

# Quick Reference Guide to Optimization with Intel® C++ and Fortran Compilers v16

For IA-32 processors, Intel® 64 processors, compatible non-Intel processors, Intel® Xeon Phi™ coprocessors and Intel® Graphics Technology.

## Contents

Application Performance .....	2
General Optimization Options** .....	3
Parallel Performance** .....	4
Parallel Performance Using Intel® Cilk™ Plus.....	5
Recommended Processor-Specific Optimization Options** .....	6
Compiling for Offload.....	7
Compiling for Intel® Graphics Technology <sup>§§</sup> .....	7
Environment Variables for Intel® Graphics Technology <sup>§§</sup> .....	7
Optimizing for the Intel® Xeon Phi™ Coprocessor x100 product family <sup>§</sup> .....	8
Environment Variables for the Intel® Xeon Phi™ Coprocessor <sup>§</sup> .....	8
Interprocedural Optimization (IPO) and Profile-Guided Optimization (PGO) Options.....	9
Floating-Point Arithmetic Options .....	10
Fine-Tuning (All Processors).....	11
Debug Options.....	12
‡ Optimization Notice .....	12

**For product and purchase information, visit the Intel® Software Development Tools site at: <http://intel.ly/sw-dev-tools>.**

<sup>§§</sup> Compiler support for Intel® Graphics Technology depends on operating system support.

<sup>§</sup> Intel® MIC architecture and Intel® Xeon Phi™ coprocessors are supported by compilers within Intel® Parallel Studio XE, but not within Intel® System Studio.

\*\* Several of these options are available for both Intel® and non-Intel microprocessors but they may perform more optimizations for Intel microprocessors than they perform for non-Intel microprocessors.‡

# Application Performance

## A Step-by-Step Approach to Application Tuning with Intel Compilers

Before you begin performance tuning, check correctness of your application by building it without optimization using **/Od (-O0)**.

1. Measure performance using the general optimization options (Windows\* **/O1, /O2 or /O3**; Linux\* and OS X\* **-O1, -O2, or -O3**) to see which one works best for your application. Most users should start at **/O2 (-O2)**, the default, before trying more advanced optimizations. Next, try **/O3 (-O3)** for loop-intensive applications.\*\*
2. Fine-tune performance using processor-specific options such as **/Qx (-x)** or **/arch (-m)**. Examples are **/QxCORE-AVX2 (-xcore-avx2)** for the 4<sup>th</sup> Generation Intel® Core™ processor family and **/arch:SSE3 (-msse3)** for compatible, non-Intel processors that support at least the Intel® SSE3 instruction set. Or, **/QxHOST (-xhost)** will use the most advanced instruction set for the processor on which you compiled.\*\*
3. Add interprocedural optimization (IPO), **/Qipo (-ipo)** and/or profile-guided optimization (PGO), **/Qprof-gen** and **/Qprof-use (-prof-gen and -prof-use)**, then measure performance again to determine whether your application benefits from one or both of them.
4. Use Intel® Advisor and Intel® VTune™ Amplifier<sup>††</sup> to help you identify serial and parallel performance “hotspots” within your application that could benefit from further performance tuning. Use the compiler optimization report **/Qopt-report (-qopt-report)** to help identify individual optimization opportunities.
5. Further optimize your application for SIMD through explicit vector programming using the Intel® Cilk™ Plus language extensions for C/C++ or the SIMD features of OpenMP\* 4.0 with **/Qopenmp-simd (-qopenmp-simd)**.<sup>†</sup>
6. Optimize for parallel execution on multi-threaded, multi-core and multi-processor systems using: the auto-parallelization option **/Qparallel (-parallel)**; the Intel Cilk Plus language extensions for C/C++; OpenMP pragmas or directives along with the option **/Qopenmp (-qopenmp)**<sup>†</sup>; or by using the Intel® Performance Libraries included with the product.\*\* Use Intel® Inspector to reduce the time to market for threaded applications by diagnosing memory and threading errors and speeding up the development process.

For more details, please consult the main product documentation at <https://software.intel.com/intel-software-technical-documentation>. The Intel® Compiler User and Reference Guide includes dedicated sections on compiling applications for Intel® MIC Architecture and for Intel® Graphics Technology.

\*\* Several of these options are available for both Intel® and non-Intel microprocessors but they may perform more optimizations for Intel microprocessors than they perform for non-Intel microprocessors.‡

†OpenMP is supported by compilers within Intel® Parallel Studio XE, but not within Intel® System Studio.

††Some features of this product cannot be used on non-Intel microprocessors.

## General Optimization Options\*\*

Windows*	Linux* & OS X*	Comment
<code>/Od</code>	<code>-O0</code>	<b>No optimization.</b> Used during the early stages of application development and debugging.
<code>/Os</code> <code>/O1</code>	<code>-Os</code> <code>-O1</code>	<b>Optimize for size.</b> Omits optimizations that tend to increase object size. Creates the smallest optimized code in most cases. May be useful in large server/database applications where memory paging due to larger code size is an issue.
<code>/O2</code>	<code>-O2</code>	<b>Maximize speed.</b> Default setting. Enables many optimizations, including vectorization. Creates faster code than <code>/O1 (-O1)</code> in most cases.
<code>/O3</code>	<code>-O3</code>	Enables <code>/O2 (-O2)</code> optimizations plus more aggressive loop and memory-access optimizations, such as scalar replacement, loop unrolling, loop blocking to allow more efficient use of cache and additional data prefetching.  The <code>/O3 (-O3)</code> option is particularly recommended for applications that have loops that do many floating-point calculations or process large data sets. These aggressive optimizations may occasionally slow down other types of applications compared to <code>/O2 (-O2)</code> .
<code>/Qopt-report [n]</code>	<code>-qopt-report [n]</code>	Generates an optimization report, by default written to a file with extension <code>.optprt</code> . <code>n</code> specifies the level of detail, from <code>0</code> (no report) to <code>5</code> (maximum detail). Default is <code>2</code> .
<code>/Qopt-report-file:name</code>	<code>-qopt-report-file=name</code>	Writes an optimization report to <code>stderr</code> , <code>stdout</code> or to the file <code>name</code> .
<code>/Qopt-report-phase:name1, name2, ...</code>	<code>-qopt-report-phase=name1, name2,...</code>	Optimization reports are generated for optimization phases <code>name1</code> , <code>name2</code> , etc. Possible phases include: <b>all</b> – Optimization reports for all phases (default) <b>loop</b> – Loop nest and memory optimizations <b>vec</b> – auto-vectorization and explicit vector programming <b>par</b> – auto-parallelization <b>openmp</b> – threading using OpenMP <b>cg</b> – code generation <b>ipo</b> – Interprocedural Optimization, including inlining <b>pgo</b> – Profile Guided Optimization <b>offload</b> – offload of data to and/or execution on Intel® MIC Architecture or Intel® Graphics Technology.
<code>/Qopt-report-help</code>	<code>-qopt-report-help</code>	Displays all possible values of <code>name</code> for <code>/Qopt-report-phase (-qopt-report-phase)</code> above. No compilation is performed.
<code>/Qopt-report-routine: substring</code>	<code>-qopt-report-routine= substring</code>	Generates reports only for functions or subroutines whose names contain <b>substring</b> . By default, reports are generated for all functions and subroutines.
<code>/Qopt-report-filter:"string"</code>	<code>-qopt-report-filter="string"</code>	Restricts reports to the file, function or subroutine and/or ranges of line numbers specified by <b>"string"</b> , e.g. "myfile,myfun,line1-line2".

\*\* Several of these options are available for both Intel® and non-Intel microprocessors but they may perform more optimizations for Intel microprocessors than they perform for non-Intel microprocessors.‡

## Parallel Performance\*\*

Windows*	Linux* & OS X*	Comment
/Qopenmp†	-qopenmp†	Multi-threaded code and/or SIMD code is generated when OpenMP* directives are present. For Fortran only, makes local arrays automatic and may require an increased stack size.
/Qopenmp-simd†	-qopenmp-simd†	SIMD code is generated when OpenMP SIMD directives are present.
/Qopenmp-stubs	-qopenmp-stubs	Ignores OpenMP directives and links references to OpenMP run-time library functions to stub (dummy) functions assuming single-threaded operation.
/Qparallel	-parallel	The auto-parallelizer detects simply structured loops that may be safely executed in parallel, including the DO CONCURRENT construct and loops implied by Intel® Cilk™ Plus array notation, and automatically generates multi-threaded code for these loops.
/Qpar-threshold[:n]	-par-threshold[n]	Sets a threshold for the auto-parallelization of loops based on the likelihood of a performance benefit. <b>n=0</b> (parallelize loops regardless of computational work volume ) to <b>n=100</b> (default; parallelize loops only if a performance benefit is highly likely). Must be used in conjunction with <b>/Qparallel (-parallel)</b> .
/Qpar-affinity: <i>name</i>	-par-affinity= <i>name</i>	Specifies thread-processor affinity for OpenMP or auto-parallelized applications. Typical values of <b>name</b> are <b>none</b> (default), <b>scatter</b> and <b>compact</b> . Has effect only when compiling the main program. See the compiler user and reference guide for more settings and details.
/Qopt-matmul[-]	-q[no-]opt-matmul	This option enables [disables] a compiler-generated Matrix Multiply (matmul) library call by identifying matrix multiplication loop nests, if any, and replacing them with a matmul library call for improved performance. This option is enabled by default if options <b>/O3 (-O3)</b> and <b>/Qparallel (-parallel)</b> are specified. This option has no effect unless option <b>/O2 (-O2)</b> or higher is set.
/Qcilk-serialize	-cilk-serialize	This option includes a header file, cilk_stubs.h, that causes the compiler to ignore all Intel® Cilk™ Plus threading keywords, resulting in a serial executable. (C/C++ only). See the "Using Intel Cilk Plus" section of the user and reference guide for more detail.
/Qcoarray	-coarray	Enables the coarray feature of Fortran 2008 (Fortran only). Options are <b>shared</b> , <b>distributed</b> , <b>coprocessor</b> and <b>single</b> . See the compiler reference guide for more detail.
/Qmkl: <i>name</i>	-mkl= <i>name</i>	Requests that the Intel® Math Kernel Library (Intel® MKL) be linked. Off by default. Possible values of <b>name</b> are: <b>parallel</b> Links the threaded part of Intel MKL (default) <b>sequential</b> Links the non-threaded part of Intel MKL <b>cluster</b> Links cluster and sequential parts of Intel MKL

\*\* Several of these options are available for both Intel® and non-Intel microprocessors but they may perform more optimizations for Intel microprocessors than they perform for non-Intel microprocessors.†

## Parallel Performance Using Intel® Cilk™ Plus

Threading Keywords	Description (C/C++ only)
<code>cilk_spawn</code>	Allows (but does not require) a spawned function to be run in parallel with the caller, subject to dynamic scheduling by the Intel® Cilk™ Plus runtime.
<code>cilk_sync</code>	Introduces a barrier: function cannot continue until all spawned children are complete.
<code>cilk_for</code>	Introduces a for loop whose iterations are allowed (but not required) to run in parallel.

**Reducers** allow reduction operations, such as accumulation of a sum, to be executed safely in parallel.

E.g. `cilk::reducer< cilk::op_add<unsigned int> >` declares a reducer to sum unsigned ints.

**Holders:** The `cilk::holder` template class provides a convenient form of task local storage that is thread safe.

**Array Notation:** a readable, explicitly data-parallel C/C++ language extension that facilitates generation of SIMD parallel code by the vectorizer at optimization level `-O2` or higher and asserts absence of dependencies.

Syntax: `array[<lower bound>:<length>:<stride>].`

Examples: `bb[:][:] = 0` zeros the entire two-dimensional array `bb` (size and shape must be known to compiler).

`c[j:len] = sqrt(c[k:len:2])` takes the square root of alternate elements of `c` starting at `c[k]`, and asserts that this is safe to vectorize (e.g., `j<k`).

**Reduction functions** are available, e.g. `__sec_reduce_add(a[:])` sums the elements of array `a`.

**SIMD-enabled Functions:** a language extension that permits functions to be called in either scalar or SIMD mode, allowing loops containing function calls to be vectorized efficiently. The compiler generates an alternate function version where one or more scalar arguments may be replaced by vectors.

C/C++ declaration syntax: `__declspec (vector(clauses)) func_name(arguments)` (or `__attribute__`)

Fortran equivalent: `!DIR$ ATTRIBUTES VECTOR: (clauses) :: func_name`

Optional clauses include *uniform*, *linear*, *mask*, *processor* and *vectorlength*.

The vector version of the function may be invoked directly using array notation, or indirectly via a loop, e.g.:

```
a[:] = func_name(b[:],c[:],d,...);
for (int i=0; i<n; i++) a[i] = func_name(b[i],c[i],d,...);
DO J=1,N; A(J) = FUNC_NAME(B(J),C(J),D,...); ENDDO
```

Similar functionality is supported via the **DECLARE SIMD** feature of OpenMP\* 4.0<sup>†</sup>. In certain cases, a SIMD pragma or directive may be needed to ensure vectorization of a loop containing a SIMD-enabled function.

### Explicit vector programming using the SIMD pragma or directive

This tells the compiler to vectorize a loop using SIMD instructions. The programmer is responsible for correctness, e.g. by explicitly specifying private variables and reductions. Semantics are similar to those for OpenMP worksharing directives.

The compiler also supports similar functionality via the SIMD clause of OpenMP 4.0<sup>†</sup>.

C/C++ syntax: `#pragma simd [clauses]` Fortran syntax: `!DIR$ SIMD [clauses]`

Clause	Description
<code>private(var1,var2,...)</code>	Specifies which variables need to be private to each iteration of the loop, to avoid conflicts.
<code>reduction(operator:var1,var2,...)</code>	Instructs the compiler to accumulate a vector reduction into <code>var1</code> , <code>var2</code> , ... under operator <code>oper</code> .
<code>linear(var1:step1,...)</code>	<code>var1</code> is incremented by <code>step1</code> for each iteration of the scalar loop.

Other supported clauses: `firstprivate`, `lastprivate`, `[no]assert`, `vectorlength`, `vectorlengthfor`, `vectorremainder`

The `_Simd` and `_Reduction` keywords provide an alternative to `#pragma simd reduction(...)`

For more information, see [www.cilk.com](http://www.cilk.com) and the Intel C++ Compiler User and Reference Guide

## Recommended Processor-Specific Optimization Options\*\*

Windows*	Linux* OS X*	Comment
<b>/Qx</b> <i>target</i>	<b>-x</b> <i>target</i>	<p>Generates specialized code for any Intel® processor that supports the instruction set specified by <i>target</i>. The executable will not run on non-Intel processors or on Intel processors that support only lower instruction sets. Possible values of <i>target</i>, from highest to lowest instruction set:</p> <p><b>CORE-AVX512, MIC-AVX512, COMMON-AVX512, CORE-AVX2, AVX, SSE4.2, ATOM_SSE4.2, SSE4.1, ATOM_SSSE3, SSSE3, SSE3, SSE2</b></p> <p><b>Note:</b> This option enables additional optimizations that are not enabled by the <i>/arch</i> or <i>-m</i> options. On 64 bit OS X, options <b>SSE3</b> and <b>SSE2</b> are not supported.</p>
<b>/arch:</b> <i>target</i>	<b>-m</b> <i>target</i>	<p>Generates specialized code for any Intel processor or compatible, non-Intel processor that supports the instruction set specified by <i>target</i>. Running the executable on an Intel processor or compatible, non-Intel processor that does not support the specified instruction set may result in a run-time error. Possible values of <i>target</i> : <b>AVX, SSE4.2, SSE4.1, SSSE3, SSE3, SSE2, IA32</b></p> <p><b>Note:</b> Option <b>IA32</b> generates non-specialized, generic x86/x87 code. It is supported on IA-32 architecture only. It is not supported on OS X.</p>
<b>/Qx</b> HOST	<b>-x</b> host	<p>Generates instruction sets up to the highest that is supported by the compilation host. On Intel processors, this corresponds to the most suitable <b>/Qx (-x)</b> option; on compatible, non-Intel processors, this corresponds to the most suitable of the <b>/arch (-m)</b> options <b>IA32, SSE2 or SSE3</b>. This option may result in additional optimizations for Intel® microprocessors that are not performed for non-Intel microprocessors.‡</p>
<b>/Qax</b> <i>target</i>	<b>-ax</b> <i>target</i>	<p>May generate specialized code for any Intel processor that supports the instruction set specified by <i>target</i>, while also generating a default code path. Possible values of <i>target</i> : <b>CORE-AVX512, MIC-AVX512, COMMON-AVX512, CORE-AVX2, AVX, SSE4.2, SSE4.1, SSSE3, SSE3, SSE2</b></p> <p>Multiple values, separated by commas, may be used to tune for additional Intel processors in the same executable, e.g. <b>/QaxAVX,SSE4.2</b>. The default code path will run on any Intel or compatible, non-Intel processor that supports at least <b>SSE2</b>, but may be modified by using in addition a <b>/Qx (-x)</b> or <b>/arch (-m)</b> switch. For example, to generate a specialized code path optimized for the 4<sup>th</sup> generation Intel® Core™ processor family and a default code path optimized for Intel processors or compatible, non-Intel processors that support at least <b>SSE3</b>, use <b>/QaxCORE-AVX2 /arch:SSE3 (-axcore-avx2 -msse3</b> on Linux*).</p> <p>At runtime, the application automatically detects whether it is running on an Intel processor, and if so, selects the most appropriate code path. If an Intel processor is not detected, the default code path is selected.</p> <p><b>Note:</b> On 64 bit OS X*, options <b>sse3</b> and <b>sse2</b> are not supported.</p> <p>This option may result in additional optimizations for Intel microprocessors that are not performed for non-Intel microprocessors.‡</p>

Please see the online article "[Intel® compiler options for Intel® SSE and Intel® AVX generation and processor-specific optimizations](#)" to view the latest recommendations for processor-specific optimization options.

\*\* Several of these options are available for both Intel® and non-Intel microprocessors but they may perform more optimizations for Intel microprocessors than they perform for non-Intel microprocessors.‡

## Compiling for Offload

Windows*	Linux*	Comment
<code>/Qoffload[-]</code> <code>/Qoffload[:<i>kywd</i>]</code>	<code>-q[no-]offload</code> <code>-qoffload=<i>kywd</i></code>	Controls whether the compiler honors language constructs for offloading to Intel® MIC architecture or Intel® Graphics Technology. <i>kywd</i> can take values: <b>none</b> : offload constructs are ignored and all code is compiled for execution on the host only. Equivalent to <code>/Qoffload- (-qno-offload)</code> <b>mandatory</b> : if target is not available, program fails unless status clause is present, in which case offload code is skipped. <b>optional</b> : if target is not available, all code is executed on host cpu. Default is <code>/Qoffload:mandatory (-qoffload=mandatory)</code>
<code>/Qoffload-option, target, tool, "option-list"</code>	<code>-qoffload-option, target, tool, "option-list"</code>	Specifies options to be used for the target compilation but not for the host. <b>target</b> may be <i>mic</i> (for Intel® MIC architecture) or <i>gfx</i> (for Intel® Graphics Technology). <b>tool</b> may be <i>compiler, ld, link</i> or <i>as</i> .
<code>/Qoffload-attribute-target: target-name</code>	<code>-qoffload-attribute-target =target-name</code>	Flags every file-scope function or data object with the offload attribute <code>target(target-name)</code> , where <b>target-name</b> is <i>mic</i> (for Intel MIC architecture) or <i>gfx</i> (for Intel Graphics Technology).
<code>/Qopt-report-phase:offload</code>	<code>-qopt-report-phase offload</code>	Generates a compile time report of variables that will be copied to or from the host and the coprocessor or processor graphics.
<code>__INTEL_OFFLOAD</code>	<code>__INTEL_OFFLOAD</code>	Predefined macro for use in offload programming on the host

## Compiling for Intel® Graphics Technology <sup>§§</sup>

Windows* (32 and 64 bits)	Linux* (64 bits only)	Comment
<code>/Qgpu-arch:arch</code>	<code>-mgpu-arch=arch</code>	The compiler generates native instructions for the graphics processor that is on the Intel® microarchitecture codenamed <i>ivybridge, haswell</i> or <i>broadwell</i> , as specified by <b>arch</b> . Default is virtual instructions translated by the jit engine.
<code>__GFX__</code>	<code>__GFX__</code>	Predefined macro for use when programming for Intel Graphics Technology

## Environment Variables for Intel® Graphics Technology <sup>§§</sup>

Variable	Comment
<code>GFX_CPU_BACKUP=1</code>	Offload code is executed on host when target is not available (default). If = <b>0</b> , application fails if target is not available.
<code>GFX_MAX_THREAD_COUNT</code>	Maximum number of target threads when parallelizing loop nests. Default= <b>-1</b> (system default).
<code>GFX_OFFLOAD_TIMEOUT=n</code>	Offload tasks time out after <b>n</b> seconds (default= <b>60</b> ). The system recovery timeout may need to be disabled or increased for this option to be effective.
<code>GFX_SHOW_TIME=1</code>	Prints offload timing information at end of execution. Default= <b>0</b> (no printing)
<code>GFX_LOG_OFFLOAD=n</code>	Generates an offload log. <b>n</b> specifies the level of detail, from <b>0</b> (no log) to <b>3</b> (maximum detail). Default is <b>0</b> .

<sup>§§</sup> Compiler support for Intel® Graphics Technology depends on operating system support.

For more information, see the Getting Started guide at <https://software.intel.com/articles/getting-started-with-compute-offload-to-intelr-graphics-technology> and the Intel® Compiler User and Reference Guides.

## Optimizing for the Intel® Xeon Phi™ Coprocessor x100 product family<sup>§</sup> (formerly code name Knights Corner)

Windows*	Linux*	Comment
/Qmic	-mmic	Builds an application that runs natively on the Intel® Xeon Phi™ coprocessor x100 product family. (Off by default).
/Qopt-streaming-cache-evict:n	-qopt-streaming-cache-evict=n	Controls whether compiler generates a cache line evict instruction after a streaming store. <i>n=0</i> no cleveict; <i>n=1</i> L1 cleveict only; <i>n=2</i> L2 cleveict only (default); <i>n=3</i> L1 and L2 cleveict generated.
/Qopt-assume-safe-padding	-qopt-assume-safe-padding	Asserts that compiler may safely access up to 64 bytes beyond the end of array or dynamically allocated objects as accessed by the user program. User is responsible for padding. Off by default.
/Qopt-threads-per-core:n	-qopt-threads-per-core=n	Hint to the compiler to optimize for <i>n</i> threads per physical core, where <i>n=1, 2, 3</i> or <i>4</i> .
/Qopt-prefetch:n	-qopt-prefetch=n	Enables increasing levels of software prefetching for <i>n=0</i> to <i>4</i> . Default is <i>n=3</i> at optimization levels of <b>-O2</b> or higher.
/Qimf-domain-exclusion:n	-fimf-domain-exclusion=n	Specifies special case arguments for which math functions need not conform to IEEE standard. The bits of <i>n</i> correspond to the domains: <b>0</b> – extreme values (e.g. very large; very small; close to singularities); <b>1</b> – NaNs; <b>2</b> – infinities; <b>3</b> – denormals; <b>4</b> – zeros.
/Qopt-gather-scatter-unroll	-qopt-gather-scatter-unroll	Specifies an alternative loop unroll sequence for gather and scatter loops
/align:array64byte	-align array64byte	Seek to align the start of arrays at a memory address that is divisible by 64, to enable aligned loads and help vectorization. (Fortran only)
__MIC__	__MIC__	Predefined macro for use in programming for Intel® MIC architecture

## Environment Variables for the Intel® Xeon Phi™ Coprocessor<sup>§</sup>

Variable	Comment
OFFLOAD_REPORT=<n>	Provides a run-time report for offload applications <i>n=1</i> reports execution times on host and on coprocessor <i>n=2</i> also reports on data transfers between host and coprocessor <i>n=3</i> detail on device initialization and individual variable transfers
OFFLOAD_DEVICES=<n1,n2,...>	Restricts the process on the host to use only the physical coprocessors <i>n1, n2</i> , etc., numbered from <b>0</b> .
MIC_STACKSIZE=<n>M	Sets the maximum stack size on the coprocessor for offload applications. In this example, <i>n</i> is in Megabytes.
MIC_ENV_PREFIX=<name>	Specify prefix to distinguish environment variables on the coprocessor from ones on the host, for offload applications. E.g. if <b>name</b> =MIC, then MIC_OMP_NUM_THREADS controls the number of OpenMP* threads on the coprocessor.
MIC_USE_2MB_BUFFERS=<n>M	Offloaded pointer variables whose runtime data length exceeds <i>n</i> MB will be allocated in large, 2MB pages.

For more optimization detail, see <https://software.intel.com/articles/advanced-optimizations-for-intel-mic-architecture> ; for building for Intel® MIC architecture in general, see <http://software.intel.com/mic-developer>, <https://software.intel.com/articles/programming-and-compiling-for-intel-many-integrated-core-architecture> and the Intel® Compiler User and Reference Guides at <https://software.intel.com/intel-cplusplus-compiler-16.0-user-and-reference-guide> and <https://software.intel.com/intel-fortran-compiler-16.0-user-and-reference-guide>.

<sup>§</sup> Intel® MIC architecture and Intel Xeon Phi coprocessors are supported by compilers within Intel® Parallel Studio XE, but not within Intel® System Studio.

## Interprocedural Optimization (IPO) and Profile-Guided Optimization (PGO) Options

Windows*	Linux* OS X*	Comment
<b>/Qip</b>	<b>-ip</b>	Single file interprocedural optimizations, including selective inlining, within the current source file.
<b>/Qipo[n]</b>	<b>-ipo[n]</b>	Permits inlining and other interprocedural optimizations among multiple source files. The optional argument <b>n</b> controls the maximum number of link-time compilations (or number of object files) spawned. Default for <b>n</b> is <b>0</b> (the compiler chooses).  Caution: This option can in some cases significantly increase compile time and code size.
<b>/Qipo-jobs[n]</b>	<b>-ipo-jobs[n]</b>	Specifies the number of commands (jobs) to be executed simultaneously during the link phase of Interprocedural Optimization (IPO). The default is <b>1</b> job.
<b>/Ob2</b>	<b>-finline- functions -finline-level=2</b>	This option enables function inlining within the current source file at the compiler's discretion. This option is enabled by default at <b>/O2</b> and <b>/O3</b> ( <b>-O2</b> and <b>-O3</b> ).  Caution: For large files, this option may sometimes significantly increase compile time and code size. It can be disabled by <b>/Ob0</b> ( <b>-fno-inline-functions</b> on Linux* and OS X*)
<b>/Qinline- factor:n</b>	<b>-finline-factor=n</b>	This option scales the total and maximum sizes of functions that can be inlined. The default value of <b>n</b> is 100, i.e., 100% or a scale factor of one.
<b>/Qprof-gen [:kywd]</b>	<b>-prof-gen [=kywd]</b>	Instruments a program for profile generation. <b>kywd= threadsafe</b> allows profile generation for threaded applications. <b>kywd=srcpos</b> and <b>globdata</b> collect additional data useful for function and data ordering.
<b>/Qprof-use</b>	<b>-prof-use</b>	Enables the use of profiling information during optimization.
<b>/Qprof-dir dir</b>	<b>-prof-dir dir</b>	Specifies a directory for profiling output files, *.dyn and *.dpi.
<b>PROF_DIR</b>	<b>PROF_DIR</b>	Environment variable to specify a directory for profiling output files (alternative to <b>/Qprof-dir</b> or <b>-prof-dir</b> ).
<b>/Qprofile- functions</b>	<b>-profile- functions</b>	Instruments functions so that a profile of execution time spent in each function may be generated.
<b>/Qprofile-loops</b>	<b>-profile-loops</b>	Instruments functions to generate a profile of each loop or loop nest in serial code. See "Profile Function or Loop Execution Time" in the main compiler documentation for additional detail and how to view profiles.

## Floating-Point Arithmetic Options

Windows*	Linux* & OS X*	Comment
<b>/fp:name</b>	<b>-fp-model name</b>	<p>Controls tradeoffs between performance, accuracy and reproducibility of floating-point results at a high level.</p> <p>Possible values of <b>name</b>:</p> <p><b>Fast[=1 =2]</b> – Allows more aggressive optimizations at a slight cost in accuracy or reproducibility. (default <b>fast=1</b>)</p> <p><b>precise</b> – Disallows optimizations that might produce slight variations in floating point results.</p> <p><b>double/extended/source</b> – Requires intermediate results to be computed in the corresponding precision. Implies <b>precise</b> unless overridden. <b>double</b> and <b>extended</b> are not available for Fortran.</p> <p><b>except</b> – Enforces floating point exception semantics.</p> <p><b>strict</b> – enables both the <b>precise</b> and <b>except</b> options and does not assume the default floating-point environment. Suppresses generation of fused multiply-add (FMA) instructions by the compiler.</p> <p><b>Recommendation:</b> <b>/fp:precise /fp:source (-fp-model precise -fp-model source)</b> is the recommended form for the majority of situations where enhanced floating point consistency and reproducibility are needed.</p>
<b>/Qopt-dynamic-align[-]</b>	<b>-q[no]-opt-dynamic-align</b>	<p>Allows [disables] certain optimizations that depend on data alignment at run-time, and that could cause small variations in floating-point results when the same, serial application is run repeatedly on the same input data. On by default unless <b>/fp:precise (-fp-model precise)</b> is set.</p>
<b>/Qftz[-]</b>	<b>-ftz[-]</b>	<p>When the main program or dll main is compiled with this option, denormals (resulting from Intel® SSE or Intel® AVX instructions) at run time are flushed to zero for the whole program (dll). Default is on except at <b>/Od (-O0)</b>.</p>
<b>/Qimf-precision: name</b>	<b>-fimf-precision: name</b>	<p>Sets the accuracy for math library functions. Default is OFF (compiler uses default heuristics). Possible values of <b>name</b> are <b>high</b>, <b>medium</b> and <b>low</b>. Reduced precision may lead to increased performance and vice versa, particularly for vectorized code.</p>
<b>/Qimf-arch-consistency: true</b>	<b>-fimf-arch-consistency=true</b>	<p>Ensures that math library functions produce consistent results across different Intel or compatible, non-Intel processors of the same architecture. May decrease run-time performance. The default is "<b>false</b>" (off).</p>
<b>/Qprec-div[-]</b>	<b>-[no]-prec-div</b>	<p>Improves [reduces] precision of floating point divides. This may slightly degrade [improve] performance.</p>
<b>/Qprec-sqrt[-]</b>	<b>-[no]-prec-sqrt</b>	<p>Improves [reduces] precision of square root computations. This may slightly degrade [improve] performance.</p>
<b>/Qprotect-parens[-]</b> <b>/assume: [no]protect _parens</b>	<b>-fprotect-parens[-]</b> <b>-assume [no]protect _parens</b>	<p>(C/C++ option) Expressions are evaluated in the order specified by parentheses (Fortran option) Default is off unless <b>/fp:precise (-fp-model precise)</b> is specified.</p>
<b>/Qfma[-]</b>	<b>-[no]fma</b>	<p>Suppresses generation of fused multiply-add (FMA) instructions by the compiler (may still be present in run-time libraries).</p>

See also <http://software.intel.com/articles/consistency-of-floating-point-results-using-the-intel-compiler>

## Fine-Tuning (All Processors)

Windows*	Linux* & OS X*	Comment
/Qunroll[n]	-unroll[n]	Sets the maximum number of times to unroll loops. <b>n=0</b> disables loop unrolling. The default is <b>/Qunroll (-unroll)</b> , which uses default heuristics.
/Qopt-prefetch:n	-qopt-prefetch=n	Enables increasing levels of software prefetching from <b>n=0</b> (no prefetching) to <b>n=4</b> (aggressive prefetching). The default value is <b>2</b> when <b>/O3 (-O3)</b> is enabled. Warning: excessive prefetching may result in resource conflicts that degrade performance.
/Qopt-block-factor:n	-qopt-block-factor=n	Specifies a preferred loop blocking factor <b>n</b> , the number of loop iterations in a block, overriding default heuristics. Loop blocking, enabled at <b>/O3 (-O3)</b> , is designed to increase the reuse of data in cache.
/Qopt-streaming-stores:mode	-qopt-streaming-stores mode	Controls streaming store generation. Values for <b>mode</b> : <b>always</b> Encourages generation of streaming stores that bypass cache, assuming application is memory bound with little data reuse <b>never</b> Disables generation of streaming stores <b>auto</b> Uses default compiler heuristics
/Qrestrict[-]	-[no]restrict	Enables [disables] pointer disambiguation with the <b>restrict</b> keyword. Off by default. (C/C++ only)
/Oa	-fno-alias	Assumes no aliasing in the program. Off by default.
/Ow	-fno-fnalias	Assumes no aliasing within functions. Off by default.
/Qalias-args[-]	-fargument-[no]alias	Implies function arguments may be aliased [are not aliased]. On by default. (C/C++ only). <b>-fargument-noalias</b> often helps the compiler to vectorize loops involving function array arguments.
/Qansi-alias[-]	-[no]-ansi-alias	Enables [disables] ANSI and ISO C Standard aliasability rules. Defaults: disabled on Windows*; enabled on Linux* and OS X*.
/Qopt-class-analysis[-]	-q[no]-opt-class-analysis	C++ class hierarchy information is used to analyze and resolve C++ virtual function calls at compile time. If a C++ application contains non-standard C++ constructs, such as pointer downcasting, it may result in different behavior. Off by default, but enabled with <b>/Qipo</b> (Windows) or <b>-ipo</b> (Linux and OS X). (C++ only)
/Qvec-threshold:n	-vec-threshold=n	Sets a threshold <b>n</b> for the auto-vectorization of loops based on the likelihood of a performance benefit. Default is <b>n=100</b> . <b>n=0</b> (vectorize loops regardless of amount of computational work) to <b>n=100</b> (vectorize loops only if a performance benefit is highly likely).
/Qvec[-]	-[no]-vec	Enables [disables] auto-vectorization. On by default at <b>/O2 (-O2)</b>
/align:arraynnbyte	-align arraynnbyte	Seek to align the start of arrays at a memory address that is divisible by <b>nn</b> , to facilitate aligned loads and help vectorization. (Fortran only)
/assume [no] buffered_io	-assume [no] buffered_io	Accumulates data for successive sequential reads or writes into a buffer for more efficient I/O. Default is unbuffered. (Fortran only)
	-f[no]-exceptions	<b>-fexceptions</b> enables exception handling table generation (default for C++). <b>-fno-exceptions</b> , default for C or Fortran, may result in smaller code. For C++, it causes exception specifications to be parsed but ignored. Any use of exception handling constructs (such as try blocks and throw statements) will produce an error if any function in the call chain has been compiled with <b>-fno-exceptions</b> .

## Debug Options

Windows*	Linux* OS X*	Comment
<b>/Zi</b> <b>/debug</b> <b>/debug:full</b> <b>/debug:all</b>	<b>-g</b> <b>-debug</b> <b>-debug full</b> <b>-debug all</b>	Produces debug information for use with any of the common platform debuggers, for full symbolic debugging of unoptimized code. Turns off optimization unless <b>/O2 (-O2)</b> (or another <b>O</b> option) is specified. Debug symbols will generally increase the size of object modules and may slightly degrade performance of optimized code.
<b>/debug:none</b>	<b>-debug none</b>	No debugging information is generated. (default)
<b>/debug:minimal</b>	<b>-debug minimal</b>	Generates line number information for debugging, but not local symbols.
<b>/debug:inline</b> <b>-debug-info</b>	<b>-debug inline-debug-info</b>	This option causes symbols for inlined functions to be associated with the source of the called function, instead of with the caller. Not enabled by <b>/debug:full (-debug full)</b> unless <b>-O2</b> is specified.
	<b>-debug extended</b>	produces additional information for improved symbolic debugging of optimized code. (Linux* only; not enabled by <b>-debug full</b> ).
	<b>-debug parallel</b>	generates additional symbols and instrumentation for debugging threaded code. (Linux only; not enabled by <b>-debug full</b> ).
<b>/Qsox[-]</b>	<b>-[no-]sox</b> (Linux only)	Embeds the compiler version and options used as strings in the object file (Windows* and Linux) and in the executable (Linux). Off by default.
<b>/Qtraceback</b>	<b>-traceback</b>	Compiler includes slight extra information in the object file to provide source file traceback information when a severe error occurs at run time. May be used with optimized code. (Fortran applications only)
<b>/Qinit:snan, arrays</b>	<b>-init=snan, arrays</b>	Initializes certain floating-point variables to a signaling NaN to facilitate run-time detection of uninitialized floating-point variables. (Fortran only). Sets <b>/fpe:0 (-fpe0)</b> to unmask floating-point exceptions.

### ‡ Optimization Notice

Intel's compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include SSE2, SSE3, and SSSE3 instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel. Microprocessor-dependent optimizations in this product are intended for use with Intel microprocessors. Certain optimizations not specific to Intel microarchitecture are reserved for Intel microprocessors. Please refer to the applicable product User and Reference Guides for more information regarding the specific instruction sets covered by this notice.

Notice revision #20110804

**For product and purchase information, visit the Intel® Software Development Tools site at: <http://intel.ly/sw-dev-tools>.**

Intel, the Intel logo, Intel VTune, Intel Core, Intel Cilk and Intel Xeon Phi are trademarks of Intel Corporation in the U.S. and other countries.  
\* Other names and brands may be claimed as the property of others.

© 2015, Intel Corporation. All rights reserved.