

Intel® Software  
Development Products



# Quick Reference Guide to Optimization with Intel® Compilers Version 13

For IA-32 processors, Intel® 64 processors,  
and compatible, non-Intel processors



## Application Performance

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

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

1. Use the general optimization options (Windows\* `/O1`, `/O2`, or `/O3`; Linux\* and OS X\* `-O1`, `-O2`, or `-O3`) and determine which one works best for your application by measuring performance with each. Most users should start at `/O2 (-O2)` (default) before trying more advanced optimizations. Next, try `/O3 (-O3)` for loop-intensive applications.<sup>1</sup>
2. Fine-tune performance to target IA-32 and Intel® 64-based systems with processor-specific options. Examples are `/QxAVX (-xavx)` for the 2<sup>nd</sup>- and 3<sup>rd</sup>-Generation Intel® Core™ processor families, and `/arch:SSE3 (-msse3)` for compatible, non-Intel processors that support at least the SSE3 instruction set. Alternatively, you can use `/QxHOST (-xhost)`, which will use the most advanced instruction set for the processor on which you compiled.<sup>1</sup> For a more extensive list of options that optimize for specific processors or instruction sets, see the table “Recommended Processor-Specific Optimization Options.”
3. Add interprocedural optimization (IPO), `/Qipo (-ipo)` and/or profile-guided optimization (PGO), `/Qprof-gen` and `/Qprof-use (-prof-gen` and `-prof-use)`<sup>1</sup>, then measure performance again to determine whether your application benefits from one or both of them.
4. Optimize your application for vector and parallel execution on multithreaded, multicore, and multiprocessor systems using: advice from the Guided Auto-Parallelism (GAP) feature, `/Qguide (-guide)`; the Intel® Cilk™ Plus language extensions for C/C++; the parallel performance options `/Qparallel (-parallel)` or `/Qopenmp (-openmp)`; or by using the Intel® Performance Libraries included with the product.<sup>1</sup>
5. Use Intel® VTune™ Amplifier XE to help you identify serial and parallel performance hotspots, so that you know which specific parts of your application could benefit from further tuning. Use Intel® Inspector XE to reduce the time to market for threaded applications by diagnosing memory and threading errors and speeding up the development process. These products cannot be used on non-Intel microprocessors.

Please consult the main product documentation for more details. For tuning applications for Intel® Xeon Phi™ coprocessors, see the references in the section “Options for the Intel® Xeon Phi™ Coprocessor.”

<sup>1</sup>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.

## Intel® Software Development Products

### General Optimization Options

These options are available for both Intel and non-Intel microprocessors, but they may result in more optimizations for Intel microprocessors than for non-Intel microprocessors.

Windows*	Linux*/OS X*	Comment
<b>/Od</b>	<b>-O0</b>	<b>No optimization.</b> Used during the early stages of application development and debugging. Use a higher setting when the application is working correctly.
<b>/O1</b>	<b>-O1</b>	<b>Optimize for size.</b> Omits optimizations that tend to increase object size. Creates the smallest optimized code in most cases.  This option is useful in many large server/database applications where memory paging due to larger code size is an issue.
<b>/O2</b>	<b>-O2</b>	<b>Maximize speed.</b> Default setting. Enables many optimizations, including vectorization. Creates faster code than <b>/O1 (-O1)</b> in most cases.
<b>/O3</b>	<b>-O3</b>	Enables <b>/O2 (-O2)</b> optimizations, plus more aggressive loop and memory-access optimizations, such as scalar replacement, loop unrolling, code replication to eliminate branches, loop blocking to allow more efficient use of cache, and additional data prefetching.  The <b>/O3 (-O3)</b> 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 <b>/O2 (-O2)</b> .
<b>/Qopt-report[:<i>n</i>]</b>	<b>-opt-report [<i>n</i>]</b>	Generates an optimization report directed to stderr. <i>n</i> specifies the level of detail, from 0 (no report) to 3 (maximum detail). Default is 2.
<b>/Qopt-report-file:<i>name</i></b>	<b>-opt-report-file=<i>name</i></b>	Generates an optimization report directed to the file <i>name</i> .

Windows*	Linux*/OS X*	Comment
<b>/Qopt-report-phase:<i>name</i></b>	<b>-opt-report-phase=<i>name</i></b>	Optimization reports are generated for phase <i>name</i> . The option can be used multiple times in the same compilation to get output from multiple phases. Some commonly used <i>name</i> arguments are as follows: <b>all</b> —All possible optimization reports for all phases (default) <b>ipo_inl</b> —Inlining report from the Interprocedural Optimizer <b>hlo</b> —High-Level Optimizer (includes loop and memory optimizations) <b>hpo</b> —High-Performance Optimizer (includes vectorizer and parallelizer) <b>pgo</b> —Profile Guided Optimizer
<b>/Qopt-report-help</b>	<b>-opt-report-help</b>	Displays all possible values of <i>name</i> for /Qopt-report-phase (-opt-report-phase) above. No compilation is performed.
<b>/Qopt-report-routine:<i>string</i></b>	<b>-opt-report-routine=<i>string</i></b>	Generates reports only for functions or subroutines whose names contain <i>string</i> . By default, reports are generated for all functions and subroutines.

## Parallel Performance

Options that use OpenMP\* or auto-parallelization are available for both Intel and non-Intel microprocessors, but these options may result in additional optimizations on Intel microprocessors that do not occur on non-Intel microprocessors.

Windows*	Linux*/OS X*	Comment
<b>/Qopenmp</b>	<b>-openmp</b>	Causes multithreaded code to be generated when OpenMP* directives are present. For Fortran only: makes local arrays automatic and may require an increased stack size. See <a href="http://www.openmp.org">http://www.openmp.org</a> for the OpenMP API specification.
<b>/Qparallel</b>	<b>-parallel</b>	The auto-parallelizer detects simply structured loops that may be safely executed in parallel, including loops implied by Intel® Cilk™ Plus array notation, and automatically generates multithreaded code for these loops.

<b>/Qpar-report[:n]</b>	<b>-par-report[n]</b>	Controls the auto-parallelizer's diagnostic level. <i>n</i> specifies the level of detail, from 0 (no report) to 3 (maximum detail). Default is 0.
<b>/Qpar-threshold[:n]</b>	<b>-par-threshold[n]</b>	Sets a threshold for the auto-parallelization of loops based on the likelihood of a performance benefit. <i>n</i> =0 to 100, default 100. 0: Parallelize loops regardless of computation work volume. 100: Parallelize loops only if a performance benefit is highly likely. Must be used in conjunction with <b>/Qparallel (-parallel)</b> .
<b>/Qpar-affinity: name</b>	<b>-par-affinity= name</b>	Specifies thread-processor affinity for OpenMP or auto-parallelized applications. Typical values of <i>name</i> are <i>none</i> (default), <i>scatter</i> , and <i>compact</i> . Has effect only when compiling the main program. See the compiler User and Reference Guide for more settings and details.

Windows*	Linux*/OS X*	Comment
<b>/Qguide[:n]</b>	<b>-guide[=n]</b>	Guided auto-parallelization. Causes the compiler to suggest ways to help loops to vectorize or auto-parallelize, without producing any objects or executables. Auto-parallelization advice is given only if the option <b>/Qparallel (-parallel)</b> is also specified. <i>n</i> is an optional value from 1 to 4 specifying increasing levels of guidance to be provided; level 4 being the most advanced and aggressive. If <i>n</i> is omitted, the default is 4.
<b>/Qopt-matmul[-]</b>	<b>-[no-]opt-matmul</b>	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.
<b>/Qcilk-serialize</b>	<b>-cilk-serialize</b>	This option includes a header file, <code>cilk_stubs.h</code> , which 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.
<b>/Qcoarray:shared</b>	<b>-coarray=shared</b>	Enables coarrays from the Fortran 2008 standard on shared memory systems (Fortran only). See the Compiler Reference Guide for more coarray options and detail.

<code>/Qmkl:name</code>	<code>-mkl=name</code>	Requests that the Intel® Math Kernel Library (Intel® MKL) be linked. Off by default. Possible values of <i>name</i> are: <i>parallel</i> —Links the threaded part of Intel MKL (default) <i>sequential</i> —Links the non-threaded part of Intel MKL <i>cluster</i> —Links the cluster and sequential parts of Intel MKL
-------------------------	------------------------	---

## Parallel Performance Using Intel® Cilk™ Plus

Threading Keywords (C/C++)	Description
<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_opadd<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 threadsafe.

**Array Notation:** A readable, explicitly data-parallel C and C++ language extension that facilitates generation of SIMD parallel code by the compiler 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 of the array object must be known to the 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`).

**Elemental Functions:** A language extension that allows functions to be called in either scalar or SIMD parallel 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)`

Fortran equivalent: `IDIR$ ATTRIBUTES VECTOR: (clauses) :: func_name`

Optional clauses include *procname*, *vectorlength*, *vectorlengthfor*, *linear*, *uniform*, and *mask*.

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

```
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
```

In certain cases, a SIMD pragma (C/C++) or directive (Fortran) may be needed to ensure vectorization of a loop containing an elemental function.

#### Explicit vector programming using the SIMD pragma (C/C++) and directive (Fortran)

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 the OpenMP\* directives:

#pragma omp parallel for (!\$OMP PARALLEL DO).

C/C++ syntax: #pragma simd (clauses) Fortran syntax: !DIR\$ SIMD (clauses)

Clause	Description
<b>private(var1,var2,...)</b>	Specifies which variables need to be private to each iteration of the loop, to avoid conflicts.
<b>reduction(oper:var1,var2,...)</b>	Instructs the compiler to accumulate a vector reduction into <i>var1</i> , <i>var2</i> , ... under operator <i>oper</i> .
<b>[no]assert</b>	Compiler emits an error if it is unable to vectorize the loop. (default noassert).
<b>linear(var1:step1,...)</b>	<i>var1</i> is incremented by <i>step1</i> for each iteration of the scalar loop.

Other supported clauses: vectorlength, vectorlengthfor, firstprivate, lastprivate, vecremainder.

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<sup>‡</sup>

Windows*	Linux*/OS X*	Comment
<b>/Qxtarget</b>	<b>-xtarget</b>	<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: CORE-AVX2, AVX, SSE4.2, SSE4.1, SSSE3, SSE3, SSE2</p> <p><b>Note:</b> On OS X*, options SSE3 and SSE2 are not supported.</p> <p>This option enables additional optimizations that are not enabled by the /arch or -m options.</p>
<b>/arch:target</b>	<b>-mtarget</b>	<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 runtime error.</p> <p>Possible values of <i>target</i>: AVX, SSE4.2, SSE4.1, SSSE3, SSE3, SSE2, IA32</p> <p><b>Note:</b> Option IA32 generates non-specialized, generic x86/x87 code. It is supported on IA-32 architecture only. On OS X, options SSE3, SSE2 and IA32 are not supported.</p>
<b>/QxHOST</b>	<b>-xhost</b>	<p>Generates instruction sets up to the highest that is supported by the compilation host. On Intel processors, this corresponds to the most suitable /Qx (-x) option; on compatible, non-Intel processors, this corresponds to the most suitable of the /arch (-m) options IA32, SSE2, or SSE3. This option may result in additional optimizations for Intel® microprocessors that are not performed for non-Intel microprocessors.<sup>‡</sup></p>



Windows*	Linux*/OS X*	Comment
<b><i>/Qaxtarget</i></b>	<b><i>-axtarget</i></b>	<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>: CORE-AVX2, AVX, SSE4.2, SSE4.1, SSSE3, SSE3, SSE2.</p> <p>Multiple values, separated by commas, may be used to tune for additional Intel processors in the same executable (e.g., <i>/QaxAVX,SSE4.2</i>). The default code path will run on any Intel or compatible, non-Intel processor that supports at least SSE2, but may be modified by using, in addition, a <i>/Qx (-x)</i> or <i>/arch (-m)</i> switch.</p> <p>For example, to generate a specialized code path optimized for the 2<sup>nd</sup> Generation Intel® Core™ processor family and a default code path optimized for Intel processors or compatible, non-Intel processors that support at least SSE3, use <i>/QaxAVX /arch:SSE3</i> (<i>-axavx -msse3</i> 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 OS X, options <i>sse3</i> and <i>sse2</i> 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 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. These options are described in greater detail in the Intel® Compiler User and Reference Guides.

## Options for the Intel® Xeon Phi™ Coprocessor (Linux\* only)

<b><i>-mmic</i></b>	Builds an application that runs natively on Intel® Xeon Phi™ coprocessors. Off by default.
<b><i>-no-offload</i></b>	The compiler ignores language constructs for offloading to Intel Xeon Phi coprocessors and builds for the host only. By default, offload constructs are honored and compiled for execution on the coprocessor.
<b><i>-offload-option, mic, tool, "option-list"</i></b>	Specifies options to be used for the target compilation but not for the host. <i>tool</i> may be <i>compiler, ld, or as</i> .

<b>-opt-report-phase offload</b>	Generates a report at compile time of variables that will be copied from the host to the coprocessor and vice versa.
<b>-opt-cache-evict=<i>n</i></b>	Controls whether compiler generates a cache line evict instruction after a streaming store. <i>n=0</i> no clevict; <i>n=1</i> L1 clevict only; <i>n=2</i> L2 clevict only (default); <i>n=3</i> L1 and L2 clevict generated.
<b>-opt-assume-safe-padding</b>	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.
<b>-opt-threads-per-core=<i>n</i></b>	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> .
<b>-opt-prefetch=<i>n</i></b>	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 <code>-O2</code> or higher.
<b>-fimf-domain-exclusion=<i>n</i></b>	Specifies special case arguments for which math functions need not conform to IEEE standard. The bits of <i>n</i> correspond to these domains:  <i>0</i> : extreme values (e.g., very large; very small; close to singularities); <i>1</i> : NaNs; <i>2</i> : infinities; <i>3</i> : denormals; <i>4</i> : zeros.
<b>-align array64byte</b> (Fortran only)	Where possible, align the start of arrays at a memory address that is divisible by 64, to enable aligned loads and help vectorization.

## Environment Variables for the Intel® Xeon Phi™ Coprocessor

Variable	Comment
<b>OFFLOAD_REPORT=&lt;<i>n</i>&gt;</b>	Provides a runtime 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
<b>OFFLOAD_DEVICES= &lt;<i>n1,n2,...</i>&gt;</b>	Restricts the process on the host to use only the physical coprocessors <i>n1</i> , <i>n2</i> , etc., numbered from 0.
<b>MIC_STACKSIZE=&lt;<i>n</i>&gt;M</b>	Sets the maximum stack size on the coprocessor for offload applications. In this example, <i>n</i> is in megabytes.
<b>MIC_ENV_PREFIX=&lt;name&gt;</b>	Specifies prefix to distinguish environment variables on the coprocessor from ones on the host, for offload applications.  For example, if <i>name</i> =MIC, then MIC_OMP_NUM_THREADS controls the number of OpenMP* threads on the coprocessor.

<b>MIC_USE_2MB_BUFFERS=&lt;n&gt;M</b>	Offloaded pointer variables whose runtime data length exceeds <i>n</i> MB will be allocated in large, 2MB pages.
---------------------------------------	--

For more optimization detail, see <http://software.intel.com/en-us/articles/advanced-optimizations-for-intel-mic-architecture>.

For building for Intel® Many Integrated Core (MIC) architecture in general, see <http://software.intel.com/en-us/articles/programming-and-compiling-for-intel-many-integrated-core-architecture>, <http://software.intel.com/mic-developer> and the Intel® Compiler User and Reference Guides.

## 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 <i>n</i> controls the maximum number of link-time compilations (or number of object files) spawned. Default for <i>n</i> is 0 (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 1 job.
<b>/Ob2</b>	<b>-finline-functions</b> <b>-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 /O2 and /O3 (-O2 and -O3).  Caution: For large files, this option may sometimes significantly increase compile time and code size. It can be disabled by /Ob0 (-fno-inline-functions 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 <i>n</i> is 100 (i.e., 100% or a scale factor of one).
<b>/Qprof-gen</b>	<b>-prof-gen</b>	Instruments a program for profile generation.

Windows*	Linux*/OS X*	Comment
<code>/Qprof-use</code>	<code>-prof-use</code>	Enables the use of profiling information during optimization.
<code>/Qprof-dir dir</code>	<code>-prof-dir dir</code>	Specifies a directory for the profiling output files, *.dyn and *.dpi.
<code>/Qprofile-functions</code>	<code>-profile-functions</code>	Instruments functions so that a profile of execution time spent in each function may be generated.
<code>/Qprofile-loops</code>	<code>-profile-loops</code>	Instruments functions to generate a profile of each loop or loop nest. See "Profile Function or Loop Execution Time" in the main compiler documentation for additional detail and information on how to view profiles.

## Floating-Point Arithmetic Options

Windows*	Linux*/OS X*	Comment
<code>/fp:name</code>	<code>-fp-model name</code>	<p>May enhance the consistency of floating-point results by restricting certain optimizations. Possible values of <i>name</i>:</p> <p><b>fast=[1 2]</b>—Allows more aggressive optimizations at a slight cost in accuracy or consistency. (<b>fast=1</b> is the default). This may include some additional optimizations that are performed on Intel® microprocessors but not on non-Intel microprocessors.</p> <p><b>precise</b>—Allows only value-safe optimizations on floating point code.</p> <p><b>double/extended/source</b>—Intermediate results are computed in double, extended, or source precision. Implies <b>precise</b> unless overridden.</p> <p>The <b>double</b> and <b>extended</b> options are not available for the Intel® Fortran compiler.</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.</p> <p><b>Recommendation:</b> <code>/fp:precise /fp:source (-fp-model precise -fp-model source)</code> is the recommended form for the majority of situations where enhanced floating-point consistency and reproducibility are needed.</p>
<code>/Qftz[-]</code>	<code>-ftz[-]</code>	When the main program or dll main is compiled with this option, denormals (resulting from Intel® SSE or AVX instructions) at runtime are flushed to zero for the whole program (dll). The default is on except at <code>/Od (-OO)</code> .

Windows*	Linux*/OS X*	Comment
<b>/Qimf-precision:<i>name</i></b>	<b>-fimf-precision:<i>name</i></b>	This option defines the accuracy for math library functions. The default is OFF (compiler uses default heuristics). Possible values of <i>name</i> are <i>high</i> , <i>medium</i> , and <i>low</i> . Reduced precision may lead to increased performance and vice versa. Many routines in the math library are more highly optimized for Intel microprocessors than for non-Intel microprocessors.
<b>/Qimf-arch-consistency:<i>true</i></b>	<b>-fimf-arch-consistency=<i>true</i></b>	Ensures that math library functions produce consistent results across different Intel* or compatible, non-Intel processors of the same architecture. May decrease runtime performance. The default is " <i>false</i> " (off).
<b>/Qprec-div[-]</b>	<b>-[no-]prec-div</b>	Improves [reduces] precision of floating-point divides. This may slightly degrade [improve] performance.
<b>/Qprec-sqrt[-]</b>	<b>-[no-]prec-sqrt</b>	Improves [reduces] precision of square root computations. This may slightly degrade [improve] performance.

For more information, go to: <http://software.intel.com/en-us/articles/consistency-of-floating-point-results-using-the-intel-compiler>.

## Fine-Tuning (All Processors)

Windows*	Linux*/OS X*	Comment
<b>/Qunroll[<i>n</i>]</b>	<b>-unroll[<i>n</i>]</b>	Sets the maximum number of times to unroll loops. <b>/Qunroll0</b> (-unroll0) disables loop unrolling. The default is <b>/Qunroll</b> (-unroll), which uses default heuristics.
<b>/Qopt-prefetch:<i>n</i></b>	<b>-opt-prefetch=<i>n</i></b>	Controls the level of software prefetching. <i>n</i> is an optional value between 0 (no prefetching) and 4 (aggressive prefetching), with a default value of 2 when high-level optimization is enabled. Warning: excessive prefetching may result in resource conflicts that degrade performance.
<b>/Qopt-block-factor:<i>n</i></b>	<b>-opt-block-factor=<i>n</i></b>	Specifies preferred loop blocking factor <i>n</i> , the number of loop iterations in a block, overriding default heuristics. Loop blocking is enabled at <b>/O3</b> (-O3) and is designed to increase the reuse of data in cache.

<b>/Qopt-streaming-stores:mode</b>	<b>-opt-streaming-stores mode</b>	Specifies whether streaming stores may be generated. Values for <i>mode</i> :  <b>always</b> —Encourages compiler to generate 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> —Default compiler heuristics for streaming store generation
<b>/Qrestrict[-]</b>	<b>-[no]restrict</b>	Enables [disables] pointer disambiguation with the <b>restrict</b> keyword. Off by default. (C/C++ only)
<b>/Oa</b>	<b>-fno-alias</b>	Assumes no aliasing in the program. Off by default.
<b>/Ow</b>	<b>-fno-fnalias</b>	Assumes no aliasing within functions. Off by default.
<b>/Qalias-args[-]</b>	<b>-fargument-[no]alias</b>	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.

Windows*	Linux*/OS X*	Comment
<b>/Qopt-class-analysis[-]</b>	<b>-[no-]opt-class-analysis</b>	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. Default is off, but it is turned on by default with the <b>/Qipo</b> (Windows*) or <b>-ipo</b> (Linux* and OS X*) compiler option, enabling improved C++ optimization. (C++ only).
	<b>-f[no-]exceptions</b>	<b>-fexceptions</b> , default for C++, enables exception handling table generation.  <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> .
<b>/Qvec-threshold:n</b>	<b>-vec-threshold n</b>	Sets a threshold <i>n</i> for the vectorization of loops based on the probability of performance gain. $0 \leq n \leq 100$ , default $n=100$ .  <b>0</b> —Vectorize loops regardless of amount of computational work.  <b>100</b> —Vectorize loops only if a performance benefit is almost certain.

<b>/Qvec-report:n</b>	<b>-vec-report n</b>	Controls the vectorizer's diagnostic levels. Values <i>n</i> = 0, 1, 2, 3, and 6 specify increasing levels of detail. Default is 0 (no report).
-----------------------	----------------------	---

## 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 /O2 (-O2) and makes /Od (-O0) the default unless /O2 (-O2) (or another O 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- 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 /debug:full (-debug full) unless -O2 is specified.
	<b>-debug extended</b>	Produces additional information for improved symbolic debugging of optimized code. Not enabled by /debug:full (-debug full).
	<b>-debug parallel</b>	Generates additional symbols and instrumentation for debugging threaded code. (Linux* only; not enabled by -debug full).
<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 runtime. May be used with optimized code. (Fortran applications only)

## ‡ 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 Products site at:  
[www.intel.com/software/products/compilers](http://www.intel.com/software/products/compilers).

© 2013, Intel Corporation. All rights reserved. Intel, the Intel logo, Intel Cilk, Intel Core, VTune, and Intel Xeon Phi are trademarks of Intel Corporation in the U.S. and/or other countries. \*Other names and brands may be claimed as the property of others.