with that version. When multiple versions of Visual Studio are installed, the compiler installation lets you select which version you want to use. In this case, *Qvc* defaults to the version you choose.

Description

This option specifies compatibility with Visual C++ or Visual Studio.

Option	Description
<i>/Qvc7.1</i>	Specifies compatibility with Microsoft* Visual Studio .NET 2003.
<i>/Qvc8</i>	Specifies compatibility with Microsoft* Visual Studio 2005.

Option	Description
/Qvc9	Specifies compatibility with Microsoft* Visual Studio 2008.

Alternate Options

None

vec, Qvec

Enables or disables vectorization and transformations enabled for vectorization.

IDE Equivalent

None

Architectures

IA-32, Intel® 64 architectures

Syntax

Linux and Mac OS X: `-vec`
`-no-vec`

Windows: `/Qvec`
`/Qvec-`

Arguments

None

Default

`-vec` Vectorization is enabled.
or `/Qvec`

Description

This option enables or disables vectorization and transformations enabled for vectorization.

To disable vectorization and transformations enabled for vectorization, specify `-no-vec` (Linux and Mac OS X) or `/Qvec-` (Windows).

Alternate Options

None

vec-guard-write, Qvec-guard-write

Tells the compiler to perform a conditional check in a vectorized loop.

IDE Equivalent

None

Architectures

IA-32, Intel® 64 architectures

Syntax

Linux and Mac OS X: `-vec-guard-write`
`-no-vec-guard-write`

Windows: `/Qvec-guard-write`
`/Qvec-guard-write-`

Arguments

None

Default

`-no-vec-guard-write` or `/Qvec-guard-write-` The compiler uses default heuristics when checking vectorized loops.

Description

This option tells the compiler to perform a conditional check in a vectorized loop. This checking avoids unnecessary stores and may improve performance.

Alternate Options

None

vec-report, Qvec-report

Controls the diagnostic information reported by the vectorizer.

IDE Equivalent

Windows: **Compilation Diagnostics > Vectorizer Diagnostic Level**

Linux: None

Mac OS X: **Diagnostics > Vectorizer Diagnostic Report**

Architectures

IA-32, Intel® 64 architectures

Syntax

Linux and Mac OS X: `-vec-report[n]`

Windows: `/Qvec-report[n]`

Arguments

<i>n</i>	Is a value denoting which diagnostic messages to report. Possible values are:
0	Tells the vectorizer to report no diagnostic information.
1	Tells the vectorizer to report on vectorized loops.
2	Tells the vectorizer to report on vectorized and non-vectorized

- loops.
- | | |
|---|---|
| 3 | Tells the vectorizer to report on vectorized and non-vectorized loops and any proven or assumed data dependences. |
| 4 | Tells the vectorizer to report on non-vectorized loops. |
| 5 | Tells the vectorizer to report on non-vectorized loops and the reason why they were not vectorized. |

Default

`-vec-report1` If the vectorizer has been enabled, it reports diagnostics on vectorized loops.
 or `/Qvec-report1`

Description

This option controls the diagnostic information reported by the vectorizer. The vectorizer report is sent to stdout.

If you do not specify *n*, it is the same as specifying `-vec-report1` (Linux and Mac OS X) or `/Qvec-report1` (Windows).

The vectorizer is enabled when certain compiler options are specified, such as option `-ax` or `-x` (Linux and Mac OS X), option `/Qax` or `/Qx` (Windows), option `-arch SSE` or `-arch SSE2` (Linux and Mac OS X), option `/architecture:SSE` or `/architecture:SSE2` (Windows).

If this option is specified from within the IDE, the report is included in the build log if the Generate Build Logs option is selected.

Alternate Options

None

x, Qx

Tells the compiler to generate optimized code specialized for the Intel processor that executes your program.

IDE Equivalent

Windows: **Code Generation > Intel Processor-Specific Optimization Optimization > Use Intel(R) Processor Extensions**

Linux: None

Mac OS X: **Code Generation > Intel Processor-Specific Optimization**

Architectures

IA-32, Intel® 64 architectures

Syntax

Linux and Mac OS X: *-xprocessor*

Windows: */Qxprocessor*

Arguments

processor

Indicates the processor for which code is generated. Many of the following descriptions refer to Intel® Streaming SIMD Extensions (Intel® SSE) and Supplemental Streaming SIMD Extensions (Intel® SSSE). Possible values are:

Host	Can generate instructions for the highest instruction set and processor available on the compilation host.
------	--

- SSE4 . 2 Can generate Intel® SSE4 Efficient Accelerated String and Text Processing instructions supported by Intel® Core™ i7 processors. Can generate Intel® SSE4 Vectorizing Compiler and Media Accelerator, Intel® SSSE3, SSE3, SSE2, and SSE instructions and it can optimize for the Intel® Core™ processor family.
- SSE4 . 1 Can generate Intel® SSE4 Vectorizing Compiler and Media Accelerator instructions for Intel processors. Can generate Intel® SSSE3, SSE3, SSE2, and SSE instructions and it can optimize for Intel® 45nm Hi-k next generation Intel® Core™ microarchitecture. This replaces value S, which is [deprecated](#).
- SSE3_ATOM Can generate MOVBE instructions for Intel processors and it can optimize for the Intel® Atom™ processor and Intel® Centrino® Atom™ Processor Technology.

- SSSE3 Can generate Intel® SSSE3, SSE3, SSE2, and SSE instructions for Intel processors and it can optimize for the Intel® Core™2 Duo processor family. This replaces value T, which is [deprecated](#).
- SSE3 Can generate Intel® SSE3, SSE2, and SSE instructions for Intel processors and it can optimize for processors based on Intel® Core™ microarchitecture and Intel NetBurst® microarchitecture. This replaces value P, which is [deprecated](#).
- SSE2 Can generate Intel® SSE2 and SSE instructions for Intel processors, and it can optimize for Intel® Pentium® 4 processors, Intel® Pentium® M processors, and Intel® Xeon® processors with Intel® SSE2. This value is not available on Mac OS X systems. This replaces value N, which is [deprecated](#).

Default

Windows systems: For more information on the default values, see
Arguments above, option `m` (Linux and Mac OS X) and
option `arch` (Windows).
`/arch:SSE2`

Linux systems: –

`msse2`

Mac OS X systems

using IA-32

architecture: `SSE3`

Mac OS X systems

using Intel® 64

architecture: `SSSE3`

Description

This option tells the compiler to generate optimized code specialized for the Intel processor that executes your program. It also enables optimizations in addition to Intel processor-specific optimizations. The specialized code generated by this option may run only on a subset of Intel processors.

This option can enable optimizations depending on the argument specified. For example, it may enable Intel® Streaming SIMD Extensions 4 (Intel® SSE4), Intel® Supplemental Streaming SIMD Extensions 3 (Intel® SSSE3), Intel® Streaming SIMD Extensions 3 (Intel® SSE3), Intel® Streaming SIMD Extensions 2 (Intel® SSE2), or Intel® Streaming SIMD Extensions (Intel® SSE) instructions. The binaries produced by these values will run on Intel processors that support all of the features for the targeted processor. For example, binaries produced with SSE3 will run on an Intel® Core™ 2 Duo processor, because that processor completely supports all of the capabilities of the Intel® Pentium® 4 processor, which the SSE3 value targets. Specifying the SSSE3 value has the potential of using more features and optimizations available to the Intel® Core™ 2 Duo processor.

Do not use *processor* values to create binaries that will execute on a processor that is not compatible with the targeted processor. The resulting program may fail with an illegal instruction exception or display other unexpected behavior. For

example, binaries produced with SSE3 may produce code that will not run on Intel® Pentium® III processors or earlier processors that do not support SSE2 instructions.

Compiling the function `main()` with any of the *processor* values produces binaries that display a fatal run-time error if they are executed on unsupported processors. For more information, see *Optimizing Applications*.

If you specify more than one *processor* value, code is generated for only the highest-performing processor specified. The highest-performing to lowest-performing *processor* values are: SSE4.2, SSE4.1, SSSE3, SSE3, SSE2. Note that *processor* value SSE3_ATOM does not fit within this group.

Compiler options `m` and `arch` produce binaries that should run on processors not made by Intel that implement the same capabilities as the corresponding Intel processors.

Previous value `O` is deprecated and has been replaced by option `-msse3` (Linux and Mac OS X) and option `/arch:SSE3` (Windows).

Previous values `W` and `K` are deprecated. The details on replacements are as follows:

- Mac OS X systems: On these systems, there is no exact replacement for `W` or `K`. You can upgrade to the default option `-msse3` (IA-32 architecture) or option `-mssse3` (Intel® 64 architecture).
- Windows and Linux systems: The replacement for `W` is `-msse2` (Linux) or `/arch:SSE2` (Windows). There is no exact replacement for `K`. However, on Windows systems, `/QxK` is interpreted as `/arch:IA32`; on Linux systems, `-xK` is interpreted as `-mia32`. You can also do one of the following:
 - Upgrade to option `-msse2` (Linux) or option `/arch:SSE2` (Windows). This will produce one code path that is specialized for Intel® SSE2. It will not run on earlier processors
 - Specify the two option combination `-mia32 -axSSE2` (Linux) or `/arch:IA32 /QaxSSE2` (Windows). This combination will produce an executable that runs on any processor with IA-32 architecture but with an additional specialized Intel® SSE2 code path.

The `-x` and `/Qx` options enable additional optimizations not enabled with option `-m` or option `/arch`.

Alternate Options

None

See Also

[ax, Qax](#) compiler option

[m](#) compiler option

[arch](#) compiler option

zero, Qzero

Initializes to zero all local scalar variables of intrinsic type INTEGER, REAL, COMPLEX, or LOGICAL that are saved but not yet initialized.

IDE Equivalent

Windows: **Data > Initialize Local Saved Scalars to Zero**

Linux: None

Mac OS X: **Data > Initialize Local Saved Scalars to Zero**

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-zero`
`-nozero`

Windows: `/Qzero`
`/Qzero-`

Arguments

None

Default

`-nozero` Local scalar variables are not initialized to zero.

or

`/Qzero-`

Description

This option initializes to zero all local scalar variables of intrinsic type INTEGER, REAL, COMPLEX, or LOGICAL that are saved but not yet initialized.

Use `-save` (Linux and Mac OS X) or `/Qsave` (Windows) on the command line to make all local variables specifically marked as SAVE.

Alternate Options

None

See Also

[save](#) compiler option

r8, r16

See [real-size](#).

rcd, Qrcd

Enables fast float-to-integer conversions.

IDE Equivalent

None

Architectures

IA-32, Intel® 64 architectures

Syntax

Linux and Mac OS X: `-rcd`

Windows: `/Qrct`

Arguments

None

Default

OFF Floating-point values are truncated when a conversion to an integer is involved. On Windows, this is the same as specifying `/QIfist-`.

Description

This option enables fast float-to-integer conversions. It can improve the performance of code that requires floating-point-to-integer conversions. The system default floating-point rounding mode is round-to-nearest. However, the Fortran language requires floating-point values to be truncated when a conversion to an integer is involved. To do this, the compiler must change the rounding mode to truncation before each floating-point-to-integer conversion and change it back afterwards.

This option disables the change to truncation of the rounding mode for all floating-point calculations, including floating point-to-integer conversions. This option can improve performance, but floating-point conversions to integer will not conform to Fortran semantics.

Alternate Options

Linux and Mac OS X: None

Windows: `/QIfist`

rct, Qrct

Sets the internal FPU rounding control to Truncate.

IDE Equivalent

None

Architectures

IA-32, Intel® 64 architectures

Syntax

Linux and Mac OS X: `-rct`

Windows: `/Qrct`

Arguments

None

Default

OFF The compiler uses the default setting for the FPU rounding control.

Description

This option sets the internal FPU rounding control to Truncate.

Alternate Options

Linux and Mac OS X: None

Windows: `/rounding-mode:chopped`

real-size

Specifies the default KIND for real and complex variables.

IDE Equivalent

Windows: **Data > Default Real KIND**

Linux: None

Mac OS X: **Data > Default Real KIND**

Architectures

IA-32, Intel® 64, IA-64 architectures

SyntaxLinux and Mac OS X: `-real-size size`Windows: `/real-size:size`**Arguments**

size Is the size for real and complex variables. Possible values are:
32, 64, or 128.

Default

`real-` Default real and complex variables are 4 bytes long

`size` (REAL(KIND=4) and COMPLEX(KIND=4)).

32

Description

This option specifies the default size (in bits) for real and complex variables.

Option	Description
<code>real-size</code> 32	Makes default real and complex variables 4 bytes long. REAL declarations are treated as single precision REAL (REAL(KIND=4)) and COMPLEX declarations are treated as COMPLEX (COMPLEX(KIND=4)).
<code>real-size</code> 64	Makes default real and complex variables 8 bytes long. REAL declarations are treated as DOUBLE PRECISION (REAL(KIND=8)) and COMPLEX declarations are treated as DOUBLE COMPLEX (COMPLEX(KIND=8)).
<code>real-size</code> 128	Makes default real and complex variables 16 bytes long. REAL declarations are treated as extended precision REAL (REAL(KIND=16)); COMPLEX and DOUBLE COMPLEX

Option	Description
--------	-------------

declarations are treated as extended precision COMPLEX (COMPLEX(KIND=16)).

These compiler options can affect the result type of intrinsic procedures, such as CMPLX, FLOAT, REAL, SNGL, and AIMAG, which normally produce single-precision REAL or COMPLEX results. To prevent this effect, you must explicitly declare the kind type for arguments of such intrinsic procedures.

For example, if `real-size 64` is specified, the CMPLX intrinsic will produce a result of type DOUBLE COMPLEX (COMPLEX(KIND=8)). To prevent this, you must explicitly declare any real argument to be REAL(KIND=4), and any complex argument to be COMPLEX(KIND=4).

Alternate Options

`real-size 64` Linux and Mac OS X: `-r8, -autodouble`
 Windows: `/4R8, /Qautodouble`

`real-size 128` Linux and Mac OS X: `-r16`
 Windows: `/4R16`

recursive

Tells the compiler that all routines should be compiled for possible recursive execution.

IDE Equivalent

Windows: **Code Generation > Enable Recursive Routines**

Linux: None

Mac OS X: **Code Generation > Enable Recursive Routines**

Architectures

IA-32, Intel® 64, IA-64 architectures

Systems: Windows, Linux

Syntax

Linux and Mac OS X: `-recursive`
`-norecursive`
Windows: `/recursive`
`/norecursive`

Arguments

None

Default

`norecursive` Routines are not compiled for possible recursive execution.

Description

This option tells the compiler that all routines should be compiled for possible recursive execution. It sets the `automatic` option.

Alternate Options

None

See Also

[automatic](#) compiler option

reentrancy

Tells the compiler to generate reentrant code to support a multithreaded application.

IDE Equivalent

Windows: **Code Generation > Generate Reentrant Code**

Linux: None

Mac OS X: **Code Generation > Generate Reentrant Code****Architectures**

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-reentrancy keyword`
`-noreentrancy`

Windows: `/reentrancy:keyword`
`/noreentrancy`

Arguments

keyword Specifies details about the program. Possible values are:

<code>none</code>	Tells the run-time library (RTL) that the program does not rely on threaded or asynchronous reentrancy. The RTL will not guard against such interrupts inside its own critical regions. This is the same as specifying <code>noreentrancy</code> .
<code>async</code>	Tells the run-time library (RTL) that the program may contain asynchronous (AST) handlers that could call the RTL. This causes the RTL to guard against AST interrupts inside its own critical regions.
<code>threaded</code>	Tells the run-time library (RTL) that the program is multithreaded, such as programs using the POSIX threads library. This causes the RTL to use thread locking to guard its own critical

regions.

Default

`noreentrancy` The compiler does not generate reentrant code for applications.

Description

This option tells the compiler to generate reentrant code to support a multithreaded application.

If you do not specify a keyword for reentrancy, it is the same as specifying `reentrancy threaded`.

Note that if option `threads` is specified, it sets option `reentrancy threaded`, since multithreaded code must be reentrant.

Alternate Options

None

See Also

[threads](#) compiler option

RTCu

See [check](#).

S

Causes the compiler to compile to an assembly file only and not link.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-S`

Windows: `/S`

Arguments

None

Default

OFF Normal compilation and linking occur.

Description

This option causes the compiler to compile to an assembly file only and not link. On Linux and Mac OS X systems, the assembly file name has a `.s` suffix. On Windows systems, the assembly file name has an `.asm` suffix.

Alternate Options

Linux and Mac OS X: None

Windows: `/Fa`, `/asmfile`

See Also

[Fa](#) compiler option

safe-cray-ptr, Qsafe-cray-ptr

Tells the compiler that Cray* pointers do not alias other variables.

IDE Equivalent

Windows: **Data > Assume Cray Pointers Do Not Share Memory Locations**

Linux: None

Mac OS X: **Data > Assume Cray Pointers Do Not Share Memory Locations**

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-safe-cray-ptr`

Windows: `/Qsafe-cray-ptr`

Arguments

None

Default

OFF The compiler assumes that Cray pointers alias other variables.

Description

This option tells the compiler that Cray pointers do not alias (that is, do not specify sharing memory with) other variables.

Alternate Options

None

Example

Consider the following:

```
pointer (pb, b)
pb = getstorage()
do i = 1, n
b(i) = a(i) + 1
enddo
```

By default, the compiler assumes that `b` and `a` are aliased. To prevent such an assumption, specify the `-safe-cray-ptr` (Linux and Mac OS X) or `/Qsafe-cray-ptr` (Windows) option, and the compiler will treat `b(i)` and `a(i)` as independent of each other.

However, if the variables are intended to be aliased with Cray pointers, using the option produces incorrect results. In the following example, you should not use the option:

```
pointer (pb, b)
```

```
pb = loc(a(2))
do i=1, n
b(i) = a(i) +1
enddo
```

save, Qsave

Causes variables to be placed in static memory.

IDE Equivalent

Windows: **Data > Local Variable Storage**

Linux: None

Mac OS X: **Data > Local Variable Storage**

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-save`

Windows: `/Qsave`

Arguments

None

Default

`-auto- scalar` Scalar variables of intrinsic types INTEGER, REAL, COMPLEX, and LOGICAL are allocated to the run-time stack.
 or Note that if option `recursive`, `-openmp` (Linux and Mac OS X), or `/Qopenmp` (Windows) is specified, the default is `-auto- scalar` automatic (Linux) or `/Qauto` (Windows).

Description

This option saves all variables in static allocation except local variables within a recursive routine and variables declared as AUTOMATIC.

If you want all local, non-SAVED variables to be allocated to the run-time stack, specify option `automatic`.

Alternate Options

Linux and Mac OS X: `-noautomatic`, `-noauto`

Windows: `/noautomatic`, `/noauto`, `/4Na`

See Also

[automatic](#) compiler option

[auto_scalar](#) compiler option

save-temps, Qsave-temps

Tells the compiler to save intermediate files created during compilation.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-save-temps`
`-no-save-temps`

Windows: `/Qsave-temps`
`/Qsave-temps-`

Arguments

None

Default

Linux and Mac OS X: `-no-save-temps` On Linux and Mac OS X systems, the compiler deletes intermediate

Windows: .obj files are saved files after compilation is completed.
 On Windows systems, the compiler saves only intermediate object files after compilation is completed.

Description

This option tells the compiler to save intermediate files created during compilation. The names of the files saved are based on the name of the source file; the files are saved in the current working directory.

If `-save-temps` or `/Qsave-temps` is specified, the following occurs:

- The object .o file (Linux and Mac OS X) or .obj file (Windows) is saved.
- The assembler .s file (Linux and Mac OS X) or .asm file (Windows) is saved if you specified `-use-asm` (Linux or Mac OS X) or `/Quse-asm` (Windows).
- The .i or .i90 file is saved if the fpp preprocessor is invoked.

If `-no-save-temps` is specified on Linux or Mac OS X systems, the following occurs:

- The .o file is put into `/tmp` and deleted after calling `ld`.
- The preprocessed file is not saved after it has been used by the compiler.

If `/Qsave-temps-` is specified on Windows systems, the following occurs:

- The .obj file is not saved after the linker step.
- The preprocessed file is not saved after it has been used by the compiler.



Note

This option only saves intermediate files that are normally created during compilation.

Alternate Options

None

Example

If you compile program *my_foo.F* on a Linux or Mac OS X system and you specify option `-save-temps` and option `-use-asm`, the compilation will produce files *my_foo.o*, *my_foo.s*, and *my_foo.i*.

If you compile program *my_foo.fpp* on a Windows system and you specify option `/Qsave-temps` and option `/Quse-asm`, the compilation will produce files *my_foo.obj*, *my_foo.asm*, and *my_foo.i*.

scalar-rep, Qscalar-rep

Enables scalar replacement performed during loop transformation.

IDE Equivalent

None

Architectures

IA-32 architecture

Syntax

Linux and Mac OS X: `-scalar-rep`
`-no-scalar-rep`

Windows: `/Qscalar-rep`
`/Qscalar-rep-`

Arguments

None

Default

`-no-scalar-rep` Scalar replacement is not performed during loop transformation.

`or/Qscalar-rep-`

Description

This option enables scalar replacement performed during loop transformation. To use this option, you must also specify `o3`.

Alternate Options

None

See Also

[O](#) compiler option

shared

Tells the compiler to produce a dynamic shared object instead of an executable.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux: `-shared`

Mac OS X: `None`

Windows: `None`

Arguments

None

Default

OFF The compiler produces an executable.

Description

This option tells the compiler to produce a dynamic shared object (DSO) instead of an executable. This includes linking in all libraries dynamically and passing `-shared` to the linker.

On systems using IA-32 architecture and Intel® 64 architecture, you must specify option `fpic` for the compilation of each object file you want to include in the shared library.

Alternate Options

None

See Also

[dynamiclib](#) compiler option

[fpic](#) compiler option

[xlinker](#) compiler option

shared-intel

Causes Intel-provided libraries to be linked in dynamically.

IDE Equivalent

Windows: None

Linux: None

Mac OS X: **Run-Time > Intel Runtime Libraries**

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-shared-intel`

Windows: None

Arguments

None

Default

OFF Intel libraries are linked in statically, with the exception of libguide on Linux* and Mac OS* X systems, where it is linked in dynamically.

Description

This option causes Intel-provided libraries to be linked in dynamically. It is the opposite of `-static-intel`.



Note

On Mac OS X systems, when you set "Intel Runtime Libraries" to "Dynamic", you must also set the DYLD_LIBRARY_PATH environment variable within Xcode or an error will be displayed.

Alternate Options

Linux and Mac OS X: `-i-dynamic` (this is a [deprecated](#) option)

Windows: None

See Also

[static-intel](#) compiler option

shared-libgcc

Links the GNU `libgcc` library dynamically.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux: `-shared-libgcc`
Mac OS X: `None`
Windows: `None`

Arguments

None

Default

`-shared-libgcc` The compiler links the `libgcc` library dynamically.

Description

This option links the GNU `libgcc` library dynamically. It is the opposite of option `static-libgcc`.

This option is useful when you want to override the default behavior of the `static` option, which causes all libraries to be linked statically.

Alternate Options

None

See Also

[static-libgcc](#) compiler option

source

Tells the compiler to compile the file as a Fortran source file.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: None

Windows: `/source:file`

Arguments

file Is the name of the file.

Default

OFF Files that do not end in standard Fortran file extensions are not compiled as Fortran files.

Description

This option tells the compiler to compile the file as a Fortran source file.

This option is useful when you have a Fortran file with a nonstandard file extension (that is, not one of .F, .FOR, or .F90).

This option assumes the file specified uses fixed source form. If the file uses free source form, you must also specify option *free*.

Alternate Options

Linux and Mac OS X: `-Tf file`

Windows: `/Tf file`

See Also

[extfor](#) compiler option

[free](#) compiler option

sox, Qsox

Tells the compiler to save the compiler options and version number in the executable.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-sox`
`-no-sox`

Windows: `/Qsox`
`/Qsox-`

Arguments

None

Default

`-no-sox` The compiler does not save the compiler options and
or `/Qsox-` version number in the executable.

Description

This option tells the compiler to save the compiler options and version number in the executable. The size of the executable on disk is increased slightly by the inclusion of these information strings.

This option forces the compiler to embed in each object file or assembly output a string that contains information about the compiler version and compilation options for each source file that has been compiled. When you link the object files into an executable file, the linker places each of the information strings into the header of the executable. It is then possible to use a tool, such as a strings utility, to determine what options were used to build the executable file.

If `-no-sox` or `/Qsox-` is specified, this extra information is not put into the object or assembly output generated by the compiler.

Alternate Options

None

stand

Tells the compiler to issue compile-time messages for nonstandard language elements.

IDE Equivalent

Windows: **Compilation Diagnostics > Warn For Nonstandard Fortran**

Linux: None

Mac OS X: **Compiler Diagnostics > Warn For Nonstandard Fortran**

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-stand [keyword]`

`-nostand`

Windows: `/stand[:keyword]`

`/nostand`

Arguments

keyword Specifies the language to use as the standard. Possible values are:

<code>none</code>	Issue no messages for nonstandard language elements.
<code>f90</code>	Issue messages for language elements that are not standard in Fortran 90.
<code>f95</code>	Issue messages for language elements that are not standard in

Fortran 95.

`f03` Issue messages for language elements that are not standard in Fortran 2003.

Default

`nostand` The compiler issues no messages for nonstandard language elements.

Description

This option tells the compiler to issue compile-time messages for nonstandard language elements.

If you do not specify a keyword for `stand`, it is the same as specifying `stand f95`.

Option	Description
<code>stand none</code>	Tells the compiler to issue no messages for nonstandard language elements. This is the same as specifying <code>nostand</code> .
<code>stand f90</code>	Tells the compiler to issue messages for language elements that are not standard in Fortran 90.
<code>stand f95</code>	Tells the compiler to issue messages for language elements that are not standard in Fortran 95.
<code>stand f03</code>	Tells the compiler to issue messages for language elements that are not standard in Fortran 2003. This option is set if you specify <code>warn stderrs</code> .

Alternate Options

`stand none` Linux and Mac OS X: `-nostand`
Windows: `/nostand, /4Ns`

<code>stand f90</code>	Linux and Mac OS X: <code>-std90</code> Windows: <code>/4Ys</code>
<code>stand f95</code>	Linux and Mac OS X: <code>-std95</code> Windows: None
<code>stand f03</code>	Linux and Mac OS X: <code>-std03, -stand, -std</code> Windows: <code>/stand</code>

See Also

[warn](#) compiler option

static

Prevents linking with shared libraries.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux:	<code>-static</code>
Mac OS X:	None
Windows:	<code>/static</code>

Arguments

None

Default

`static` The compiler does not link with shared libraries.

Description

This option prevents linking with shared libraries. It causes the executable to link all libraries statically.

Alternate Options

None

staticlib

Invokes the `libtool` command to generate static libraries.

IDE Equivalent

None

Architectures

IA-32, Intel® 64 architectures

Syntax

Linux:	None
Mac OS X:	<code>-staticlib</code>
Windows:	None

Arguments

None

Default

OFF The compiler produces an executable.

Description

This option invokes the `libtool` command to generate static libraries. When passed this option, the compiler uses the `libtool` command to produce a static library instead of an executable when linking.

To build dynamic libraries, you should specify option `-dynamiclib` or `libtool -dynamic <objects>`.

Alternate Options

None

See Also

[dynamiclib](#) compiler option

static-intel

Causes Intel-provided libraries to be linked in statically.

IDE Equivalent

Windows: None

Linux: None

Mac OS X: **Run-Time > Intel Runtime Libraries**

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-static-intel`

Windows: None

Arguments

None

Default

OFF Intel libraries are linked in statically, with the exception of `libguide`, which is linked in dynamically.

Description

This option causes Intel-provided libraries to be linked in statically. It is the opposite of `-shared-intel`.

Alternate Options

Linux and Mac OS X: `i-static` (this is a [deprecated](#) option)

Windows: None

See Also

[shared-intel](#) compiler option

static-libgcc

Links the GNU `libgcc` library statically.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux: `-static-libgcc`

Mac OS X: `None`

Windows: `None`

Arguments

None

Default

OFF `DEFAULT_DESC`

Description

This option links the GNU `libgcc` library statically. It is the opposite of option `libgcc`.

This option is useful when you want to override the default behavior of the `libgcc` option, which causes all libraries to be linked statically.

Alternate Options

None

See Also

[shared-libgcc](#) compiler option

std, std90, std95, std03

See [stand](#).

syntax-only

Tells the compiler to check only for correct syntax.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-syntax-only`

Windows: `/syntax-only`

Arguments

None

Default

OFF Normal compilation is performed.

Description

This option tells the compiler to check only for correct syntax. It lets you do a quick syntax check of your source file.

Compilation stops after the source file has been parsed. No code is generated, no object file is produced, and some error checking done by the optimizer is bypassed.

Warnings and messages appear on `stderr`.

Alternate Options

Linux: `-y`, `-fsyntax-only`, `-syntax` (this is a [deprecated](#) option)

Mac OS X: `-y`, `-fsyntax-only`

Windows: `/Zs`

T

Tells the linker to read link commands from a file.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux: -T*file*
Mac OS X: None
Windows: None

Arguments

file Is the name of the file.

Default

OFF The linker does not read link commands from a file.

Description

This option tells the linker to read link commands from a file.

Alternate Options

None

tcheck, Qtcheck

Enables analysis of threaded applications.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux: -tcheck
Mac OS X: None
Windows: /Qtcheck

Arguments

None

Default

OFF Threaded applications are not instrumented by the compiler for analysis by Intel® Thread Checker.

Description

This option enables analysis of threaded applications.

To use this option, you must have Intel® Thread Checker installed, which is one of the Intel® Threading Analysis Tools. If you do not have this tool installed, the compilation will fail. Remove the `-tcheck` (Linux) or `/Qtcheck` (Windows) option from the command line and recompile.

For more information about Intel® Thread Checker (including an evaluation copy), open the page associated with threading tools at [Intel® Software Development Products](#).

Alternate Options

None

tcollect, Qtcollect

Inserts instrumentation probes calling the Intel® Trace Collector API.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux: `-tcollect[lib]`

Mac OS X: None
 Windows: /`Qtcollect[:lib]`

Arguments

lib Is one of the Intel® Trace Collector libraries; for example, VT, VTcs, VTmc, or VTfs. If you do not specify *lib*, the default library is VT.

Default

OFF Instrumentation probes are not inserted into compiled applications.

Description

This option inserts instrumentation probes calling the Intel® Trace Collector API. To use this option, you must have the Intel® Trace Collector installed and set up through one of its set-up scripts. This tool is a component of the Intel® Trace Analyzer and Collector.

This option provides a flexible and convenient way of instrumenting functions of a compiled application. For every function, the entry and exit points are instrumented at compile time to let the Intel® Trace Collector record functions beyond the default MPI calls. For non-MPI applications (for example, threaded or serial), you must ensure that the Intel® Trace Collector is properly initialized (VT_initialize/VT_init).



Caution

Be careful with full instrumentation because this feature can produce very large trace files.

For more details, see the Intel® Trace Collector User Guide.

Alternate Options

None

See Also

[tcollect-filter](#), [Qtcollect-filter](#) compiler option

tcollect-filter, Qtcollect-filter

Lets you enable or disable the instrumentation of specified functions.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux: `-tcollect-filter file`

Mac OS X: `None`

Windows: `/Qtcollect-filter:file`

Arguments

file Is a configuration file that lists filters, one per line. Each filter consists of a regular expression string and a switch. Strings with leading or trailing white spaces must be quoted. Other strings do not have to be quoted. The switch value can be ON, on, OFF, or off.

Default

OFF Functions are not instrumented. However, if option `-tcollect` (Linux) or `/Qtcollect` (Windows) is specified, the filter setting is `.* ON` and all functions get instrumented.

Description

This option lets you enable or disable the instrumentation of specified functions. During instrumentation, the regular expressions in the file are matched against the function names. The switch specifies whether matching functions are to be instrumented or not. Multiple filters are evaluated from top to bottom with increasing precedence.

The names of the functions to match against are formatted as follows:

- The source file name is followed by a colon-separated function name. Source file names should contain the full path, if available. For example:

```
/home/joe/src/file.f:FOO_bar
```

- Classes and function names are separated by double colons. For example:

```
/home/joe/src/file.fpp:app::foo::bar
```

You can use option `-opt-report` (Linux) or `/Qopt-report` (Windows) to get a full list of file and function names that the compiler recognizes from the compilation unit. This list can be used as the basis for filtering in the configuration file.

To use this option, you must have the Intel® Trace Collector installed and set up through one of its set-up scripts. This tool is a component of the Intel® Trace Analyzer and Collector.

For more details, see the Intel® Trace Collector User Guide.

Alternate Options

None

Consider the following filters in a configuration file:

```
'.*' OFF '.*vector.*' ON
```

The above will cause instrumentation of only those functions having the string 'vector' in their names. No other function will be instrumented. Note that reversing the order of the two lines will prevent instrumentation of all functions.

To get a list of the file or routine strings that can be matched by the regular expression filters, generate an optimization report with `tcollect` information. For example:

```
Windows OS: ifort /Qtcollect /Qopt-report /Qopt-report-phase tcollect
Linux OS: ifort -tcollect -opt-report -opt-report-phase tcollect
```

See Also

[tcollect, Qtcollect](#) compiler option

Tf

See [source](#).

threads

Tells the linker to search for unresolved references in a multithreaded run-time library.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-threads`
`-nothreads`
Windows: `/threads`
`/nothreads`

Arguments

None

Default

Systems On systems using IA-32 architecture and IA-64 architecture, using Intel® the linker does not search for unresolved references in a 64 mutithreaded run-time library. On systems using Intel® 64 architecture: architectures, it does.

`threads`

Systems
using IA-32

architecture
and IA-64
architecture:
nothreads

Description

This option tells the linker to search for unresolved references in a multithreaded run-time library.

This option sets option `reentrancy threaded`.

Windows systems: The following table shows which options to specify for a multithreaded run-time library.

Type of Library	Options Required	Alternate Option
Multithreaded	<code>/libs:static</code> <code>/threads</code>	<code>/MT</code>
Debug multithreaded	<code>/libs:static</code> <code>/threads</code> <code>/dbglibs</code>	<code>/MTd</code>
Multithreaded DLLs	<code>/libs:dll</code> <code>/threads</code>	<code>/MD</code>
Multithreaded debug DLLs	<code>/libs:dll</code> <code>/threads</code> <code>/dbglibs</code>	<code>/MDd</code>

Alternate Options

None

See Also

Building Applications: Specifying Consistent Library Types; Programming with Mixed Languages Overview

tprofile, Qtprofile

Intel® Fortran Compiler User and Reference Guides

Generates instrumentation to analyze multi-threading performance.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux: `-tprofile`

Mac OS X: `None`

Windows: `/Qtprofile`

Arguments

None

Default

OFF Instrumentation is not generated by the compiler for analysis by Intel® Thread Profiler.

Description

This option generates instrumentation to analyze multi-threading performance. To use this option, you must have Intel® Thread Profiler installed, which is one of the Intel® Threading Analysis Tools. If you do not have this tool installed, the compilation will fail. Remove the `-tprofile` (Linux) or `/Qtprofile` (Windows) option from the command line and recompile.

For more information about Intel® Thread Profiler (including an evaluation copy), open the page associated with threading tools at [Intel® Software Development Products](#).

Alternate Options

None

traceback

Tells the compiler to generate extra information in the object file to provide source file traceback information when a severe error occurs at run time.

IDE Equivalent

Windows: **Run-time > Generate Traceback Information**

Linux: None

Mac OS X: **Run-time > Generate Traceback Information**

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-traceback`
`-notraceback`

Windows: `/traceback`
`/notraceback`

Arguments

None

Default

`notraceback` No extra information is generated in the object file to produce traceback information.

Description

This option tells the compiler to generate extra information in the object file to provide source file traceback information when a severe error occurs at run time. When the severe error occurs, source file, routine name, and line number correlation information is displayed along with call stack hexadecimal addresses (program counter trace).

Note that when a severe error occurs, advanced users can also locate the cause of the error using a map file and the hexadecimal addresses of the stack displayed when the error occurs.

This option increases the size of the executable program, but has no impact on run-time execution speeds.

It functions independently of the debug option.

On Windows systems, `traceback` sets the `/Oy-` option, which forces the compiler to use EBP as the stack frame pointer.

On Windows systems, the linker places the traceback information in the executable image, in a section named ".trace". To see which sections are in an image, use the command:

```
link -dump -summary your_app_name.exe
```

To see more detailed information, use the command:

```
link -dump -headers your_app_name.exe
```

On Windows systems, when requesting traceback, you must set Enable Incremental Linking in the VS .NET* IDE Linker Options to No. On systems using IA-32 architecture and Intel® 64 architecture, you must also set Omit Frame Pointers (the `/Oy` option) in the Optimization Options to "No."

On Linux systems, to display the section headers in the image (including the header for the .trace section, if any), use the command:

```
objdump -h your_app_name.exe
```

On Mac OS X systems, to display the section headers in the image, use the command:

```
otool -l your_app_name.exe
```

Alternate Options

None

tune

Determines the version of the architecture for which the compiler generates instructions.

IDE Equivalent

None

Architectures

IA-32, Intel® 64 architectures

Syntax

Linux and Mac OS X: `-tune keyword`

Windows: `/tune:keyword`

Arguments

keyword Specifies the processor type. Possible values are:

<code>pn1</code>	Optimizes for the Intel® Pentium® processor.
<code>pn2</code>	Optimizes for the Intel® Pentium® Pro, Intel® Pentium® II, and Intel® Pentium® III processors.
<code>pn3</code>	Optimizes for the Intel® Pentium® Pro, Intel® Pentium® II, and Intel® Pentium® III processors. This is the same as specifying <code>pn2</code> .
<code>pn4</code>	Optimizes for the Intel® Pentium® 4 processor.

Default

`pn4` The compiler optimizes for the Intel® Pentium® 4 processor.

Description

This option determines the version of the architecture for which the compiler generates instructions.

On systems using Intel® 64 architecture, only *keyword* `pn4` is valid.

Alternate Options

None

u (Linux* and Mac OS* X)

See [warn](#).

u (Windows*)

Undefines all previously defined preprocessor values.

IDE Equivalent

Windows: **Preprocessor > Undefine All Preprocessor Definitions**

Linux: None

Mac OS X: None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: None

Windows: /u

Arguments

None

Default

OFF Defined preprocessor values are in effect until they are undefined.

Description

This option undefines all previously defined preprocessor values. To undefine specific preprocessor values, use the /U option.

Alternate Options

None

See Also

[U](#) compiler option

U

Undefines any definition currently in effect for the specified symbol.

IDE Equivalent

Windows: **Preprocessor > Undefine Preprocessor Definitions**

Linux: None

Mac OS X: **Preprocessor > Undefine Preprocessor Definitions**

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-Uname`

Windows: `/Uname`

Arguments

name Is the name of the symbol to be undefined.

Default

OFF Symbol definitions are in effect until they are undefined.

Description

This option undefines any definition currently in effect for the specified symbol. On Windows systems, use the `/u` option to undefine all previously defined preprocessor values.

Alternate Options

Linux and Mac OS X: None

Windows: `/undefine:name`

See Also

[u \(Windows\)](#) compiler option

undefine

See [U](#).

unroll, Qunroll

Tells the compiler the maximum number of times to unroll loops.

IDE Equivalent

Windows: **Optimization > Loop Unroll Count**

Linux: None

Mac OS X: **Optimization > Loop Unroll Count**

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-unroll[=n]`

Windows: `/Qunroll[:n]`

Arguments

n Is the maximum number of times a loop can be unrolled. To disable loop enrolling, specify 0. On systems using IA-64 architecture, you can only specify a value of 0.

Default

`-unroll` The compiler uses default heuristics when unrolling loops.
or `/Qunroll`

Description

This option tells the compiler the maximum number of times to unroll loops. If you do not specify *n*, the optimizer determines how many times loops can be unrolled.

Alternate Options

Linux and Mac OS X: `-funroll-loops`

Windows: `/unroll`

See Also

Optimizing Applications: Loop Unrolling

unroll-aggressive, Qunroll-aggressive

Determines whether the compiler uses more aggressive unrolling for certain loops.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-unroll-aggressive`
 `-no-unroll-aggressive`

Windows: `/Qunroll-aggressive`
 `/Qunroll-aggressive-`

Arguments

None

Default

`-no-unroll-aggressive` The compiler uses default heuristics when
`or/Quroll-` unrolling loops.
`aggressive-`

Description

This option determines whether the compiler uses more aggressive unrolling for certain loops. The positive form of the option may improve performance.

On IA-32 architecture and Intel® 64 architecture, this option enables aggressive, complete unrolling for loops with small constant trip counts.

On IA-64 architecture, this option enables additional complete unrolling for loops that have multiple exits or outer loops that have a small constant trip count.

Alternate Options

None

uppercase, Quppercase

See [names](#).

us

See [assume](#).

use-asm, Quse-asm

Tells the compiler to produce objects through the assembler.

IDE Equivalent

None

Architectures

`-use-asm`: IA-32 architecture, Intel® 64 architecture, IA-64 architecture

/Quse-asm: IA-64 architecture

Syntax

Linux and Mac OS X: -use-asm
 -no-use-asm

Windows: /Quse-asm
 /Quse-asm-

Arguments

None

Default

-no-use-asm The compiler produces objects directly.
or /Quse-asm-

Description

This option tells the compiler to produce objects through the assembler.

Alternate Options

None

v

Specifies that driver tool commands should be displayed and executed.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-v[file]`

Windows: `None`

Arguments

file Is the name of a file.

Default

OFF No tool commands are shown.

Description

This option specifies that driver tool commands should be displayed and executed.

If you use this option without specifying a file name, the compiler displays only the version of the compiler.

If you want to display processing information (pass information and source file names), specify option `watch:all`.

Alternate Options

Linux and Mac OS X: `-watch cmd`

Windows: `/watch:cmd`

See Also

[dryrun](#) compiler option

[watch](#) compiler option

V (Linux* and Mac OS* X)

See [logo](#)

V (Windows*)

See [bintext](#).

vec, Qvec

Enables or disables vectorization and transformations enabled for vectorization.

IDE Equivalent

None

Architectures

IA-32, Intel® 64 architectures

Syntax

Linux and Mac OS X: `-vec`
`-no-vec`

Windows: `/Qvec`
`/Qvec-`

Arguments

None

Default

`-vec` Vectorization is enabled.
or `/Qvec`

Description

This option enables or disables vectorization and transformations enabled for vectorization.

To disable vectorization and transformations enabled for vectorization, specify `-no-vec` (Linux and Mac OS X) or `/Qvec-` (Windows).

Alternate Options

None

vec-guard-write, Qvec-guard-write

Tells the compiler to perform a conditional check in a vectorized loop.

IDE Equivalent

None

Architectures

IA-32, Intel® 64 architectures

Syntax

Linux and Mac OS X: `-vec-guard-write`
`-no-vec-guard-write`

Windows: `/Qvec-guard-write`
`/Qvec-guard-write-`

Arguments

None

Default

`-no-vec-guard-write` The compiler uses default heuristics when
or `/Qvec-guard-write-` checking vectorized loops.

Description

This option tells the compiler to perform a conditional check in a vectorized loop. This checking avoids unnecessary stores and may improve performance.

Alternate Options

None

vec-report, Qvec-report

Controls the diagnostic information reported by the vectorizer.

IDE Equivalent

Windows: **Compilation Diagnostics > Vectorizer Diagnostic Level**

Linux: None

Mac OS X: **Diagnostics > Vectorizer Diagnostic Report**

Architectures

IA-32, Intel® 64 architectures

Syntax

Linux and Mac OS X: `-vec-report[n]`

Windows: `/Qvec-report[n]`

Arguments

<i>n</i>	Is a value denoting which diagnostic messages to report. Possible values are:
0	Tells the vectorizer to report no diagnostic information.
1	Tells the vectorizer to report on vectorized loops.
2	Tells the vectorizer to report on vectorized and non-vectorized loops.
3	Tells the vectorizer to report on vectorized and non-vectorized loops and any proven or assumed data dependences.
4	Tells the vectorizer to report on non-vectorized loops.
5	Tells the vectorizer to report on non-vectorized loops and the

reason why they were not
vectorized.

Default

`-vec-report1` If the vectorizer has been enabled, it reports diagnostics
on vectorized loops.
or `/Qvec-report1`

Description

This option controls the diagnostic information reported by the vectorizer. The vectorizer report is sent to stdout.

If you do not specify *n*, it is the same as specifying `-vec-report1` (Linux and Mac OS X) or `/Qvec-report1` (Windows).

The vectorizer is enabled when certain compiler options are specified, such as option `-ax` or `-x` (Linux and Mac OS X), option `/Qax` or `/Qx` (Windows), option `-arch SSE` or `-arch SSE2` (Linux and Mac OS X), option `/architecture:SSE` or `/architecture:SSE2` (Windows).

If this option is specified from within the IDE, the report is included in the build log if the Generate Build Logs option is selected.

Alternate Options

None

vms

Causes the run-time system to behave like HP* Fortran on OpenVMS* Alpha systems and VAX* systems (VAX FORTRAN*).

IDE Equivalent

Windows: **Compatibility > Enable VMS Compatibility**

Linux: None

Mac OS X: **Compatibility > Enable VMS Compatibility**

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-vms`

`-novms`

Windows: `/vms`

`/novms`

Arguments

None

Default

`novms` The run-time system follows default Intel® Fortran behavior.

Description

This option causes the run-time system to behave like HP* Fortran on OpenVMS* Alpha systems and VAX* systems (VAX FORTRAN*).

It affects the following language features:

- Certain defaults
In the absence of other options, `vms` sets the defaults as `check format` and `check output_conversion`.
- Alignment
Option `vms` does not affect the alignment of fields in records or items in common blocks. For compatibility with HP Fortran on OpenVMS systems, use `align norecords` to pack fields of records on the next byte boundary.
- Carriage control default

If option `vms` and option `ccdefault default` are specified, carriage control defaults to FORTRAN if the file is formatted and the unit is connected to a terminal.

- INCLUDE qualifiers

`/LIST` and `/NOLIST` are recognized at the end of the file name in an INCLUDE statement at compile time. If the file name in the INCLUDE statement does not specify the complete path, the path used is the current directory. Note that if `vms` is not specified, the path used is the directory where the file that contains the INCLUDE statement resides.

- Quotation mark character

A quotation mark (") character is recognized as starting an octal constant ("0..7) instead of a character literal ("...").

- Deleted records in relative files

When a record in a relative file is deleted, the first byte of that record is set to a known character (currently '@'). Attempts to read that record later result in ATTACCNON errors. The rest of the record (the whole record, if `vms` is not specified) is set to nulls for unformatted files and spaces for formatted files.

- ENDFILE records

When an ENDFILE is performed on a sequential unit, an actual 1-byte record containing a Ctrl/Z is written to the file. If `vms` is not specified, an internal ENDFILE flag is set and the file is truncated. The `vms` option does not affect ENDFILE on relative files: these files are truncated.

- Implied logical unit numbers

The `vms` option enables Intel Fortran to recognize certain environment variables at run time for ACCEPT, PRINT, and TYPE statements and for READ and WRITE statements that do not specify a unit number (such as READ (*,1000)).

- Treatment of blanks in input

The `vms` option causes the defaults for the keyword BLANK in OPEN statements to become 'NULL' for an explicit OPEN and 'ZERO' for an implicit OPEN of an external or internal file.

- **OPEN statement effects**

Carriage control defaults to FORTRAN if the file is formatted, and the unit is connected to a terminal. Otherwise, carriage control defaults to LIST. The `vms` option affects the record length for direct access and relative organization files. The buffer size is increased by 1 to accommodate the deleted record character.

- **Reading deleted records and ENDFILE records**

The run-time direct access READ routine checks the first byte of the retrieved record. If this byte is '@' or NULL ("\0"), then an ATTACCNON error is returned. The run-time sequential access READ routine checks to see if the record it just read is one byte long and contains a Ctrl/Z. If this is true, it returns EOF.

Alternate Options

Linux and Mac OS X: None

Windows: `/Qvms`

See Also

[align](#) compiler option

[ccdefault](#) compiler option

[check](#) compiler option

w

See keywords `none` and `nogeneral` in [warn](#)

W0, W1

See [warn](#).

W0, W1

See [warn](#).

Wa

Passes options to the assembler for processing.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-Wa,option1[,option2,...]`

Windows: None

Arguments

<i>option</i>	Is an assembler option. This option is not processed by the driver and is directly passed to the assembler.
---------------	---

Default

OFF No options are passed to the assembler.

Description

This option passes one or more options to the assembler for processing. If the assembler is not invoked, these options are ignored.

Alternate Options

None

warn

Specifies diagnostic messages to be issued by the compiler.

IDE Equivalent

Windows: **General > Compile Time Diagnostics** (/warn:all, /warn:none)

Compilation Diagnostics > Treat Warnings as Errors (/warn:[no]errors)

Compilation Diagnostics > Treat Fortran Standard Warnings as Errors

(/warn:[no]stderrs)

Compilation Diagnostics > Compile Time Diagnostics (/warn:all,

/warn:none)

Compilation Diagnostics > Warn for Undeclared Symbols

(/warn:[no]declarations)

Compilation Diagnostics > Warn for Unused Variables

(/warn:[no]unused)

Compilation Diagnostics > Warn When Removing %LOC

(/warn:[no]ignore_loc)

Compilation Diagnostics > Warn When Truncating Source Line

(/warn:[no]truncated_source)

Compilation Diagnostics > Warn for Unaligned Data

(/warn:[no]alignments)

Compilation Diagnostics > Warn for Uncalled Routine

(/warn:[no]uncalled)

Compilation Diagnostics > Suppress Usage Messages (/warn:[no]usage)

Compilation Diagnostics > Check Routine Interfaces

(/warn:[no]interfaces)

Linux: None

Mac OS X: **General > Compile Time Diagnostics** (-warn all, -warn

none)

Compiler Diagnostics > Warn For Unaligned Data (-warn

[no]alignments)

Compiler Diagnostics > Warn For Undeclared Symbols (-warn

[no]declarations)

Compiler Diagnostics > Treat Warnings as Errors (-warn error)

Compiler Diagnostics > Warn When Removing %LOC (-warn
[no]ignore_loc)

Compiler Diagnostics > Check Routine Interfaces (-warn
[no]interfaces)

Compiler Diagnostics > Treat Fortran Standard Warnings As Errors (-warn
[no]stderrs)

Compiler Diagnostics > Warn When Truncating Source Line (-warn
[no]truncated_source)

Compiler Diagnostics > Warn For Uncalled Routine (-warn [no]uncalled)

Compiler Diagnostics > Warn For Unused Variables (-warn [no]unused)

Compiler Diagnostics > Suppress Usage Messages (-warn [no]usage)

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: -warn [*keyword*]
-nowarn

Windows: /warn[:*keyword*]
/nowarn

Arguments

keyword Specifies the diagnostic messages to be issued. Possible values are:

none	Disables all warning messages.
------	--------------------------------

[no]alignments	Determines whether warnings occur for data that is not naturally aligned.
----------------	---

<code>[no]declarations</code>	Determines whether warnings occur for any undeclared symbols.
<code>[no]errors</code>	Determines whether warnings are changed to errors.
<code>[no]general</code>	Determines whether warning messages and informational messages are issued by the compiler.
<code>[no]ignore_loc</code>	Determines whether warnings occur when <code>%LOC</code> is stripped from an actual argument.
<code>[no]interfaces</code>	Determines whether the compiler checks the interfaces of all SUBROUTINEs called and FUNCTIONs invoked in your compilation against an external set of interface blocks.
<code>[no]stderrs</code>	Determines whether warnings about Fortran standard violations are changed to errors.
<code>[no]truncated_source</code>	Determines whether warnings occur when

	source exceeds the maximum column width in fixed-format files.
[no]uncalled	Determines whether warnings occur when a statement function is never called
[no]unused	Determines whether warnings occur for declared variables that are never used.
[no]usage	Determines whether warnings occur for questionable programming practices.
all	Enables all warning messages.

Default

alignments	Warnings are issued about data that is not naturally aligned.
general	All information-level and warning-level messages are enabled.
usage	Warnings are issued for questionable programming practices.
nodeclarations	No errors are issued for undeclared symbols.
noerrors	Warning-level messages are not changed to error-level messages.

<code>noignore_loc</code>	No warnings are issued when <code>%LOC</code> is stripped from an argument.
<code>nointerfaces</code>	The compiler does not check interfaces of SUBROUTINES called and FUNCTIONS invoked in your compilation against an external set of interface blocks.
<code>nostderrors</code>	Warning-level messages about Fortran standards violations are not changed to error-level messages.
<code>notruncated_source</code>	No warnings are issued when source exceeds the maximum column width in fixed-format files.
<code>nouncalled</code>	No warnings are issued when a statement function is not called.
<code>nounused</code>	No warnings are issued for variables that are declared but never used.

Description

This option specifies the diagnostic messages to be issued by the compiler.

Option	Description
<code>warn none</code>	Disables all warning messages. This is the same as specifying <code>nowarn</code> .
<code>warn noalignments</code>	Disables warnings about data that is not naturally aligned.
<code>warn undeclared_symbols</code>	Enables error messages about any undeclared symbols.
<code>declarations</code>	This option makes the default data type of a variable undefined (IMPLICIT NONE) rather than using the implicit Fortran rules.

Option	Description
<code>warn errors</code>	Tells the compiler to change all warning-level messages to error-level messages; this includes warnings about Fortran standards violations.
<code>warn nogeneral</code>	Disables all informational-level and warning-level diagnostic messages.
<code>warn ignore_loc</code>	Enables warnings when %LOC is stripped from an actual argument.
<code>warn interfaces</code>	Tells the compiler to check the interfaces of all SUBROUTINEs called and FUNCTIONs invoked in your compilation against a set of interface blocks stored separately from the source being compiled. The compiler generates a compile-time message if the interface used to invoke a routine does not match the interface defined in a .mod file external to the source (that is, in a .mod generated by option <code>gen-interfaces</code> as opposed to a .mod file USED in the source). The compiler looks for these .mods in the current directory or in the directory specified by the <code>include (-I)</code> or <code>-module</code> option.
<code>warn stderrs</code>	Tells the compiler to change all warning-level messages about Fortran standards violations to error-level messages. This option sets the <code>std03</code> option (Fortran 2003 standard). If you want Fortran 95 standards violations to become errors, you must specify options <code>warn stderrs</code> and <code>std95</code> .
<code>warn truncated_source</code>	Enables warnings when a source line exceeds the maximum column width in fixed-format source files. The maximum column width for fixed-format files is 72, 80, or

Option	Description
	132, depending on the setting of the <code>extend-source</code> option. The <code>warn truncated_source</code> option has no effect on truncation; lines that exceed the maximum column width are always truncated. This option does not apply to free-format source files.
<code>warn uncalled</code>	Enables warnings when a statement function is never called.
<code>warn unused</code>	Enables warnings for variables that are declared but never used.
<code>warn nousage</code>	Disables warnings about questionable programming practices. Questionable programming practices, although allowed, often are the result of programming errors; for example: a continued character or Hollerith literal whose first part ends before the statement field and appears to end with trailing spaces. Note that the <code>/pad-source</code> option can prevent this error.
<code>warn all</code>	Enables all warning messages. This is the same as specifying <code>warn</code> . This option does not set options <code>warn errors</code> or <code>warn stderrs</code> . To enable all the additional checking to be performed and force the severity of the diagnostic messages to be severe enough to not generate an object file, specify <code>warn all warn errors</code> or <code>warn all warn stderrs</code> . On Windows systems: In the Property Pages, Custom means that diagnostics will be specified on an individual basis.

Alternate Options

Intel® Fortran Compiler User and Reference Guides

warn none	Linux and Mac OS X: -nowarn, -w, -W0, -warn nogeneral Windows: /nowarn,/w, /W0, /warn:nogeneral
warn declarations	Linux and Mac OS X: -implicitnone, -u Windows: /4Yd
warn nodeclarations	Linux and Mac OS X: None Windows: /4Nd
warn general	Linux and Mac OS X: -w1 Windows: /w1
warn nogeneral	Linux and Mac OS X: -W0, -w, -nowarn, -warn none Windows: /W0, /w, /nowarn, /warn:none
warn stderrs	Linux and Mac OS X: -e90, -e95, -e03 Windows: None
warn nousage	Linux and Mac OS X: -cm Windows: /cm
warn all	Linux and Mac OS X: -warn Windows: /warn

watch

Tells the compiler to display certain information to the console output window.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-watch [keyword]`
`-nowatch`

Windows: `/watch[:keyword]`
`/nowatch`

Arguments

keyword Determines what information is displayed. Possible values are:

<code>none</code>	Disables <code>cmd</code> and <code>source</code> .
<code>[no]cmd</code>	Determines whether driver tool commands are displayed and executed.
<code>[no]source</code>	Determines whether the name of the file being compiled is displayed.
<code>all</code>	Enables <code>cmd</code> and <code>source</code> .

Default

`nowatch` Pass information and source file names are not displayed to the console output window.

Description

Tells the compiler to display processing information (pass information and source file names) to the console output window.

Option	Description
<code>watch none</code>	Tells the compiler to not display pass information and source file names to the console output window. This is the same as specifying <code>nowatch</code> .
<code>watch cmd</code>	Tells the compiler to display and execute driver tool commands.
<code>watch</code>	Tells the compiler to display the name of the file being compiled.

Option	Description
--------	-------------

source

`watch all` Tells the compiler to display pass information and source file names to the console output window. This is the same as specifying `watch` with no *keyword*.

Alternate Options

`watch` Linux and Mac OS X: `-v`

`cmd` Windows: None

See Also

[v](#) compiler option

WB

Turns a compile-time bounds check into a warning.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-WB`

Windows: `/WB`

Arguments

None

Default

OFF Compile-time bounds checks are errors.

Description

This option turns a compile-time bounds check into a warning.

Alternate Options

None

what

Tells the compiler to display its detailed version string.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-what`

Windows: `/what`

Arguments

None

Default

OFF The version strings are not displayed.

Description

This option tells the compiler to display its detailed version string.

Alternate Options

None

winapp

Tells the compiler to create a graphics or Fortran Windows application and link against the most commonly used libraries.

IDE Equivalent

Windows: **Libraries > Use Common Windows Libraries**

Linux: None

Mac OS X: None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: None

Windows: /winapp

Arguments

None

Default

OFF No graphics or Fortran Windows application is created.

Description

This option tells the compiler to create a graphics or Fortran Windows application and link against the most commonly used libraries.

Alternate Options

Linux and Mac OS X: None

Windows: /MG

Winline

Enables diagnostics about what is inlined and what is not inlined.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-winline`

Windows: None

Arguments

None

Default

OFF No diagnostics are produced about what is inlined and what is not inlined.

Description

This option enables diagnostics about what is inlined and what is not inlined. The diagnostics depend on what interprocedural functionality is available.

Alternate Options

None

WI

Passes options to the linker for processing.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-wl,option1[,option2,...]`

Windows: None

Arguments

option Is a linker option. This option is not processed by the driver and is directly passed to the linker.

Default

OFF No options are passed to the linker.

Description

This option passes one or more options to the linker for processing. If the linker is not invoked, these options are ignored.

This option is equivalent to specifying option `-Qoption,link,options`.

Alternate Options

None

See Also

[Option](#) compiler option

Wp

Passes options to the preprocessor.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-Wp,option1[,option2,...]`

Windows: None

Arguments

option Is a preprocessor option. This option is not processed by the driver and is directly passed to the preprocessor.

Default

OFF No options are passed to the preprocessor.

Description

This option passes one or more options to the preprocessor. If the preprocessor is not invoked, these options are ignored.

This option is equivalent to specifying option `-Qoption, fpp, options`.

Alternate Options

None

See Also

[Qoption](#) compiler option

x, Qx

Tells the compiler to generate optimized code specialized for the Intel processor that executes your program.

IDE Equivalent

Windows: **Code Generation > Intel Processor-Specific Optimization**

Optimization > Use Intel(R) Processor Extensions

Linux: None

Mac OS X: **Code Generation > Intel Processor-Specific Optimization**

Architectures

IA-32, Intel® 64 architectures

Syntax

Linux and Mac OS X: `-xprocessor`

Windows: `/Qxprocessor`

Arguments

processor Indicates the processor for which code is generated. Many of the following descriptions refer to Intel® Streaming SIMD Extensions (Intel® SSE) and Supplemental Streaming SIMD Extensions (Intel® SSSE). Possible values are:

Host	Can generate instructions for the highest instruction set and processor available on the compilation host.
SSE4.2	Can generate Intel® SSE4 Efficient Accelerated String and Text Processing instructions supported by Intel® Core™ i7 processors. Can generate Intel® SSE4 Vectorizing Compiler and Media Accelerator, Intel® SSSE3, SSE3, SSE2, and SSE instructions and it can optimize for the Intel® Core™ processor

family.

- SSE4 . 1 Can generate Intel® SSE4 Vectorizing Compiler and Media Accelerator instructions for Intel processors. Can generate Intel® SSSE3, SSE3, SSE2, and SSE instructions and it can optimize for Intel® 45nm Hi-k next generation Intel® Core™ microarchitecture. This replaces value S, which is [deprecated](#).
- SSE3_ATOM Can generate MOVBE instructions for Intel processors and it can optimize for the Intel® Atom™ processor and Intel® Centrino® Atom™ Processor Technology.
- SSSE3 Can generate Intel® SSSE3, SSE3, SSE2, and SSE instructions for Intel processors and it can optimize for the Intel® Core™2 Duo processor family. This replaces value T, which is [deprecated](#).
- SSE3 Can generate Intel® SSE3, SSE2, and SSE instructions for Intel processors and it can

optimize for processors based on Intel® Core™ microarchitecture and Intel NetBurst® microarchitecture. This replaces value P, which is [deprecated](#).

SSE2

Can generate Intel® SSE2 and SSE instructions for Intel processors, and it can optimize for Intel® Pentium® 4 processors, Intel® Pentium® M processors, and Intel® Xeon® processors with Intel® SSE2. This value is not available on Mac OS X systems. This replaces value N, which is [deprecated](#).

Default

Windows systems: For more information on the default values, see Arguments above, option `m` (Linux and Mac OS X) and option `arch` (Windows).
`/arch:SSE2`

Linux systems: -

`msse2`

Mac OS X systems
using IA-32

architecture: `SSE3`

Mac OS X systems
using Intel® 64

architecture: `SSSE3`

Description

This option tells the compiler to generate optimized code specialized for the Intel processor that executes your program. It also enables optimizations in addition to Intel processor-specific optimizations. The specialized code generated by this option may run only on a subset of Intel processors.

This option can enable optimizations depending on the argument specified. For example, it may enable Intel® Streaming SIMD Extensions 4 (Intel® SSE4), Intel® Supplemental Streaming SIMD Extensions 3 (Intel® SSSE3), Intel® Streaming SIMD Extensions 3 (Intel® SSE3), Intel® Streaming SIMD Extensions 2 (Intel® SSE2), or Intel® Streaming SIMD Extensions (Intel® SSE) instructions. The binaries produced by these values will run on Intel processors that support all of the features for the targeted processor. For example, binaries produced with SSE3 will run on an Intel® Core™ 2 Duo processor, because that processor completely supports all of the capabilities of the Intel® Pentium® 4 processor, which the SSE3 value targets. Specifying the SSSE3 value has the potential of using more features and optimizations available to the Intel® Core™ 2 Duo processor.

Do not use *processor* values to create binaries that will execute on a processor that is not compatible with the targeted processor. The resulting program may fail with an illegal instruction exception or display other unexpected behavior. For example, binaries produced with SSE3 may produce code that will not run on Intel® Pentium® III processors or earlier processors that do not support SSE2 instructions.

Compiling the function `main()` with any of the *processor* values produces binaries that display a fatal run-time error if they are executed on unsupported processors. For more information, see *Optimizing Applications*.

If you specify more than one *processor* value, code is generated for only the highest-performing processor specified. The highest-performing to lowest-performing *processor* values are: SSE4.2, SSE4.1, SSSE3, SSE3, SSE2. Note that *processor* value SSE3_ATOM does not fit within this group.

Compiler options `m` and `arch` produce binaries that should run on processors not made by Intel that implement the same capabilities as the corresponding Intel processors.

Previous value `O` is deprecated and has been replaced by option `-msse3` (Linux and Mac OS X) and option `/arch:SSE3` (Windows).

Previous values `W` and `K` are deprecated. The details on replacements are as follows:

- Mac OS X systems: On these systems, there is no exact replacement for `W` or `K`. You can upgrade to the default option `-msse3` (IA-32 architecture) or option `-mssse3` (Intel® 64 architecture).
- Windows and Linux systems: The replacement for `W` is `-msse2` (Linux) or `/arch:SSE2` (Windows). There is no exact replacement for `K`. However, on Windows systems, `/QxK` is interpreted as `/arch:IA32`; on Linux systems, `-xK` is interpreted as `-mia32`. You can also do one of the following:
 - Upgrade to option `-msse2` (Linux) or option `/arch:SSE2` (Windows). This will produce one code path that is specialized for Intel® SSE2. It will not run on earlier processors
 - Specify the two option combination `-mia32 -axSSE2` (Linux) or `/arch:IA32 /QaxSSE2` (Windows). This combination will produce an executable that runs on any processor with IA-32 architecture but with an additional specialized Intel® SSE2 code path.

The `-x` and `/Qx` options enable additional optimizations not enabled with option `-m` or option `/arch`.

Alternate Options

None

See Also

[ax, Qax](#) compiler option

[m](#) compiler option

[arch](#) compiler option

X

Removes standard directories from the include file search path.

IDE Equivalent

Windows: **Preprocessor > Ignore Standard Include Path** (/noinclude)

Linux: None

Mac OS X: **Preprocessor > Ignore Standard Include Path** (/noinclude)

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-X`

Windows: `/X`

Arguments

None

Default

OFF Standard directories are in the include file search path.

Description

This option removes standard directories from the include file search path. It prevents the compiler from searching the default path specified by the FPATH environment variable.

On Linux and Mac OS X systems, specifying `-X` (or `-noinclude`) prevents the compiler from searching in `/usr/include` for files specified in an `INCLUDE` statement.

You can use this option with the `I` option to prevent the compiler from searching the default path for include files and direct it to use an alternate path.

This option affects fpp preprocessor behavior and the `USE` statement.

Alternate Options

Linux and Mac OS X: `-nostdinc`

Windows: `/noinclude`

See Also

[I](#) compiler option

Xlinker

Passes a linker option directly to the linker.

IDE Equivalent

None

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-Xlinkeroption`

Windows: None

Arguments

option Is a linker option.

Default

OFF No options are passed directly to the linker.

Description

This option passes a linker option directly to the linker.

If `-Xlinker -shared` is specified, only `-shared` is passed to the linker and no special work is done to ensure proper linkage for generating a shared object. –

`xlinker` just takes whatever arguments are supplied and passes them directly to the linker.

If you want to pass compound options to the linker, for example `"-L $HOME/lib"`, you must use one of the following methods:

```
-Xlinker -L -Xlinker $HOME/lib
-Xlinker "-L $HOME/lib"
-Xlinker -L\ $HOME/lib
```

Alternate Options

None

See Also

[shared](#) compiler option

[link](#) compiler option

y

See [syntax-only](#).

g, Zi, Z7

Tells the compiler to generate full debugging information in the object file.

IDE Equivalent

Windows: **General > Debug Information Format** (/Z7, /Zd, /Zi)

Linux: None

Mac OS X: **General > Generate Debug Information** (-g)

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-g`

Windows: `/Zi`
 `/Z7`

Arguments

None

Default

OFF No debugging information is produced in the object file.

Description

This option tells the compiler to generate symbolic debugging information in the object file for use by debuggers.

The compiler does not support the generation of debugging information in assemblable files. If you specify this option, the resulting object file will contain debugging information, but the assemblable file will not.

This option turns off `O2` and makes `O0` (Linux and Mac OS X) or `Od` (Windows) the default unless `O2` (or another `O` option) is explicitly specified in the same command line.

On Linux systems using Intel® 64 architecture and Linux and Mac OS X systems using IA-32 architecture, specifying the `-g` or `-O0` option sets the `-fno-omit-frame-pointer` option.

For more information on `Zi` and `Z7`, see keyword `full` in [debug \(Windows*\)](#).

Alternate Options

Linux and Mac OS X: None

Windows: `/debug:full` (or `/debug`)

See Also

[Zd](#) compiler option

Zd

This option has been deprecated. Use keyword `minimal` in [debug \(Windows*\)](#).

zero, Qzero

Initializes to zero all local scalar variables of intrinsic type INTEGER, REAL, COMPLEX, or LOGICAL that are saved but not yet initialized.

IDE Equivalent

Windows: **Data > Initialize Local Saved Scalars to Zero**

Linux: None

Mac OS X: **Data > Initialize Local Saved Scalars to Zero**

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: `-zero`
`-nozero`

Windows: `/Qzero`
`/Qzero-`

Arguments

None

Default

`-nozero` Local scalar variables are not initialized to zero.

or

`/Qzero-`

Description

This option initializes to zero all local scalar variables of intrinsic type INTEGER, REAL, COMPLEX, or LOGICAL that are saved but not yet initialized.

Use `-save` (Linux and Mac OS X) or `/Qsave` (Windows) on the command line to make all local variables specifically marked as SAVE.

Alternate Options

None

See Also

[save](#) compiler option

g, Zi, Z7

Tells the compiler to generate full debugging information in the object file.

IDE Equivalent

Windows: **General > Debug Information Format** (/Z7, /Zd, /Zi)

Linux: None

Mac OS X: **General > Generate Debug Information** (-g)

Architectures

IA-32, Intel® 64, IA-64 architectures

Syntax

Linux and Mac OS X: -g

Windows: /Zi

 /Z7

Arguments

None

Default

OFF No debugging information is produced in the object file.

Description

This option tells the compiler to generate symbolic debugging information in the object file for use by debuggers.

The compiler does not support the generation of debugging information in assemblable files. If you specify this option, the resulting object file will contain debugging information, but the assemblable file will not.

This option turns off `o2` and makes `o0` (Linux and Mac OS X) or `od` (Windows) the default unless `o2` (or another `o` option) is explicitly specified in the same command line.

On Linux systems using Intel® 64 architecture and Linux and Mac OS X systems using IA-32 architecture, specifying the `-g` or `-O0` option sets the `-fno-omit-frame-pointer` option.

For more information on `zi` and `z7`, see keyword `full` in [debug \(Windows*\)](#).

Alternate Options

Linux and Mac OS X: None

Windows: `/debug:full` (or `/debug`)

See Also

[zd](#) compiler option

ZI

See keyword `none` in [libdir](#)

Zp

See keyword `recnbyte` in [align](#).

Zs

See [syntax-only](#).

Quick Reference Guides and Cross References

Quick Reference Guides and Cross References

The topic summarizes Intel® Fortran compiler options used on Windows* OS, Linux* OS and Mac OS* X.

If you want to see the summarized Windows* OS options, see [this topic](#).

If you want to see the summarized Linux* OS and Mac OS* X options, see [this topic](#).

Windows* OS Quick Reference Guide and Cross Reference

The table in this section summarizes Intel® Fortran compiler options used on Windows* OS . Each summary also shows the equivalent compiler options on Linux* OS and Mac OS* X.

If you want to see the summarized Linux* OS and Mac OS* X options, see [this table](#).

Some compiler options are only available on systems using certain architectures, as indicated by these labels:

Label	Meaning
i32	The option is available on systems using IA-32 architecture.
i64em	The option is available on systems using Intel® 64 architecture.
i64	The option is available on systems using IA-64 architecture.

If "only" appears in the label, the option is only available on the identified system or architecture.

If no label appears, the option is available on all supported systems and architectures.

For more details on the options, refer to the [Alphabetical Compiler Options](#) section.

The Intel® Fortran Compiler includes the Intel® Compiler Option Mapping tool. This tool lets you find equivalent options by specifying compiler option `-map-opts` (Linux and Mac OS X) or `/Qmap-opts` (Windows).

For information on conventions used in this table, see [Conventions](#).

Quick Reference of Windows* OS Options

The following table summarizes all supported Windows* OS options. It also shows equivalent Linux* OS and Mac OS* X options, if any.

Option	Description	Default	Equivalent Linux* OS
/1	Executes at least one iteration of DO loops.	OFF	-1
/4I{2 4 8}	Specifies the default KIND for integer and logical variables; same as the <code>/integer-size</code> option.	/4I4	-i{2 4 8}
/4L{72 80 132}	Treats the statement field of each fixed-form source line as ending in column 72, 80, or 132; same as the <code>/extend-source</code> option.	/4L72	-72, -80
/4Na, /4Ya	Determines where local variables are stored. /4Na is the same as <code>/save</code> . /4Ya is the same as <code>/automatic</code> .	/4Ya	None
/4Naltparam, /4Yaltparam	Determines whether alternate syntax is allowed for PARAMETER statements; same as the <code>/altparam</code> option).	/4Yaltparam	None
/4Nb, /4Yb	Determines whether checking is performed for run-time failures (same as the <code>/check</code> option).	/4Nb	None
/4Nd, /4Yd	Determines whether error	/4Nd	-warn no

Option	Description	Default	Equivalent Linux* OS
	messages are issued for undeclared symbols. /4Nd is the same as /warn:nodeclarations. /4Yd is the same as /warn:declarations.		warn dec
/4Nf, /4Yf	Specifies the format for source files. /4Nf is the same as /fixed. /4Yf is the same as /free.	/4Nf	-fixed, -
/4Nportlib, /4Yportlib	Determines whether the compiler links to the library of portability routines.	/4Yportlib	None
/4Ns, /4Ys	Determines whether the compiler changes warning messages about Fortran standards violations to error messages. /4Ns is the same as /stand:none. /4Ys is the same as /stand:f90.	/4Ns	-stand n
/4R8, /4R16	Specifies the default KIND for real and complex variables. /4R8 is the same as /real-size:64. /4R16 is the same as	OFF	-real-si size 128

Option	Description	Default	Equivalent Linux* OS
	<code>/real-size:128.</code>		
<code>/align[:keyword]</code>	Tells the compiler how to align certain data items.	keywords: nocommons nodcommons records nosequence	<code>-align [k</code>
<code>/allow:keyword</code>	Determines how the fpp preprocessor treats Fortran end-of-line comments in preprocessor directive lines.	keyword: <code>fpp_comments</code>	<code>/allow ke</code>
<code>/[no]altparam</code>	Allows alternate syntax (without parentheses) for PARAMETER statements.	<code>/altparam</code>	<code>-[no]altp</code>
<code>/arch:processor</code> (i32, i64em)	Tells the compiler to generate optimized code specialized for the processor that executes your program; same as option <code>/architecture</code> .	varies; see option description	<code>-arch pro</code> <code>processo</code> (i32, i64em)
<code>/asmattr:keyword</code>	Specifies the contents of an assembly listing file.	<code>/noasmattr</code>	None
<code>/asmfile[:file dir]</code>	Specifies that an assembly listing file should be generated.	<code>/noasmfile</code>	None
<code>/assume:keyword</code>	Tells the compiler to make	keywords:	<code>-assume k</code>

Option	Description	Default	Equivalent Linux* OS
	certain assumptions.	nobsc nobuffered_io nobyterecl nocc_omp nodummy_aliases nominus0 noold_boz old_unit_star old_xor protect_constants noprotect_parens norealloc_lhs source_include nostd_mod_proc_name nounderscore nowriteable-strings	
/[no]automatic	Causes all variables to be allocated to the run-time stack; same as the /auto option.	/Qauto-scalar	-[no]aut
/bigobj	Increases the number of sections that an object file can contain.	OFF	None
/bintext:string	Places a text string into the object file (.obj) being generated by the compiler.	OFF	None
/c	Prevents linking.	OFF	-c

Option	Description	Default	Equivalent Linux* OS
<code>/C</code>	Performs checking for all run-time failures; same as the <code>/check:all</code> option.	OFF	<code>-C</code>
<code>/CB</code>	Performs run-time checking on array subscript and character substring expressions; same as the <code>/check:bounds</code> option.	OFF	<code>-CB</code>
<code>/ccdefault:keyword</code>	Specifies the type of carriage control used when a file is displayed at a terminal screen.	<code>/ccdefault:default</code>	<code>-ccdefault</code>
<code>/check[:keyword]</code>	Checks for certain conditions at run time.	<code>/nocheck</code>	<code>-check [k]</code>
<code>/cm</code>	Disables all messages about questionable programming practices; same as specifying option <code>/warn:nousage</code> .	OFF	<code>-cm</code>
<code>/compile-only</code>	Causes the compiler to compile to an object file only and not link; same as the <code>/c</code> option.	OFF	None
<code>/Qcomplex-limited-range[-]</code>	Enables the use of basic algebraic expansions of some arithmetic operations	<code>/Qcomplex-limited-range-</code>	<code>-[no-]complex-range</code>

Option	Description	Default	Equivalent Linux* OS
	involving data of type COMPLEX.		
<code>/convert:keyword</code>	Specifies the format of unformatted files containing numeric data.	<code>/convert:native</code>	<code>-convert</code>
<code>/CU</code>	Enables run-time checking for uninitialized variables. This option is the same as <code>/check:uninit</code> and <code>/RTCu</code> .	OFF	<code>-CU</code>
<code>/Dname [=value]</code>	Defines a symbol name that can be associated with an optional value.	<code>/noD</code>	<code>-Dname [=value]</code>
<code>/[no]d-lines</code>	Compiles debugging statements indicated by the letter D in column 1 of the source code.	<code>/nod-lines</code>	<code>-[no]d-lines</code>
<code>/[no]dbglibs</code>	Tells the linker to search for unresolved references in a debug run-time library.	<code>/nodbglibs</code>	None
<code>/debug:keyword</code>	Specifies the type of debugging information generated by the compiler in the object file.	<code>/debug:full (IDE)</code> <code>/debug:none (command line)</code>	<code>-debug keyword</code> Note: the L X option takes <code>keyword</code> s
<code>/debug-parameters[:keyword]</code>	Tells the compiler to generate debug information	<code>/nodebug-parameters</code>	<code>-debug-p [keyword]</code>

Option	Description	Default	Equivalent Linux* OS
	for PARAMETERS used in a program.		
<code>/define:name [=value]</code>	Defines a symbol name that can be associated with an optional value; same as the <code>/D<name> [=value]</code> option.	OFF	None
<code>/dll</code>	Specifies that a program should be linked as a dynamic-link (DLL) library.	OFF	None
<code>/double-size:size</code>	Defines the default KIND for DOUBLE PRECISION and DOUBLE COMPLEX variables.	<code>/double-size:64</code>	<code>-double-</code>
<code>/E</code>	Causes the Fortran preprocessor to send output to stdout.	OFF	<code>-E</code>
<code>/EP</code>	Causes the Fortran preprocessor to send output to stdout, omitting #line directives.	OFF	<code>-EP</code>
<code>/error-limit:n</code>	Specifies the maximum number of error-level or fatal-level compiler errors allowed for a file specified on the command line.	<code>/error-limit:30</code>	<code>-error-l</code>

Option	Description	Default	Equivalent Linux* OS
<code>/exe:{file dir}</code>	Specifies the name for a built program or dynamic-link library.	OFF	-o
<code>/extend-source[:size]</code>	Specifies the length of the statement field in a fixed-form source file.	<code>/extend-source:72</code>	-extend-
<code>/extfor:ext</code>	Specifies file extensions to be processed by the compiler as Fortran files.	OFF	None
<code>/extfpp:ext</code>	Specifies file extensions to be recognized as a file to be preprocessed by the Fortran preprocessor.	OFF	None
<code>/extlnk:ext</code>	Specifies file extensions to be passed directly to the linker.	OFF	None
<code>/Fn</code>	Specifies the stack reserve amount for the program.	OFF	None
<code>/f66</code>	Tells the compiler to apply FORTRAN 66 semantics.	OFF	-f66
<code>/f77rtl</code>	Tells the compiler to use the run-time behavior of FORTRAN 77.	OFF	-f77rtl
<code>/Fa[:file dir]</code>	Specifies that an assembly listing file should be generated; same as option	OFF	-S

Option	Description	Default	Equivalent Linux* OS
	<i>/asmfile</i> and <i>/S</i> .		
<i>/FAc</i> , <i>/FAs</i> , <i>/FAcs</i>	Specifies the contents of an assembly listing file. <i>/FAc</i> is the same as the <i>/asmattr:machine</i> option. <i>/FAs</i> is the same as the <i>/asmattr:source</i> option. <i>/FAcs</i> is the same as the <i>/asmattr:all</i> option.	OFF	None
<i>/fast</i>	Maximizes speed across the entire program.	OFF	<i>-fast</i>
<i>/Fefile</i>	Specifies the name for a built program or dynamic-link library; same as the <i>/exe</i> option.	OFF	<i>-o</i>
<i>/FI</i>	Specifies source files are in determined by file suffix fixed format; same as the <i>/fixed</i> option.		<i>-FI</i>
<i>/[no]fixed</i>	Specifies source files are in determined by file suffix fixed format.		<i>-[no]fix</i>
<i>/[no]fltconsistency</i>	Enables improved floating-point consistency.	<i>/nofltconsistency</i>	<i>-[no]flt</i>
<i>/Fm[file]</i>	Tells the linker to generate a link map file; same as the <i>/map</i> option.	OFF	None

Option	Description	Default	Equivalent Linux* OS
<code>/Fofile</code>	Specifies the name for an object file; same as the <code>/object</code> option.	OFF	None
<code>/fp:keyword</code>	Controls the semantics of floating-point calculations.	<code>/fp:fast=1</code>	<code>-fp-mode</code>
<code>/[no]fpconstant</code>	Tells the compiler that single-precision constants assigned to double-precision variables should be evaluated in double precision.	<code>/nofpconstant</code>	<code>-[no]fpco</code>
<code>/fpe:n</code>	Specifies floating-point exception handling for the main program at run-time.	<code>/fpe:3</code>	<code>-fpen</code>
<code>/fpp [n]</code> <code>/fpp["option"]</code>	Runs the Fortran preprocessor on source files before compilation.	<code>/nofpp</code>	<code>- fpp[n]</code> <code>- fpp["o</code>
<code>/fpscomp[:keyword]</code>	Specifies compatibility with Microsoft* Fortran PowerStation or Intel® Fortran.	<code>/fpscomp:libs</code>	<code>-fpscomp</code>
<code>/FR</code>	Specifies source files are in free format; same as the <code>/free</code> option.	determined by file suffix	<code>-FR</code>
<code>/[no]free</code>	Specifies source files are in free format.	determined by file suffix	<code>-[no]fre</code>

Option	Description	Default	Equivalent Linux* OS
/G2 (i64 only)	Optimizes application performance for systems using IA-64 architecture.	OFF	None
/G2-p9000 (i64 only)	Optimizes for Dual-Core Intel® Itanium® 2 Processor 9000 series.	ON	-mtune i
/G{5 6 7} (i32, i64em)	Optimizes application performance for systems using IA-32 architecture and Intel® 64 architecture. These options have been deprecated.	/G7	None
/GB	Optimizes for Intel® Pentium® Pro, Pentium® II and Pentium® III processors; same as the /G6 option.	OFF	None
/Ge	Enables stack-checking for all functions. Deprecated.	OFF	None
/gen-interfaces[:[no]source]	Tells the compiler to generate an interface block for each routine in a source file.	/nogen-interfaces	-gen-int [[no]sou
/Gm	Tells the compiler to use calling convention CVF; same as the /iface:cvf	OFF	None

Option	Description	Default	Equivalent Linux* OS
	option.		
/Gs[n]	Disables stack-checking for routines with a specified number of bytes of local variables and compiler temporaries.	/Gs4096	None
/GS[-]	Determines whether the compiler generates code that detects some buffer overruns.	/GS-	-f[no-]scheck
/Gz	Tells the compiler to use STDCALL; same as the /iface:stdcall option.	OFF	None
/heap-arrays[:size]	Puts automatic arrays and arrays created for temporary computations on the heap instead of the stack.	/heap-arrays-	-heap-ar
/help[category]	Displays all available compiler options or a category of compiler options; same as the /? option.	OFF	-help
/homeparams	Tells the compiler to store parameters passed in	OFF	None

Option	Description	Default	Equivalent Linux* OS
	registers to the stack.		
<code>/I:dir</code>	Specifies a directory to add to the include path.	OFF	<code>-Idir</code>
<code>/iface:keyword</code>	Specifies the default calling convention and argument-passing convention for an application.	<code>/iface:default</code>	None
<code>/include</code>	Specifies a directory to add to the include path; same as the <code>/I</code> option.	OFF	None
<code>/inline[:keyword]</code>	Specifies the level of inline function expansion.	OFF	None
<code>/[no]intconstant</code>	Tells the compiler to use FORTRAN 77 semantics to determine the kind parameter for integer constants.	<code>/nointconstant</code>	<code>-[no]int</code>
<code>/integer-size:size</code>	Specifies the default KIND for integer and logical variables.	<code>/integer-size:32</code>	<code>-integer</code>
<code>/LD</code>	Specifies that a program should be linked as a dynamic-link (DLL) library.	OFF	None
<code>/libdir[:keyword]</code>	Controls whether linker options for search libraries are included in object files	<code>/libdir:all</code>	None

Option	Description	Default	Equivalent Linux* OS
	generated by the compiler.		
<code>/libs:keyword</code>	Tells the linker to search for unresolved references in a specific run-time library.	<code>/libs:static</code>	None
<code>/link</code>	Passes options to the linker at compile time.	OFF	None
<code>/[no]logo</code>	Displays the compiler version information.	<code>/logo</code>	<code>-[no]log</code>
<code>/map[:file]</code>	Tells the linker to generate a link map file.	<code>/nomap</code>	None
<code>/MD and /MDd</code>	Tells the linker to search for unresolved references in a multithreaded, debug, dynamic-link run-time library.	OFF	None
<code>/MDs and /MDsd</code>	Tells the linker to search for unresolved references in a single-threaded, dynamic-link run-time library.	OFF	None
<code>/MG</code>	Tells the compiler to create a graphics or Fortran Windows application and link against the most commonly used libraries.	OFF	None
<code>/ML and /MLd</code>	Tells the linker to search for unresolved references in a	<code>/ML i32, i64: /ML i64em: OFF</code>	None

Option	Description	Default	Equivalent Linux* OS
	single-threaded, static run-time library.		
<i>/module:path</i>	Specifies the directory where module files should be placed when created and where they should be searched for.	OFF	<i>-module p</i>
<i>/MP[:n]</i>	Creates multiple processes that can be used to compile large numbers of source files at the same time.	OFF	<i>-multipl</i>
<i>/MT and /MTd</i>	Tells the linker to search for unresolved references in a multithreaded, static run-time library.	i32, i64: OFF i64em: <i>/MT/noreentrancy</i>	None
<i>/MW</i>	Tells the linker to search for unresolved references in a Fortran QuickWin library; same as <i>/libs:qwin</i> .	OFF	None
<i>/MWs</i>	Tells the linker to search for unresolved references in a Fortran standard graphics library; same as <i>/libs:qwins</i> .	OFF	None
<i>/names:keyword</i>	Specifies how source code identifiers and external	<i>/names:uppercase</i>	<i>-names ke</i>

Option	Description	Default	Equivalent Linux* OS
	names are interpreted.		
/nbs	Tells the compiler to treat the backslash character (\) as a normal character in character literals; same as the /assume:nobscc option.	/nbs	-nbs
/noinclude	Removes standard directories from the include file search path; same as the /x option.	OFF	None
/O[<i>n</i>]	Specifies the code optimization for applications.	/O2	-O[<i>n</i>]
/Ob <i>n</i>	Specifies the level of inline function expansion. <i>n</i> = 0, 1, or 2.	/Ob2 if /O2 is in effect or /O3 is specified /Ob0 if /Od is specified	-inline-
/object: <i>file</i>	Specifies the name for an object file.	OFF	None
/Od	Disables optimizations.	OFF	-O0
/Og	Enables global optimizations.	/Og	None
/Op	Enables improved floating-point consistency.	OFF	-mp
/optimize: <i>n</i>	Affects optimizations	/optimize:3 or	-On

Option	Description	Default	Equivalent Linux* OS
	performed by the compiler; <code>/optimize:4</code> <code>n = 1, 2, 3, or 4.</code>		
<code>/Os</code>	Enables optimizations that do not increase code size and produces smaller code size than O2.	OFF (unless <code>/O1</code> is specified)	<code>-Os</code>
<code>/Ot</code>	Enables all speed optimizations.	<code>/Ot</code> (unless <code>/O1</code> is specified)	None
<code>/Ox</code>	Same as the <code>/O2</code> option.	<code>/Ox</code>	<code>-O2</code>
<code>/Oy[-]</code> (i32 only)	Determines whether EBP is used as a general-purpose register in optimizations.	<code>/Oy</code> (unless <code>/Od</code> is specified)	<code>-f[no-]pointer</code> (i32, i64em)
<code>/P</code>	Causes the Fortran preprocessor to send output to a file, which is named by default; same as the <code>-preprocess-only</code> option.	OFF	<code>-P</code>
<code>/[no]pad-source</code>	Specifies padding for fixed-form source records.	<code>/nopad-source</code>	<code>-[no]pad</code>
<code>/pdbfile[:file]</code>	Specifies that any debug information generated by the compiler should be saved to a program database file.	<code>/nopdbfile</code>	None
<code>/preprocess-only</code>	Causes the Fortran	<code>/nopreprocess-only</code>	<code>-preproc</code>

Option	Description	Default	Equivalent Linux* OS
	preprocessor to send output to a file.		
<code>/Qansi-alias</code>	Tells the compiler to assume the program adheres to the Fortran Standard type aliasability rules.	<code>/Qansi-alias</code>	<code>-ansi-al</code>
<code>/Qauto</code>	Causes all variables to be allocated on the stack, rather than in local static storage.	<code>/Qauto-scalar</code>	<code>-auto</code>
<code>/Qauto-scalar</code>	Causes allocation of scalar variables of intrinsic types INTEGER, REAL, COMPLEX, and LOGICAL to the run-time stack.	<code>/Qauto-scalar</code>	<code>-auto-sc</code>
<code>/Qautodouble</code>	Makes default real and complex variables 8 bytes long; same as the <code>/real-size:64</code> option.	OFF	<code>-autodou</code>
<code>/Qaxp</code> (i32, i64em)	Tells the compiler to generate multiple, processor-specific auto-dispatch code paths for Intel processors if there is a performance benefit.	OFF	<code>-axp</code> (i32, i64em)

Option	Description	Default	Equivalent Linux* OS
<code>/Qchkstk[-]</code> (i64 only)	Enables stack probing when the stack is dynamically expanded at run-time.	<code>/Qchkstk</code>	None
<code>/Qcommon-args</code>	Tells the compiler that dummy (formal) arguments to procedures share memory locations with other dummy arguments or with COMMON variables that are assigned; same as <code>/assume:dummy_aliases</code> .	OFF	<code>-common-</code>
<code>/Qcomplex-limited-range[-]</code>	Enables the use of basic algebraic expansions of some arithmetic operations involving data of type COMPLEX.	<code>/Qcomplex-limited-range-</code>	<code>-[no-]com-range</code>
<code>/Qcpp</code>	Runs the Fortran preprocessor on source files before compilation; same as the <code>/fpp</code> option.	OFF	<code>-cpp</code>
<code>/Qd-lines[-]</code>	Compiles debugging statements indicated by the letter D in column 1 of the source code; can also be specified as <code>/d-lines</code> .	OFF	<code>-[no]d-l</code>
<code>/Qdiag-type:diag-list</code>	Controls the display of	OFF	<code>-diag-ty</code>

Option	Description	Default	Equivalent Linux* OS
	diagnostic information.		
/Qdiag-dump	Tells the compiler to print all enabled diagnostic messages and stop compilation.	OFF	-diag-dump
/Qdiag-enable:sv-include	Tells the Static Verifier to analyze include files and source files when issuing diagnostic message.	OFF	-diag-enable include
/Qdiag-error-limit:n	Specifies the maximum number of errors allowed before compilation stops.	/Qdiag-error-limit:30	-diag-error-limit:n
/Qdiag-file[:file]	Causes the results of diagnostic analysis to be output to a file.	OFF	-diag-file[:file]
/Qdiag-file-append[:file]	Causes the results of diagnostic analysis to be appended to a file.	OFF	-diag-file-append[:file]
/Qdiag-id-numbers[-]	Tells the compiler to display diagnostic messages by using their ID number values.	/Qdiag-id-numbers	-[no-]diag-id-numbers
/Qdiag-once:id[,id,...]	Tells the compiler to issue one or more diagnostic messages only once	OFF	-diag-once:id[,id,...]
/Qdps	Specifies that the alternate	/Qdps	-dps

Option	Description	Default	Equivalent Linux* OS
	syntax for PARAMETER statements is allowed; same as the /altparam option.		
/Qdyncom "common1,common2,..."	Enables dynamic allocation of common blocks at run time.	OFF	-dyncom "common1
/Qextend-source	This is the same as specifying option /extend-source:132.	OFF	-extend-
/Qfast-transcendentals[-]	Enables the compiler to replace calls to transcendental functions with faster but less precise implementations.	/Qfast-transcendentals	-[no-]fast-transcend
/Qfma[-] (i64 only)	Enables the combining of floating-point multiplies and add/subtract operations.	/Qfma	-[no-]fma (i64 only; L
/Qfnalign[:n] (i32, i64em)	Tells the compiler to align functions on an optimal byte boundary.	/Qfnalign-	-falign- (i32, i64em
/Qfnsplit[-] (i32, i64)	Enables function splitting.	/Qfnsplit-	-[no-]fn (i64 only; L
/Qfp-port[-] (i32, i64em)	Rounds floating-point results after floating-point operations.	/Qfp-port-	-[no-]fp (i32, i64em

Option	Description	Default	Equivalent Linux* OS
<code>/Qfp-relaxed[-]</code> (i64 only)	Enables use of faster but slightly less accurate code sequences for math functions, such as divide and sqrt.	<code>/Qfp-relaxed-</code>	<code>-[no-]fp</code> (i64 only; L
<code>/Qfp-speculation:mode</code>	Tells the compiler the mode in which to speculate on floating-point operations.	<code>/Qfp-speculation:fast</code>	<code>-fp-spec</code>
<code>/Qfp-stack-check</code> (i32, i64em)	Generates extra code after every function call to ensure that the FP (floating-point) stack is in the expected state.	OFF	<code>-fp-stack</code> (i32, i64em)
<code>/Qfpp[n]</code>	Runs the Fortran preprocessor on source files before compilation.	<code>/nofpp</code>	<code>-fpp[n]</code>
<code>/Qftz[-]</code>	Flushes denormal results to zero.	i64: <code>/Qftz-</code> i32, i64em: <code>/Qftz</code>	<code>-[no-]ftz</code>
<code>/Qglobal-hoist[-]</code>	Enables certain optimizations that can move memory loads to a point earlier in the program execution than where they appear in the source.	<code>/Qglobal-hoist-</code>	<code>-[no-]gl</code>
<code>/QIA64-fr32</code> (i64 only)	Disables use of high floating-point registers.	OFF	None

Option	Description	Default	Equivalent Linux* OS
<code>/QIfist</code> (i32 only)	Enables fast float-to-integer conversions; same as option <code>/Qrcd</code> .	OFF	None
<code>/Qinline-debug-info</code>	Produces enhanced source position information for inlined code.	OFF	<code>-inline-</code>
<code>/Qinline-dllimport[-]</code>	Determines whether <code>dllimport</code> functions are inlined.	<code>/Qinline-dllimport</code>	None
<code>/Qinline-factor==n</code>	Specifies the percentage multiplier that should be applied to all inlining options that define upper limits.	<code>/Qinline-factor-</code>	<code>-inline-</code>
<code>/Qinline-forceinline</code>	Specifies that an inline routine should be inlined whenever the compiler can do so.	OFF	<code>-inline-</code>
<code>/Qinline-max-per-compile=n</code>	Specifies the maximum number of times inlining may be applied to an entire compilation unit.	<code>/Qinline-max-per-compile-</code>	<code>-inline-compile=</code>
<code>/Qinline-max-per-routine=n</code>	Specifies the maximum number of times the inliner may inline into a particular routine.	<code>/Qinline-max-per-routine-</code>	<code>-inline-routine=</code>

Option	Description	Default	Equivalent Linux* OS
<code>/Qinline-max-size=<i>n</i></code>	Specifies the lower limit for the size of what the inliner considers to be a large routine.	<code>/Qinline-max-size-</code>	<code>-inline-</code>
<code>/Qinline-max-total-size=<i>n</i></code>	Specifies how much larger a routine can normally grow when inline expansion is performed.	<code>/Qinline-max-total-size-</code>	<code>-inline-size=<i>n</i></code>
<code>/Qinline-min-size=<i>n</i></code>	Specifies the upper limit for the size of what the inliner considers to be a small routine.	<code>/Qinline-min-size-</code>	<code>-inline-</code>
<code>/Qinstruction=[no]movbe</code>	Determines whether MOVBE instructions are generated for Intel® processors.	<code>/Qinstruction:movbe</code>	<code>-minstruc</code>
<code>/Qinstrument-functions[-]</code>	Determines whether function entry and exit points are instrumented.	<code>/Qinstrument-functions-</code>	<code>-f[no-]in</code> <code>function</code>
<code>/Qip[-]</code>	Enables additional single-file interprocedural optimizations.	OFF	<code>-[no-]ip</code>
<code>/Qip-no-inlining</code>	Disables full and partial inlining enabled by <code>-ip</code> .	OFF	<code>-ip-no-i</code>
<code>/Qip-no-pinlining</code> (i32, i64em)	Disables partial inlining.	OFF	<code>-ip-no-p</code> (i32, i64em)

Option	Description	Default	Equivalent Linux* OS
<code>/QIPF-flt-eval-method0</code> (i64 only)	Tells the compiler to evaluate the expressions involving floating-point operands in the precision indicated by the variable types declared in the program. Deprecated.	OFF	<code>-[no-]IPF-method0</code> (i64 only; L
<code>/QIPF-fltacc[-]</code> (i64 only)	Tells the compiler to apply optimizations that affect floating-point accuracy. Deprecated.	<code>/QIPF-fltacc-</code>	<code>-IPF-fltacc-</code> (i64 only; L
<code>/QIPF-fma</code> (i64 only)	Enables the combining of floating-point multiplies and add/subtract operations. Deprecated; use <code>/Qfma</code> .	<code>/QIPF-fma</code>	<code>-IPF-fma</code> (i64 only; L
<code>/QIPF-fp-relaxed</code> (i64 only)	Enables use of faster but slightly less accurate code sequences for math functions, such as divide and sqrt. Deprecated; use <code>/Qfp-relaxed</code> .	<code>/QIPF-fp-relaxed-</code>	<code>-IPF-fp-relaxed-</code> (i64 only; L
<code>/Qipo[n]</code>	Enables multifile IP optimizations between files.	OFF	<code>-ipo[n]</code>
<code>/Qipo-c</code>	Generates a multifile object file that can be used in further link steps.	OFF	<code>-ipo-c</code>

Option	Description	Default	Equivalent Linux* OS
/Qipo-jobs: <i>n</i>	Specifies the number of commands to be executed simultaneously during the link phase of Interprocedural Optimization (IPO).	/Qipo-jobs:1	-ipo-jobs
/Qipo-S	Generates a multifile assembly file that can be used in further link steps.	OFF	-ipo-S
/Qipo-separate	Generates one object file per source file.	OFF	-ipo-sep (Linux only)
/Qivdep-parallel (i64 only)	Tells the compiler that there is no loop-carried memory dependency in any loop following an IVDEP directive.	OFF	-ivdep-p (i64 only; L
/Qkeep-static-consts[-]	Tells the compiler to preserve allocation of variables that are not referenced in the source.	/Qkeep-static- consts-	-f[no-]k consts
/Qlocation, <i>string,dir</i>	Specifies a directory as the location of the specified tool in <i>string</i> .	OFF	-Qlocati
/Qlowercase	Causes the compiler to ignore case differences in identifiers and to convert	OFF	-lowerca

Option	Description	Default	Equivalent Linux* OS
	external names to lowercase; same as the <code>/names:lowercase</code> option.		
<code>/Qmap-opts</code>	Maps one or more Windows* compiler options to their equivalent on a Linux* system (or vice versa).	OFF	<code>-map-opts</code>
<code>/Qnobss-init</code>	Places any variables that are explicitly initialized with zeros in the DATA section.	OFF	<code>-no-bss-</code>
<code>/Qonetrip</code>	This is the same as specifying option <code>/onetrip</code> .	OFF	<code>-onetrip</code>
<code>/Qopenmp</code>	Enables the parallelizer to generate multithreaded code based on OpenMP* directives.	OFF	<code>-openmp</code>
<code>/Qopenmp-lib:type</code>	Lets you specify an OpenMP* run-time library to use for linking.	<code>/Qopenmp-lib:legacy</code>	<code>-openmp-</code> (Linux only)
<code>/Qopenmp-link:library</code>	Lets you specify an OpenMP* run-time library to use for linking.	<code>/Qopenmp-lib:legacy</code>	<code>-openmp-</code> (Linux only)
<code>/Qopenmp-profile</code>	Enables analysis of	OFF	<code>-openmp-</code>

Option	Description	Default	Equivalent Linux* OS
	OpenMP* applications.		(Linux only)
<code>/Qopenmp-report[n]</code>	Controls the OpenMP parallelizer's level of diagnostic messages.	<code>/Qopenmp-report1</code>	<code>-openmp-</code>
<code>/Qopenmp-stubs</code>	Enables compilation of OpenMP programs in sequential mode.	OFF	<code>-openmp-</code>
<code>/Qopenmp-threadprivate:type</code>	Lets you specify an OpenMP* threadprivate implementation.	<code>/Qopenmp-threadprivate:legacy type</code>	<code>-openmp-</code> (Linux only)
<code>/Qopt-block-factor:n</code>	Lets you specify a loop blocking factor.	OFF	<code>-opt-blo</code>
<code>/Qopt-jump-tables:keyword</code>	Enables or disables generation of jump tables for switch statements.	<code>/Qopt-jump-tables:default</code>	<code>-opt-jump-tables=k</code>
<code>/Qopt-loadpair[-]</code> (i64 only)	Enables or disables loadpair optimization.	<code>/Qopt-loadpair-</code>	<code>-[no-]op</code> (i64 only; L
<code>/Qopt-mem-bandwidthn</code> (i64 only)	Enables performance tuning and heuristics that control memory bandwidth use among processors.	<code>/Qopt-mem-bandwidth0</code> <code>/Qopt-mem-bandwidth1</code>	<code>-opt-mem</code> (i64 only; L
<code>/Qopt-mod-versioning[-]</code> (i64 only)	Enables or disables versioning of modulo operations for certain types of operands.	<code>/Qopt-mod-versioning-</code>	<code>-[no-]op</code> versioni (i64 only; L
<code>/Qopt-multi-version-</code>	Tells the compiler to use	<code>/Qopt-multi-version-</code>	<code>-[no-]op</code>

Option	Description	Default	Equivalent Linux* OS
<code>aggressive[-]</code> (i32, i64em)	aggressive multi-versioning to check for pointer aliasing and scalar replacement.	<code>aggressive-</code>	<code>version-</code> (i32, i64em)
<code>/Qopt-prefetch[:n]</code>	Enables prefetch insertion optimization.	i64: <code>/Qopt-prefetch</code> i32, i64em: <code>/Qopt-prefetch-</code>	<code>-opt-pre</code>
<code>/Qopt-prefetch-initial-values[-]</code> (i64 only)	Enables or disables prefetches that are issued before a loop is entered.	<code>/Qopt-prefetch-initial-values</code>	<code>-[no-]op</code> <code>initial-</code> (i64 only; L
<code>/Qopt-prefetch-issue-excl-hint[-]</code> (i64 only)	Determines whether the compiler issues prefetches for stores with exclusive hint.	<code>/Qopt-prefetch-issue-excl-hint-</code>	<code>-[no-]op</code> <code>issue-ex</code> (i64 only; L
<code>/Qopt-prefetch-next-iteration[-][:n]</code> (i64 only)	Enables or disables prefetches for a memory access in the next iteration of a loop.	<code>-opt-prefetch-next-iteration</code>	<code>/Qopt-pr</code> <code>iteration</code> (i64 only; L
<code>/Qopt-ra-region-strategy[:keyword]</code> (i32, i64em)	Selects the method that the register allocator uses to partition each routine into regions.	<code>/Qopt-ra-region-strategy:default</code>	<code>-opt-ra-</code> <code>strategy</code> (i32, i64em)
<code>/Qopt-report:n</code>	Tells the compiler to generate an optimization report to <i>stderr</i> .	<code>/Qopt-report:2</code>	<code>-opt-rep</code>
<code>/Qopt-report-file:file</code>	Tells the compiler to generate an optimization	OFF	<code>-opt-rep</code>

Option	Description	Default	Equivalent Linux* OS
	report named <i>file</i> .		
/Qopt-report-help	Displays the optimizer phases available for report generation.	OFF	-opt-rep
/Qopt-report-phase: <i>phase</i>	Specifies an optimizer phase to use when optimization reports are generated.	OFF	-opt-rep phase= <i>ph</i>
/Qopt-report-routine: <i>string</i>	Generates a report on all routines or the routines containing the specified <i>string</i> .	OFF	-opt-rep routine= <i>st</i>
/Qopt-streaming-stores: <i>keyword</i> (i32, i64em)	Enables generation of streaming stores for optimization.	/Qopt-streaming-stores:auto	-opt-str keyword (i32, i64em)
/Qopt-subscript-in-range[-] (i32, i64em)	Determines whether the compiler assumes no overflows in the intermediate computation of subscript expressions in loops.	/Qopt-subscript-in-range-	-[no-]op in-range (i32, i64em)
/Qoption, <i>string,options</i>	Passes <i>options</i> to the tool specified in <i>string</i> .	OFF	-Qoption,
/Qpad	Enables the changing of the variable and array memory layout.	/Qpad-	-pad

Option	Description	Default	Equivalent Linux* OS
<code>/Qpad-source</code>	This is the same as specifying option <code>/pad-source</code> .	<code>/Qpad-source-</code>	<code>-pad-sou</code>
<code>/Qpar-adjust-stack:n</code> (i32, i64em)	Tells the compiler to generate code to adjust the stack size for a fiber-based main thread.	<code>/Qpar-adjust-stack:0</code>	None
<code>/Qpar-report[n]</code>	Controls the diagnostic information reported by the auto-parallelizer.	<code>/Qpar-report1</code>	<code>-par-rep</code>
<code>/Qpar-runtime-control[-]</code>	Generates code to perform run-time checks for loops that have symbolic loop bounds.	<code>/Qpar-runtime-control-</code>	<code>-[no-]pa control</code>
<code>/Qpar-schedule-keyword</code> <code>[[:]n]</code>	Specifies a scheduling algorithm for DO loop iterations.	OFF	<code>-par-sch keyword[=</code>
<code>/Qpar-threshold[[:]n]</code>	Sets a threshold for the auto-parallelization of loops.	<code>/Qpar-threshold:100</code>	<code>-par-thr</code>
<code>/Qparallel</code>	Tells the auto-parallelizer to generate multithreaded code for loops that can be safely executed in parallel.	OFF	<code>-paralle</code>
<code>/Qpcn</code> (i32, i64em)	Enables control of floating-point significand precision.	<code>/Qpc64</code>	<code>-pcn</code> (i32, i64em)

Option	Description	Default	Equivalent Linux* OS
/Qprec	Improves floating-point precision and consistency.	OFF	-mp1
/Qprec-div[-]	Improves precision of floating-point divides.	/Qprec-div-	-[no-]pr
/Qprec-sqrt (i32, i64em)	Improves precision of square root implementations.	/Qprec-sqrt-	-prec-sq (i32, i64em)
/Qprefetch	Enables prefetch insertion optimization. Deprecated; use /Qopt-prefetch	i64: /Qprefetch i32, i64em: /Qprefetch-	-prefetch
/Qprof-data-order[-]	Enables or disables data ordering if profiling information is enabled.	/Qprof-data-order-	-[no-]pr (Linux only)
/Qprof-dir <i>dir</i>	Specifies a directory for profiling information output files.	OFF	-prof-di
/Qprof-file <i>file</i>	Specifies a file name for the profiling summary file.	OFF	-prof-fi
/Qprof-func-order[-]	Enables or disables function ordering if profiling information is enabled.	/Qprof-func-order-	-[no-]pr (Linux only)
/Qprof-gen[: <i>keyword</i>]	Produces an instrumented object file that can be used in profile-guided optimization.	/Qprof-gen-	-prof-gen
/Qprof-genx	Produces an instrumented	OFF	-prof-gen

Option	Description	Default	Equivalent Linux* OS
	object file that includes extra source position information. Deprecated; use <code>/Qprof-gen:srcpos</code> .		
<code>/Qprof-hotness-threshold:n</code>	Lets you set the hotness threshold for function grouping and function ordering.	OFF	<code>-prof-hotness-threshold:n</code> (Linux only)
<code>/Qprof-src-dir[-]</code>	Determines whether directory information of the source file under compilation is considered when looking up profile data records.	<code>/Qprof-src-dir</code>	<code>-[no-]prof-src-dir</code>
<code>/Qprof-src-root:dir</code>	Lets you use relative directory paths when looking up profile data and specifies a directory as the base.	OFF	<code>-prof-src-root:dir</code>
<code>/Qprof-src-root-cwd</code>	Lets you use relative directory paths when looking up profile data and specifies the current working directory as the base.	OFF	<code>-prof-src-root-cwd</code>
<code>/Qprof-use[:arg]</code>	Enables the use of profiling	OFF	<code>-prof-use[:arg]</code>

Option	Description	Default	Equivalent Linux* OS
	information during optimization.		
<code>/Qrcd</code> (i32, i64em)	Enables fast float-to-integer conversions.	OFF	<code>-rcd</code> (i32, i64em)
<code>/Qrct</code> (i32, i64em)	Sets the internal FPU rounding control to Truncate.	OFF	<code>-rct</code> (i32, i64em)
<code>/Qsafe-cray-ptr</code>	Tells the compiler that Cray* pointers do not alias other variables.	OFF	<code>-safe-cray</code>
<code>/Qsave</code>	Causes variables to be placed in static memory.	<code>/Qauto-scalar</code>	<code>-save</code>
<code>/Qsave-temps[-]</code>	Tells the compiler to save intermediate files created during compilation.	<code>.obj</code> files are saved	<code>-[no-]save</code>
<code>/Qscalar-rep[-]</code> (i32 only)	Enables scalar replacement performed during loop transformation (requires <code>/O3</code>).	<code>/Qscalar-rep-</code>	<code>-[no-]scalar-rep</code> (i32 only)
<code>/Qsfalign[n]</code> (i32 only)	Specifies stack alignment for functions. <i>n</i> is 8 or 16.	<code>/Qsfalign8</code>	None
<code>/Qsox[-]</code>	Tells the compiler to save the compiler options and version number in the executable.	<code>/Qsox-</code>	<code>-[no-]sox</code>
<code>/Qtcheck</code>	Enables analysis of	OFF	<code>-tcheck</code>

Option	Description	Default	Equivalent Linux* OS
	threaded applications.		(Linux only)
/Qtcollect[: <i>lib</i>]	Inserts instrumentation probes calling the Intel(R) Trace Collector API.	OFF	-tcollect (Linux only)
/Qtcollect-filter[= <i>file</i>]	Lets you enable or disable the instrumentation of specified functions.	OFF	-tcollect (Linux only)
/Qtprofile	Generates instrumentation to analyze multi-threading performance.	OFF	-tprofile (Linux only)
/Qtrapuv	Initializes stack local variables to an unusual value.	OFF	-ftrapuv
/Qunroll[: <i>n</i>]	Tells the compiler the maximum number of times to unroll loops; same as the /unroll[: <i>n</i>] option.	/Qunroll	-unroll[=
/Qunroll-aggressive[-] (i32, i64em)	Determines whether the compiler uses more aggressive unrolling for certain loops.	/Qunroll-aggressive-	-[no-]unroll-aggressive (i32, i64em)
/Quppercase	Causes the compiler to ignore case differences in identifiers and to convert external names to uppercase; same as the	/Quppercase	-uppercase

Option	Description	Default	Equivalent Linux* OS
	/names:uppercase option.		
/Quse-asm[-] (i64 only)	Tells the compiler to produce objects through the assembler.	/Quse-asm-	-[no-]use-asm
/Quse-msasm-symbols (i32,i64em)	Tells the compiler to use a dollar sign ("\$\$") when producing symbol names.	OFF	None
/Quse-vcdebug (i32 only)	Tells the compiler to issue debug information compatible with the Visual C++ debugger.	OFF	None
/Qvc7.1 /Qvc8 /Qvc9 (i32, i64em)	Specifies compatibility with Microsoft* Visual C++ or Microsoft* Visual Studio.	varies; see option description	None
/Qvec[-] (i32, i64em)	Enables or disables vectorization and transformations enabled for vectorization.	/Qvec	-[no-]vectorize (i32, i64em)
/Qvec-guard-write[-] (i32, i64em)	Tells the compiler to perform a conditional check in a vectorized loop.	/Qvec-guard-write-	-[no-]vectorize-guard-write (i32, i64em)
/Qvec-report[n] (i32, i64em)	Controls the diagnostic information reported by the vectorizer.	/Qvec-report1	-vec-report (i32, i64em)

Option	Description	Default	Equivalent Linux* OS
<code>/Qvms</code>	Causes the run-time system to behave like HP Fortran for OpenVMS* Alpha systems and VAX* systems (VAX FORTRAN*) in certain ways; same as the <code>/vms</code> option.	<code>/novms</code>	<code>-vms</code>
<code>/Qxprocessor</code> (i32, i64em)	Tells the compiler to generate optimized code specialized for the Intel processor that executes your program.	varies; see the option description	<code>-xproces</code> (i32, i64em)
<code>/Qzero[-]</code>	Initializes to zero all local scalar variables of intrinsic type INTEGER, REAL, COMPLEX, or LOGICAL that are saved but not yet initialized.	OFF	<code>-[no]zer</code>
<code>/real-size: size</code>	Specifies the default KIND for real variables.	<code>/real-size:32</code>	<code>-real-si</code>
<code>/[no]recursive</code>	Tells the compiler that all routines should be compiled for possible recursive execution.	<code>/norecursive</code>	<code>-[no]rec</code>
<code>/reentrancy: keyword</code>	Tells the compiler to generate reentrant code to support a multithreaded	<code>/noreentrancy</code>	<code>-reentra</code>

Option	Description	Default	Equivalent Linux* OS
	application.		
/RTCu	Enables run-time checking for uninitialized variables; same as option /check:uninit.	OFF	-check un
/S	Causes the compiler to compile to an assembly file only and not link.	OFF	-S
/source:file	Tells the compiler to compile the file as a Fortran source file.	OFF	None
/stand:keyword	Tells the compiler to issue compile-time messages for nonstandard language elements.	/nostand	-stand ke
/static	Prevents linking with shared libraries.	/static	-static (Linux only)
/syntax-only	Tells the compiler to check only for correct syntax.	OFF	-syntax-
/Tf file	Tells the compiler to compile the file as a Fortran source file; same as the /source option.	OFF	-Tf file
/[no]threads	Tells the linker to search for unresolved references in a multithreaded run-time	i32, i64: /nothreads i64em: /threads	-[no]thr

Option	Description	Default	Equivalent Linux* OS
	library.		
<code>/[no]traceback</code>	Tells the compiler to generate extra information in the object file to provide source file traceback information when a severe error occurs at run time.	<code>/notraceback</code>	<code>-[no]tra</code>
<code>/tune:keyword</code> (i32, i64em)	Determines the version of the architecture for which the compiler generates instructions.	<code>/tune:pn4</code>	<code>-tune key</code> (i32, i64em)
<code>/u</code>	Undefines all previously defined preprocessor values.	OFF	None Note: the L X option <code>-u</code>
<code>/Uname</code>	Undefines any definition currently in effect for the specified symbol; same as the <code>/undefine</code> option.	OFF	<code>-Uname</code>
<code>/undefine:name</code>	Undefines any definition currently in effect for the specified symbol; same as option <code>/U</code> .	OFF	None
<code>/unroll[:n]</code>	Tells the compiler the maximum number of times to unroll loops. This is the same as <code>/Qunroll</code> , which	<code>/unroll</code>	<code>-unroll[=</code>

Option	Description	Default	Equivalent Linux* OS
	is the recommended option to use.		
<code>/us</code>	Tells the compiler to append an underscore character to external user-defined names; same as the <code>/assume:underscore</code> option.	OFF	<code>-us</code>
<code>/Vstring</code>	Places the text string specified into the object file (.obj) being generated by the compiler; same as the <code>/bintext</code> option.	OFF	None
<code>/[no]vms</code>	Causes the run-time system to behave like HP* Fortran on OpenVMS* Alpha systems and VAX* systems (VAX FORTRAN*).	<code>/novms</code>	<code>-[no]vms</code>
<code>/w</code>	Disables all warning messages; same as specifying option <code>/warn:none</code> or <code>/warn:nogeneral</code> .	OFF	<code>-w</code>
<code>/Wn</code>	Disables (<code>n=0</code>) or enables (<code>n=1</code>) all warning	<code>/W1</code>	<code>-Wn</code>

Option	Description	Default	Equivalent Linux* OS
	messages.		
<code>/warn:keyword</code>	Specifies diagnostic messages to be issued by the compiler.	keywords: alignments general usage nodeclarations noerrors noignore_loc nointerfaces nostderrors notruncated_source nuncalled nounused	<code>-warn key</code>
<code>/watch[:keyword]</code>	Tells the compiler to display certain information to the console output window.	<code>/nowatch</code>	<code>-watch [k</code>
<code>/WB</code>	Turns a compile-time bounds check error into a warning.	OFF	<code>-WB</code>
<code>/what</code>	Tells the compiler to display its detailed version string.	OFF	<code>-what</code>
<code>/winapp</code>	Tells the compiler to create a graphics or Fortran Windows application and link against the most	OFF	None

Option	Description	Default	Equivalent Linux* OS
	commonly used libraries.		
/X	Removes standard directories from the include file search path.	OFF	-X
/zd	Tells the compiler to generate line numbers and minimal debugging information. Deprecated; use option <code>/debug:minimal</code> .	OFF	None
/zi or /z7	Tells the compiler to generate full debugging information in the object file; same as the <code>/debug:full</code> or <code>/debug</code> option.	OFF	-g
/zl	Prevents any linker search options from being included into the object file; same as the <code>/libdir:none</code> or <code>/nolibdir</code> option.	OFF	None
/Zp[n]	Aligns fields of records and components of derived types on the smaller of the size boundary specified or the boundary that will naturally align them; same	/Zp16	-Zp[n]

Option	Description	Default	Equivalent Linux* OS
	as the <code>/align:recn byte</code> option.		
<code>/Zs</code>	Tells the compiler to perform syntax checking only; same as the <code>/syntax-only</code> option.	OFF	None

See Also

[-map-opts, /Qmap-opts](#) compiler option

Linux* OS and Mac OS* X Quick Reference Guide and Cross Reference

The table in this section summarizes Intel® Fortran compiler options used on Linux* OS and Mac OS* X. Each summary also shows the equivalent compiler options on Windows* OS.

If you want to see the summarized Windows* options, see [this table](#).

Some compiler options are only available on systems using certain architectures, as indicated by these labels:

Label	Meaning
i32	The option is available on systems using IA-32 architecture.
i64em	The option is available on systems using Intel® 64 architecture.
i64	The option is available on systems using IA-64 architecture.

If "only" appears in the label, the option is only available on the identified system or architecture.

If no label appears, the option is available on all supported systems and architectures.

For more details on the options, refer to the [Alphabetical Compiler Options](#) section.

The Intel® Fortran Compiler includes the Intel® Compiler Option Mapping tool. This tool lets you find equivalent options by specifying compiler option `-map-opts` (Linux and Mac OS X) or `/Qmap-opts` (Windows).

For information on conventions used in this table, see [Conventions](#).

Quick Reference of Linux OS and Mac OS X Options

The following table summarizes all supported Linux OS and Mac OS X options. It also shows equivalent Windows OS options, if any.

Option	Description	Default	Equivalent Option Windows* OS
-1	Executes at least one iteration of DO loops.	OFF	/1
-66	Tells the compiler to use FORTRAN 66 semantics.	OFF	None
-72, -80, -132	Treats the statement field of each fixed-form source line as ending in column 72, 80, or 132; same as the <code>-extend-source</code> option.	-72	/4L{72 80 132}
-align [keyword]	Tells the compiler how to align certain data items.	keywords: nocommons nodcommons records nosequence	/align[:keyword]
-allow [no]fpp_comments	Determines how the fpp preprocessor	-allow fpp_comments	/allow:[no]fpp_

Option	Description	Default	Equivalent Option Windows* OS
	treats Fortran end-of-line comments in preprocessor directive lines.		
-[no]altparam	Allows alternate syntax (without parentheses) for PARAMETER statements.	-altparam	/[no]altparam
-[no-]ansi-alias	Tells the compiler to assume the program adheres to the Fortran Standard type aliasability rules.	-ansi-alias	/Qansi-alias[-]
-arch <i>keyword</i> (i32, i64em)	Tells the compiler to generate optimized code specialized for the processor that executes your program.	varies; see the option description	/arch: <i>keyword</i> (i32, i64em)
-assume <i>keyword</i>	Tells the compiler to make certain assumptions.	keywords: nobsc nobuffered_io nobyterecl nocc_omp nodummy_aliases nominus0	/assume: <i>keyword</i>

Option	Description	Default	Equivalent Option Windows* OS
		noold_boz old_unit_star old_xor protect_constants nprotect_parens norealloc_lhs source_include nostd_mod_proc_name underscore no2underscores nowriteable-strings	
-[no]automatic	Causes all variables to be allocated to the run-time stack; same as the <code>-auto</code> option.	-auto-scalar	/[no]automatic
-auto-scalar	Causes allocation of scalar variables of intrinsic types INTEGER, REAL, COMPLEX, and LOGICAL to the run-time stack.	-auto-scalar	/Qauto-scalar
-autodouble	Makes default real and complex variables 8 bytes long; same as the <code>-real-size 64</code>	OFF	/Qautodouble

Option	Description	Default	Equivalent Option Windows* OS
	option.		
<i>-axprocessor</i> (i32, i64em)	Tells the compiler to generate multiple, processor-specific auto-dispatch code paths for Intel processors if there is a performance benefit.	OFF	<i>/Qaxprocessor</i> (i32, i64em)
<i>-Bdir</i>	Specifies a directory that can be used to find include files, libraries, and executables.	OFF	None
<i>-Bdynamic</i> (Linux only)	Enables dynamic linking of libraries at run time.	OFF	None
<i>-Bstatic</i> (Linux only)	Enables static linking of a user's library.	OFF	None
<i>-c</i>	Causes the compiler to compile to an object file only and not link.	OFF	<i>/c</i>
<i>-CB</i>	Performs run-time checks on whether array subscript and	OFF	<i>/CB</i>

Option	Description	Default	Equivalent Option Windows* OS
<i>-ccdefault keyword</i>	substring references are within declared bounds; same as the <code>-check bounds</code> option. Specifies the type of carriage control used when a file is displayed at a terminal screen.	<code>-ccdefault default</code>	<code>/ccdefault:keyword</code>
<code>-check [keyword]</code>	Checks for certain conditions at run time.	<code>-nocheck</code>	<code>/check[:keyword]</code>
<code>-cm</code>	Disables all messages about questionable programming practices; same as specifying option <code>-warn nousage</code> .	OFF	<code>/cm</code>
<code>-common-args</code>	Tells the compiler that dummy (formal) arguments to procedures share memory locations with other dummy arguments or with	OFF	<code>/Qcommon-args</code>

Option	Description	Default	Equivalent Option Windows* OS
	COMMON variables that are assigned. This option is the same as <code>-assume dummy_aliases</code> .		
<code>-[no-]complex-limited-range</code>	Enables the use of basic algebraic expansions of some arithmetic operations involving data of type COMPLEX.	<code>-no-complex-limited-range</code>	<code>/Qcomplex-limited-range[-]</code>
<code>-convert keyword</code>	Specifies the format of unformatted files containing numeric data.	<code>-convert native</code>	<code>/convert:keyword</code>
<code>-cpp</code>	Runs the Fortran preprocessor on source files prior to compilation.	OFF	<code>/Qcpp</code>
<code>-CU</code>	Enables run-time checking for uninitialized variables. This option is the same as <code>-check uninit</code> .	OFF	<code>/CU</code>
<code>-cxxlib[=dir]</code>	Tells the compiler to link using the C++	<code>-no-cxxlib</code>	None

Option	Description	Default	Equivalent Option Windows* OS
	run-time libraries provided by gcc.		
-cxxlib-nostd	Prevents the compiler from linking with the standard C++ library.	-no-cxxlib	None
-Dname [=value]	Defines a symbol name that can be associated with an optional value.	-noD	/Dname [=value]
-[no]d-lines	Compiles debugging statements indicated by the letter D in column 1 of the source code.	-nod-lines	/[no]d-lines or /
-DD	Compiles debugging statements indicated by the letter D in column 1 of the source code; same as the -d-lines option.	-nod-lines	/d-lines
-debug keyword	Specifies settings that enhance debugging.	-debug none	/debug:keyword Note: the Windows different keyword s
-debug-parameters	Tells the compiler to	-nodebug-parameters	/debug-

Option	Description	Default	Equivalent Option Windows* OS
[<i>keyword</i>]	generate debug information for PARAMETERs used in a program.		parameters[: <i>keyw</i>
-diag-type <i>diag-list</i>	Controls the display of diagnostic information.	OFF	/Qdiag-type : <i>di</i>
-diag-dump	Tells the compiler to print all enabled diagnostic messages and stop compilation.	OFF	/Qdiag-dump
-diag-enable sv-include	Tells the Static Verifier to analyze include files and source files when issuing diagnostic message.	OFF	/Qdiag-enable : s
-diag-error-limit <i>n</i>	Specifies the maximum number of errors allowed before compilation stops.	-diag-error-limit 30	/Qdiag-error-li
-diag-file[= <i>file</i>]	Causes the results of diagnostic analysis to be output to a file.	OFF	/Qdiag-file[: <i>fi</i>
-diag-file-append[= <i>file</i>]	Causes the results of diagnostic analysis to	OFF	/Qdiag-file- append[: <i>file</i>]

Option	Description	Default	Equivalent Option Windows* OS
	be appended to a file.		
-[no-]diag-id-numbers	Tells the compiler to display diagnostic messages by using their ID number values.	-diag-id-numbers	/Qdiag-id-numbe
-diag-once <i>id</i> [, <i>id</i> ,...]	Tells the compiler to issue one or more diagnostic messages only once	OFF	/Qdiag-once: <i>id</i> [,
-double-size <i>size</i>	Defines the default KIND for DOUBLE PRECISION and DOUBLE COMPLEX variables.	-double-size 64	/double-size: <i>s</i>
-dps	Specifies that the alternate syntax for PARAMETER statements is allowed; same as the -altparam option.	-dps	/Qdps
-dryrun	Specifies that driver tool commands should be shown but not executed.	OFF	None
-dumpmachine	Displays the target	OFF	None

Option	Description	Default	Equivalent Option Windows* OS
	machine and operating system configuration.		
<code>-dynamic-linkerfile</code> (Linux only)	Specifies a dynamic linker in <i>file</i> other than the default.	OFF	None
<code>-dynamiclib</code> (i32, i64em; Mac OS X only)	Invokes the <i>libtool</i> command to generate dynamic libraries.	OFF	None
<code>-dyncom</code> "common1,common2,..."	Enables dynamic allocation of common blocks at run time.	OFF	/Qdyncom "common1,common2,..."
<code>-E</code>	Causes the Fortran preprocessor to send output to stdout.	OFF	/E
<code>-e03, -e95, -e90</code>	Causes the compiler to issue errors instead of warnings for nonstandard Fortran; same as the <code>-warn stderrors</code> option.	OFF	None
<code>-EP</code>	Causes the Fortran preprocessor to send output to stdout,	OFF	/EP

Option	Description	Default	Equivalent Option Windows* OS
	omitting #line directives.		
<code>-error-limit <i>n</i></code>	Specifies the maximum number of error-level or fatal-level compiler errors allowed for a file specified on the command line; same as <code>-diag-error-limit</code> .	<code>-error-limit 30</code>	<code>/error-limit:<i>n</i></code>
<code>-extend-source <i>size</i></code>	Specifies the length of the statement field in a fixed-form source file.	<code>-extend-source 72</code>	<code>/extend-source:</code>
<code>-f66</code>	Tells the compiler to use FORTRAN 66 semantics.	OFF	<code>/f66</code>
<code>-[no]f77rtl</code>	Tells the compiler to use FORTRAN 77 run-time behavior.	OFF	<code>/[no]f77rtl</code>
<code>-f[no-]alias</code>	Determines whether aliasing should be assumed in the program.	ON	None
<code>-falign-functions[=<i>n</i>]</code>	Tells the compiler to	<code>-no-falign-</code>	<code>/Qfnalign[:<i>n</i>]</code>

Option	Description	Default	Equivalent Option Windows* OS
(i32, i64em)	align functions on an optimal byte boundary.	functions	(i32, i64em)
-falign-stack[= <i>mode</i>] (i32 only)	Tells the compiler to align functions on an optimal byte boundary.	-falign-stack=default	None
-fast	Maximizes speed across the entire program.	OFF	/fast
-[no-]fast-transcendentals	Enables the compiler to replace calls to transcendental functions with faster but less precise implementations.	ON	/Qfast-transcendentals]
-fcode-asm	Produces an assembly file with optional machine code annotations.	OFF	/FAc
-f[no-]exceptions	Enables exception handling table generation.	-fno-exceptions	None
-f[no-]fnalias	Specifies that aliasing should be assumed within	OFF	-ffnalias

Option	Description	Default	Equivalent Option Windows* OS
	functions.		
-FI	Specifies source files are in fixed format; same as the <code>-fixed</code> option.	determined by file suffix	/FI
-f[no-]inline	Tells the compiler to inline functions declared with <code>cDEC\$ ATTRIBUTES FORCEINLINE</code> .	-fno-inline	None
-f[no-]inline-functions	Enables function inlining for single file compilation.	-finline-functions	/Ob2
-finline-limit= <i>n</i>	Lets you specify the maximum size of a function to be inlined.	OFF	None
-f[no-]instrument-functions	Determines whether function entry and exit points are instrumented.	-fno-instrument-functions	/Qinstrument-fu]
-[no]fixed	Specifies source files are in fixed format.	determined by file suffix	/[no]fixed
-f[no-]keep-static-consts	Tells the compiler to preserve allocation of variables that are not referenced in the	-fno-keep-static-consts	/Qkeep-static-c

Option	Description	Default	Equivalent Option Windows* OS
	source.		
-[no]fltconsistency	Enables improved floating-point consistency.	OFF	/[no]fltconsist
-[no-]fma (i64 only; Linux only)	Enables the combining of floating-point multiplies and add/subtract operations.	-fma	/Qfma[-] (i64 only)
-f[no-]math-errno	Tells the compiler that <i>errno</i> can be reliably tested after calls to standard math library functions.	-fno-math-errno	None
-fminshared	Tells the compiler to treat a compilation unit as a component of a main program and not to link it as a shareable object.	OFF	None
-[no-]fnsplit (i64 only; Linux only)	Enables function splitting.	OFF	/Qfnsplit[-] (i32, i64)
-f[no-]omit-frame- pointer (i32, i64em)	Determines whether EBP is used as a general-purpose	-fomit-frame- pointer (unless option -O0 or -g is specified)	/Oy[-] (i32 only)

Option	Description	Default	Equivalent Option Windows* OS
	register in optimizations. This is the same as specifying option <code>-fp</code> , which is deprecated.		
<code>-fp-model <i>keyword</i></code>	Controls the semantics of floating-point calculations.	<code>-fp-model fast</code>	<code>/fp:<i>keyword</i></code>
<code>-[no-]fp-port</code>	Rounds floating-point results after floating-point operations, so rounding to user-declared precision happens at assignments and type conversions (some impact on speed).	<code>-no-fp-port</code>	<code>/Qfp-port[-]</code>
<code>-[no-]fp-relaxed</code> (i64 only; Linux only)	Enables use of faster but slightly less accurate code sequences for math functions, such as divide and sqrt.	<code>-no-fp-relaxed</code>	<code>/Qfp-relaxed[-]</code> (i64 only)
<code>-fp-speculation=<i>mode</i></code>	Tells the compiler the mode in which to	<code>-fp-speculation=fast</code>	<code>/Qfp-speculatio</code>

Option	Description	Default	Equivalent Option Windows* OS
	speculate on floating-point operations.		
-fp-stack-check (i32, i64em)	Generates extra code after every function call to ensure that the FP (floating-point) stack is in the expected state.	OFF	/Qfp-stack-check (i32, i64em)
-[no]fpconstant	Tells the compiler that single-precision constants assigned to double-precision variables should be evaluated in double precision.	OFF	/[no]fpconstant
-fpen	Specifies floating-point exception handling at run time for the main program.	-fpe3	/fpe:n
-f[no-]pic, -fPIC	Generates position-independent code.	-fno-pic	None
-fpie (Linux only)	Tells the compiler to generate position-independent code.	OFF	None
-fpp[n] or -fpp[="option"]	Runs the Fortran preprocessor on	-nofpp	/fpp[n] or /fpp[="

Option	Description	Default	Equivalent Option Windows* OS
	source files prior to compilation.		
<code>-fpscomp [keyword]</code>	Specifies compatibility with Microsoft* Fortran PowerStation or Intel® Fortran.	<code>-fpscomp libs</code>	<code>/fpscomp[:keyword]</code>
<code>-FR</code>	Specifies source files are in free format; same as the <code>-free</code> option.	determined by file suffix	<code>/FR</code>
<code>-fr32</code> (i64 only; Linux only)	Disables use of high floating-point registers.	OFF	None
<code>-[no]free</code>	Specifies source files are in free format.	determined by file suffix	<code>/[no]free</code>
<code>-fsource-asm</code>	Produces an assembly file with optional source code annotations.	OFF	<code>/FAs</code>
<code>-f[no-]stack-security-check</code>	Determines whether the compiler generates code that detects some buffer overruns; same as <code>-f[no-]stack-</code>	<code>-fno-stack-security-check</code>	<code>/GS[-]</code>

Option	Description	Default	Equivalent Option Windows* OS
<code>-fsyntax-only</code>	protector. Specifies that the source file should be checked only for correct syntax; same as <code>-syntax-only</code> .	OFF	None
<code>-ftrapuv</code>	Initializes stack local variables to an unusual value.	OFF	/Qtrapuv
<code>-[no-]ftz</code>	Flushes denormal results to zero.	i64: <code>-no-ftz</code> i32, i64em: <code>-ftz</code>	/Qftz[-]
<code>-[no-]func-groups</code> (i32, i64em; Linux only)	Enables or disables function grouping if profiling information is enabled. This option is deprecated, use <code>-prof-func-groups</code> .	<code>-no-func-groups</code>	None
<code>-funroll-loops</code>	Tells the compiler to unroll user loops based on the default optimization heuristics; same as <code>-unroll</code> , which is the recommended option.	<code>-funroll-loops</code>	/Qunroll

Option	Description	Default	Equivalent Option Windows* OS
<code>-fverbose-asm</code>	Produces an assembly file with compiler comments, including options and version information.	<code>-fno-verbose-asm</code>	None
<code>-fvisibility=keyword</code> <code>-fvisibility-keyword=</code> <code>file</code>	Specifies the default visibility for global symbols; the 2nd form indicates symbols in a file.	<code>-fvisibility=default</code>	None
<code>-g</code>	Produces symbolic debug information in the object file.	OFF	<code>/zi, /z7</code>
<code>-gdwarf2</code>	Enables generation of debug information using the DWARF2 format.	OFF	None
<code>-gen-interfaces</code> <code>[[no]source]</code>	Tells the compiler to generate an interface block for each routine in a source file.	<code>-nogen-interfaces</code>	<code>/gen-interfaces:[no]</code>
<code>-[no-]global-hoist</code>	Enables certain optimizations that can move memory loads to a point earlier in the program execution than where	<code>-global-hoist</code>	<code>/Qglobal-hoist[</code>

Option	Description	Default	Equivalent Option Windows* OS
	they appear in the source.		
<code>-heap-arrays [size]</code>	Puts automatic arrays and arrays created for temporary computations on the heap instead of the stack.	<code>-no-heap-arrays</code>	<code>/heap-arrays[:size]</code>
<code>-help [category]</code>	Displays the list of compiler options.	OFF	<code>/help [category]</code>
<code>-Idir</code>	Specifies a directory to add to the include path.	OFF	<code>/Idir</code>
<code>-i-dynamic</code>	Links Intel-provided libraries dynamically. This is a deprecated option; use <code>-shared-intel</code> .	OFF	None
<code>-i-static</code>	Links Intel-provided libraries statically. This is a deprecated option; use <code>-static-intel</code> .	OFF	None
<code>-i{2 4 8}</code>	Specifies the default KIND for integer and logical variables;	<code>-i4</code>	<code>/4I{2 4 8}</code>

Option	Description	Default	Equivalent Option Windows* OS
	same as the <code>-integer-size</code> option.		
<code>-idirafterdir</code>	Adds a directory to the second include file search path.	OFF	None
<code>-implicitnone</code>	Sets the default type of a variable to undefined; same as option <code>warn</code> declarations.	OFF	/4Yd
<code>-inline-debug-info</code>	Produces enhanced source position information for inlined code.	OFF	/Qinline-debug-
<code>-inline-factor=<i>n</i></code>	Specifies the percentage multiplier that should be applied to all inlining options that define upper limits.	<code>-no-inline-factor</code>	/Qinline-factor
<code>-inline-forceinline</code>	Specifies that an inline routine should be inlined whenever the compiler can do so.	OFF	/Qinline-forcei

Option	Description	Default	Equivalent Option Windows* OS
<code>-inline-level=<i>n</i></code>	Specifies the level of inline function expansion. <i>n</i> = 0, 1, or 2.	<code>-inline-level=2</code> if <code>-O2</code> is in effect <code>-inline-level=0</code> if <code>-O0</code> is specified	<code>/Ob<i>n</i></code>
<code>-inline-max-per-compile=<i>n</i></code>	Specifies the maximum number of times inlining may be applied to an entire compilation unit.	<code>-no-inline-max-per-compile</code>	<code>/Qinline-max-per-compile=<i>n</i></code>
<code>-inline-max-per-routine=<i>n</i></code>	Specifies the maximum number of times the inliner may inline into a particular routine.	<code>-no-inline-max-per-routine</code>	<code>/Qinline-max-per-routine=<i>n</i></code>
<code>-inline-max-size=<i>n</i></code>	Specifies the lower limit for the size of what the inliner considers to be a large routine.	<code>-no-inline-max-size</code>	<code>/Qinline-max-size=<i>n</i></code>
<code>-inline-max-total-size=<i>n</i></code>	Specifies how much larger a routine can normally grow when inline expansion is performed.	<code>-no-inline-max-total-size</code>	<code>/Qinline-max-total-size=<i>n</i></code>
<code>-inline-min-size=<i>n</i></code>	Specifies the upper limit for the size of what the inliner	<code>-no-inline-min-size</code>	<code>/Qinline-min-size=<i>n</i></code>

Option	Description	Default	Equivalent Option Windows* OS
	considers to be a small routine.		
-[no]intconstant	Tells the compiler to use FORTRAN 77 semantics to determine the KIND for integer constants.	-nointconstant	/[no]intconstant
-integer-size <i>size</i>	Specifies the default KIND for integer and logical variables.	-integer-size 32	/integer-size:s
-[no-]ip	Enables additional single-file interprocedural optimizations.	OFF	/Qip[-]
-ip-no-inlining	Disables full and partial inlining enabled by -ip.	OFF	/Qip-no-inlinin
-ip-no-pinlining	Disables partial inlining.	OFF	/Qip-no-pinlini
-IPF-flt-eval-method0 (i64 only; Linux only)	Tells the compiler to evaluate the expressions involving floating-point operands in the precision indicated by the variable types	OFF	/QIPF-flt-eval- (i64 only)

Option	Description	Default	Equivalent Option Windows* OS
	declared in the program. Deprecated.		
-IPF-fltacc (i64 only; Linux only)	Tells the compiler to apply optimizations that affect floating-point accuracy. Deprecated.	-no-IPF-fltacc	/QIPF-fltacc (i64 only)
-IPF-fma (i64 only; Linux only)	Enables the combining of floating-point multiplies and add/subtract operations. Deprecated; use -fma.	-IPF-fma	/QIPF-fma (i64 only)
-IPF-fp-relaxed (i64 only; Linux only)	Enables use of faster but slightly less accurate code sequences for math functions, such as divide and sqrt. Deprecated; use -fp-relaxed.	-no-IPF-fp-relaxed	/QIPF-fp-relaxed (i64 only)
-ipo[n]	Enables multifile IP optimizations between files.	OFF	/Qipo[n]
-ipo-c	Generates a multifile	OFF	/Qipo-c

Option	Description	Default	Equivalent Option Windows* OS
	object file that can be used in further link steps.		
<code>-ipo-jobsn</code>	Specifies the number of commands to be executed simultaneously during the link phase of Interprocedural Optimization (IPO).	<code>-ipo-jobs1</code>	/Qipo-jobs: n
<code>-ipo-S</code>	Generates a multifile assembly file that can be used in further link steps.	OFF	/Qipo-S
<code>-ipo-separate</code> (Linux only)	Generates one object file per source file.	OFF	/Qipo-separate
<code>-isystemdir</code>	Specifies a directory to add to the start of the system include path.	OFF	None
<code>-ivdep-parallel</code> (i64 only; Linux only)	Tells the compiler that there is no loop-carried memory dependency in any loop following an IVDEP directive.	OFF	/Qivdep-parallel (i64 only)

Option	Description	Default	Equivalent Option Windows* OS
<code>-lstring</code>	Tells the linker to search for a specified library when linking.	OFF	None
<code>-Ldir</code>	Tells the linker where to search for libraries before searching the standard directories.	OFF	None
<code>-[no]logo</code>	Displays compiler version information.	<code>-nologo</code>	<code>/[no]logo</code>
<code>-lowercase</code>	Causes the compiler to ignore case differences in identifiers and to convert external names to lowercase; same as the <code>-names lowercase</code> option.	<code>-lowercase</code>	<code>/Qlowercase</code>
<code>-m[processor]</code> (i32, i64em)	Tells the compiler to generate optimized code specialized for the processor that executes your program.	varies; see option description	<code>/arch</code>
<code>-m32, -m64</code> (i32, i64em)	Tells the compiler to generate code for IA-32 architecture or Intel® 64	OFF	None

Option	Description	Default	Equivalent Option Windows* OS
	architecture, respectively.		
<code>-map-opts</code> (Linux only)	Maps one or more Linux* compiler options to their equivalent on a Windows* system (or vice versa).	OFF	<code>/Qmap-opts</code>
<code>-march=processor</code> (i32, i64em; Linux only)	Tells the compiler to generate code for a specified processor.	i32: OFF i64em: - march=pentium4	None
<code>-mmodel=mem_model</code> (i64em only; Linux only)	Tells the compiler to use a specific memory model to generate code and store data.	<code>-mmodel=small</code>	None
<code>-mdynamic-no-pic</code> (i32 only; Mac OS X only)	Generates code that is not position-independent but has position-independent external references.	OFF	None
<code>-mieee-fp</code>	Tells the compiler to use IEEE floating point comparisons. This is the same as specifying option - <code>fltconsistency</code> or	OFF	<code>/fltconsistency</code>

Option	Description	Default	Equivalent Option Windows* OS
-	<p data-bbox="662 331 724 359">-mp.</p> <p data-bbox="662 401 943 596">Determines whether MOVBE instructions are generated for Intel® processors.</p>	-minstruction=movbe /Qinstruction=	
-mixed_str_len_arg	<p data-bbox="662 638 964 1106">Tells the compiler that the hidden length passed for a character argument is to be placed immediately after its corresponding character argument in the argument list.</p>	OFF	/iface:mixed_st
-module <i>path</i>	<p data-bbox="662 1148 932 1505">Specifies the directory where module files should be placed when created and where they should be searched for.</p>	OFF	/module: <i>path</i>
-mp	<p data-bbox="662 1547 915 1688">Enables improved floating-point consistency.</p>	OFF	/Op
-mp1	<p data-bbox="662 1730 927 1871">Improves floating-point precision and consistency.</p>	OFF	/Qprec

Option	Description	Default	Equivalent Option Windows* OS
<code>-m[no-]relax</code> (i64 only)	Determines whether the compiler passes linker option <code>-relax</code> to the linker.	<code>-mno-relax</code>	None
<code>-mtune=<i>processor</i></code>	Performs optimizations for a particular processor. <code>-mtune=itanium2-p9000</code> is equivalent to <code>/G2</code> .	i32: <code>-mtune=pentium4</code> i64: <code>-mtune=itanium2-p9000</code>	None
<code>-multiple-processes=<i>n</i></code>	Creates multiple processes that can be used to compile large numbers of source files at the same time.	OFF	<code>/MP: <i>n</i></code>
<code>-names <i>keyword</i></code>	Specifies how source code identifiers and external names are interpreted.	<code>-names lowercase</code>	<code>/names: <i>keyword</i></code>
<code>-nbs</code>	Tells the compiler to treat the backslash character (<code>\</code>) as a normal character in character literals; same as the <code>-assume nobsc</code> option.	<code>-nbs</code>	<code>/nbs</code>

Option	Description	Default	Equivalent Option Windows* OS
-no-bss-init	Tells the compiler to place in the DATA section any variables explicitly initialized with zeros.	OFF	/Qno-bss-init
-nodefaultlibs	Prevents the compiler from using standard libraries when linking.	OFF	None
-nodefine	Specifies that all preprocessor definitions apply only to fpp and not to Intel® Fortran conditional compilation directives.	OFF	/nodefine
-nofor-main	Specifies the main program is not written in Fortran, and prevents the compiler from linking <i>for_main.o</i> into applications.	OFF	None
-noinclude	Prevents the compiler from searching in a	OFF	/noinclude

Option	Description	Default	Equivalent Option Windows* OS
	directory previously added to the include path for files specified in an INCLUDE statement.		
-nolib-inline	Disables inline expansion of standard library or intrinsic functions.	OFF	None
-nostartfiles	Prevents the compiler from using standard startup files when linking.	OFF	None
-nostdinc	Removes standard directories from the include file search path; same as the -x option.	OFF	None
-nostdlib	Prevents the compiler from using standard libraries and startup files when linking.	OFF	None
-nus	Disables appending an underscore to external user-defined names; same as the	OFF	None

Option	Description	Default	Equivalent Option Windows* OS
	-assume nounderscore option.		
-ofile	Specifies the name for an output file.	OFF	None
-O[<i>n</i>]	Specifies the code optimization for applications.	-O2	/O[<i>n</i>]
-O0	Disables all optimizations.	OFF	/Od
-onetrip	Executes at least one iteration of DO loops.	OFF	/onetrip
-openmp	Enables the parallelizer to generate multithreaded code based on OpenMP* directives.	OFF	/Qopenmp
-openmp-lib <i>type</i> (Linux only)	Lets you specify an OpenMP* run-time library to use for linking.	-openmp-lib legacy	/Qopenmp-lib:ty
-openmp-link <i>library</i>	Controls whether the compiler links to static or dynamic OpenMP run-time	-openmp-link dynamic	/Qopenmp-link:l

Option	Description	Default	Equivalent Option Windows* OS
	libraries.		
-openmp-profile (Linux only)	Enables analysis of OpenMP* applications.	OFF	/Qopenmp-profile
-openmp-report[<i>n</i>]	Controls the OpenMP parallelizer's level of diagnostic messages.	-openmp-report1	/Qopenmp-report
-openmp-stubs	Enables compilation of OpenMP programs in sequential mode.	OFF	/Qopenmp-stubs
-openmp-threadprivate <i>type</i> (Linux only)	Lets you specify an OpenMP* threadprivate implementation.	-openmp-threadprivate legacy	/Qopenmp-threadprivate: <i>t</i>
-opt-block-factor= <i>n</i>	Lets you specify a loop blocking factor.	OFF	/Qopt-block-fac
-opt-jump- tables= <i>keyword</i>	Enables or disables generation of jump tables for switch statements.	-opt-jump- tables=default	/Qopt-jump- tables: <i>keyword</i>
-[no-]opt-loadpair (i64 only; Linux only)	Enables or disables loadpair optimization.	-no-opt-loadpair	/Qopt-loadpair[(i64 only)
-opt-malloc-options= <i>n</i> (i32, i64em)	Lets you specify an alternate algorithm for malloc().	-opt-malloc- options=0	None
-opt-mem-bandwidth <i>n</i>	Enables performance	-opt-mem-bandwidth0	/Qopt-mem-bandw

Option	Description	Default	Equivalent Option Windows* OS
(i64 only; Linux only)	tuning and heuristics that control memory bandwidth use among processors.	for serial compilation; -opt-mem-bandwidth1 for parallel compilation	(i64 only)
-[no-]opt-mod-versioning (i64 only; Linux only)	Enables or disables versioning of modulo operations for certain types of operands.	-no-opt-mod-versioning	/Qopt-mod-versi (i64 only)
-[no-]opt-multi-version-aggressive (i32, i64em)	Tells the compiler to use aggressive multi-versioning to check for pointer aliasing and scalar replacement.	-no-opt-multi-version-aggressive	/Qopt-multi-ver aggressive[-] (i32, i64em)
-opt-prefetch[=n]	Enables prefetch insertion optimization.	i64: -opt-prefetch i32, i64em: -no-opt-prefetch	/Qopt-prefetch[
-[no-]opt-prefetch-initial-values (i64 only; Linux only)	Enables or disables prefetches that are issued before a loop is entered.	-opt-prefetch-initial-values	/Qopt-prefetch- values[-] (i64 only)
-[no-]opt-prefetch-issue-excl-hint (i64 only; Linux only)	Determines whether the compiler issues prefetches for stores with exclusive hint.	-no-opt-prefetch-issue-excl-hint	/Qopt-prefetch- excl-hint[-] (i64 only)
-[no-]opt-prefetch-	Enables or disables	-opt-prefetch-next-	/Qopt-prefetch-

Option	Description	Default	Equivalent Option Windows* OS
<code>next-iteration</code> (i64 only; Linux only)	prefetches for a memory access in the next iteration of a loop.	<code>iteration</code>	<code>iteration[-][:n]</code> (i64 only)
<code>-opt-ra-region-strategy[=keyword]</code> (i32, i64em)	Selects the method that the register allocator uses to partition each routine into regions.	<code>-opt-ra-region-strategy=default</code>	<code>/Qopt-ra-region-strategy[:keyword]</code> (i32, i64em)
<code>-opt-report [n]</code>	Tells the compiler to generate an optimization report to <code>stderr</code> .	<code>-opt-report 2</code>	<code>/Qopt-report[:n]</code>
<code>-opt-report-file=file</code>	Specifies the name for an optimization report.	OFF	<code>/Qopt-report-fi</code>
<code>-opt-report-help</code>	Displays the optimizer phases available for report generation.	OFF	<code>/Qopt-report-he</code>
<code>-opt-report-phase=phase</code>	Specifies an optimizer phase to use when optimization reports are generated.	OFF	<code>/Qopt-report-ph</code>
<code>-opt-report-</code>	Tells the compiler to	OFF	<code>/Qopt-report-</code>

Option	Description	Default	Equivalent Option Windows* OS
<code>routine=<i>string</i></code>	generate reports on the routines containing specified text.		<code>routine:<i>string</i></code>
<code>-opt-streaming-stores <i>keyword</i></code> (i32, i64em)	Enables generation of streaming stores for optimization.	<code>-opt-streaming-stores auto</code>	<code>/Qopt-streaming-stores:<i>keyword</i></code> (i32, i64em)
<code>-[no-]opt-subscript-in-range</code> (i32, i64em)	Determines whether the compiler assumes no overflows in the intermediate computation of subscript expressions in loops.	<code>-no-opt-subscript-in-range</code>	<code>/Qopt-subscript-range[-]</code> (i32, i64em)
<code>-p</code>	Compiles and links for function profiling with gprof(1).	OFF	None
<code>-P</code>	Causes the Fortran preprocessor to send output to a file, which is named by default; same as the <code>-preprocess-only</code> option.	OFF	<code>/P</code>
<code>-[no]pad</code>	Enables the changing of the variable and	OFF	<code>/Qpad[-]</code>

Option	Description	Default	Equivalent Option Windows* OS
	array memory layout.		
-[no]pad-source	Specifies padding for fixed-form source records.	OFF	/[no]pad-source source[-]
-par-report[n]	Controls the diagnostic information reported by the auto-parallelizer.	-par-report1	/Qpar-report[n]
-[no-]par-runtime-control	Generates code to perform run-time checks for loops that have symbolic loop bounds.	-no-par-runtime-control	/Qpar-runtime-c
-par-schedule-keyword [=n]	Specifies a scheduling algorithm for DO loop iterations.	OFF	/Qpar-schedule- [[:]n]
-par-threshold[n]	Sets a threshold for the auto-parallelization of loops.	-par-threshold100	/Qpar-threshold
-parallel	Tells the auto-parallelizer to generate multithreaded code	OFF	/Qparallel

Option	Description	Default	Equivalent Option Windows* OS
	for loops that can be safely executed in parallel.		
<code>-pcn</code> (i32, i64em)	Enables control of floating-point significand precision.	<code>-pc80</code>	<code>/Qpcn</code> (i32, i64em)
<code>-pg</code>	Compiles and links for function profiling with <code>gprof(1)</code> ; same as the <code>-p</code> option.	OFF	None
<code>-pie</code> (Linux only)	Produces a position-independent executable on processors that support it.	OFF	None
<code>-[no-]prec-div</code>	Improves precision of floating-point divides.	<code>-prec-div</code>	<code>/Qprec-div[-]</code>
<code>-[no-]prec-sqrt</code> (i32, i64em)	Improves precision of square root implementations.	<code>-no-prec-sqrt</code>	<code>/Qprec-sqrt[-]</code> (i32, i64em)
<code>-prefetch</code>	Enables prefetch insertion optimization. Deprecated; use <code>-opt-prefetch</code> .	i64: <code>-prefetch</code> i32, i64em: <code>-no-prefetch</code>	<code>/Qprefetch</code>
<code>-preprocess-only</code>	Causes the Fortran	OFF	<code>/preprocess-onl</code>

Option	Description	Default	Equivalent Option Windows* OS
	preprocessor to send output to a file, which is named by default; same as the <code>-P</code> option.		
<code>-print-multi-lib</code>	Prints information about where system libraries should be found.	OFF	None
<code>-[no-]prof-data-order</code> (Linux only)	Enables or disables data ordering if profiling information is enabled.	<code>-no-prof-data-order</code> / <code>/Qprof-data-ord</code>	
<code>-prof-dir dir</code>	Specifies a directory for profiling information output files.	OFF	<code>/Qprof-dir dir</code>
<code>-prof-file file</code>	Specifies a file name for the profiling summary file.	OFF	<code>/Qprof-file fil</code>
<code>-[no-]prof-func-groups</code> (i32, i64em; Linux only)	Enables or disables function grouping if profiling information is enabled.	<code>-no-prof-func-groups</code>	None
<code>-[no-]prof-func-order</code> (Linux only)	Enables or disables function ordering if	<code>-no-prof-func-order</code> / <code>/Qprof-func-ord</code>	

Option	Description	Default	Equivalent Option Windows* OS
	profiling information is enabled.		
<code>-prof-gen[=keyword]</code>	Produces an instrumented object file that can be used in profile-guided optimization.	<code>-[no-]prof-gen</code>	<code>/Qprof-gen[:keyw</code>
<code>-prof-genx</code>	Produces an instrumented object file that includes extra source position information. Deprecated; use <code>-prof-gen=srcpos</code> .	OFF	<code>/Qprof-genx</code>
<code>-prof-hotness-threshold=n</code> (Linux only)	Lets you set the hotness threshold for function grouping and function ordering.	OFF	<code>/Qprof-hotness-threshold:n</code>
<code>-[no-]prof-src-dir</code>	Determines whether directory information of the source file under compilation is considered when looking up profile data records.	<code>-prof-src-dir</code>	<code>/Qprof-src-dir[</code>
<code>-prof-src-root=dir</code>	Lets you use relative directory paths when	OFF	<code>/Qprof-src-root</code>

Option	Description	Default	Equivalent Option Windows* OS
	looking up profile data and specifies a directory as the base.		
<code>-prof-src-root-cwd</code>	Lets you use relative directory paths when looking up profile data and specifies the current working directory as the base.	OFF	<code>/Qprof-src-root</code>
<code>-prof-use[=arg]</code>	Enables the use of profiling information during optimization.	<code>-no-prof-use</code>	<code>/Qprof-use[:arg]</code>
<code>-Qinstall dir</code>	Specifies the root directory where the compiler installation was performed.	OFF	None
<code>-Qlocation,string,dir</code>	Specifies a directory as the location of the specified tool in <i>string</i> .	OFF	<code>/Qlocation,stri</code>
<code>-Qoption,string,options</code>	Passes <i>options</i> to the specified tool in <i>string</i> .	OFF	<code>/Qoption,string</code>
<code>-r8, -r16</code>	Specifies the default KIND for real and complex variables. –	OFF	<code>/4R8, /4R16</code>

Option	Description	Default	Equivalent Option Windows* OS
	r8 is the same as <code>/real-size:64.</code> r16 is the same as <code>/real-size:128.</code>		
<code>-rcd</code> (i32, i64em)	Enables fast float-to-integer conversions.	OFF	<code>/Qrcd</code> (i32, i64em)
<code>-rct</code> (i32, i64em)	Sets the internal FPU rounding control to Truncate.	OFF	<code>/Qrct</code> (i32, i64em)
<code>-real-size size</code>	Specifies the default KIND for real variables.	<code>-real-size 32</code>	<code>/real-size:size</code>
<code>-recursive</code>	Tells the compiler that all routines should be compiled for possible recursive execution.	<code>-norecursive</code>	<code>/recursive</code>
<code>-reentrancy keyword</code>	Tells the compiler to generate reentrant code to support a multithreaded application.	<code>-noreentrancy</code>	<code>/reentrancy:key</code>
<code>-S</code>	Causes the compiler to compile to an assembly file (.s) only and not link; same as	OFF	<code>/S</code>

Option	Description	Default	Equivalent Option Windows* OS
	options /Fa and /asmfile.		
-safe-cray-ptr	Tells the compiler that Cray* pointers do not alias other variables.	OFF	/Qsafe-cray-ptr
-save	Causes variables to be placed in static memory.	-auto-scalar	/Qsave
-[no-]save-temps	Tells the compiler to save intermediate files created during compilation.	-no-save-temps	/Qsave-temps[-]
-[no-]scalar-rep (i32 only)	Enables scalar replacement performed during loop transformation (requires -O3).	-no-scalar-rep	/Qscalar-rep[-] (i32 only)
-shared (Linux only)	Tells the compiler to produce a dynamic shared object instead of an executable.	OFF	None
-shared-intel	Links Intel-provided libraries dynamically.	OFF	None
-shared-libgcc (Linux only)	Links the GNU libgcc library dynamically.	-shared-libgcc	None

Option	Description	Default	Equivalent Option Windows* OS
-[no-]sox	Tells the compiler to save the compiler options and version in the executable.	-no-sox	/Qsox[-]
-stand <i>keyword</i>	Causes the compiler to issue compile-time messages for nonstandard language elements.	-nostand	/stand:keyword
-static (Linux only)	Prevents linking with shared libraries.	-static	/static
-staticlib (i32, i64em; Mac OS X only)	Invokes the libtool command to generate static libraries.	OFF	None
-static-intel	Links Intel-provided libraries statically.	OFF	None
-static-libgcc (Linux only)	Links the GNU libgcc library statically.	OFF	None
-std90	Causes the compiler to issue messages for language elements that are not standard in Fortran 90; same as -stand f90.	OFF	/stand:f90

Option	Description	Default	Equivalent Option Windows* OS
-std95	Causes the compiler to issue messages for language elements that are not standard in Fortran 95; same as <code>-stand f95</code> .	OFF	<code>/stand:f95</code>
-std03	Causes the compiler to issue messages for language elements that are not standard in Fortran 2003; same as <code>-std</code> or <code>-stand f03</code> .	OFF	<code>/stand:f03</code>
-syntax-only	Specifies that the source file should be checked only for correct syntax.	OFF	<code>/syntax-only</code>
-T <i>file</i> (Linux only)	Tells the linker to read link commands from the specified <i>file</i> .	OFF	None
-tcheck (Linux only)	Enables analysis of threaded applications.	OFF	<code>/Qtcheck</code>
-tcollect [<i>lib</i>]	Inserts instrumentation	OFF	<code>/Qtcollect[=<i>lib</i>]</code>

Option	Description	Default	Equivalent Option Windows* OS
(Linux only)	probes calling the Intel(R) Trace Collector API.		
-tcollect-filter <i>[file]</i>	Lets you enable or disable the instrumentation of specified functions.	OFF	/Qtcollect-filt
(Linux only)			
-Tf <i>file</i>	Tells the compiler to compile the file as a Fortran source file.	OFF	/Tf <i>file</i>
-[no]threads	Tells the linker to search for unresolved references in a multithreaded run-time library.	i32, i64: -nothreads i64em: -threads	/[no]threads
-tprofile	Generates instrumentation to analyze multi-threading performance.	OFF	/Qtprofile
(Linux only)			
-[no]traceback	Tells the compiler to generate extra information in the object file to provide source file traceback information when a severe error occurs	-notraceback	/[no]traceback

Option	Description	Default	Equivalent Option Windows* OS
	at run time.		
-tune <i>keyword</i> (i32, i64em)	Determines the version of the architecture for which the compiler generates instructions.	-tune pn4	/tune: <i>keyword</i> (i32, i64em)
-u	Enables error messages about any undeclared symbols; same as the <code>-warn declarations</code> option.	OFF	None Note: the Windows not the same
-U <i>name</i>	Undefines any definition currently in effect for the specified symbol.	OFF	/U <i>name</i>
-unroll[<i>n</i>]	Tells the compiler the maximum number of times to unroll loops. <code>-unroll</code> is the same as option <code>-funroll-loops</code> .	-unroll	/unroll[: <i>n</i>]
-[no-]unroll- aggressive (i32, i64em)	Determines whether the compiler uses more aggressive unrolling for certain	-no-unroll- aggressive	/Qunroll-aggress (i32, i64em)

Option	Description	Default	Equivalent Option Windows* OS
-uppercase	loops. Causes the compiler to ignore case differences in identifiers and to convert external names to uppercase; same as the <code>-names uppercase</code> option.	OFF	/Quppercase
-us	Tells the compiler to append an underscore character to external user-defined names; same as the <code>-assume underscore</code> option.	-us	/us
-[no-]use-asm	Tells the compiler to produce objects through the assembler.	-no-use-asm	/Quse-asm[-] (i32 only)
-v [<i>file</i>]	Tells the driver that tool commands should be shown and executed.	OFF	None
-V	Displays the compiler version information; same as the <code>-logo</code>	OFF	None

Option	Description	Default	Equivalent Option Windows* OS
	option.		
-[no-]vec (i32, i64em)	Enables or disables vectorization and transformations enabled for vectorization.	-no-vec	/Qvec[-] (i32, i64em)
-[no-]vec-guard-write (i32, i64em)	Tells the compiler to perform a conditional check in a vectorized loop.	-no-vec-guard-write	/Qvec-guard-wri (i32, i64em)
-vec-report[n] (i32, i64em)	Controls the diagnostic information reported by the vectorizer.	-vec-report1	/Qvec-report[n] (i32, i64em)
-[no]vms	Causes the run-time system to behave like HP* Fortran on OpenVMS* Alpha systems and VAX* systems (VAX FORTRAN*).	-novms	/[no]vms
-w	Disables all warning messages; same as specifying option -warn none or -warn no general.	OFF	/w

Option	Description	Default	Equivalent Option Windows* OS
-Wn	Disables ($n=0$) or enables ($n=1$) all warning messages.	-W1	/Wn
-Wa,o1[,o2,...]	Passes options (o1,o2, and so forth) to the assembler for processing.	OFF	None
-warn [keyword]	Specifies diagnostic messages to be issued by the compiler.	keywords: alignments nodeclarations noerrors general noignore_loc nointerfaces nostderrors notruncated_source nouncalled nounused usage	/warn[:keyword]
-[no]watch [keyword]	Tells the compiler to display certain information to the console output window.	-nowatch	/[no]watch[:keyw
-WB	Turns a compile-time bounds check error	OFF	/WB

Option	Description	Default	Equivalent Option Windows* OS
	into a warning.		
-what	Tells the compiler to display its detailed version string.	OFF	/what
-winline	Enables diagnostics about what is inlined and what is not inlined.	OFF	None
-wl,option1 [,option2,...]	Passes options (o1, o2, and so forth) to the linker for processing.	OFF	None
-wp,option1 [,option2,...]	Passes options (o1, o2, and so forth) to the preprocessor.	OFF	None
-xp (i32, i64em)	Tells the compiler to generate optimized code specialized for the Intel processor that executes your program.	varies; see the option description	/Qxp (i32, i64em)
-x	Removes standard directories from the include file search path.	OFF	/x
-xlinker <i>option</i>	Passes a linker	OFF	None

Option	Description	Default	Equivalent Option Windows* OS
	option directly to the linker		
-y	Specifies that the source file should be checked only for correct syntax; same as the <code>-syntax-only</code> option.	OFF	/Zs
-[no]zero	Initializes to zero all local scalar variables of intrinsic type INTEGER, REAL, COMPLEX, or LOGICAL that are saved but not yet initialized.	-nozero	/Qzero[-]
-Zp[n]	Aligns fields of records and components of derived types on the smaller of the size boundary specified or the boundary that will naturally align them.	-Zp16	/Zp[n]

See Also

[-map-opts, /Qmap-opts](#) compiler option

Related Options

Related Options

This topic lists related options that can be used under certain conditions.

Cluster OpenMP* Options (Linux* OS only)

The Cluster OpenMP* (CLOMP or Cluster OMP) options are available if you have a separate license for the Cluster OpenMP product.

These options can be used on Linux* operating systems running on Intel® 64 and IA-64 architectures.

Option	Description
-[no-]cluster-openmp	Lets you run an OpenMP program on a cluster.
-[no-]cluster-openmp-profile	Links a Cluster OMP program with profiling information.
-[no-]clomp-sharable-propagation	Reports variables that need to be made sharable by the user with Cluster OpenMP.
-[no-]clomp-sharable-info	Reports variables that the compiler automatically makes sharable for Cluster OpenMP.
-[no-]clomp-sharable-commons	Makes all COMMONs sharable by default for Cluster OpenMP.
-[no-]clomp-sharable-modvars	Makes all variables in modules sharable by default for Cluster OpenMP.
-[no-]clomp-sharable-localsaves	Makes all SAVE variables sharable by default for Cluster OpenMP.
-[no-]clomp-sharable-argexprs	Makes all expressions in function and subroutine call statements sharable by default for Cluster OpenMP.

For more information on these options, see the Cluster OpenMP documentation.

Linking Tools and Options

This topic describes how to use the Intel® linking tools, `xild` (Linux* OS and Mac OS* X) or `xilink` (Windows* OS).

The Intel linking tools behave differently on different platforms. The following sections summarize the primary differences between the linking behaviors.

Linux and Mac OS Linking Behavior Summary

The linking tool invokes the compiler to perform IPO if objects containing IR (intermediate representation) are found. (These are mock objects.) It invokes GNU `ld` to link the application.

The command-line syntax for `xild` is the same as that of the GNU linker:

```
xild [<options>] <normal command-line>
```

where:

- `[<options>]`: (optional) one or more options supported only by `xild`.
- `<normal command-line>`: linker command line containing a set of valid arguments for `ld`.

To create `app` using IPO, use the option `-ofile` as shown in the following example:

```
xild -qipo_fas -oapp a.o b.o c.o
```

The linking tool calls the compiler to perform IPO for objects containing IR and creates a new list of object(s) to be linked. The linker then calls `ld` to link the object files that are specified in the new list and produce the application with the name specified by the `-o` option. The linker supports the `-ipo`, `-ipoN`, and `-ipo-separate` options.

Windows Linking Behavior Summary

The linking tool invokes the Intel compiler to perform multi-file IPO if objects containing IR (intermediate representation) is found. These are mock objects. It invokes Microsoft* `link.exe` to link the application.

Windows Linking Behavior Summary

The command-line syntax for the Intel® linker is the same as that of the Microsoft linker:

```
xilink [<options>] <normal command-line>
```

where:

- [*<options>*]: (optional) one or more options supported only by `xilink`.
- *<normal command-line>*: linker command line containing a set of valid arguments for the Microsoft linker.

To place the multifile IPO executable in `ipo_file.exe`, use the linker option `/out:file`; for example:

```
xilink -qipo_fas /out:ipo_file.exe a.obj b.obj c.obj
```

The linker calls the compiler to perform IPO for objects containing `IR` and creates a new list of object(s) to be linked. The linker calls Microsoft `link.exe` to link the object files that are specified in the new list and produce the application with the name specified by the `/out:file` linker option.

Using the Linking Tools

You must use the Intel linking tools to link your application if the following conditions apply:

- Your source files were compiled with multifile IPO enabled. Multi-file IPO is enabled by specifying the `-ipo` (Linux and Mac OS X) or `/Qipo` (Windows) command-line option.
- You normally would invoke either the GNU linker (`ld`) or the Microsoft linker (`link.exe`) to link your application.

The following table lists the available, case-insensitive options supported by the Intel linking tools and briefly describes the behavior of each option:

Linking Tools Option	Description
<code>-qhelp</code>	Lists the available linking tool options. Same as passing no option.

Linking Tools Option	Description
<code>-qnoipo</code>	Disables multi-file IPO compilation.
<code>-qipo_fa[<i>{file/dir}</i>]</code>	<p data-bbox="797 338 1391 598">Produces assembly listing for the multi-file IPO compilation. You may specify an optional name for the listing file, or a directory (with the backslash) in which to place the file.</p> <p data-bbox="797 611 1391 709">The default listing name is depends on the platform:</p> <ul data-bbox="797 722 1391 821" style="list-style-type: none"> <li data-bbox="797 722 1391 766">• Linux and Mac OS X: <i>ipo_out.s</i> <li data-bbox="797 779 1391 821">• Windows: <i>ipo_out.asm</i> <p data-bbox="797 833 1391 1266">If the Intel linking tool invocation results in multi-object compilation, either because the application is big or because the user explicitly instructed the compiler to generate multiple objects, the first <i>.s</i> (Linux and Mac OS X) or <i>.asm</i> (Windows) file takes its name from the <code>-qipo_fa</code> option.</p> <p data-bbox="797 1278 1391 1665">The compiler derives the names of subsequent <i>.s</i> (Linux and Mac OS X) or <i>.asm</i> (Windows) files by appending an incrementing number to the name, for example, <i>foo.asm</i> and <i>foo1.asm</i> for <i>ipo_fafoo.asm</i>. The same is true for the <code>-qipo_fo</code> option (listed below).</p>
<code>-qipo_fo[<i>{file/dir}</i>]</code>	<p data-bbox="797 1692 1391 1841">Produces object file for the multi-file IPO compilation. You may specify an optional name for the object file, or a</p>

Linking Tools Option	Description
	directory (with the backslash) in which to place the file. The default object file name is depends on the platform: <ul style="list-style-type: none"> Linux and Mac OS X: <i>ipo_out.o</i> Windows: <i>ipo_out.obj</i>
<code>-qipo_fas</code>	Add source lines to assembly listing.
<code>-qipo_fac</code>	Adds code bytes to the assembly listing.
<code>-qipo_facs</code>	Add code bytes and source lines to assembly listing.
<code>-quseenv</code>	Disables override of existing PATH, LIB, and INCLUDE variables.
<code>-lib</code>	Invokes librarian instead of linker.
<code>-libtool</code>	Mac OS X: Invokes <code>libtool</code> to create a library instead of <code>ld</code> .
<code>-qv</code>	Displays version information.

See Also

Optimizing Applications: Using IPO

Fortran Preprocessor Options

The Fortran preprocessor (fpp) may be invoked automatically or by specifying option `fpp`.

The following options are available if `fpp` is in effect.

fpp Option	Description
<code>-B</code>	Specifies that C++-style comments should not be recognized.

fpp Option	Description
-C	Specifies that C-style comments should not be recognized. This is the same as specifying <code>-c_com=no</code> .
-c_com={yes no}	Determines whether C-style comments are recognized. If you specify <code>-c_com=no</code> or <code>-C</code> , C-style comments are not recognized. By default, C-style comments are recognized; that is, <code>-c_com=yes</code> .
-Dname	<p>Defines the preprocessor variable name as 1 (one). This is the same as if a <code>-Dname=1</code> option appeared on the fpp command line, or as if a</p> <pre data-bbox="810 1094 1032 1119">#define name 1</pre> <p>line appeared in the source file processed by fpp.</p>
-Dname=def	<p>Defines name as if by a <code>#define</code> directive. This is the same as if a</p> <pre data-bbox="810 1430 1065 1455">#define name def</pre> <p>line appeared in the source file processed by fpp. The <code>-D</code> option has lower precedence than the <code>-U</code> option. That is, if the same name is used in both a <code>-U</code> option and a <code>-D</code> option, the name will be undefined regardless of the order of the options.</p>

fpp Option	Description
-e	Tells the compiler to accept extended source lines. For fixed format, lines can contain up to 132 characters. For free format, lines can contain up to 32768 characters.
-e80	Tells the compiler to accept extended source lines. For fixed format, lines can contain up to 80 characters.
-fixed	Tells the compiler to assume fixed format in the source file.
-free	Tells the compiler to assume free format in the source file.
-f_com={yes no}	<p>Determines whether Fortran-style end-of-line comments are recognized or ignored by fpp. If you specify -f_com=no, Fortran style end-of-line comments are processed as part of the preprocessor directive. By default, Fortran style end-of-line comments are recognized by fpp on preprocessor lines and are ignored by fpp; that is, -f_com=yes. For example:</p> <pre data-bbox="810 1587 1256 1646">#define max 100 ! max number do i = 1, max + 1</pre> <p>If you specify -f_com=yes, fpp will output</p> <pre data-bbox="810 1833 1081 1856">do i = 1, 100 + 1</pre>

fpp Option	Description
	<p>If you specify <code>-f_com=no</code>, fpp will output</p> <pre data-bbox="818 436 1289 468">do i = 1, 100 ! max number + 1</pre>
-help	Displays information about fpp options.
-I<dir>	<p>Inserts directory <dir> into the search path for <code>#include</code> files with names not beginning with <code>"/</code>. The <dir> is inserted ahead of the standard list of "include" directories so that <code>#include</code> files with names enclosed in double-quotes (<code>"</code>) are searched for first in the directory of the file with the <code>#include</code> line, then in directories named with <code>-I</code> options, and lastly, in directories from the standard list. For <code>#include</code> files with names enclosed in angle-brackets (<code><></code>), the directory of the file with the <code>#include</code> line is not searched.</p>
-m	Expands macros everywhere. This is the same as <code>-macro=yes</code> .
macro={yes no_com no}	<p>Determines the behavior of macro expansion. If you specify <code>-macro=no_com</code>, macro expansion is turned off in comments. If you specify <code>-macro=no</code>, no macro expansion occurs anywhere. By default, macros are expanded everywhere; that is, <code>-</code></p>

fpp Option	Description
-noB	<p>macro=yes.</p> <p>Specifies that C++-style comments should be recognized.</p>
-noC	<p>Specifies that C-style comments should be recognized. This is the same as -c_com=yes.</p>
-noJ	<p>Specifies that F90-style comments should be recognized in a #define line. This is the same as -f_com=no.</p>
-no-fort-cont	<p>Specifies that IDL style format should be recognized. This option is only for the IDE. Note that macro arguments in IDL may have a C-like continuation character "\" which is different from the Fortran continuation character "&". Fpp should recognize the C-like continuation character and process some other non-Fortran tokens so that the IDL processor can recognize them.</p>
-P	<p>Tells the compiler that line numbering directives should not be added to the output file. This line-numbering directive appears as</p>
-Uname	<p>#line-number file-name</p> <p>Removes any initial definition of name, where name is an fpp variable that is</p>

fpp Option	Description
-undef	<p>predefined on a particular preprocessor. Here is a partial list of symbols that may be predefined, depending upon the architecture of the system:</p> <p>Operating System: __APPLE__, __unix, and __linux</p> <p>Hardware: __i386, __ia64, __x86_64</p>
-V	Removes initial definitions for all predefined symbols.
-w[0]	Displays the fpp version number.
-Xu	Prevents warnings from being output. By default, warnings are output to stderr.
-Xw	Converts uppercase letters to lowercase, except within character-string constants. The default is to leave the case as is.
-Y<dir>	Tells the compiler that in fixed-format source files, the blank or space symbol " " is insignificant. By default, the space symbol is the delimiter of tokens for this format.
-Y<dir>	Adds directory <dir> to the end of the system include paths.

For details on how to specify these options on the compiler command line, see [fpp, Qfpp](#).

Floating-point Operations

Overview: Floating-point Operations

This section introduces the floating-point support in the Intel® Fortran Compiler and provides information about using floating-point operations in your applications. The section also briefly describes the IEEE* Floating-Point Standard (IEEE 754).

The following table lists some possible starting points:

If you are trying to...	Then start with...
Understand the programming objectives of floating-point applications	Programming Objectives of Floating-Point Applications
Use the <code>-fp-model</code> (Linux* and Mac OS* X) or <code>/fp</code> (Windows*) option	Using the <code>-fp-model</code> or <code>/fp</code> Option
Set the flush-to-zero (FTZ) or denormals-are-zero (DAZ) flags	Setting the FTZ and DAZ Flags
Handle floating-point exceptions	Handling Floating-Point Exceptions
Tuning the performance of floating-point applications	Overview: Tuning Performance of Floating-Point Applications
Learn about the IEEE Floating-Point Standard	Overview: Understanding IEEE Floating-Point Standard

Floating-point Options Quick Reference

The Intel® Compiler provides various options for you to optimize floating-point calculations with varying degrees of accuracy and predictability on different Intel architectures. This topic lists these compiler options and provides information about their supported architectures and operating systems.

IA-32, Intel® 64, and IA-64 architectures

Linux* and Mac OS* X	Windows*	Description
<code>-fp-model</code>	<code>/fp</code>	Specifies semantics used in floating-point

Linux* and Mac OS* X	Windows*	Description
<code>-fp-speculation</code>	<code>/Qfp-speculation</code>	<p>calculations. Values are precise, fast [=1/2], strict, source, double, extended, [no-]except and except[-].</p> <ul style="list-style-type: none"> • -fp-model compiler option <p>Specifies the speculation mode for floating-point operations. Values are fast, safe, strict, and off.</p> <ul style="list-style-type: none"> • -fp-speculation compiler option
<code>-prec-div</code>	<code>/Qprec-div</code>	<p>Attempts to use slower but more accurate implementation of floating-point divide. Use this option to disable the divide optimizations in cases where it is important to maintain the full range and precision for floating-point division. Using this option results in greater accuracy with some loss of performance.</p> <p>Specifying <code>-no-prec-div</code></p>

Linux* and Mac OS* X	Windows*	Description
		<p>(Linux and Mac OS X) or /Qprec-div- (Windows) enables optimizations that result in slightly less precise results than full IEEE division.</p> <ul style="list-style-type: none"> • -prec-div compiler option
-complex-limited-range	/Qcomplex-limited-range	<p>Enables the use of basic algebraic expansions of some arithmetic operations involving data of type COMPLEX. This can cause performance improvements in programs that use a lot of COMPLEX arithmetic. Values at the extremes of the exponent range might not compute correctly.</p> <ul style="list-style-type: none"> • -complex-limited-range compiler option
-ftz	/Qftz	<p>The default behavior depends on the architecture. Refer to the following topic for details:</p> <ul style="list-style-type: none"> • -ftz compiler option
-fpe	/fpe	<p>By default, the Fortran</p>

Linux* and Mac OS* X	Windows*	Description
		<p>compiler disables all floating-point exceptions and floating underflow is gradual.</p> <p>This option controls which exceptions are enabled by the compiler. It also controls whether floating-point underflow is gradual or abrupt.</p> <ul style="list-style-type: none"> • -fpe compiler option

IA-32 and Intel® 64 architectures

Linux* and Mac OS* X	Windows*	Description
<code>-prec-sqrt</code>	<code>/Qprec-sqrt</code>	<p>Improves the accuracy of square root implementations, but using this option may impact speed.</p> <ul style="list-style-type: none"> • -prec-sqrt compiler option
<code>-pc</code>	<code>/Qpc</code>	<p>Changes the floating point significant precision. Use this option when compiling applications.</p> <p>The application must use <code>PROGRAM</code> as the entry point, and you must</p>

Linux* and Mac OS* X	Windows*	Description
		compile the source file containing PROGRAM with this option.
		<ul style="list-style-type: none"> • -pc compiler option
-rcd	/Qrcd	Disables rounding mode changes for floating-point-to-integer conversions.
		<ul style="list-style-type: none"> • -rcd compiler option
-fp-port	/Qfp-port	Causes floating-point values to be rounded to the source precision at assignments and casts.
		<ul style="list-style-type: none"> • -fp-port compiler option
-mp1	/Qprec	This option rounds floating-point values to the precision specified in the source program prior to comparisons. It also implies <code>-prec-div</code> and <code>-prec-sqrt</code> (Linux and Mac OS X) or <code>/Qprec-div</code> and <code>/Qprec-sqrt</code> (Windows).
		<ul style="list-style-type: none"> • -mp1 compiler option

IA-64 architecture only

Linux*	Windows*	Description
--------	----------	-------------

Linux*	Windows*	Description
<code>-IPF-fma</code>	<code>/QIPF-fma</code>	<p>Enables or disables the contraction of floating-point multiply and add/subtract operations into a single operation.</p> <ul style="list-style-type: none"> • -IPF-fma compiler option
<code>-IPF-fp-relaxed</code> (deprecated)	<code>/QIPF-fp-relaxed</code> (deprecated)	<p>Enables use of faster but slightly less accurate code sequences for math functions, such as the <code>sqrt()</code> function and the divide operation. As compared to strict IEEE* precision, using this option slightly reduces the accuracy of floating-point calculations performed by these functions, usually limited to the least significant binary digit. This option is deprecated.</p> <ul style="list-style-type: none"> • -IPF-fp-relaxed compiler option

Understanding Floating-point Operations

Using the `-fp-model (/fp)` Option

The `-fp-model` (Linux* and Mac OS* X) or `/fp` (Windows*) option allows you to control the optimizations on floating-point data. You can use this option to tune

the performance, level of accuracy, or result consistency across platforms for floating-point applications.

For applications that do not require support for denormalized numbers, the `-fp-model` or `/fp` option can be combined with the `-ftz` (Linux* and Mac OS* X) or `/Qftz` (Windows*) option to flush denormalized results to zero in order to obtain improved runtime performance on processors based on:

- IA-32 and Intel® 64 architectures
- IA-64 architecture

You can use keywords to specify the semantics to be used. Possible values of the keywords are as follows:

Keyword	Description
<code>precise</code>	Enables value-safe optimizations on floating-point data and rounds intermediate results to source-defined precision.
<code>fast[=1 2]</code>	Enables more aggressive optimizations on floating-point data.
<code>strict</code>	Enables <code>precise</code> and <code>except</code> , disables contractions, and enables the property that allows modification of the floating-point environment.
<code>source</code>	Enables value-safe optimizations on floating-point data and rounds intermediate results to source-defined precision (same as <code>precise</code> keyword).
<code>double</code>	Rounds intermediate results to 53-bit (double) precision and enables value-safe optimizations.
<code>extended</code>	Rounds intermediate results to 64-bit (extended) precision and enables value-safe optimizations.
<code>[no-]except</code> (Linux* and Mac OS* X) or	Determines whether floating-point exception semantics are used.

Keyword	Description
---------	-------------

`except [-]`

(Windows*)

The default value of the option is `-fp-model fast=1` or `/fp:fast=1`, which means that the compiler uses more aggressive optimizations on floating-point calculations.



Note

Using the default option keyword `-fp-model fast` or `/fp:fast`, you may get significant differences in your result depending on whether the compiler uses x87 or SSE2 instructions to implement floating-point operations. Results are more consistent when the other option keywords are used.

Several examples are provided to illustrate the usage of the keywords. These examples show:

- A small example of source code
Note that the same source code is considered in all the included examples.
- The semantics that are used to interpret floating-point calculations in the source code
- One or more possible ways the compiler may interpret the source code
Note that there are several ways the compiler may interpret the code; we show just some of these possibilities.

`-fp-model fast` or `/fp:fast`

Example source code:

```
REAL T0, T1, T2;
...
T0 = 4.0E + 0.1E + T1 + T2;
```

When this option is specified, the compiler applies the following semantics:

- Additions may be performed in any order
- Intermediate expressions may use single, double, or extended double precision

- The constant addition may be pre-computed, assuming the default rounding mode

Using these semantics, the following shows some possible ways the compiler may interpret the original code:

```
REAL T0, T1, T2;  
...  
T0 = (T1 + T2) + 4.1E;  
REAL T0, T1, T2;  
...  
T0 = (T1 + 4.1E) + T2;
```

-fp-model source or /fp:source

This setting is equivalent to `-fp-model precise or /fp:precise` on systems based on the Intel® 64 architecture.

Example source code:

```
REAL T0, T1, T2;  
...  
T0 = 4.0E + 0.1E + T1 + T2;
```

When this option is specified, the compiler applies the following semantics:

- Additions are performed in program order, taking into account any parentheses
- Intermediate expressions use the precision specified in the source code
- The constant addition may be pre-computed, assuming the default rounding mode

Using these semantics, the following shows a possible way the compiler may interpret the original code:

```
REAL T0, T1, T2;  
...  
T0 = ((4.1E + T1) + T2);
```

-fp-model strict or /fp:strict

Example source code:

```
REAL T0, T1, T2;  
...  
T0 = 4.0E + 0.1E + T1 + T2;
```

When this option is specified, the compiler applies the following semantics:

- Additions are performed in program order, taking into account any parentheses

- Intermediate expressions use the precision specified in the source code. In Fortran, intermediate expressions always use source precision in modes other than fast.
- The constant addition will not be pre-computed, because there is no way to tell what rounding mode will be active when the program runs.

Using these semantics, the following shows a possible way the compiler may interpret the original code:

```
REAL T0, T1, T2;
...
T0 = REAL ( ((REAL)4.0E + (REAL)0.1E) + (REAL)T1) + (REAL)T2);
```

See Also

[-fp-model compiler option](#) Controls the semantics of floating-point calculations.

[/fp compiler option](#) Controls the semantics of floating-point calculations.

Floating-point Optimizations

Application performance is an important goal of the Intel® Compilers, even at default optimization levels. A number of optimizations involve transformations, such as evaluation of constant expressions at compiler time, hoisting invariant expressions out of loops, or changes in the order of evaluation of expressions. These optimizations usually help the compiler produce most efficient code possible. However, this may not be true for floating-point applications, because some optimizations may affect accuracy, reproducibility, and performance. Some optimizations are not consistent with strict interpretation of the ANSI or ISO standards for Fortran, which can result in differences in rounding and small variants in floating-point results that may be more or less accurate than the ANSI-conformant result.

Intel Compilers provide the `-fp-model` (Linux* and Mac OS* X) or `/fp` (Windows*) option, which allows you to control the optimizations performed when you build an application. The option allows you to specify the compiler rules for:

- Value safety: Whether the compiler may perform transformations that could affect the result. For example, in the SAFE mode, the compiler won't transform x/x to 1.0. The UNSAFE mode is the default.

- Floating-point expression evaluation: How the compiler should handle the rounding of intermediate expressions.
- Floating-point contractions: Whether the compiler should generate floating-point multiply-add (FMA) on processors based on the IA-64 architecture. When enabled, the compiler may generate FMA for combined multiply/add; when disabled, the compiler must generate separate multiply/add with intermediate rounding.
- Floating-point environment access: Whether the compiler must account for the possibility that the program might access the floating-point environment, either by changing the default floating-point control settings or by reading the floating-point status flags. This is disabled by default. You can use the `-fp-model:strict` (Linux and Mac OS X) / `/fp:strict` (Windows) option to enable it.
- Precise floating-point exceptions: Whether the compiler should account for the possibility that floating-point operations might produce an exception. This is disabled by default. You can use `-fp-model:strict` (Linux and Mac OS X) or `/fp:strict` (Windows); or `-fp-model:except` (Linux and Mac OS X) or `/fp:except` (Windows) to enable it.

The following table describes the impact of different keywords of the option on compiler rules and optimizations:

Keyword	Value Safety	Floating-Point Expression Evaluation	Floating-Point Contractions	Floating-Point Environment Access	Precise Floating-Point Exceptions
<code>precise</code>	Safe	Varies	Yes	No	No
<code>source</code>		Source			
<code>strict</code>	Safe	Varies	No	Yes	Yes
<code>fast=1</code> (default)	Unsafe	Unknown	Yes	No	No

Keyword	Value Safety	Floating- Point Expression Evaluation	Floating- Point Contractions	Floating- Point Environment Access	Precise Floating- Point Exceptions
<code>fast=2</code>	Very unsafe	Unknown	Yes	No	No
<code>except</code>	Unaffected	Unaffected	Unaffected	Unaffected	Yes
<code>except-</code>	Unaffected	Unaffected	Unaffected	Unaffected	No



Note

It is illegal to specify the `except` keyword in an unsafe safety mode. Based on the objectives of an application, you can choose to use different sets of compiler options and keywords to enable or disable certain optimizations, so that you can get the desired result.

Programming Objectives of Floating-point Applications

In general, the programming objectives of the floating-point applications fall into the following categories:

- Accuracy: The application produces that results that are close to the correct result.
- Reproducibility and portability: The application produces results that are consistent across different runs, different set of build options, different compilers, different platforms, and different architectures.
- Performance: The application produces the most efficient code possible.

Based on the goal of an application, you will need to balance the tradeoffs among these objectives. For example, if you are developing a 3D graphics engine, then performance can be the most important factor to consider, and reproducibility and accuracy can be your secondary concerns.

Intel® Compiler provides appropriate compiler options, such as the `-fp-model` (Linux* and Mac OS* X) or `/fp` (Windows*) option, which allows you to tune your

applications based on specific objectives. The compiler processes the code differently when you specify different compiler options.

In most case, an application will be much more complicated. You should select appropriate compiler options by carefully considering your programming objectives and balancing the tradeoffs among these objectives.

Denormal Numbers

A normalized number is a number for which both the exponent (including offset) and the most significant bit of the mantissa are non-zero. For such numbers, all the bits of the mantissa contribute to the precision of the representation.

The smallest normalized single precision floating-point number greater than zero is about 1.1754943^{-38} . Smaller numbers are possible, but those numbers must be represented with a zero exponent and a mantissa whose leading bit(s) are zero, which leads to a loss of precision. These numbers are called denormalized numbers; denormals (newer specifications refer to these as subnormal numbers).

Denormal computations use both hardware or operating system resources to handle them, which can cost hundreds of clock cycles.

- Denormal computations take much longer to calculate on processors based on IA-32 and Intel® 64 architectures than normal computations.
- Denormals are computed in software on processors based on the IA-64 architecture, and the computation usually requires hundreds of clock cycles, which results in excessive kernel time.

There are several ways to handle denormals and increase the performance of your application:

- Scale the values into the normalized range.
- Use a higher precision data type with a larger dynamic range.
- Flush denormals to zero.

See Also

Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 1: Basic Architecture

Intel® Itanium® Architecture Software Developer's Manual, Volume 1:
Application Architecture

Institute of Electrical and Electronics Engineers, Inc*. (IEEE) web site for
information about the current floating-point standards and recommendations

Floating-point Environment

The floating-point environment is a collection of registers that control the behavior of the floating-point machine instructions and indicate the current floating-point status. The floating-point environment can include rounding mode controls, exception masks, flush-to-zero (FTZ) controls, exception status flags, and other floating-point related features.

For example, on IA-32 and Intel® 64 architectures, bit 15 of the MXCSR register enables the flush-to-zero mode, which controls the masked response to an single-instruction multiple-data (SIMD) floating-point underflow condition.

The floating-point environment affects most floating-point operations; therefore, correct configuration to meet your specific needs is important. For example, the exception mask bits define which exceptional conditions will be raised as exceptions by the processor. In general, the default floating-point environment is set by the operating system. You don't need to configure the floating-point environment unless the default floating-point environment does not suit your needs.

There are several methods available if you want to modify the default floating-point environment. For example, you can use inline assembly, compiler built-in functions, and library functions.

Changing the default floating-point environment affects runtime results only. This does not affect any calculations which are pre-computed at compile time.

See Also

Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 1: Basic Architecture

Setting the FTZ and DAZ Flags

In Intel® processors, the flush-to-zero (FTZ) and denormals-are-zero (DAZ) flags in the MXCSR register are used to control floating-point calculations. When the FTZ and DAZ flags are enabled, the Single Instructions and Multiple Data (SIMD) floating-point computation can be accelerated, thus improving the performance of the application.

You can use the `-ftz` (Linux* and Mac OS* X) or `/Qftz` (Windows*) option to flush denormal results to zero when the application is in the gradual underflow mode. This option may improve performance if the denormal values are not critical to your application's behavior.

The `-ftz` or `/Qftz` option sets or resets the FTZ and the DAZ hardware flags. The following table describes how the compiler process denormal values based on the status of the FTZ and DAZ flags:

Flag	When set to ON, the compiler...	When set to OFF, the compiler...	Supported on
FTZ	Sets denormal results from floating-point calculations to zero.	Does not change the denormal results.	IA-64 and Intel® 64 architectures
DAZ	Treats denormal values used as input to floating-point instructions as zero.	Does not change the denormal instruction inputs.	Intel® 64 architecture

- FTZ and DAZ are not supported on all IA-32 architectures. On systems based on the IA-64 architecture, FTZ always works, while on systems based on the IA-32 and Intel® 64 architectures, it only applies to SSE instructions. Hence if your application happened to generate denormals using x87 instructions, FTZ does not apply.

- DAZ and FTZ flags are not compatible with IEEE Standard 754, so you should only consider enabling them when strict compliance to the IEEE standard is not required and application performance has higher priority than application accuracy.

Options `-ftz` and `/Qftz` are performance options. Setting these options does not guarantee that all denormals in a program are flushed to zero. They only cause denormals generated at run time to be flushed to zero.

When `-ftz` or `/Qftz` is used in combination with an SSE-enabling option on systems based on the IA-32 architecture (for example, `-xW` or `/QxW`), the compiler will insert code in the main routine to set FTZ and DAZ. When `-ftz` or `/Qftz` is used without such an option, the compiler will insert code to conditionally set FTZ/DAZ based on a run-time processor check. `-no-ftz` (Linux and Mac OS X) or `/Qftz-` (Windows) will prevent the compiler from inserting any code that might set FTZ or DAZ.

The `-ftz` or `/Qftz` option only has an effect when the main program is being compiled. It sets the FTZ/DAZ mode for the process. The initial thread and any threads subsequently created by that process will operate in the FTZ/DAZ mode. On systems based on the IA-64 architecture, optimization option O3 sets `-ftz` and `/Qftz`; optimization option O2 sets `-no-ftz` (Linux) and `/Qftz-` (Windows). On systems based on the IA-32 and Intel® 64 architectures, every optimization option O level, except O0, sets `-ftz` and `/Qftz`.

If this option produces undesirable results of the numerical behavior of your program, you can turn the FTZ/DAZ mode off by using `-no-ftz` or `/Qftz-` in the command line while still benefiting from the O3 optimizations.

For some non-Intel processors, the flags can be set manually by calling the following Intel Fortran intrinsic:

Example

```
RESULT = FOR_SET_FPE (FOR_M_ABRUPT_UND)
```

See Also

Tuning Performance

Overview: Tuning Performance

This section describes several programming guidelines that can help you improve the performance of a floating-point applications:

- [Avoid exact floating-point comparisons](#)
- Avoid exceeding representable ranges during computation; handling these cases can have a performance impact.
- Use REAL variables in single precision format unless the extra precision obtained through DOUBLE or REAL*8 is required because a larger precision formation will also increase memory size and bandwidth requirements.
- [Reduce the impact of denormal exceptions for all supported architectures.](#)
- [Avoid mixed data type arithmetic expressions.](#)

Avoiding Exact Floating-point Comparison

It is unsafe for applications to rely on exact floating-point comparisons. Slight variations in rounding can change the outcome of such comparisons, leading to non-convergence or other unexpected behavior.

Tests for equality of floating-point quantities should be made within some tolerance related to the expected precision of the calculation, for example, by using the Fortran intrinsic function EPSILON. The following examples demonstrate the concept:

Example

```
if (foo() == 2.0)
```

Where `foo()` may be as close to 2.0 as can be imagined without actually exactly matching 2.0. You can improve the behavior of such codes by using inexact floating-point comparisons or fuzzy comparisons to test a value to within a certain tolerance, as shown below:

Example

```
epsilon = 1E-8;  
if (abs(foo() - 2.0) <= epsilon)
```

Handling Floating-point Array Operations in a Loop Body

Following the guidelines below will help autovectorization of the loop.

- Statements within the loop body may contain float or double operations (typically on arrays). The following arithmetic operations are supported: addition, subtraction, multiplication, division, negation, square root, MAX, MIN, and mathematical functions such as SIN and COS.
- Writing to a float scalar/array and a double scalar/array within the same loop decreases the chance of autovectorization due to the differences in the vector length (that is, the number of elements in the vector register) between float and double types. If autovectorization fails, try to avoid using mixed data types.

Reducing the Impact of Denormal Exceptions

Denormalized floating-point values are those that are too small to be represented in the normal manner; for example, the mantissa cannot be left-justified.

Denormal values require hardware or operating system interventions to handle the computation, so floating-point computations that result in denormal values may have an adverse impact on performance.

There are several ways to handle denormals to increase the performance of your application:

- Scale the values into the normalized range
- Use a higher precision data type with a larger dynamic range
- Flush denormals to zero

For example, you can translate them to normalized numbers by multiplying them using a large scalar number, doing the remaining computations in the normal space, then scaling back down to the denormal range. Consider using this method when the small denormal values benefit the program design.

If you change the declaration of a variable you might also need to change the libraries you call to use the variable. Another strategy that might result in

increased performance is to increase the amount of precision of intermediate values using the `-fp-model [double|extended]` option; however, if you increased precision as the solution to slow performance caused by denormal numbers you must verify the resulting changes actually increase performance. Finally, In many cases denormal numbers be treated safely as zero without adverse effects on program results. Depending on the target architecture, use flush-to-zero (FTZ) options.

IA-32 and Intel® 64 Architectures

These architectures take advantage of the FTZ and DAZ (denormals-are-zero) capabilities of Streaming SIMD Extensions (SSE), Streaming SIMD Extensions 2 (SSE2), and Streaming SIMD Extensions 3 (SSE3), and Supplemental Streaming SIMD Extensions 3 (SSSE3) instructions.

By default, the compiler for the IA-32 architecture generates code that will run on machines that do not support SSE instructions. The compiler implements floating-point calculations using the x87 floating-point unit, which does not benefit from the FTZ and DAZ settings. You can use the `-x` (Linux* and Mac OS* X) or `/Qx` (Windows*) option to enable the compiler to implement floating-point calculations using the SSE and SSE2 instructions. The compiler for the Intel® 64 architecture generates SSE2 instructions by default.

The FTZ and DAZ modes are enabled by default when you compile the source file containing PROGRAM using the Intel Compiler. The compiler generates a call to a library routine that performs a runtime processor check. The FTZ and DAZ modes are set provided that the modes are available for the machine on which the program is running.

IA-64 Architecture

Enable the FTZ mode by using the `-ftz` (Linux and Mac OS X) or `/Qftz` (Windows) option on the source file containing PROGRAM. The `-O3` (Linux and Mac OS X) or `/O3` (Windows) option automatically enables `-ftz` or `/Qftz`.



Note

After using flush-to-zero, ensure that your program still gives correct results when treating denormalized values as zero.

See Also

[Setting the FTZ and DAZ Flags](#)

Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 1: Basic Architecture

Avoiding Mixed Data Type Arithmetic Expressions

Avoid mixing integer and floating-point (REAL) data in the same computation. Expressing all numbers in a floating-point arithmetic expression (assignment statement) as floating-point values eliminates the need to convert data between fixed and floating-point formats. Expressing all numbers in an integer arithmetic expression as integer values also achieves this. This improves run-time performance.

For example, assuming that `I` and `J` are both INTEGER variables, expressing a constant number (2.) as an integer value (2) eliminates the need to convert the data. The following examples demonstrate inefficient and efficient code.

Example: Inefficient Code

```
INTEGER I, J
I = J / 2.
```

Example: Efficient Code

```
INTEGER I, J
I = J / 2
```

You can use different sizes of the same general data type in an expression with minimal or no effect on run-time performance. For example, using REAL, DOUBLE PRECISION, and COMPLEX floating-point numbers in the same

floating-point arithmetic expression has minimal or no effect on run-time performance. However, this practice of mixing different sizes of the same general data type in an expression can lead to unexpected results due to operations being performed in a lower precision than desired.

Using Efficient Data Types

In cases where more than one data type can be used for a variable, consider selecting the data types based on the following hierarchy, listed from most to least efficient:

- Integer
- Single-precision real, expressed explicitly as `REAL`, `REAL (KIND=4)`, or `REAL*4`
- Double-precision real, expressed explicitly as `DOUBLE PRECISION`, `REAL (KIND=8)`, or `REAL*8`
- Extended-precision real, expressed explicitly as `REAL (KIND=16)` or `REAL*16`

However, keep in mind that in an arithmetic expression, you should avoid mixing integer and floating-point data.

Handling Floating-point Exceptions

Overview: Controlling Floating-point Exceptions

When developing applications, you may need to control the exceptions that can occur during the run-time processing of floating-point numbers. These exceptions can be categorized into specific types: overflow, divide-by-zero, underflow, and invalid operations.

Overflow

Overflow is signaled whenever the destination format's largest finite number is exceeded in magnitude by what would have been the rounded floating-point result. The result computed is rounding mode specific:

- Round-to-nearest (default): +/- Infinity in specified precision

- Round-to-zero: +/- Maximum Number in specified precision
- Round-to-+Infinity: +Infinity or -(Maximum Positive Number) in specified precision
- Round-to--Infinity: (Maximum Positive Number) or -Infinity in specified precision

For example, in round-to-nearest mode $1E30 * 1E30$ overflows the single-precision floating-point range and results in a +Infinity; $-1E30 * 1E30$ results in a -Infinity.

Divide-by-zero

Divide-by-zero is signaled when the divisor is zero and the dividend is a finite nonzero number. The computed result is a correctly signed Infinity.

For example, $2.0E0/+0.0$ produces a divide-by-zero exception and results in a +Infinity; $-2.0E0/+0.0$ produces a divide-by-zero exception and results in a -Infinity.

Underflow

Underflow occurs when a computed result (of an add, subtract, multiply, divide, or math function call) falls beyond the minimum range in magnitude of normalized numbers of the floating-point data type. Each floating-point type (32-, 64-, and 128-bit) has a denormalized range where very small numbers can be represented with some loss of precision. This is called gradual underflow. For example, the lower bound for normalized single-precision floating-point is approximately $1E-38$, while the lower bound for denormalized single-precision floating-point is approximately $1E-45$. Results falling below the lower bound of the denormalized range simply become zero. $1E-30 / 1E10$ underflows the normalized range but not the denormalized range so the result is the denormal value $1E-40$. $1E-30 / 1E30$ underflows the entire range and the result is zero.

Invalid operation

Invalid occurs when operands to the basic floating-point operations or math function inputs produce an undefined (QNaN) result. Some examples include:

- SNaN operand in any floating-point operation or math function call
- Division of zeroes: $(+/-0.0)/(+/-0.0)$
- Sum of Infinities having different signs: Infinity + (-Infinity)
- Difference of Infinities having the same sign: $(+/-Infinity) - (+/-Infinity)$
- Product of signed Infinities with zero: $(+/-Inf) * 0$
- Math Function Domain Errors: $\log(\text{negative})$, $\sqrt{\text{negative}}$, $\text{asin}(|x|>1)$

With the Intel® Compiler, you can use the `-fpe` (Linux* and Mac OS* X) or `/fpe` (Windows*) option to control floating-point exceptions.

Handling Floating-point Exceptions

If a floating-point exception is disabled (its bit is set to 1 with SETCONTROLFPQQ (IA-32 architecture only), it will not generate an interrupt signal if it occurs. The floating-point process may return an appropriate special value (for example, NaN or signed infinity) or may return an acceptable value (for example, in the case of a denormal operand), and the program will continue. If a floating-point exception is enabled (its bit is set to 0), it will generate an interrupt signal (software interrupt) if it occurs.

The following table lists the floating-point exception signals:

Parameter Name	Value in Hex	Description
FPE\$INVALID	#81	Invalid result
FPE\$DENORMAL	#82	Denormal operand
FPE\$ZERODIVIDE	#83	Divide by zero
FPE\$OVERFLOW	#84	Overflow
FPE\$UNDERFLOW	#85	Underflow
FPE\$INEXACT	#86	Inexact precision

If a floating-point exception interrupt occurs and you do not have an exception handling routine, the run-time system will respond to the interrupt according to

the behavior selected by the compiler option `/fpe`. Remember, interrupts only occur if an exception is enabled (set to 0).

If you do not want the default system exception handling, you need to write your own interrupt handling routine:

- Write a function that performs whatever special behavior you require on the interrupt.
- Register that function as the procedure to be called on that interrupt with `SIGNALQQ`.

Note that your interrupt handling routine must use the `cDEC$ ATTRIBUTES option C`.

The drawback of writing your own routine is that your exception-handling routine cannot return to the process that caused the exception. This is because when your exception-handling routine is called, the floating-point processor is in an error condition, and if your routine returns, the processor is in the same state, which will cause a system termination. Your exception-handling routine can therefore either branch to another separate program unit or exit (after saving your program state and printing an appropriate message). You cannot return to a different statement in the program unit that caused the exception-handling routine, because a global `GOTO` does not exist, and you cannot reset the status word in the floating-point processor.

If you need to know when exceptions occur and also must continue if they do, you must disable exceptions so they do not cause an interrupt, then poll the floating-point status word at intervals with `GETSTATUSFPQQ` (IA-32 architecture only) to see if any exceptions occurred. To clear the status word flags, call the `CLEARSTATUSFPQQ` (IA-32 architecture only) routine.

Polling the floating-point status word at intervals creates processing overhead for your program. In general, you will want to allow the program to terminate if there is an exception. An example of an exception-handling routine follows. The comments at the beginning of the `SIGTEST.F90` file describe how to compile this example.

```
! SIGTEST.F90
!Establish the name of the exception handler as the
```

```

! function to be invoked if an exception happens.
! The exception handler hand_fpe is attached below.
USE IFPORT
INTERFACE
  FUNCTION hand_fpe (sigid, except)
    !DEC$ ATTRIBUTES C :: hand_fpe
    INTEGER(4) hand_fpe
    INTEGER(2) sigid, except
  END FUNCTION
END INTERFACE
INTEGER(4) iret
REAL(4) r1, r2
r1 = 0.0
iret = SIGNALQQ(SIG$FPE, hand_fpe)
WRITE(*,*) 'Set exception handler. Return = ', iret
! Cause divide-by-zero exception
r1 = 0.0
r2 = 3/r1
END
! Exception handler routine hand_fpe
FUNCTION hand_fpe (signum, excnum)
  !DEC$ ATTRIBUTES C :: hand_fpe
  USE IFPORT
  INTEGER(2) signum, excnum
  WRITE(*,*) 'In signal handler for SIG$FPE'
  WRITE(*,*) 'signum = ', signum
  WRITE(*,*) 'exception = ', excnum
  SELECT CASE(excnum)
    CASE( FPE$INVALID )
      STOP ' Floating point exception: Invalid number'
    CASE( FPE$DENORMAL )
      STOP ' Floating point exception: Denormalized number'
    CASE( FPE$ZERODIVIDE )
      STOP ' Floating point exception: Zero divide'
    CASE( FPE$OVERFLOW )
      STOP ' Floating point exception: Overflow'
    CASE( FPE$UNDERFLOW )
      STOP ' Floating point exception: Underflow'
    CASE( FPE$INEXACT )
      STOP ' Floating point exception: Inexact precision'
    CASE DEFAULT
      STOP ' Floating point exception: Non-IEEE type'
  END SELECT
  hand_fpe = 1
END
END

```

Checking the Floating-point Stack State

On systems based on the IA-32 architectures, when an application calls a function that returns a floating-point value, the returned floating-point value is supposed to be on the top of the floating-point stack. If the return value is not

used, the compiler must pop the value off of the floating-point stack in order to keep the floating-point stack in the correct state.

On systems based on Intel(R) 64 architectures, floating-point values are usually returned in the xmm0 register. The floating-point stack is used only when the return value is a long doublean internal 80-bit floating-point data type on Linux* and Mac OS* X systems.

If the application calls a function without defining or incorrectly defining the function's prototype, the compiler cannot determine if the function must return a floating-point value. Consequently, the return value is not popped off the floating-point stack if it is not used. This can cause the floating-point stack to overflow.

The overflow of the stack results in two undesirable situations:

- A NaN value gets involved in the floating-point calculations
- The program results become unpredictable; the point where the program starts making errors can be arbitrarily far away from the point of the actual error.

For systems based on the IA-32 and Intel® 64 architectures, the `-fp-stack-check` (Linux* and Mac OS* X) or `/Qfp-stack-check` (Windows*) option checks whether a program makes a correct call to a function that should return a floating-point value. If an incorrect call is detected, the option places a code that marks the incorrect call in the program. The `-fp-stack-check` (Linux* and Mac OS* X) or `/Qfp-stack-check` (Windows*) option marks the incorrect call and makes it easy to find the error.



Note

The `-fp-stack-check` (Linux* and Mac OS* X) and the `/Qfp-stack-check` (Windows*) option causes significant code generation after every function/subroutine call to ensure that the floating-point stack is maintained in the correct state. Therefore, using this option slows down the program being compiled. Use the option only as a debugging aid to find floating point stack underflow/overflow problems, which can be otherwise hard to find.

See Also

[-fp-stack-check, /Qfp-stack-check option](#) Tells the compiler to generate extra code after every function call to ensure that the floating-point stack is in the expected state.

File *fordef.for* and Its Usage

The parameter file *fordef.for* contains symbols and INTEGER*4 values corresponding to the classes of floating-point representations. Some of these classes are exceptional ones such as bit patterns that represent positive denormalized numbers.

With this file of symbols and with the FP_CLASS intrinsic function, you have the flexibility of identifying exceptional numbers so that, for example, you can replace positive and negative denormalized numbers with true zero.

The following is a simple example of identifying floating-point bit representations:

```
include 'fordef.for'
real*4 a
integer*4 class_of_bits
a = 57.0
class_of_bits = fp_class(a)
if ( class_of_bits .eq. for_k_fp_pos_norm .or. &
     class_of_bits .eq. for_k_fp_neg_norm ) then
  print *, a, ' is a non-zero and non-exceptional value'
else
  print *, a, ' is zero or an exceptional value'
end if
end
```

In this example, the symbol *for_k_fp_pos_norm* in the file *fordef.for* plus the REAL*4 value 57.0 to the FP_CLASS intrinsic function results in the execution of the first print statement.

The table below explains the symbols in the file *fordef.for* and their corresponding floating-point representations.

Symbols in *fordef.for*

Symbol Name	Class of Floating-Point Bit Representation
FOR_K_FP_SNAN	Signaling NaN
FOR_K_FP_QNAN	Quiet NaN
FOR_K_FP_POS_INF	Positive infinity

Symbol Name	Class of Floating-Point Bit Representation
FOR_K_FP_NEG_INF	Negative infinity
FOR_K_FP_POS_NORM	Positive normalized finite number
FOR_K_FP_NEG_NORM	Negative normalized finite number
FOR_K_FP_POS_DENORM	Positive denormalized number
FOR_K_FP_NEG_DENORM	Negative denormalized number
FOR_K_FP_POS_ZERO	Positive zero
FOR_K_FP_NEG_ZERO	Negative zero

Another example of using file *fordef.for* and intrinsic function `FP_CLASS` follows. The goals of this program are to quickly read any 32-bit pattern into a `REAL*4` number from an unformatted file with no exception reporting and to replace denormalized numbers with true zero:

```
include 'fordef.for'
real*4 a(100)
integer*4 class_of_bits
! open an unformatted file as unit 1
!
!   ...
read (1) a
do i = 1, 100
  class_of_bits = fp_class(a(i))
  if ( class_of_bits .eq. for_k_fp_pos_denorm .or. &
      class_of_bits .eq. for_k_fp_neg_denorm ) then
    a(i) = 0.0
  end if
end do
close (1)
end
```

You can compile this program with any value of `-fpem` (Linux* and Mac OS* X) or `/fpe:n` (Windows*). Intrinsic function `FP_CLASS` helps to find and replace denormalized numbers with zeroes before the program can attempt to perform calculations on the denormalized numbers.

On the other hand, if this program did not replace denormalized numbers read from unit 1 with zeroes and the program was compiled with `-fpe0` or `/fpe:0`, then the first attempted calculation on a denormalized number would result in a floating-point exception. If you compile with `/fpe:0` flush-to-zero is enabled. If

the resulting calculation creates a divide-by-zero, overflow, or invalid operation, then the application should abort with a floating-point exception. Otherwise, a program using the data will run to completion, perhaps faster and with different answers.

File *fordef.for* and intrinsic function `FP_CLASS` can work together to identify NaNs. A variation of the previous example would contain the symbols `for_k_fp_snan` and `for_k_fp_qnan` in the IF statement. A faster way to do this is based on the intrinsic `ISNAN` function. One modification of the previous example, using `ISNAN`, follows:

```
! The ISNAN function does not need file fordef.for
real*4 a(100)
! open an unformatted file as unit 1
!
!   ...
read (1) a
do i= 1, 100
  if ( isnan (a(i)) ) then
    print *, 'Element ', i, ' contains a NaN'
  end if
end do
close (1)
end
```

You can compile this program with any value of `-fpen` or `/fpe:n`.

Setting and Retrieving Floating-point Status and Control Words (IA-32)

Overview: Setting and Retrieving Floating-point Status and Control Words

The FPU (floating-point unit) on systems based on the IA-32 architecture contains eight floating-point registers the system uses for numeric calculations, status and control words, and error pointers. You normally need to consider only the status and control words, and then only when customizing your floating-point environment.

The FPU status and control words correspond to 16-bit registers whose bits hold the value of a state of the FPU or control its operation. Intel Fortran defines a set of symbolic constants to set and reset the proper bits in the status and control words.



Note

The symbolic constants and the library routines used to read and write the control and status registers only affect the x87 control and status registers. They do not affect the MXCSR register (the control and status register for the Intel(R) SSE and Intel(R) SSE2 instructions).

They do not affect the MXCSR (the control and status register for the Intel(R) SSE and Intel(R) SSE2 instructions). For example:

```

USE IFPORT
INTEGER(2) status, control, controlo, mask_all_traps
CALL GETCONTROLFPQQ(control)
WRITE (*, 9000) 'Control word: ', control
!   Save old control word
controlo = control
!   Clear the rounding control flags
control = IAND(control,NOT(FPCW$MCW_RC))
!   Set new control to round up
control = IOR(control,FPCW$UP)
CALL SETCONTROLFPQQ(control)
CALL GETCONTROLFPQQ(control)
WRITE (*, 9000) 'Control word: ', control
! Demonstrate setting and clearing exception mask flags
mask_all_traps = FPCW$INVALID + FPCW$DENORMAL + &
  FPCW$ZERODIVIDE + FPCW$OVERFLOW + &
  FPCW$UNDERFLOW + FPCW$INEXACT
!   Clear the exception mask flags
control = IAND(control,NOT(FPCW$MCW_EM))
!   Set new exception mask to disallow overflow
!   (i.e., enable overflow traps)
!   but allow (i.e., mask) all other exception conditions.
control = IOR(control,IEOR(mask_all_traps,FPCW$OVERFLOW))
CALL SETCONTROLFPQQ(control)
CALL GETCONTROLFPQQ(control)
WRITE (*, 9000) 'Control word: ', control
9000 FORMAT (1X, A, Z4)
END

```

The status and control symbolic constants (such as FPCW\$OVERFLOW and FPCW\$CHOP in the preceding example) are defined as INTEGER(2) parameters in the module IFORT.F90 in the . . . \INCLUDE folder. The status and control words are logical combinations (such as with .AND.) of different parameters for different FPU options.

The name of a symbolic constant takes the general form *name\$option*. The prefix *name* is one of the following:

Prefixes for Parameter Flags

name	Meaning
FPSW	Floating-point status word
FPCW	Floating-point control word
SIG	Signal
FPE	Floating-point exception
MTH	Math function

The suffix *option* is one of the options available for that *name*. The parameter *name\$option* corresponds either to a status or control option (for example, `FPSW$ZERODIVIDE`, a status word parameter that shows whether a zero-divide exception has occurred or not) or *name\$option* corresponds to a mask, which sets all symbolic constants to 1 for all the options of *name*. You can use the masks in logical functions (such as `IAND`, `IOR`, and `NOT`) to set or to clear all options for the specified *name*. The following sections define the *options* and illustrate their use with examples.

You can control the floating-point processor options (on systems based on the IA-32 architecture) and find out its status with the run-time library routines `GETSTATUSFPQQ` (IA-32 architecture only), `GETCONTROLFPQQ` (IA-32 architecture only), and `SETCONTROLFPQQ` (IA-32 architecture only). Examples of using these routines also appear in the following sections.

See Also

[Understanding Floating-Point Status Word \(IA-32 architecture only\)](#)

[Understanding Floating-Point Control Word \(IA-32 architecture only\)](#)

Understanding Floating-point Status Word

On systems based on the IA-32 architecture, the FPU status word includes bits that show the floating-point exception state of the processor. The status word parameters describe six exceptions: invalid result, denormalized operand, zero divide, overflow, underflow and inexact precision. These are described in the section, [Loss of Precision Errors](#). When one of the bits is set to 1, it means a past

floating-point operation produced that exception type. (Intel Fortran initially clears all status bits. It does not reset the status bits before performing additional floating-point operations after an exception occurs. The status bits accumulate.)

The following table shows the floating-point exception status parameters:

Parameter Name	Value in Hex	Description
FPSW\$MSW_EM	#003F	Status Mask (set all bits to 1)
FPSW\$INVALID	#0001	An invalid result occurred
FPSW\$DENORMAL	#0002	A denormal operand occurred
FPSW\$ZERODIVIDE	#0004	A divide by zero occurred
FPSW\$OVERFLOW	#0008	An overflow occurred
FPSW\$UNDERFLOW	#0010	>An underflow occurred
FPSW\$INEXACT	#0020	Inexact precision occurred

You can find out which exceptions have occurred by retrieving the status word and comparing it to the exception parameters. For example:

USE IFPORT

```
INTEGER(2) status
CALL GETSTATUSFPQQ(status)
IF (IAND (status, FPSW$INEXACT) > 0) THEN
    WRITE (*, *) "Inexact precision has occurred"
ELSE IF (IAND (status, FPSW$DENORMAL) > 0) THEN
    WRITE (*, *) "Denormal occurred"
END IF
```

To clear the status word flags, call the `CLEARSTATUSFPQQ` (IA-32 architecture only) routine.

Note

The `GETSTATUSFPQQ` and `CLEARSTATUSFPQQ` routines only affect the x87 status register. They do not affect the `MXCSR` register (the control and status register for the Intel(R) SSE and Intel(R) SSE2 instructions).

Floating-point Control Word Overview

On systems based on the IA-32 architecture, the FPU control word includes bits that control the FPU's precision, rounding mode, and whether exceptions

generate signals if they occur. You can read the control word value with `GETCONTROLFPQQ` (IA-32 architecture only) to find out the current control settings, and you can change the control word with `SETCONTROLFPQQ` (IA-32 architecture only).

 Note

The `GETCONTROLFPQQ` and `SETCONTROLFPQQ` routines only affect the x87 status register. They do not affect the `MXCSR` register (the control and status register for the SSE and SSE2 instructions).

Each bit in the floating-point control word corresponds to a mode of the floating-point math processor. The `IFORT.F90` module file in the `... \INCLUDE` folder contains the `INTEGER(2)` parameters defined for the control word, as shown in the following table:

Parameter Name	Value in Hex	Description
<code>FPCW\$MCW_PC</code>	<code>#0300</code>	Precision control mask
<code>FPCW\$64</code>	<code>#0300</code>	64-bit precision
<code>FPCW\$53</code>	<code>#0200</code>	53-bit precision
<code>FPCW\$24</code>	<code>#0000</code>	24-bit precision
<code>FPCW\$MCW_RC</code>	<code>#0C00</code>	Rounding control mask
<code>FPCW\$CHOP</code>	<code>#0C00</code>	Truncate
<code>FPCW\$UP</code>	<code>#0800</code>	Round up
<code>FPCW\$DOWN</code>	<code>#0400</code>	Round down
<code>FPCW\$NEAR</code>	<code>#0000</code>	Round to nearest
<code>FPCW\$MCW_EM</code>	<code>#003F</code>	Exception mask
<code>FPCW\$INVALID</code>	<code>#0001</code>	Allow invalid numbers
<code>>FPCW\$DENORMAL</code>	<code>#0002</code>	Allow denormals (very small numbers)
<code>FPCW\$ZERODIVIDE</code>	<code>#0004</code>	Allow divide by zero

Parameter Name	Value in Hex	Description
FPCW\$OVERFLOW	#0008	Allow overflow
FPCW\$UNDERFLOW	#0010	Allow underflow
FPCW\$INEXACT	#0020	Allow inexact precision

The control word defaults are:

- 53-bit precision
- Round to nearest (rounding mode)
- The denormal, underflow, overflow, divide-by-zero, invalid, and inexact precision exceptions are disabled (do not generate an exception). To change exception handling, you can use the `-fpe` (Linux* and Mac OS* X) or the `/fpe` (Windows*) compiler option or the `FOR_SET_FPE` routine.

Using Exception, Precision, and Rounding Parameters

This topic describes the exception, precision, and rounding parameters that you can use for the control word.

Exception Parameters

An exception is disabled if its bit is set to 1 and enabled if its bit is cleared to 0. If an exception is disabled (exceptions can be disabled by setting the flags to 1 with `SETCONTROLFPQQ` [IA-32 architecture only]), it will not generate an interrupt signal if it occurs. The floating-point process will return an appropriate special value (for example, NaN or signed infinity), but the program continues. You can find out which exceptions (if any) occurred by calling `GETSTATUSFPQQ` (IA-32 architecture only).

If errors on floating-point exceptions are enabled (by clearing the flags to 0 with `SETCONTROLFPQQ` [IA-32 architecture only]), the operating system generates an interrupt when the exception occurs. By default these interrupts cause run-time errors, but you can capture the interrupts with `SIGNALQQ` and branch to your own error-handling routines.

You should remember not to clear all existing settings when changing one. The values you want to change should be combined with the existing control word in an inclusive-OR operation (`IOR`) if you do not want to reset all options. For example:

```
USE IFPORT
INTEGER(2) control, newcontrol
CALL GETCONTROLFPQQ(control)
newcontrol = IOR(control, FPCW$INVALID)
! Invalid exception set (disabled).
CALL SETCONTROLFPQQ(newcontrol)
```

Note

The `GETCONTROLFPQQ`, `SETCONTROLFPQQ`, and `GETSTATUSFPQQ` routines only affect the x87 status register. They do not affect the `MXCSR` register (the control and status register for the Intel(R) SSE and Intel(R) SSE2 instructions).

Precision Parameters

On systems based on the IA-32 architecture, the precision bits control the precision to which the FPU rounds floating-point numbers. For example:

```
USE IFPORT
INTEGER(2) control, holdcontrol, newcontrol
CALL GETCONTROLFPQQ(control)
! Clear any existing precision flags.
holdcontrol = IAND(control, NOT(FPCW$MCW_PC))
newcontrol = IOR(holdcontrol, FPCW$64)
! Set precision to 64 bits.
CALL SETCONTROLFPQQ(newcontrol)
```

The precision options are mutually exclusive. If you set more than one, you may get an invalid mode or a mode other than the one you want. Therefore, you should clear the precision bits before setting a new precision mode.

Note

The `GETCONTROLFPQQ` and `SETCONTROLFPQQ` routines only affect the x87 status register. They do not affect the `MXCSR` register (the control and status register for the Intel(R) SSE and Intel(R) SSE2 instructions).

Rounding Parameters

On systems based on the IA-32 architecture, the rounding flags control the method of rounding that the FPU uses. For example:

```
USE IFPORT
INTEGER(2) status, control, controlo, mask_all_traps
CALL GETCONTROLFPQQ(control)
WRITE (*, 9000) 'Control word: ', control
>
! Save old control word
controlo = control
! Clear the rounding control flags
control = IAND(control,NOT(FPCW$MCW_RC))
! Set new control to round up
control = IOR(control,FPCW$SUP)
CALL SETCONTROLFPQQ(control)
CALL GETCONTROLFPQQ(control)
WRITE (*, 9000) 'Control word: ', control
```

The rounding options are mutually exclusive. If you set more than one, you may get an invalid mode or a mode other than the one you want. Therefore, you should clear the rounding bits before setting a new rounding mode.

Note

The `GETCONTROLFPQQ` and `SETCONTROLFPQQ` routines only affect the x87 status register. They do not affect the `MXCSR` register (the control and status register for the Intel(R) SSE and Intel(R) SSE2 instructions).

Handling Floating-point Exceptions with the `-fpe` or `/fpe` Compiler Option

Using the `-fpe` or `/fpe` Compiler Option

The `-fpen` (Linux* and Mac OS* X) or `/fpe:n` (Windows*) option allows some control over the results of floating-point exceptions.

`-fpe0` or `/fpe:0` restricts floating-point exceptions by enabling the overflow, the divide-by-zero, and the invalid floating-point exceptions. The program will print an error message and abort if any of these exceptions occurs. If a floating underflow occurs, the result is set to zero and execution continues. This is called flush-to-zero. This option sets `-IPF_fp_speculationstrict` (Linux and Mac OS X) or `/QIPF_fp_speculationstrict` (Windows) if no specific `-IPF_fp_speculation` or `/QIPF_fp_speculation` option is specified. The `-fpe0` or `/fpe:0` option sets `-ftz` (Linux and Mac OS X) `/Qftz` (Windows). To

get more detailed location information about where the exception occurred, use `-traceback` (Linux and Mac OS X) or `/traceback` (Windows).

 Note

On systems based on the IA-32 and Intel® 64 architectures , explicitly setting `-fpe0` or `/fpe:0` can degrade performance since the generated code stream must be synchronized after each floating-point instruction to allow for abrupt underflow fix-up.

`-fpe1` or `/fpe:1` restricts only floating-point underflow. Floating-point overflow, floating-point divide-by-zero, and floating-point invalid produce exceptional values (NaN and signed Infinities) and execution continues. If a floating-point underflow occurs, the result is set to zero and execution continues. The `/fpe:1` option sets `-ftz` or `/Qftz`.

 Note

On systems based on the IA-32 and Intel® 64 architectures , explicitly setting `-fpe1` or `/fpe:1` can degrade performance since the generated code stream must be synchronized after each floating-point instruction to allow for abrupt underflow fix-up.

`-fpe3` or `/fpe:3` is the default on all processors, which allows full floating-point exception behavior. Floating-point overflow, floating-point divide-by-zero, and floating-point invalid produce exceptional values (NaN and signed Infinities) and execution continues. Floating underflow is gradual: denormalized values are produced until the result becomes 0.

The `-fpe` or `/fpe` option enables exceptions in the Fortran main program only. The floating-point exception behavior set by the Fortran main program remains in effect throughout the execution of the entire program unless changed by the programmer. If the main program is not Fortran, the user can use the Fortran intrinsic `FOR_SET_FPE` to set the floating-point exception behavior.

When compiling different routines in a program separately, you should use the same value of *n* in `-fpen` or `/fpe:n`.

An example follows:

```
IMPLICIT NONE
```

```

real*4 res_uflow, res_oflow
real*4 res_dbyz, res_inv
real*4 small, big, zero, scale
small = 1.0e-30
big = 1.0e30
zero = 0.0
scale = 1.0e-10
! IEEE underflow condition (Underflow Raised)
res_uflow = small * scale
write(6,100)"Underflow: ",small, " *", scale, " = ", res_uflow
! IEEE overflow condition (Overflow Raised)
res_oflow = big * big
write(6,100)"Overflow:", big, " *", big, " = ", res_oflow
! IEEE divide-by-zero condition (Divide by Zero Raised)
res_dbyz = -big / zero
write(6,100)"Div-by-zero:", -big, " /", zero, " = ", res_dbyz
! IEEE invalid condition (Invalid Raised)
res_inv = zero / zero
write(6,100)"Invalid:", zero, " /", zero, " = ", res_inv
100 format(A14,E8.1,A2,E8.1,A2,E10.1)
end

```

Consider the following command line:

```
ifort fpe.f90 -fpe0 -fp-model strict -g (Linux and Mac OS X)
```

```
ifort fpe.f90 /fpe:0 /fp:strict /traceback (Windows)
```

Output similar to the following should result:

Windows:

```

Underflow: 0.1E-29 * 0.1E-09 = 0.0E+00
forrtl: error (72): floating overflow
Image      PC          Routine Line      Source
fpe.exe     0040115B Unknown Unknown Unknown
fpe.exe     0044DFC0 Unknown Unknown Unknown
fpe.exe     00433277 Unknown Unknown Unknown
kernel32.dll 7C816D4F Unknown Unknown Unknown

```

Linux and Mac OS X:

```

./a.out
Underflow: 0.1E-29* 0.1E-09 = 0.0E+00
forrtl: error (72): floating overflow
Image      PC          Routine Line      Source
a.out      0804A063 Unknown Unknown Unknown
a.out      08049E78 Unknown Unknown Unknown
Unknown    B746B748 Unknown Unknown Unknown
a.out      08049D31 Unknown Unknown Unknown
Aborted

```

The following command line uses /fpe1:

```
ifort fpe.f90 -fpe1 -g (Linux and Mac OS X)
```

```
ifort fpe.f90 /fpe:1 /traceback (Windows)
```

The following output is produced:

```
Underflow: 0.1E-29 * 0.1E-09 = 0.0E+00
```

```
Overflow: 0.1E+31 * 0.1E+31 = Infinity  
Div-by-zero: -0.1E+31 / 0.0E+00 = -Infinity  
Invalid: 0.0E+00 / 0.0E+00 = NaN
```

The following command line uses `/fpe3`:

```
ifort fpe.f90 -fpe3 -g (Linux and Mac OS X)
```

```
ifort fpe.f90 /fpe:3 /traceback (Windows)
```

The following output is produced:

```
Underflow: 0.1E-29 * 0.1E-09 = 0.1E-39  
Overflow: 0.1E+31 * 0.1E+31 = Infinity  
Div-by-zero: -0.1E+31 / 0.0E+00 = -Infinity  
Invalid: 0.0E+00 / 0.0E+00 = NaN
```

Understanding the Impact of Application Types

The full consequences of the `/fpe` option depend on the type of application you are building. You only get the full support for the chosen option setting in a Fortran Console or QuickWin/Standard Graphics application, assuming you do not override the default run-time exception handler. The work to achieve the full behavior is done partly by each of the default run-time handler, the Fortran compiler, the math library, the underlying hardware, and the operating system.

Floating-point Exceptions in Fortran Console, Fortran QuickWin, and Fortran Standard Graphics Applications

When you build a console application, the compiler generates a few calls at the beginning of your Fortran main program to Fortran run-time routines that initialize the environment, either with default options or in accordance with your selected compile-time options.

For floating-point exception handling, `/fpe:3` is the default. The run-time system initializes the hardware to mask all exceptions. With `/fpe:3`:

- The Intel Fortran run-time system does this initialization automatically (no call from the compiled code).
- The IA-32 architecture hardware automatically generates the default IEEE result for exceptional conditions.
- Because traps are masked with `/fpe:3`, there are no traps and you may see exceptional values like Nan's, Infinities, and denormalized numbers in your computations.

Users can poll the [floating-point status word](#) with `GETSTATUSFPQQ` to see if an exception has occurred and can clear the status register with `CLEARSTATUSFPQQ`.

If you specify `/fpe:0`, the compiler generates a call to an Intel Fortran run-time routine `FOR_SET_FPE`, with an argument that unmask all floating-point traps in the [floating-point control word](#). In this case, the hardware does not supply the default IEEE result. It traps to the operating system, which then looks for a condition handler.

In a Fortran console or Fortran QuickWin application, the Intel Fortran run-time system provides a default condition handler unless you establish your own. For all exceptions except underflow, the run-time system just prints out an error message and aborts the application. For underflow, the run-time system replaces the result with zero. This treatment of underflow with `/fpe:0` is called *abrupt underflow to 0 (zero)*, as opposed to *gradual underflow to 0* provided with `/fpe:3`.

Fixing up underflow results to zero can significantly degrade the performance of your application based on IA-32 architecture. If you are experiencing a large number of underflows, consider changing your code to avoid underflows or consider masking underflow traps and allowing the hardware to operate on denormalized numbers. The IA-32 architecture based hardware is designed to operate correctly in the denormalized range and doing so is much faster than trapping to fix up a result to zero.

Another important point to understand about selecting the `/fpe` option is that the generated code must support the trapping mode. When an IA-32 architecture based floating-point instruction generates an exceptional result, you do not necessarily get a trap. The instruction must be followed by an `wait` instruction or another floating-point operate instruction to cause the trap to occur.

The Intel Fortran compiler generates machine code to support these requirements in accordance with your selected `/fpe` option setting. In other words, the generated code must support the trapping mode. You can see this by compiling a simple test program and look at the machine code listing with

`/fpe:0` first, then `/fpe:3`. There are no `fwait` instructions in the `/fpe:3` code. Even if you replace the default run-time exception handler with your own handler, you may still want to compile with `/fpe:0` to generate code that supports trapping.

Floating-Point Exceptions in Fortran DLL Applications

In a DLL, there is no main Fortran program (unless you have written your main program in Fortran), so there is no automatic calling of run-time routines to initialize the environment. Even if you select `/fpe:0`, there is nothing that causes the run-time system to unmask traps in the hardware so you won't see traps. You will continue to see the hardware generated default IEEE results (Nan's, Infinities, and denormalized numbers in your computations). The generated code will still do its part by supplying the `fwait` instructions, and so on, but unless the traps are unmasked somehow, no traps will occur. You can use `SETCONTROLFPQQ` or `FOR_SET_FPE` to unmask traps in the [floating-point control word](#).

There is also no default exception handling in a DLL. The main application that calls the DLL must provide this, or the code in the DLL must provide something when it is called. Since underflow processing (fixup to 0, and so on) is done by the default Fortran run-time system handler, a DLL won't have that feature automatically.

A typical strategy is to compile with `/fpe:0`, but only unmask floating divide-by-zero, floating overflow, and floating invalid traps in the [floating-point control word](#). By leaving floating-point underflow traps masked, the hardware will continue to provide gradual underflow, but other floating-point exceptions will generate traps, which the user then handles as desired.

Floating-Point Exceptions in Fortran Windows Applications

In a Fortran Windows application, the situation is similar to a Fortran DLL application. You define a `WinMain` routine in your Fortran code as the main entry point for your application. The Fortran run-time system does not define the main

routine as it does in a Fortran console or Fortran QuickWin (or Standard Graphics) application. Your code is not protected by the default handler and the default run-time initialization routines are not called.

Understanding IEEE Floating-point Standard

Overview: Understanding IEEE Floating-point Standard

This version of Intel® Compiler uses a close approximation to the IEEE floating-point standard (ANSI/IEEE Std 754-1985, *IEEE Standard for Binary Floating-Point Arithmetic*, 1985). This standard is common to many microcomputer-based systems due to the availability of fast processors that implement the required characteristics.

This section outlines the characteristics of the standard and its implementation for Intel Compiler. Except as noted, the description includes both the IEEE standard and the Intel Compiler implementation.

Floating-point Formats

The IEEE Standard 754 specifies values and requirements for floating-point representation (such as base 2). The standard outlines requirements for two formats: basic and extended, and for two word-lengths within each format: single and double.

Intel Fortran supports single-precision format (REAL(4)) and double-precision format (REAL(8)) floating-point numbers. At some levels of optimization, some floating-point numbers are stored in register file format (which equals 1 bit sign, 15 bit exponent, 64 bits of significand rounded to 53 bits), rather than being stored in IEEE single or double precision format. This can affect the values produced by some computations. The compiler option `-fp-model` (Linux* and Mac OS* X) or `/fp` (Windows*) can control how floating-point expressions are evaluated, thus leading to more predictable results.

Limitations of Numeric Conversion

The Intel® Fortran floating-point conversion solution is not expected to fulfill all floating-point conversion needs.

For instance, data fields in record structure variables (specified in a `STRUCTURE` statement) and data components of derived types (`TYPE` statement) are not converted. When they are later examined as separate fields by the program, they will remain in the binary format they were stored in on disk, unless the program is modified. With `EQUIVALENCE` statements, the data type of the variable named in the I/O statement is used.

If a program reads an I/O record containing multiple format floating-point fields into a single variable (such as an array) instead of their respective variables, the fields will not be converted. When they are later examined as separate fields by the program, they will remain in the binary format they were stored in on disk, unless the program is modified.

Conversions of the following file structure types are not supported:

- Binary data (`FORM= ' BINARY '`)
- Formatted data (`FORM= ' FORMATTED '`)
- Unformatted data (`FORM= ' UNFORMATTED '`) written by Microsoft* Fortran PowerStation or by Intel Fortran with the `/fpscomp:ioformat` compiler option in effect.

Special Values

The following lists special values that the Intel® Compiler supports and provides a brief description.

Signed zero

Intel Fortran treats zero as signed by default. The sign of zero is the same as the sign of a nonzero number. If you use the intrinsic function `SIGN` with zero as the second argument, the sign of the zero will be transferred. Comparisons, however, consider `+0` to be equal to `-0`. A signed zero is useful in certain numerical analysis algorithms, but in most applications the sign of zero is invisible.

Denormalized numbers

Denormalized numbers (denormals) fill the gap between the smallest positive normalized number and the smallest negative number. Otherwise only (+/-) 0 occurs in that interval. Denormalized numbers extend the range of computable results by allowing for gradual underflow. Systems based on the IA-32 architecture support a Denormal Operand status flag, which, when set, means at least one of the input operands to a floating-point operation is a denormal. The Underflow status flag is set when a number loses precision and becomes a denormal.

Signed infinity

Infinities are the result of arithmetic in the limiting case of operands with arbitrarily large magnitude. They provide a way to continue when an overflow occurs. The sign of an infinity is simply the sign you obtain for a finite number in the same operation as the finite number approaches an infinite value. By retrieving the status flags, you can differentiate between an infinity that results from an overflow and one that results from division by zero. Intel® Compiler treats infinity as signed by default. The output value of infinity is Infinity or -Infinity.

Not a Number

Not a Number (NaN) results from an invalid operation. For instance $0/0$ and $\text{SQRT}(-1)$ result in NaN. In general, an operation involving a NaN produces another NaN. Because the fraction of a NaN is unspecified, there are many possible NaNs. Intel® Compiler treats all NaNs identically, but provides two different types:

- Signaling NaN, which has an initial fraction bit of 0 (zero). This usually raises an invalid exception when used in an operation.
- Quiet NaN, which has an initial fraction bit of 1.

The floating-point hardware changes a signaling NaN into a quiet NaN during many arithmetic operations, including the assignment operation. An invalid exception may be raised but the resulting floating-point value will be a quiet NaN.

Fortran binary and unformatted input and output do not change the internal representations of the values as they are handled. Therefore, signaling and quiet NaNs may be read into real data and output to files in binary form.

The output value of NaN is NaN.

Representing Floating-point Numbers

Floating-point Representation

The Fortran numeric environment is flexible, which helps make Fortran a strong language for intensive numerical calculations. The Fortran standard purposely leaves the precision of numeric quantities and the method of rounding numeric results unspecified. This allows Fortran to operate efficiently for diverse applications on diverse systems.

Computations on real numbers may not yield what you expect. This happens because the hardware must represent numbers in a finite number of bits.

There are several effects of using finite floating-point numbers. The hardware is not able to represent every real number exactly, but must approximate exact representations by rounding or truncating to finite length. In addition, some numbers lie outside the range of representation of the maximum and minimum exponents and can result in calculations that underflow and overflow. As an example of one consequence, finite precision produces many numbers that, although non-zero, behave in addition as zero.

You can minimize the effects of finite representation with programming techniques; for example, by not using floating-point numbers in LOGICAL comparisons or by giving them a tolerance (for example, `IF (ABS(x-10.0) <= 0.001)`), and by not attempting to combine or compare numbers that differ by more than the number of significant bits.

Floating-point numbers approximate real numbers with a finite number of bits. The bits are calculated as shown in the following formula. The representation is binary, so the base is 2. The bits b_n represent binary digits (0 or 1). The precision P is the number of bits in the nonexponential part of the number (the significand),

and E is the exponent. With these parameters, binary floating-point numbers approximate real numbers with the values:

$$(-1)^s b_0. b_1 b_2 \dots b_{P-1} \times 2^E$$

where s is 0 or 1 (+ or -), and $E_{\min} \leq E \leq E_{\max}$

The following table gives the standard values for these parameters for single, double, and quad (extended precision) formats and the resulting bit widths for the sign, the exponent, and the full number.

Parameters for IEEE* Floating-Point Formats

Parameter	Single	Double	Quad or Extended Precision (IEEE_X)*
Sign width in bits	1	1	1
P	24	53	113
E_{\max}	+127	+1023	+16383
E_{\min}	-126	-1022	-16382
Exponent <i>bias</i>	+127	+1023	+16383
Exponent width in bits	8	11	15
Format width in bits	32	64	128

* This type is emulated in software.

The actual number of bits needed to represent the precisions 24, 53, and 113 is therefore 23, 52, and 112, respectively, because b_0 is chosen to be 1 implicitly.

A *bias* is added to all exponents so that only positive integer exponents occur.

This expedites comparisons of exponent values. The stored exponent is actually:

$$e = E + \textit{bias}$$

See Also

[Native IEEE Floating-Point Representations](#)
[Rounding Errors](#)

Retrieving Parameters of Numeric Representations

Intel Fortran includes several intrinsic functions that return details about the numeric representation. These are listed in the following table and described fully in the *Language Reference*.

Functions that Return Numeric Parameters

Name	Description	Argument/Function Type
DIGITS	DIGITS(<i>x</i>). Returns number of significant digits for data of the same type as <i>x</i> .	<i>x</i> : Integer or Real result: INTEGER(4)
EPSILON	EPSILON(<i>x</i>). Returns the smallest positive number that when added to one produces a number greater than one for data of the same type as <i>x</i> .	<i>x</i> : Real result: same type as <i>x</i>
EXPONENT	EXPONENT(<i>x</i>). Returns the exponent part of the representation of <i>x</i> .	<i>x</i> : Real result: INTEGER(4)
FRACTION	FRACTION(<i>x</i>). Returns the fractional part (significand) of the representation of <i>x</i> .	<i>x</i> : Real result: same type as <i>x</i>
HUGE	HUGE(<i>x</i>). Returns largest number that can be represented by data of type <i>x</i> .	<i>x</i> : Integer or Real result: same type as <i>x</i> .
MAXEXPONENT	MAXEXPONENT(<i>x</i>). Returns the largest positive decimal exponent for data of the same type as <i>x</i> .	<i>x</i> : Real result: INTEGER(4)
MINEXPONENT	MINEXPONENT(<i>x</i>). Returns the largest negative decimal exponent for data of the same type as <i>x</i> .	<i>x</i> : Real result: INTEGER(4)

Name	Description	Argument/Function Type
NEAREST	NEAREST(x , s). Returns the nearest different machine representable number to x in the direction of the sign of s .	x : Real s : Real and not zero result: same type as x .
PRECISION	PRECISION(x). Returns the number of significant digits for data of the same type as x .	x : Real or Complex result: INTEGER(4)
RADIX	RADIX(x). Returns the base for data of the same type as x .	x : Integer or Real result: INTEGER(4)
RANGE	RANGE(x). Returns the decimal exponent range for data of the same type as x .	x : Integer, Real or Complex result: INTEGER(4)
RRSPACING	RRSPACING(x). Returns the reciprocal of the relative spacing of numbers near x .	x : Real result: same type as x
SCALE	SCALE(x , i). Multiplies x by 2 raised to the power of i .	x : Real i : Integer result: same type as x
SET_EXPONENT	SET_EXPONENT(x , i). Returns a number whose fractional part is x and whose exponential part is i .	x : Real i : Integer result: same type as x
SPACING	SPACING(x). Returns the absolute spacing of numbers near x .	x : Real result: same type as x
TINY	TINY(x). Returns smallest positive number that can be represented by data of type x .	x : Real result: same type as x

Native IEEE Floating-point Representation

Native IEEE* Floating-point Representations Overview

The REAL(4) (IEEE* S_floating), REAL(8) (IEEE T_floating), and REAL(16) (IEEE-style X_floating) formats are stored in standard little endian IEEE binary floating-point notation. (See IEEE Standard 754 for additional information about IEEE binary floating point notation.) COMPLEX() formats use a pair of REAL values to denote the real and imaginary parts of the data.

All floating-point formats represent fractions in sign-magnitude notation, with the binary radix point to the right of the most-significant bit. Fractions are assumed to be normalized, and therefore the most-significant bit is not stored (this is called "hidden bit normalization"). This bit is assumed to be 1 unless the exponent is 0. If the exponent equals 0, then the value represented is denormalized (subnormal) or plus or minus zero.

Intrinsic REAL kinds are 4 (single precision), 8 (double precision), and 16 (extended precision), such as REAL(KIND=4) for single-precision floating-point data. Intrinsic COMPLEX kinds are also 4 (single precision), 8 (double precision), and 16 (extended precision).

To obtain the kind of a variable, use the KIND intrinsic function. You can also use a size specifier, such as REAL*4, but be aware this is an extension to the Fortran 95 standard.

If you omit certain compiler options, the default sizes for REAL and COMPLEX data declarations are as follows:

- For REAL data declarations without a kind parameter (or size specifier), the default size is REAL (KIND=4) (same as REAL*4).
- For COMPLEX data declarations without a kind parameter (or size specifier), the default data size is COMPLEX (KIND=4) (same as COMPLEX*8).

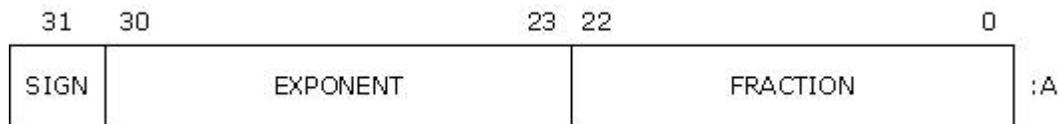
To control the size of all REAL or COMPLEX declarations without a kind parameter, use the `/real_size:64` (or `/4R8`) or `/real_size:128` (or `/4R16`) options; the default is `/real_size:32` (or `/4R4`).

You can explicitly declare the length of a REAL or a COMPLEX declaration using a kind parameter, or specify DOUBLE PRECISION or DOUBLE COMPLEX. To control the size of all DOUBLE PRECISION and DOUBLE COMPLEX declarations, use the `/double_size:128` (or `/Qautodouble`) option; the default is `/double_size:64`.

REAL(KIND=4) (REAL) Representation

REAL(4) (same as REAL(KIND=4)) data occupies 4 contiguous bytes stored in IEEE S_floating format. Bits are labeled from the right, 0 through 31, as shown below.

REAL(4) Floating-Point Data Representation



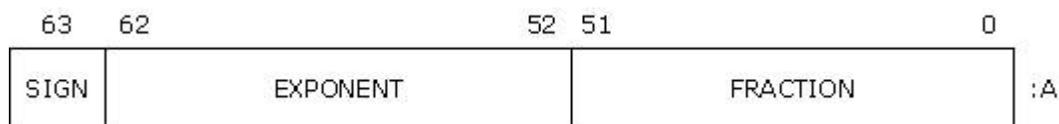
The form of REAL(4) data is sign magnitude, with bit 31 the sign bit (0 for positive numbers, 1 for negative numbers), bits 30:23 a binary exponent in excess 127 notation, and bits 22:0 a normalized 24-bit fraction including the redundant most-significant fraction bit not represented.

The value of data is in the approximate range: 1.17549435E-38 (normalized) to 3.40282347E38. The IEEE denormalized (subnormal) limit is 1.40129846E-45. The precision is approximately one part in 2^{23} ; typically 7 decimal digits.

REAL(KIND=8) (DOUBLE PRECISION) Representation

REAL(8) (same as REAL(KIND=8)) data occupies 8 contiguous bytes stored in IEEE T_floating format. Bits are labeled from the right, 0 through 63, as shown below.

REAL(8) Floating-Point Data Representation



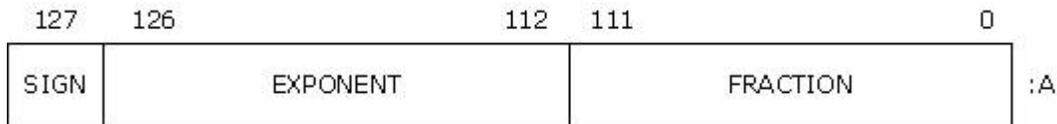
The form of REAL(8) data is sign magnitude, with bit 63 the sign bit (0 for positive numbers, 1 for negative numbers), bits 62:52 a binary exponent in excess 1023

notation, and bits 51:0 a normalized 53-bit fraction including the redundant most-significant fraction bit not represented.

The value of data is in the approximate range: 2.2250738585072013D-308 (normalized) to 1.7976931348623158D308. The IEEE denormalized (subnormal) limit is 4.94065645841246544D-324. The precision is approximately one part in 2^{52} ; typically 15 decimal digits.

REAL(KIND=16) Representation

REAL(16) (same as REAL(KIND=16)) data occupies 16 contiguous bytes stored in IEEE-style X_floating format. Bits are labeled from the right, 0 through 127, as shown below.



The form of REAL(16) data is sign magnitude, with bit 127 the sign bit (0 for positive numbers, 1 for negative numbers), bits 126:112 a binary exponent in excess 16383 notation, and bits 111:0 a normalized 113-bit fraction including the redundant most-significant fraction bit not represented.

The value of data is in the approximate range:

6.4751751194380251109244389582276465524996Q-4966 to

1.189731495357231765085759326628007016196477Q4932. Unlike other floating-point formats, there is little if any performance penalty from using denormalized extended-precision numbers. This is because accessing denormalized REAL (KIND=16) numbers does not result in an arithmetic trap (the extended-precision format is emulated in software). The smallest normalized number is 3.362103143112093506262677817321753Q-4932.

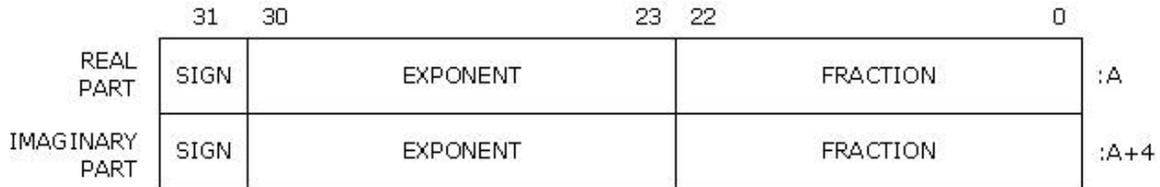
The precision is approximately one part in 2^{112} or typically 33 decimal digits.

COMPLEX(KIND=4) (COMPLEX) Representation

COMPLEX(4) (same as COMPLEX(KIND=4) and COMPLEX*8) data is 8 contiguous bytes containing a pair of REAL(4) values stored in IEEE S_floating format. The low-order 4 bytes contain REAL(4) data that represents the real part

of the complex number. The high-order 4 bytes contain REAL(4) data that represents the imaginary part of the complex number, as shown below.

COMPLEX(4) Floating-Point Data Representation

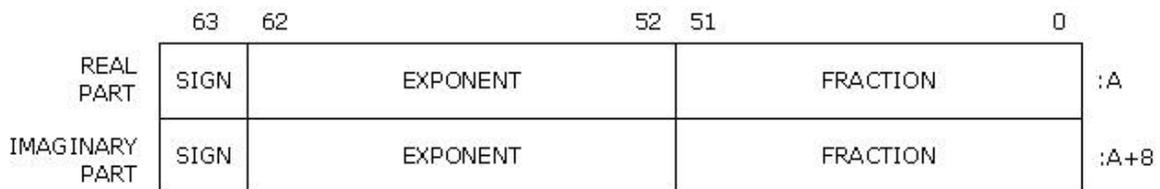


The limits and underflow characteristics for [REAL\(4\)](#) apply to the two separate real and imaginary parts of a COMPLEX(4) number. Like REAL(4) numbers, the sign bit representation is 0 (zero) for positive numbers and 1 for negative numbers.

COMPLEX(KIND=8) (DOUBLE COMPLEX) Representation

COMPLEX(8) (same as COMPLEX(KIND=8) and COMPLEX*16) data is 16 contiguous bytes containing a pair of REAL(8) values stored in IEEE T_floating format. The low-order 8 bytes contain REAL(8) data that represents the real part of the complex data. The high-order 8 bytes contain REAL(8) data that represents the imaginary part of the complex data, as shown below.

COMPLEX(8) Floating-Point Data Representation

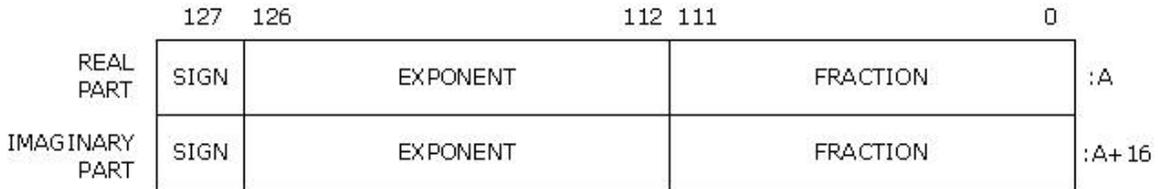


The limits and underflow characteristics for [REAL\(8\)](#) apply to the two separate real and imaginary parts of a COMPLEX(8) number. Like REAL(8) numbers, the sign bit representation is 0 (zero) for positive numbers and 1 for negative numbers.

COMPLEX(KIND=16) Representation

COMPLEX(16) (same as COMPLEX(KIND=16) or COMPLEX*32) data is 32 contiguous bytes containing a pair of REAL(16) values stored in IEEE-style

X_floating format. The low-order 16 bytes contain REAL(16) data that represents the real part of the complex data. The high-order 16 bytes contain REAL(8) data that represents the imaginary part of the complex data, as shown below.



The limits and underflow characteristics for [REAL\(16\)](#) apply to the two separate real and imaginary parts of a COMPLEX(16) number. Like REAL(16) numbers, the sign bit representation is 0 (zero) for positive numbers and 1 for negative numbers.

Handling Exceptions and Errors

Loss of Precision Errors

If a real number is not exactly one of the representable floating-point numbers, then the nearest floating-point number must represent it. The rounding error is the difference between the exact real number and its nearest floating-point representation. If the rounding error is non-zero, the rounded floating-point number is called *inexact*.

Normally, calculations proceed when an inexact value results. Almost any floating-point operation can produce an inexact result. The rounding mode (round up, round down, round nearest, truncate) is determined by the floating-point control word.

If an arithmetic operation results in a floating-point number that cannot be represented in a specific data type, the operation may produce a special value: signed zero, signed infinity, NaN, or a denormal. Numbers that have been rounded to an exactly representable floating-point number also result in a special value. Special-value results are a limiting case of the arithmetic operation involved. Special values can propagate through your arithmetic operations without causing your program to fail, and often provide usable results.

If an arithmetic operation results in an exception, the operation can cause an underflow or overflow:

- Underflow occurs when an arithmetic result is too small for the math processor to handle. Depending on the setting of the `/fpe` compiler option, underflows are set to zero (they are usually harmless) or they are left as is (denormalized).
- Overflow occurs when an arithmetic result is too large for the math processor to handle. Overflows are more serious than underflows, and may indicate an error in the formulation of a problem (for example, unintended exponentiation of a large number by a large number). Overflows generally produce an appropriately signed infinity value. (This depends on the rounding mode as per the IEEE standard.)

An arithmetic operation can also throw the following exceptions: divide-by-zero exception, an invalid exception, and an inexact exception.

You can select how exceptions are handled by setting the floating-point control word.

On Windows* systems, you can select how exceptions are handled by setting the floating-point control word.

See Also

[Special Values](#)

Rounding Errors

Although the rounding error for one real number might be acceptably small in your calculations, at least two problems can arise because of it. If you test for exact equality between what you consider to be two exact numbers, the rounding error of either or both floating-point representations of those numbers may prevent a successful comparison and produce spurious results. Also, when you calculate with floating-point numbers the rounding errors may accumulate to a meaningful loss of numerical significance.

Carefully consider the numerics of your solution to minimize rounding errors or their effects. You might benefit from using double-precision arithmetic or

restructuring your algorithm, or both. For instance, if your calculations involve arrays of linear data items, you might reduce the loss of numerical significance by subtracting the mean value of each array from each array element and by normalizing each element of such an array to the standard deviation of the array elements.

The following code segment can execute differently on various systems and produce varying results for `n`, `x`, and `s`. It also produces different results if you use the `-fp-model source` (Linux* and Mac OS* X) or `/fp:source` (Windows*; systems based on the IA-32 architecture only), or `-fp-model fast` (Linux and Mac OS X) or `/fp:fast` (Windows) compiler options. Rounding error accumulates in `x` because the floating-point representation of 0.2 is inexact, then accumulates in `s`, and affects the final value for `n`:

```

INTEGER n
REAL s, x
n = 0
s = 0.0
x = 0.0
1 n = n + 1
  x = x + 0.2
  s = s + x
IF ( x .LE. 10. ) GOTO 1 ! Will you get 51 cycles?
WRITE(*,*) 'n = ', n, '; x = ', x, '; s = ', s

```

This example illustrates a common coding problem: carrying a floating-point variable through many successive cycles and then using it to perform an `IF` test. This process is common in numerical integration. There are several remedies. You can compute `x` and `s` as multiples of an integer index, for example, replacing the statement that increments `x` with `x = n * 0.2` to avoid round-off accumulation. You might test for completion on the integer index, such as `IF (n <= 50) GOTO 1`, or use a `DO` loop, such as `DO n= 1,51`. If you must test on the real variable that is being cycled, use a realistic tolerance, such as `IF (x <= 10.001)`.

Floating-point arithmetic does not always obey the standard rules of algebra exactly. Addition is not precisely associative when round-off errors are considered. You can use parentheses to express the exact evaluation you require to compute a correct, accurate answer. This is recommended when you

specify optimization for your generated code, since associativity may otherwise be unpredictable.

The expressions $(x + y) + z$ and $x + (y + z)$ can give unequal results in some cases.

The compiler uses the default rounding mode (round-to-nearest) during compilation. The compiler performs more compile-time operations that eliminate runtime operations as the optimization level increases. If you set rounding mode to a different setting (other than round-to-nearest), that rounding mode is used only if that computation is performed at runtime. If you want to force computations to be performed at runtime, use the `-fp-model strict` (Linux or MacOS) or `/fp:strict` (Windows) option.

See Also

[OPT-NAME](#)

ULPs, Relative Error, and Machine Epsilon

ULP, Relative Error, and Machine Epsilon are terms that describe the magnitude of rounding error. A floating-point approximation to a real constant or to a computed result may err by as much as $1/2$ unit in the last place (the b_{p-1} bit). The abbreviation *ULP* represents the measure "unit in the last place." Another measure of the rounding error uses the relative error, which is the difference between the exact number and its approximation divided by the exact number. The relative error that corresponds to $1/2$ *ULP* is bounded by:

$$1/2 \cdot 2^{-P} \leq 1/2 \text{ ULP} \leq 2^{-P}$$

The upper bound $EPS = 2^{-P}$, the machine epsilon, is commonly used in discussions of rounding errors because it expresses the smallest floating-point number that you can add to 1.0 with a result that does not round to 1.0.