

# OpenMP Technical Report 3 on OpenMP 4.0 enhancements

---

This Technical Report specifies OpenMP 4.0 enhancements that are candidates for a future OpenMP 4.1: (e.g. for asynchronous execution on and data transfer to offload devices, task loop construct that helps in simplifying task creation, etc)

**All members of the OpenMP Language Working Group**

**November 13, 2014**

**Expires: November 13, 2017**

We actively solicit comments. Please provide feedback on this document either to the Editor directly or in the OpenMP Forum at [openmp.org](http://openmp.org)

**End of Public Comment Period: January 12, 2015**

This technical report describes possible future directions or extensions to the [OpenMP Specification](#).

The goal of this technical report is to build more widespread existing practice for an expanded [OpenMP](#). It gives advice on extensions or future directions to those vendors who wish to provide them possibly for trial implementation, allows [OpenMP](#) to gather early feedback, support timing and scheduling differences between official [OpenMP](#) releases, and offers a preview to users of the future directions of [OpenMP](#) with the provision stated in the next paragraph.

This technical report is non-normative. Some of the components in this technical report may be considered for standardization in a future version of [OpenMP](#), but they are not currently part of any [OpenMP Specification](#). Some of the components in this technical report may never be standardized, others may be standardized in a substantially changed form, or it may be standardized as is in its entirety.



# OpenMP Application Programming Interface

**Version 4.1rev0, November, 2014**

Copyright © 1997-2014 OpenMP Architecture Review Board.

Permission to copy without fee all or part of this material is granted, provided the OpenMP Architecture Review Board copyright notice and the title of this document appear. Notice is given that copying is by permission of OpenMP Architecture Review Board.

This page intentionally left blank in published version.

This is Revision 0-TR3 (1 Nov 2014) and includes the following tickets applied to the 4.0 LaTeX sources: 267, 268, 271, 279, 281, 282, 284, 285, 288, 289, 290, 297, 298, 300, 302, 304, 308-311, 314, 316, 318, 323-326, 328, 330, 332-336, 338, 339, 341, 343-345, 347-352, 355, 357.

This is a draft - DO NOT DISTRIBUTE

# Contents

---

<b>1</b>	<b>Introduction</b>	<b>1</b>
1.1	Scope . . . . .	1
1.2	Glossary . . . . .	2
1.2.1	Threading Concepts . . . . .	2
1.2.2	OpenMP Language Terminology . . . . .	2
1.2.3	Loop Terminology . . . . .	8
1.2.4	Synchronization Terminology . . . . .	8
1.2.5	Tasking Terminology . . . . .	9
1.2.6	Data Terminology . . . . .	10
1.2.7	Implementation Terminology . . . . .	12
1.3	Execution Model . . . . .	13
1.4	Memory Model . . . . .	17
1.4.1	Structure of the OpenMP Memory Model . . . . .	17
1.4.2	Device Data Environments . . . . .	18
1.4.3	The Flush Operation . . . . .	18
1.4.4	OpenMP Memory Consistency . . . . .	20
1.5	OpenMP Compliance . . . . .	21
1.6	Normative References . . . . .	21
1.7	Organization of this document . . . . .	23
<b>2</b>	<b>Directives</b>	<b>25</b>
2.1	Directive Format . . . . .	26
2.1.1	Fixed Source Form Directives . . . . .	28
2.1.2	Free Source Form Directives . . . . .	29
2.1.3	Stand-Alone Directives . . . . .	32
2.2	Conditional Compilation . . . . .	32
2.2.1	Fixed Source Form Conditional Compilation Sentinels . . . . .	33

2.2.2	Free Source Form Conditional Compilation Sentinel . . . . .	33
2.3	Internal Control Variables . . . . .	35
2.3.1	ICV Descriptions . . . . .	35
2.3.2	ICV Initialization . . . . .	36
2.3.3	Modifying and Retrieving ICV Values . . . . .	38
2.3.4	How ICVs are Scoped . . . . .	40
2.3.4.1	How the Per-Data Environment ICVs Work . . . . .	40
2.3.5	ICV Override Relationships . . . . .	41
2.4	Array Sections . . . . .	43
2.5	<b>parallel</b> Construct . . . . .	44
2.5.1	Determining the Number of Threads for a <b>parallel</b> Region . . . . .	48
2.5.2	Controlling OpenMP Thread Affinity . . . . .	50
2.6	Canonical Loop Form . . . . .	52
2.7	Worksharing Constructs . . . . .	54
2.7.1	Loop Construct . . . . .	55
2.7.1.1	Determining the Schedule of a Worksharing Loop . . . . .	62
2.7.2	<b>sections</b> Construct . . . . .	63
2.7.3	<b>single</b> Construct . . . . .	65
2.7.4	<b>workshare</b> Construct . . . . .	67
2.8	SIMD Constructs . . . . .	70
2.8.1	<b>simd</b> Construct . . . . .	70
2.8.2	<b>declare simd</b> Construct . . . . .	74
2.8.3	Loop SIMD Construct . . . . .	78
2.9	Tasking Constructs . . . . .	80
2.9.1	<b>task</b> Construct . . . . .	80
2.9.2	<b>taskloop</b> Construct . . . . .	83
2.9.3	<b>taskloop simd</b> Construct . . . . .	88
2.9.4	<b>taskyield</b> Construct . . . . .	89
2.9.5	Task Scheduling . . . . .	90
2.10	Device Constructs . . . . .	92
2.10.1	<b>target data</b> Construct . . . . .	92
2.10.2	<b>target</b> Construct . . . . .	93
2.10.3	<b>target update</b> Construct . . . . .	96

2.10.4	<b>declare target</b> Directive	99
2.10.5	<b>teams</b> Construct	102
2.10.6	<b>distribute</b> Construct	104
2.10.7	<b>distribute simd</b> Construct	107
2.10.8	Distribute Parallel Loop Construct	109
2.10.9	Distribute Parallel Loop SIMD Construct	110
2.10.10	<b>target enter data</b> Construct	112
2.10.11	<b>target exit data</b> Construct	114
2.11	Combined Constructs	116
2.11.1	Parallel Loop Construct	117
2.11.2	<b>parallel sections</b> Construct	118
2.11.3	<b>parallel workshare</b> Construct	120
2.11.4	Parallel Loop SIMD Construct	121
2.11.5	<b>target teams</b> construct	123
2.11.6	<b>teams distribute</b> Construct	125
2.11.7	<b>teams distribute simd</b> Construct	126
2.11.8	<b>target teams distribute</b> construct	127
2.11.9	<b>target teams distribute simd</b> Construct	128
2.11.10	Teams Distribute Parallel Loop Construct	130
2.11.11	Target Teams Distribute Parallel Loop Construct	131
2.11.12	Teams Distribute Parallel Loop SIMD Construct	132
2.11.13	Target Teams Distribute Parallel Loop SIMD Construct	133
2.12	Master and Synchronization Constructs and Clauses	135
2.12.1	<b>master</b> Construct	135
2.12.2	<b>critical</b> Construct	136
2.12.3	<b>barrier</b> Construct	138
2.12.4	<b>taskwait</b> Construct	139
2.12.5	<b>taskgroup</b> Construct	140
2.12.6	<b>atomic</b> Construct	141
2.12.7	<b>flush</b> Construct	148
2.12.8	<b>ordered</b> Construct	152
2.12.9	<b>depend</b> Clause	154

2.13	Cancellation Constructs . . . . .	156
2.13.1	<b>cancel</b> Construct . . . . .	156
2.13.2	<b>cancellation point</b> Construct . . . . .	160
2.14	Data Environment . . . . .	162
2.14.1	Data-sharing Attribute Rules . . . . .	162
2.14.1.1	Data-sharing Attribute Rules for Variables Referenced in a Construct	162
2.14.1.2	Data-sharing Attribute Rules for Variables Referenced in a Region but not in a Construct . . . . .	166
2.14.2	<b>threadprivate</b> Directive . . . . .	166
2.14.3	Data-Sharing Attribute Clauses . . . . .	172
2.14.3.1	<b>default</b> clause . . . . .	173
2.14.3.2	<b>shared</b> clause . . . . .	174
2.14.3.3	<b>private</b> clause . . . . .	176
2.14.3.4	<b>firstprivate</b> clause . . . . .	179
2.14.3.5	<b>lastprivate</b> clause . . . . .	182
2.14.3.6	<b>reduction</b> clause . . . . .	184
2.14.3.7	<b>linear</b> clause . . . . .	190
2.14.4	Data Copying Clauses . . . . .	191
2.14.4.1	<b>copyin</b> clause . . . . .	192
2.14.4.2	<b>copyprivate</b> clause . . . . .	193
2.14.5	<b>map</b> Clause . . . . .	195
2.15	<b>declare reduction</b> Directive . . . . .	199
2.16	Nesting of Regions . . . . .	205
<b>3</b>	<b>Runtime Library Routines</b>	<b>206</b>
3.1	Runtime Library Definitions . . . . .	207
3.2	Execution Environment Routines . . . . .	208
3.2.1	<b>omp_set_num_threads</b> . . . . .	208
3.2.2	<b>omp_get_num_threads</b> . . . . .	209
3.2.3	<b>omp_get_max_threads</b> . . . . .	210
3.2.4	<b>omp_get_thread_num</b> . . . . .	212
3.2.5	<b>omp_get_num_procs</b> . . . . .	213
3.2.6	<b>omp_in_parallel</b> . . . . .	213
3.2.7	<b>omp_set_dynamic</b> . . . . .	214



3.2.8	<code>omp_get_dynamic</code>	216
3.2.9	<code>omp_get_cancellation</code>	217
3.2.10	<code>omp_set_nested</code>	217
3.2.11	<code>omp_get_nested</code>	219
3.2.12	<code>omp_set_schedule</code>	220
3.2.13	<code>omp_get_schedule</code>	222
3.2.14	<code>omp_get_thread_limit</code>	223
3.2.15	<code>omp_set_max_active_levels</code>	223
3.2.16	<code>omp_get_max_active_levels</code>	225
3.2.17	<code>omp_get_level</code>	226
3.2.18	<code>omp_get_ancestor_thread_num</code>	227
3.2.19	<code>omp_get_team_size</code>	228
3.2.20	<code>omp_get_active_level</code>	229
3.2.21	<code>omp_in_final</code>	230
3.2.22	<code>omp_get_proc_bind</code>	231
3.2.23	<code>omp_set_default_device</code>	233
3.2.24	<code>omp_get_default_device</code>	234
3.2.25	<code>omp_get_num_devices</code>	235
3.2.26	<code>omp_get_num_teams</code>	235
3.2.27	<code>omp_get_team_num</code>	237
3.2.28	<code>omp_is_initial_device</code>	238
3.3	Lock Routines	239
3.3.1	<code>omp_init_lock</code> and <code>omp_init_nest_lock</code>	241
3.3.2	<code>omp_destroy_lock</code> and <code>omp_destroy_nest_lock</code>	241
3.3.3	<code>omp_set_lock</code> and <code>omp_set_nest_lock</code>	242
3.3.4	<code>omp_unset_lock</code> and <code>omp_unset_nest_lock</code>	243
3.3.5	<code>omp_test_lock</code> and <code>omp_test_nest_lock</code>	244
3.4	Timing Routines	245
3.4.1	<code>omp_get_wtime</code>	246
3.4.2	<code>omp_get_wtick</code>	248
<b>4</b>	<b>Environment Variables</b>	<b>249</b>
4.1	<code>OMP_SCHEDULE</code>	251
4.2	<code>OMP_NUM_THREADS</code>	252

4.3	<b>OMP_DYNAMIC</b> . . . . .	253
4.4	<b>OMP_PROC_BIND</b> . . . . .	253
4.5	<b>OMP_PLACES</b> . . . . .	254
4.6	<b>OMP_NESTED</b> . . . . .	256
4.7	<b>OMP_STACKSIZE</b> . . . . .	256
4.8	<b>OMP_WAIT_POLICY</b> . . . . .	257
4.9	<b>OMP_MAX_ACTIVE_LEVELS</b> . . . . .	258
4.10	<b>OMP_THREAD_LIMIT</b> . . . . .	258
4.11	<b>OMP_CANCELLATION</b> . . . . .	259
4.12	<b>OMP_DISPLAY_ENV</b> . . . . .	259
4.13	<b>OMP_DEFAULT_DEVICE</b> . . . . .	260
<b>A</b>	<b>Stubs for Runtime Library Routines</b>	<b>261</b>
A.1	C/C++ Stub Routines . . . . .	262
A.2	Fortran Stub Routines . . . . .	269
<b>B</b>	<b>OpenMP C and C++ Grammar</b>	<b>276</b>
B.1	Notation . . . . .	276
B.2	Rules . . . . .	277
<b>C</b>	<b>Interface Declarations</b>	<b>296</b>
C.1	Example of the <b>omp.h</b> Header File . . . . .	297
C.2	Example of an Interface Declaration <b>include</b> File . . . . .	299
C.3	Example of a Fortran Interface Declaration <b>module</b> . . . . .	302
C.4	Example of a Generic Interface for a Library Routine . . . . .	307
<b>D</b>	<b>OpenMP Implementation-Defined Behaviors</b>	<b>308</b>
<b>E</b>	<b>Features History</b>	<b>312</b>
E.1	Version 4.0 to 4.1 Differences . . . . .	312
E.2	Version 3.1 to 4.0 Differences . . . . .	312
E.3	Version 3.0 to 3.1 Differences . . . . .	313
E.4	Version 2.5 to 3.0 Differences . . . . .	314
	<b>Index</b>	<b>317</b>

# 1 CHAPTER 1

## 2 Introduction

---

3 The collection of compiler directives, library routines, and environment variables described in this  
4 document collectively define the specification of the OpenMP Application Program Interface  
5 (OpenMP API) for shared-memory parallelism in C, C++ and Fortran programs.

6 This specification provides a model for parallel programming that is portable across shared  
7 memory architectures from different vendors. Compilers from numerous vendors support the  
8 OpenMP API. More information about the OpenMP API can be found at the following web site

9 **`http://www.openmp.org`**

10 The directives, library routines, and environment variables defined in this document allow users to  
11 create and manage parallel programs while permitting portability. The directives extend the C, C++  
12 and Fortran base languages with single program multiple data (SPMD) constructs, tasking  
13 constructs, device constructs, worksharing constructs, and synchronization constructs, and they  
14 provide support for sharing and privatizing data. The functionality to control the runtime  
15 environment is provided by library routines and environment variables. Compilers that support the  
16 OpenMP API often include a command line option to the compiler that activates and allows  
17 interpretation of all OpenMP directives.

### 18 1.1 Scope

19 The OpenMP API covers only user-directed parallelization, wherein the programmer explicitly  
20 specifies the actions to be taken by the compiler and runtime system in order to execute the program  
21 in parallel. OpenMP-compliant implementations are not required to check for data dependencies,  
22 data conflicts, race conditions, or deadlocks, any of which may occur in conforming programs. In  
23 addition, compliant implementations are not required to check for code sequences that cause a

1 program to be classified as non- conforming. Application developers are responsible for correctly  
2 using the OpenMP API to produce a conforming program. The OpenMP API does not cover  
3 compiler-generated automatic parallelization and directives to the compiler to assist such  
4 parallelization.

## 5 1.2 Glossary

### 6 1.2.1 Threading Concepts

7 **thread** An execution entity with a stack and associated static memory, called *threadprivate*  
8 *memory*.

9 **OpenMP thread** A *thread* that is managed by the OpenMP runtime system.

10 **thread-safe routine** A routine that performs the intended function even when executed concurrently (by  
11 more than one *thread*).

12 **processor** Implementation defined hardware unit on which one or more *OpenMP threads* can  
13 execute.

14 **device** An implementation defined logical execution engine.

15 COMMENT: A *device* could have one or more *processors*.

16 **host device** The *device* on which the *OpenMP program* begins execution

17 **target device** A device onto which code and data may be offloaded from the *host device*.

### 18 1.2.2 OpenMP Language Terminology

19 **base language** A programming language that serves as the foundation of the OpenMP specification.

20 COMMENT: See Section 1.6 on page 21 for a listing of current *base*  
21 *languages* for the OpenMP API.

22 **base program** A program written in a *base language*.

1	<b>structured block</b>	For C/C++, an executable statement, possibly compound, with a single entry at the
2		top and a single exit at the bottom, or an OpenMP <i>construct</i> .
3		For Fortran, a block of executable statements with a single entry at the top and a
4		single exit at the bottom, or an OpenMP <i>construct</i> .
5		COMMENTS:
6		For all <i>base languages</i> ,
7		• Access to the <i>structured block</i> must not be the result of a branch.
8		• The point of exit cannot be a branch out of the <i>structured block</i> .
9		For C/C++:
10		• The point of entry must not be a call to <b>setjmp()</b> .
11		• <b>longjmp()</b> and <b>throw()</b> must not violate the entry/exit criteria.
12		• Calls to <b>exit()</b> are allowed in a <i>structured block</i> .
13		• An expression statement, iteration statement, selection statement, or try
14		block is considered to be a <i>structured block</i> if the corresponding
15		compound statement obtained by enclosing it in { and } would be a
16		<i>structured block</i> .
17		For Fortran:
18		• <b>STOP</b> statements are allowed in a <i>structured block</i> .
19	<b>enclosing context</b>	In C/C++, the innermost scope enclosing an OpenMP <i>directive</i> .
20		In Fortran, the innermost scoping unit enclosing an OpenMP <i>directive</i> .
21	<b>directive</b>	In C/C++, a <b>#pragma</b> , and in Fortran, a comment, that specifies <i>OpenMP program</i>
22		behavior.
23		COMMENT: See Section 2.1 on page 26 for a description of OpenMP
24		<i>directive</i> syntax.
25	<b>white space</b>	A non-empty sequence of space and/or horizontal tab characters
26	<b>OpenMP program</b>	A program that consists of a <i>base program</i> , annotated with OpenMP <i>directives</i> and
27		runtime library routines.
28	<b>conforming program</b>	An <i>OpenMP program</i> that follows all the rules and restrictions of the OpenMP
29		specification.
30	<b>declarative directive</b>	An OpenMP <i>directive</i> that may only be placed in a declarative context. A <i>declarative</i>
31		<i>directive</i> results in one or more declarations only; it is not associated with the
32		immediate execution of any user code.

1	<b>executable directive</b>	An OpenMP <i>directive</i> that is not declarative. That is, it may be placed in an
2		executable context.
3	<b>stand-alone directive</b>	An OpenMP <i>executable directive</i> that has no associated executable user code.
4	<b>construct</b>	An OpenMP <i>executable directive</i> (and for Fortran, the paired <b>end directive</b> , if any)
5		and the associated statement, loop or <i>structured block</i> , if any, not including the code
6		in any called routines. That is, in the lexical extent of an <i>executable directive</i> .
7	<b>combined construct</b>	A construct that is a shortcut for specifying one construct immediately nested inside
8		another construct. A combined construct is semantically identical to that of explicitly
9		specifying the first construct containing one instance of the second construct and no
10		other statements.
11	<b>composite construct</b>	A construct that is composed of two constructs but does not have identical semantics
12		to specifying one of the constructs immediately nested inside the other. A composite
13		construct either adds semantics not included in the constructs from which it is
14		composed or the nesting of the one construct inside the other is not conforming.
15	<b>region</b>	All code encountered during a specific instance of the execution of a given <i>construct</i>
16		or of an OpenMP library routine. A <i>region</i> includes any code in called routines as
17		well as any implicit code introduced by the OpenMP implementation. The generation
18		of a <i>task</i> at the point where a <b>task directive</b> is encountered is a part of the <i>region</i> of
19		the <i>encountering thread</i> , but the <i>explicit task region</i> associated with the <b>task</b>
20		<i>directive</i> is not. The point where a <b>target</b> or <b>teams</b> directive is encountered is a
21		part of the <i>region</i> of the <i>encountering thread</i> , but the <i>region</i> associated with the
22		<b>target</b> or <b>teams</b> directive is not.
23		COMMENTS:
24		<i>A region may also be thought of as the dynamic or runtime extent of a</i>
25		<i>construct</i> or of an OpenMP library routine.
26		<i>During the execution of an OpenMP program, a construct may give rise to</i>
27		<i>many regions.</i>
28	<b>active parallel region</b>	A <b>parallel region</b> that is executed by a <i>team</i> consisting of more than one <i>thread</i> .
29	<b>inactive parallel region</b>	A <b>parallel region</b> that is executed by a <i>team</i> of only one <i>thread</i> .
30	<b>sequential part</b>	All code encountered during the execution of an <i>initial task region</i> that is not part of
31		a <b>parallel region</b> corresponding to a <b>parallel construct</b> or a <b>task region</b>
32		corresponding to a <b>task construct</b> .
33		COMMENTS:
34		<i>A sequential part is enclosed by an implicit parallel region.</i>

1		Executable statements in called routines may be in both a <i>sequential part</i>
2		and any number of explicit <b>parallel regions</b> at different points in the
3		program execution.
4	<b>master thread</b>	The <i>thread</i> that encounters a <b>parallel construct</b> , creates a <i>team</i> , generates a set of
5		<i>implicit tasks</i> , then executes one of those <i>tasks</i> as <i>thread</i> number 0.
6	<b>parent thread</b>	The <i>thread</i> that encountered the <b>parallel construct</b> and generated a <b>parallel</b>
7		<i>region</i> is the <i>parent thread</i> of each of the <i>threads</i> in the <i>team</i> of that <b>parallel</b>
8		<i>region</i> . The <i>master thread</i> of a <b>parallel region</b> is the same <i>thread</i> as its <i>parent</i>
9		<i>thread</i> with respect to any resources associated with an <i>OpenMP thread</i> .
10	<b>child thread</b>	When a <i>thread</i> encounters a <b>parallel</b> construct, each of the <i>threads</i> in the
11		generated <b>parallel region</b> 's <i>team</i> are <i>child threads</i> of the encountering <i>thread</i> .
12		The <b>target</b> or <b>teams</b> region's <i>initial thread</i> is not a <i>child thread</i> of the <i>thread</i> that
13		encountered the <b>target</b> or <b>teams</b> construct.
14	<b>ancestor thread</b>	For a given <i>thread</i> , its <i>parent thread</i> or one of its <i>parent thread</i> 's <i>ancestor threads</i> .
15	<b>descendent thread</b>	For a given <i>thread</i> , one of its <i>child threads</i> or one of its <i>child threads</i> ' <i>descendent</i>
16		<i>threads</i> .
17	<b>team</b>	A set of one or more <i>threads</i> participating in the execution of a <b>parallel region</b> .
18		COMMENTS:
19		For an <i>active parallel region</i> , the <i>team</i> comprises the <i>master thread</i> and at
20		least one additional <i>thread</i> .
21		For an <i>inactive parallel region</i> , the <i>team</i> comprises only the <i>master thread</i> .
22	<b>league</b>	The set of <i>thread teams</i> created by a <b>target</b> construct or a <b>teams</b> construct.
23	<b>contention group</b>	An <i>initial thread</i> and its <i>descendent threads</i> .
24	<b>implicit parallel region</b>	An <i>inactive parallel region</i> that generates an <i>initial task region</i> . <i>Implicit parallel</i>
25		<i>regions</i> surround the whole OpenMP program, all <b>target</b> regions, and all <b>teams</b>
26		regions
27	<b>initial thread</b>	A <i>thread</i> that executes an <i>implicit parallel region</i> .
28	<b>nested construct</b>	A <i>construct</i> (lexically) enclosed by another <i>construct</i> .
29	<b>closely nested construct</b>	A <i>construct</i> nested inside another <i>construct</i> with no other <i>construct</i> nested between
30		them.
31	<b>nested region</b>	A <i>region</i> (dynamically) enclosed by another <i>region</i> . That is, a <i>region</i> encountered
32		during the execution of another <i>region</i> .
33		COMMENT: Some nestings are <i>conforming</i> and some are not. See
34		Section 2.16 on page 205 for the restrictions on nesting.

1	<b>closely nested region</b>	A <i>region nested</i> inside another <i>region</i> with no <b>parallel</b> <i>region nested</i> between
2		them.
3	<b>all threads</b>	All OpenMP <i>threads</i> participating in the <i>OpenMP program</i> .
4	<b>current team</b>	All <i>threads</i> in the <i>team</i> executing the innermost enclosing <b>parallel</b> <i>region</i> .
5	<b>encountering thread</b>	For a given <i>region</i> , the <i>thread</i> that encounters the corresponding <i>construct</i> .
6	<b>all tasks</b>	All <i>tasks</i> participating in the <i>OpenMP program</i> .
7	<b>current team tasks</b>	All <i>tasks</i> encountered by the corresponding <i>team</i> . Note that the <i>implicit tasks</i>
8		constituting the <b>parallel</b> <i>region</i> and any <i>descendent tasks</i> encountered during the
9		execution of these <i>implicit tasks</i> are included in this set of tasks.
10	<b>generating task</b>	For a given <i>region</i> , the task whose execution by a <i>thread</i> generated the <i>region</i> .
11	<b>binding thread set</b>	The set of <i>threads</i> that are affected by, or provide the context for, the execution of a
12		<i>region</i> .
13		The <i>binding thread</i> set for a given <i>region</i> can be <i>all threads</i> on a <i>device</i> , <i>all threads</i>
14		in a <i>contention group</i> , the <i>current team</i> , or the <i>encountering thread</i> .
15		COMMENT: The <i>binding thread</i> set for a particular <i>region</i> is described in
16		its corresponding subsection of this specification.
17	<b>binding task set</b>	The set of <i>tasks</i> that are affected by, or provide the context for, the execution of a
18		<i>region</i> .
19		The <i>binding task</i> set for a given <i>region</i> can be <i>all tasks</i> , the <i>current team tasks</i> , or the
20		<i>generating task</i> .
21		COMMENT: The <i>binding task</i> set for a particular <i>region</i> (if applicable) is
22		described in its corresponding subsection of this specification.



1	<b>binding region</b>	The enclosing <i>region</i> that determines the execution context and limits the scope of
2		the effects of the bound <i>region</i> is called the <i>binding region</i> .
3		<i>Binding region</i> is not defined for <i>regions</i> whose <i>binding thread</i> set is <i>all threads</i> or
4		the <i>encountering thread</i> , nor is it defined for <i>regions</i> whose <i>binding task set</i> is <i>all</i>
5		<i>tasks</i> .
6		COMMENTS:
7		The <i>binding region</i> for an <b>ordered</b> <i>region</i> is the innermost enclosing
8		<i>loop region</i> .
9		The <i>binding region</i> for a <b>taskwait</b> <i>region</i> is the innermost enclosing
10		<i>task region</i> .
11		For all other <i>regions</i> for which the <i>binding thread set</i> is the <i>current team</i>
12		or the <i>binding task set</i> is the <i>current team tasks</i> , the <i>binding region</i> is the
13		innermost enclosing <b>parallel</b> <i>region</i> .
14		For <i>regions</i> for which the <i>binding task set</i> is the <i>generating task</i> , the
15		<i>binding region</i> is the <i>region</i> of the <i>generating task</i> .
16		A <b>parallel</b> <i>region</i> need not be <i>active</i> nor explicit to be a <i>binding</i>
17		<i>region</i> .
18		A <i>task region</i> need not be explicit to be a <i>binding region</i> .
19		A <i>region</i> never binds to any <i>region</i> outside of the innermost enclosing
20		<b>parallel</b> <i>region</i> .
21	<b>orphaned construct</b>	A <i>construct</i> that gives rise to a <i>region</i> whose <i>binding thread set</i> is the <i>current team</i> ,
22		but is not nested within another <i>construct</i> giving rise to the <i>binding region</i> .
23	<b>worksharing construct</b>	A <i>construct</i> that defines units of work, each of which is executed exactly once by one
24		of the <i>threads</i> in the <i>team</i> executing the <i>construct</i> .
25		For C/C++, <i>worksharing constructs</i> are <b>for</b> , <b>sections</b> , and <b>single</b> .
26		For Fortran, <i>worksharing constructs</i> are <b>do</b> , <b>sections</b> , <b>single</b> and
27		<b>workshare</b> .
28	<b>place</b>	Unordered set of <i>processors</i> that is treated by the execution environment as a location
29		unit when dealing with OpenMP thread affinity.
30	<b>place list</b>	The ordered list that describes all OpenMP <i>places</i> available to the execution
31		environment.
32	<b>place partition</b>	An ordered list that corresponds to a contiguous interval in the OpenMP <i>place list</i> . It
33		describes the <i>places</i> currently available to the execution environment for a given
34		parallel region.

1	<b>SIMD instruction</b>	A single machine instruction that can operate on multiple data elements.
2	<b>SIMD lane</b>	A software or hardware mechanism capable of processing one data element from a
3		<i>SIMD instruction</i> .
4	<b>SIMD chunk</b>	A set of iterations executed concurrently, each by a <i>SIMD lane</i> , by a single <i>thread</i> by
5		means of <i>SIMD instructions</i> .

### 6 1.2.3 Loop Terminology

7	<b>loop directive</b>	An OpenMP <i>executable</i> directive whose associated user code must be a loop nest
8		that is a <i>structured block</i> .
9	<b>associated loop(s)</b>	The loop(s) controlled by a <i>loop directive</i> .
10		COMMENT: If the <i>loop directive</i> contains a <b>collapse</b> or an
11		<b>ordered</b> ( <i>n</i> ) clause then there may be more than one <i>associated loop</i> .
12	<b>sequential loop</b>	A loop that is not associated with any OpenMP <i>loop directive</i> .
13	<b>SIMD loop</b>	A loop that includes at least one <i>SIMD chunk</i> .
14	<b>doacross loop nest</b>	A loop nest that has cross-iteration dependence. An iteration is dependent on one or
15		more lexicographically earlier iterations.
16		COMMENT: The loop directive with an <b>ordered</b> clause with the
17		parameter identifies the loop(s) associated for the <i>doacross loop nest</i> .

### 18 1.2.4 Synchronization Terminology

19	<b>barrier</b>	A point in the execution of a program encountered by a <i>team of threads</i> , beyond
20		which no <i>thread</i> in the team may execute until all <i>threads</i> in the <i>team</i> have reached
21		the barrier and all <i>explicit tasks</i> generated by the <i>team</i> have executed to completion.
22		If <i>cancellation</i> has been requested, threads may proceed to the end of the canceled
23		<i>region</i> even if some threads in the team have not reached the <i>barrier</i> .
24	<b>cancellation</b>	An action that cancels (that is, aborts) an OpenMP <i>region</i> and causes executing
25		<i>implicit</i> or <i>explicit</i> tasks to proceed to the end of the canceled <i>region</i> .
26	<b>cancellation point</b>	A point at which implicit and explicit tasks check if cancellation has been requested.
27		If cancellation has been observed, they perform the <i>cancellation</i> .
28		COMMENT: For a list of cancellation points, see Section 2.13.1 on
29		page <a href="#">156</a>

## 1 1.2.5 Tasking Terminology

2	<b>task</b>	A specific instance of executable code and its <i>data environment</i> , generated when a <i>thread</i> encounters a <b>task construct</b> or a <b>parallel construct</b> .
3		
4	<b>task region</b>	A <i>region</i> consisting of all code encountered during the execution of a <i>task</i> .
5		COMMENT: A <b>parallel region</b> consists of one or more implicit <i>task regions</i> .
6		
7	<b>explicit task</b>	A <i>task</i> generated when a <b>task construct</b> is encountered during execution.
8	<b>implicit task</b>	A <i>task</i> generated by an <i>implicit parallel region</i> or generated when a <b>parallel construct</b> is encountered during execution.
9		
10	<b>initial task</b>	An <i>implicit task</i> associated with an <i>implicit parallel region</i> .
11	<b>current task</b>	For a given <i>thread</i> , the <i>task</i> corresponding to the <i>task region</i> in which it is executing.
12	<b>child task</b>	A <i>task</i> is a <i>child task</i> of its generating <i>task region</i> . A <i>child task region</i> is not part of its generating <i>task region</i> .
13		
14	<b>sibling tasks</b>	<i>Tasks</i> that are <i>child tasks</i> of the same <i>task region</i> .
15	<b>descendent task</b>	A <i>task</i> that is the <i>child task</i> of a <i>task region</i> or of one of its <i>descendent task regions</i> .
16	<b>task completion</b>	<i>Task completion</i> occurs when the end of the <i>structured block</i> associated with the <i>construct</i> that generated the <i>task</i> is reached.
17		
18		COMMENT: Completion of the <i>initial task</i> occurs at program exit.
19	<b>task scheduling point</b>	A point during the execution of the current <i>task region</i> at which it can be suspended to be resumed later; or the point of <i>task completion</i> , after which the executing <i>thread</i> may switch to a different <i>task region</i> .
20		
21		
22		COMMENT: For a list of task scheduling points, see Section <a href="#">2.9.5</a> on
23		page <a href="#">90</a> .
24	<b>task switching</b>	The act of a <i>thread</i> switching from the execution of one <i>task</i> to another <i>task</i> .
25	<b>tied task</b>	A <i>task</i> that, when its <i>task region</i> is suspended, can be resumed only by the same <i>thread</i> that suspended it. That is, the <i>task</i> is tied to that <i>thread</i> .
26		
27	<b>untied task</b>	A <i>task</i> that, when its <i>task region</i> is suspended, can be resumed by any <i>thread</i> in the team. That is, the <i>task</i> is not tied to any <i>thread</i> .
28		
29	<b>undeferrred task</b>	A <i>task</i> for which execution is not deferred with respect to its generating <i>task region</i> . That is, its generating <i>task region</i> is suspended until execution of the <i>undeferrred task</i> is completed.
30		
31		

1	<b>included task</b>	A <i>task</i> for which execution is sequentially included in the generating <i>task region</i> . That is, an <i>included task</i> is <i>undeferred</i> and executed immediately by the <i>encountering thread</i> .
2		
3		
4	<b>merged task</b>	A <i>task</i> whose <i>data environment</i> , inclusive of ICVs, is the same as that of its generating <i>task region</i> .
5		
6	<b>final task</b>	A <i>task</i> that forces all of its <i>child tasks</i> to become <i>final</i> and <i>included tasks</i> .
7	<b>task dependence</b>	An ordering relation between two <i>sibling tasks</i> : the <i>dependent task</i> and a previously generated <i>predecessor task</i> . The <i>task dependence</i> is fulfilled when the <i>predecessor task</i> has completed.
8		
9		
10	<b>dependent task</b>	A <i>task</i> that because of a <i>task dependence</i> cannot be executed until its <i>predecessor tasks</i> have completed.
11		
12	<b>predecessor task</b>	A <i>task</i> that must complete before its <i>dependent tasks</i> can be executed.
13	<b>task synchronization construct</b>	A <b>taskwait</b> , <b>taskgroup</b> , or a <b>barrier</b> <i>construct</i> .
14	<b>target task</b>	A <i>merged task</i> that is executed immediately.

## 15 1.2.6 Data Terminology

16 **variable** A named data storage block, whose value can be defined and redefined during the  
17 execution of a program.

18 Note – An array or structure element is a variable that is part of another variable.

19 **array section** A designated subset of the elements of an array.

20 **array item** An array, an array section or an array element.

21 **private variable** With respect to a given set of *task regions* or *SIMD lanes* that bind to the same  
22 **parallel region**, a *variable* whose name provides access to a different block of  
23 storage for each *task region* or *SIMD lane*.

24 A *variable* that is part of another variable (as an array or structure element) cannot be  
25 made private independently of other components.

1	<b>shared variable</b>	With respect to a given set of <i>task regions</i> that bind to the same <b>parallel region</b> , a <i>variable</i> whose name provides access to the same block of storage for each <i>task region</i> .
2		
3		
4		A <i>variable</i> that is part of another variable (as an array or structure element) cannot be
5		<i>shared</i> independently of the other components, except for static data members of
6		C++ classes.
7	<b>threadprivate variable</b>	A <i>variable</i> that is replicated, one instance per <i>thread</i> , by the OpenMP
8		implementation. Its name then provides access to a different block of storage for each
9		<i>thread</i> .
10		A <i>variable</i> that is part of another variable (as an array or structure element) cannot be
11		made <i>threadprivate</i> independently of the other components, except for static data
12		members of C++ classes.
13	<b>threadprivate memory</b>	The set of <i>threadprivate variables</i> associated with each <i>thread</i> .
14	<b>data environment</b>	The <i>variables</i> associated with the execution of a given <i>region</i> .
15	<b>device data environment</b>	The initial <i>data environment</i> associated with a device.
16	<b>mapped variable</b>	An original <i>variable</i> in a <i>data environment</i> with a corresponding <i>variable</i> in a device
17		<i>data environment</i> .
18		COMMENT: The original and corresponding <i>variables</i> may share storage.
19	<b>mappable type</b>	A type that is valid for a <i>mapped variable</i> . If a type is composed from other types
20		(such as the type of an array or structure element) and any of the other types are not
21		mappable then the type is not mappable.
22		COMMENT: Pointer types are <i>mappable</i> but the memory block to which
23		the pointer refers is not <i>mapped</i> .
24		For C: The type must be a complete type.
25		For C++: The type must be a complete type.
26		In addition, for class types:
27		• All member functions accessed in any <b>target</b> region must appear in a
28		<b>declare target</b> directive.
29		• All data members must be non-static.
30		• A <i>mappable type</i> cannot contain virtual members.
31		For Fortran: The type must be definable.
32		In addition, for derived types:

- 1                   • All type-bound procedures accessed in any target region must appear in a declare
- 2                   target directive.
- 3           **defined** For *variables*, the property of having a valid value.
- 4                   For C: For the contents of *variables*, the property of having a valid value.
- 5                   For C++: For the contents of *variables* of POD (plain old data) type, the property of
- 6                   having a valid value.
- 7                   For *variables* of non-POD class type, the property of having been constructed but not
- 8                   subsequently destructed.
- 9                   For Fortran: For the contents of *variables*, the property of having a valid value. For
- 10                  the allocation or association status of *variables*, the property of having a valid status.
- 11                                COMMENT: Programs that rely upon *variables* that are not *defined* are
- 12                                *non-conforming programs*.
- 13           **class type** For C++: *Variables* declared with one of the **class**, **struct**, or **union** keywords
- 14 **sequentially consistent** An **atomic** construct for which the **seq\_cst** clause is specified.
- 15           **atomic construct**
- non-sequentially** An **atomic** construct for which the **seq\_cst** clause is not specified
- consistent atomic**
- construct**

## 16 1.2.7 Implementation Terminology

- 17 **supporting *n* levels of** Implies allowing an *active parallel region* to be enclosed by *n-1 active parallel*
- 18 **parallelism** *regions*.
- 19           **supporting the** Supporting at least one level of parallelism.
- OpenMP API**
- 20 **supporting nested** Supporting more than one level of parallelism.
- parallelism**
- 21 **internal control** A conceptual variable that specifies runtime behavior of a set of *threads* or *tasks* in
- 22 **variable** an *OpenMP program*.
- 23                                COMMENT: The acronym ICV is used interchangeably with the term
- 24                                *internal control variable* in the remainder of this specification.
- 25           **compliant** An implementation of the OpenMP specification that compiles and executes any
- 26 **implementation** *conforming program* as defined by the specification.
- 27                                COMMENT: A *compliant implementation* may exhibit *unspecified*
- 28                                *behavior* when compiling or executing a *non-conforming program*.

1	<b>unspecified behavior</b>	A behavior or result that is not specified by the OpenMP specification or not known
2		prior to the compilation or execution of an <i>OpenMP program</i> .
3		Such <i>unspecified behavior</i> may result from:
4		• Issues documented by the OpenMP specification as having <i>unspecified behavior</i> .
5		• A <i>non-conforming program</i> .
6		• A <i>conforming program</i> exhibiting an <i>implementation defined</i> behavior.
7	<b>implementation defined</b>	Behavior that must be documented by the implementation, and is allowed to vary
8		among different <i>compliant implementations</i> . An implementation is allowed to define
9		this behavior as <i>unspecified</i> .
10		COMMENT: All features that have <i>implementation defined</i> behavior are
11		documented in Appendix D.
12	<b>deprecated</b>	Implies a construct, clause or other feature is normative in the current specification
13		but is considered obsolescent and will be removed in the future.

## 14 1.3 Execution Model

15 The OpenMP API uses the fork-join model of parallel execution. Multiple threads of execution  
 16 perform tasks defined implicitly or explicitly by OpenMP directives. The OpenMP API is intended  
 17 to support programs that will execute correctly both as parallel programs (multiple threads of  
 18 execution and a full OpenMP support library) and as sequential programs (directives ignored and a  
 19 simple OpenMP stubs library). However, it is possible and permitted to develop a program that  
 20 executes correctly as a parallel program but not as a sequential program, or that produces different  
 21 results when executed as a parallel program compared to when it is executed as a sequential  
 22 program. Furthermore, using different numbers of threads may result in different numeric results  
 23 because of changes in the association of numeric operations. For example, a serial addition  
 24 reduction may have a different pattern of addition associations than a parallel reduction. These  
 25 different associations may change the results of floating-point addition.

26 An OpenMP program begins as a single thread of execution, called an initial thread. An initial  
 27 thread executes sequentially, as if enclosed in an implicit task region, called an initial task region,  
 28 that is defined by the implicit parallel region surrounding the whole program.

29 The thread that executes the implicit parallel region that surrounds the whole program executes on  
 30 the *host device*. An implementation may support other *target devices*. If supported, one or more  
 31 devices are available to the host device for offloading code and data. Each device has its own  
 32 threads that are distinct from threads that execute on another device. Threads cannot migrate from

1 one device to another device. The execution model is host-centric such that the host device offloads  
2 **target** regions to target devices.

3 The initial thread that executes the implicit parallel region that surrounds the **target** region may  
4 execute on a *target device*. An initial thread executes sequentially, as if enclosed in an implicit task  
5 region, called an initial task region, that is defined by an implicit inactive **parallel** region that  
6 surrounds the entire **target** region.

7 When a **target** construct is encountered, the **target** region is executed by the implicit device  
8 task. The task that encounters the **target** construct waits at the end of the construct until  
9 execution of the region completes. If the target device does not exist or the implementation does not  
10 support the target device, all **target** regions associated with that device are executed by the host  
11 device.

12 The implementation must ensure that the **target** region executes as if it were executed in the data  
13 environment of the target device unless an **if** clause is present and the **if** clause expression  
14 evaluates to *false*.

15 The **teams** construct creates a *league of thread teams* where the master thread of each team  
16 executes the region. Each of these master threads is an initial thread, and executes sequentially, as if  
17 enclosed in an implicit task region that is defined by an implicit parallel region that surrounds the  
18 entire **teams** region.

19 If a construct creates a data environment, the data environment is created at the time the construct is  
20 encountered. Whether a construct creates a data environment is defined in the description of the  
21 construct.

22 When any thread encounters a **parallel** construct, the thread creates a team of itself and zero or  
23 more additional threads and becomes the master of the new team. A set of implicit tasks, one per  
24 thread, is generated. The code for each task is defined by the code inside the **parallel** construct.  
25 Each task is assigned to a different thread in the team and becomes tied; that is, it is always  
26 executed by the thread to which it is initially assigned. The task region of the task being executed  
27 by the encountering thread is suspended, and each member of the new team executes its implicit  
28 task. There is an implicit barrier at the end of the **parallel** construct. Only the master thread  
29 resumes execution beyond the end of the **parallel** construct, resuming the task region that was  
30 suspended upon encountering the **parallel** construct. Any number of **parallel** constructs  
31 can be specified in a single program.

32 **parallel** regions may be arbitrarily nested inside each other. If nested parallelism is disabled, or  
33 is not supported by the OpenMP implementation, then the new team that is created by a thread  
34 encountering a **parallel** construct inside a **parallel** region will consist only of the  
35 encountering thread. However, if nested parallelism is supported and enabled, then the new team  
36 can consist of more than one thread. A **parallel** construct may include a **proc\_bind** clause to  
37 specify the places to use for the threads in the team within the **parallel** region.

38 When any team encounters a worksharing construct, the work inside the construct is divided among  
39 the members of the team, and executed cooperatively instead of being executed by every thread.



1           There is a default barrier at the end of each worksharing construct unless the **nowait** clause is  
2 present. Redundant execution of code by every thread in the team resumes after the end of the  
3 worksharing construct.

4           When any thread encounters a **task** construct, a new explicit task is generated. Execution of  
5 explicitly generated tasks is assigned to one of the threads in the current team, subject to the  
6 thread's availability to execute work. Thus, execution of the new task could be immediate, or  
7 deferred until later according to task scheduling constraints and thread availability. Threads are  
8 allowed to suspend the current task region at a task scheduling point in order to execute a different  
9 task. If the suspended task region is for a tied task, the initially assigned thread later resumes  
10 execution of the suspended task region. If the suspended task region is for an untied task, then any  
11 thread may resume its execution. Completion of all explicit tasks bound to a given parallel region is  
12 guaranteed before the master thread leaves the implicit barrier at the end of the region. Completion  
13 of a subset of all explicit tasks bound to a given parallel region may be specified through the use of  
14 task synchronization constructs. Completion of all explicit tasks bound to the implicit parallel  
15 region is guaranteed by the time the program exits.

16           When any thread encounters a **simd** construct, the iterations of the loop associated with the  
17 construct may be executed concurrently using the SIMD lanes that are available to the thread.

18           The **cancel** construct can alter the previously described flow of execution in an OpenMP region.  
19 The effect of the **cancel** construct depends on its *construct-type-clause*. If a task encounters a  
20 **cancel** construct with a **taskgroup** *construct-type-clause*, then the task activates cancellation  
21 and continues execution at the end of its **task** region, which implies completion of that task. Any  
22 other task in that **taskgroup** that has begun executing completes execution unless it encounters a  
23 **cancellation point** construct, in which case it continues execution at the end of its **task**  
24 region, which implies its completion. Other tasks in that **taskgroup** region that have not begun  
25 execution are aborted, which implies their completion.

26           For all other *construct-type-clause* values, if a thread encounters a **cancel** construct, it activates  
27 cancellation of the innermost enclosing region of the type specified and the thread continues  
28 execution at the end of that region. Threads check if cancellation has been activated for their region  
29 at cancellation points and, if so, also resume execution at the end of the canceled region.

30           If cancellation has been activated regardless of *construct-type-clause*, threads that are waiting  
31 inside a barrier other than an implicit barrier at the end of the canceled region exit the barrier and  
32 resume execution at the end of the canceled region. This action can occur before the other threads  
33 reach that barrier.

34           Synchronization constructs and library routines are available in the OpenMP API to coordinate  
35 tasks and data access in **parallel** regions. In addition, library routines and environment  
36 variables are available to control or to query the runtime environment of OpenMP programs.

37           The OpenMP specification makes no guarantee that input or output to the same file is synchronous  
38 when executed in parallel. In this case, the programmer is responsible for synchronizing input and  
39 output statements (or routines) using the provided synchronization constructs or library routines.

1 For the case where each thread accesses a different file, no synchronization by the programmer is  
2 necessary.

## 1 1.4 Memory Model

### 2 1.4.1 Structure of the OpenMP Memory Model

3 The OpenMP API provides a relaxed-consistency, shared-memory model. All OpenMP threads  
4 have access to a place to store and to retrieve variables, called the *memory*. In addition, each thread  
5 is allowed to have its own *temporary view* of the memory. The temporary view of memory for each  
6 thread is not a required part of the OpenMP memory model, but can represent any kind of  
7 intervening structure, such as machine registers, cache, or other local storage, between the thread  
8 and the memory. The temporary view of memory allows the thread to cache variables and thereby  
9 to avoid going to memory for every reference to a variable. Each thread also has access to another  
10 type of memory that must not be accessed by other threads, called *threadprivate memory*.

11 A directive that accepts data-sharing attribute clauses determines two kinds of access to variables  
12 used in the directive's associated structured block: shared and private. Each variable referenced in  
13 the structured block has an original variable, which is the variable by the same name that exists in  
14 the program immediately outside the construct. Each reference to a shared variable in the structured  
15 block becomes a reference to the original variable. For each private variable referenced in the  
16 structured block, a new version of the original variable (of the same type and size) is created in  
17 memory for each task or SIMD lane that contains code associated with the directive. Creation of  
18 the new version does not alter the value of the original variable. However, the impact of attempts to  
19 access the original variable during the region associated with the directive is unspecified; see  
20 Section 2.14.3.3 on page 176 for additional details. References to a private variable in the  
21 structured block refer to the private version of the original variable for the current task or SIMD  
22 lane. The relationship between the value of the original variable and the initial or final value of the  
23 private version depends on the exact clause that specifies it. Details of this issue, as well as other  
24 issues with privatization, are provided in Section 2.14 on page 162.

25 The minimum size at which a memory update may also read and write back adjacent variables that  
26 are part of another variable (as array or structure elements) is implementation defined but is no  
27 larger than required by the base language.

28 A single access to a variable may be implemented with multiple load or store instructions, and  
29 hence is not guaranteed to be atomic with respect to other accesses to the same variable. Accesses  
30 to variables smaller than the implementation defined minimum size or to C or C++ bit-fields may  
31 be implemented by reading, modifying, and rewriting a larger unit of memory, and may thus  
32 interfere with updates of variables or fields in the same unit of memory.

33 If multiple threads write without synchronization to the same memory unit, including cases due to  
34 atomicity considerations as described above, then a data race occurs. Similarly, if at least one  
35 thread reads from a memory unit and at least one thread writes without synchronization to that  
36 same memory unit, including cases due to atomicity considerations as described above, then a data  
37 race occurs. If a data race occurs then the result of the program is unspecified.

1 A private variable in a task region that eventually generates an inner nested **parallel** region is  
2 permitted to be made shared by implicit tasks in the inner **parallel** region. A private variable in  
3 a task region can be shared by an explicit **task** region generated during its execution. However, it  
4 is the programmer's responsibility to ensure through synchronization that the lifetime of the  
5 variable does not end before completion of the explicit **task** region sharing it. Any other access by  
6 one task to the private variables of another task results in unspecified behavior.

## 7 1.4.2 Device Data Environments

8 When an OpenMP program begins, an implicit **target data** region for each device surrounds  
9 the whole program. Each device has a device data environment that is defined by its implicit  
10 **target data** region. Any **declare target** directives and the directives that accept  
11 data-mapping attribute clauses determine how an original variable in a data environment is mapped  
12 to a corresponding variable in a device data environment.

13 When an original variable is mapped to a device data environment and the associated  
14 corresponding variable is not present in the device data environment, a new corresponding variable  
15 (of the same type and size as the original variable) is created in the device data environment. The  
16 initial value of the new corresponding variable is determined from the clauses and the data  
17 environment of the encountering thread.

18 The corresponding variable in the device data environment may share storage with the original  
19 variable. Writes to the corresponding variable may alter the value of the original variable. The  
20 impact of this on memory consistency is discussed in Section 1.4.4 on page 20. When a task  
21 executes in the context of a device data environment, references to the original variable refer to the  
22 corresponding variable in the device data environment.

23 The relationship between the value of the original variable and the initial or final value of the  
24 corresponding variable depends on the *map-type*. Details of this issue, as well as other issues with  
25 mapping a variable, are provided in Section 2.14.5 on page 195.

26 The original variable in a data environment and the corresponding variable(s) in one or more device  
27 data environments may share storage. Without intervening synchronization data races can occur.

## 28 1.4.3 The Flush Operation

29 The memory model has relaxed-consistency because a thread's temporary view of memory is not  
30 required to be consistent with memory at all times. A value written to a variable can remain in the  
31 thread's temporary view until it is forced to memory at a later time. Likewise, a read from a variable

1 may retrieve the value from the thread's temporary view, unless it is forced to read from memory.  
2 The OpenMP flush operation enforces consistency between the temporary view and memory.

3 The flush operation is applied to a set of variables called the *flush-set*. The flush operation restricts  
4 reordering of memory operations that an implementation might otherwise do. Implementations  
5 must not reorder the code for a memory operation for a given variable, or the code for a flush  
6 operation for the variable, with respect to a flush operation that refers to the same variable.

7 If a thread has performed a write to its temporary view of a shared variable since its last flush of  
8 that variable, then when it executes another flush of the variable, the flush does not complete until  
9 the value of the variable has been written to the variable in memory. If a thread performs multiple  
10 writes to the same variable between two flushes of that variable, the flush ensures that the value of  
11 the last write is written to the variable in memory. A flush of a variable executed by a thread also  
12 causes its temporary view of the variable to be discarded, so that if its next memory operation for  
13 that variable is a read, then the thread will read from memory when it may again capture the value  
14 in the temporary view. When a thread executes a flush, no later memory operation by that thread for  
15 a variable involved in that flush is allowed to start until the flush completes. The completion of a  
16 flush of a set of variables executed by a thread is defined as the point at which all writes to those  
17 variables performed by the thread before the flush are visible in memory to all other threads and  
18 that thread's temporary view of all variables involved is discarded.

19 The flush operation provides a guarantee of consistency between a thread's temporary view and  
20 memory. Therefore, the flush operation can be used to guarantee that a value written to a variable  
21 by one thread may be read by a second thread. To accomplish this, the programmer must ensure  
22 that the second thread has not written to the variable since its last flush of the variable, and that the  
23 following sequence of events happens in the specified order:

- 24 1. The value is written to the variable by the first thread.
- 25 2. The variable is flushed by the first thread.
- 26 3. The variable is flushed by the second thread.
- 27 4. The value is read from the variable by the second thread.

28 **Note** – OpenMP synchronization operations, described in Section 2.12 on page 135 and in  
29 Section 3.3 on page 239, are recommended for enforcing this order. Synchronization through  
30 variables is possible but is not recommended because the proper timing of flushes is difficult.

## 1 1.4.4 OpenMP Memory Consistency

2 The restrictions in Section 1.4.3 on page 18 on reordering with respect to flush operations  
3 guarantee the following:

- 4 • If the intersection of the flush-sets of two flushes performed by two different threads is  
5 non-empty, then the two flushes must be completed as if in some sequential order, seen by all  
6 threads.
- 7 • If two operations performed by the same thread either access, modify, or flush the same variable,  
8 then they must be completed as if in that thread's program order, as seen by all threads.
- 9 • If the intersection of the flush-sets of two flushes is empty, the threads can observe these flushes  
10 in any order.

11 The flush operation can be specified using the **flush** directive, and is also implied at various  
12 locations in an OpenMP program: see Section 2.12.7 on page 148 for details.

13 **Note** – Since flush operations by themselves cannot prevent data races, explicit flush operations are  
14 only useful in combination with non-sequentially consistent atomic directives.

15 OpenMP programs that:

- 16 • do not use non-sequentially consistent atomic directives,
- 17 • do not rely on the accuracy of a *false* result from **omp\_test\_lock** and  
18 **omp\_test\_nest\_lock**, and
- 19 • correctly avoid data races as required in Section 1.4.1 on page 17

20 behave as though operations on shared variables were simply interleaved in an order consistent with  
21 the order in which they are performed by each thread. The relaxed consistency model is invisible  
22 for such programs, and any explicit flush operations in such programs are redundant.

23 Implementations are allowed to relax the ordering imposed by implicit flush operations when the  
24 result is only visible to programs using non-sequentially consistent atomic directives.

## 1 1.5 OpenMP Compliance

2 An implementation of the OpenMP API is compliant if and only if it compiles and executes all  
3 conforming programs according to the syntax and semantics laid out in Chapters 1, 2, 3 and 4.  
4 Appendices A, B, C, D, E and F and sections designated as Notes (see Section 1.7 on page 23) are  
5 for information purposes only and are not part of the specification.

6 The OpenMP API defines constructs that operate in the context of the base language that is  
7 supported by an implementation. If the base language does not support a language construct that  
8 appears in this document, a compliant OpenMP implementation is not required to support it, with  
9 the exception that for Fortran, the implementation must allow case insensitivity for directive and  
10 API routines names, and must allow identifiers of more than six characters

11 All library, intrinsic and built-in routines provided by the base language must be thread-safe in a  
12 compliant implementation. In addition, the implementation of the base language must also be  
13 thread-safe. For example, **ALLOCATE** and **DEALLOCATE** statements must be thread-safe in  
14 Fortran. Unsynchronized concurrent use of such routines by different threads must produce correct  
15 results (although not necessarily the same as serial execution results, as in the case of random  
16 number generation routines).

17 Starting with Fortran 90, variables with explicit initialization have the **SAVE** attribute implicitly.  
18 This is not the case in Fortran 77. However, a compliant OpenMP Fortran implementation must  
19 give such a variable the **SAVE** attribute, regardless of the underlying base language version.

20 Appendix D lists certain aspects of the OpenMP API that are implementation defined. A compliant  
21 implementation is required to define and document its behavior for each of the items in Appendix D.

## 22 1.6 Normative References

- 23 • ISO/IEC 9899:1990, *Information Technology - Programming Languages - C*.  
24 This OpenMP API specification refers to ISO/IEC 9899:1990 as C90.
- 25 • ISO/IEC 9899:1999, *Information Technology - Programming Languages - C*.  
26 This OpenMP API specification refers to ISO/IEC 9899:1999 as C99.
- 27 • ISO/IEC 14882:1998, *Information Technology - Programming Languages - C++*.  
28 This OpenMP API specification refers to ISO/IEC 14882:1998 as C++.
- 29 • ISO/IEC 1539:1980, *Information Technology - Programming Languages - Fortran*.  
30 This OpenMP API specification refers to ISO/IEC 1539:1980 as Fortran 77.

- 1           ● ISO/IEC 1539:1991, *Information Technology - Programming Languages - Fortran*.
- 2           This OpenMP API specification refers to ISO/IEC 1539:1991 as Fortran 90.
- 3           ● ISO/IEC 1539-1:1997, *Information Technology - Programming Languages - Fortran*.
- 4           This OpenMP API specification refers to ISO/IEC 1539-1:1997 as Fortran 95.
- 5           ● ISO/IEC 1539-1:2004, *Information Technology - Programming Languages - Fortran*.
- 6           This OpenMP API specification refers to ISO/IEC 1539-1:2004 as Fortran 2003. The following
- 7           features are not supported:
- 8           – IEEE Arithmetic issues covered in Fortran 2003 Section 14
- 9           – Parameterized derived types
- 10          – The **PASS** attribute
- 11          – Procedures bound to a type as operators
- 12          – Overriding a type-bound procedure
- 13          – Polymorphic entities
- 14          – **SELECT TYPE** construct
- 15          – Deferred bindings and abstract types
- 16          – Controlling IEEE underflow
- 17          – Another IEEE class value

18          Where this OpenMP API specification refers to C, C++ or Fortran, reference is made to the base  
19          language supported by the implementation.



# 1.7 Organization of this document

The remainder of this document is structured as follows:

- Chapter 2 “Directives”
- Chapter 3 “Runtime Library Routines”
- Chapter 4 “Environment Variables”
- Appendix A “Stubs for Runtime Library Routines”
- Appendix B “OpenMP C and C++ Grammar”
- Appendix C “Interface Declarations”
- Appendix D “OpenMP Implementation-Defined Behaviors”
- Appendix E “Features History”

Some sections of this document only apply to programs written in a certain base language. Text that applies only to programs whose base language is C or C++ is shown as follows:

▼ C / C++ ▼

C/C++ specific text...

▲ C / C++ ▲

Text that applies only to programs whose base language is C only is shown as follows:

▼ C ▼

C specific text...

▲ C ▲

Text that applies only to programs whose base language is C90 only is shown as follows:

▼ C90 ▼

C90 specific text...

▲ C90 ▲

Text that applies only to programs whose base language is C99 only is shown as follows:

▼ C99 ▼

C99 specific text...

▲ C99 ▲

Text that applies only to programs whose base language is C++ only is shown as follows:

1 C++ specific text... C++

2 Text that applies only to programs whose base language is Fortran is shown as follows:  
3 Fortran specific text..... Fortran

4 Where an entire page consists of, for example, Fortran specific text, a marker is shown at the top of  
5 the page like this:

Fortran (cont.)

6 Some text is for information only, and is not part of the normative specification. Such text is  
7 designated as a note, like this:

8 Note – Non-normative text....

2 **Directives**

---

3 This chapter describes the syntax and behavior of OpenMP directives, and is divided into the  
4 following sections:

- 5 • The language-specific directive format (Section 2.1 on page 26)
- 6 • Mechanisms to control conditional compilation (Section 2.2 on page 32)
- 7 • How to specify and to use array sections for all base languages (Section 2.4 on page 43)
- 8 • Control of OpenMP API ICVs (Section 2.3 on page 35)
- 9 • Details of each OpenMP directive (Section 2.5 on page 44 to Section 2.16 on page 205)

▼ C / C++ ▼

10 In C/C++, OpenMP directives are specified by using the **#pragma** mechanism provided by the C  
11 and C++ standards.

▲ C / C++ ▲  
▼ Fortran ▼

12 In Fortran, OpenMP directives are specified by using special comments that are identified by  
13 unique sentinels. Also, a special comment form is available for conditional compilation.

▲ Fortran ▲

14 Compilers can therefore ignore OpenMP directives and conditionally compiled code if support of  
15 the OpenMP API is not provided or enabled. A compliant implementation must provide an option  
16 or interface that ensures that underlying support of all OpenMP directives and OpenMP conditional  
17 compilation mechanisms is enabled. In the remainder of this document, the phrase *OpenMP*  
18 *compilation* is used to mean a compilation with these OpenMP features enabled.

## 1 Restrictions

2 The following restriction applies to all OpenMP directives:

- 3 • OpenMP directives may not appear in **PURE** or **ELEMENTAL** procedures.

## 4 2.1 Directive Format

5 OpenMP directives for C/C++ are specified with the **pragma omp** preprocessing directive. The syntax  
6 of an OpenMP directive is formally specified by the grammar in Appendix B, and informally as  
7 follows:

```
#pragma omp directive-name [clause[ [, ] clause] ... ] new-line
```

8 Each directive starts with **pragma omp**. The remainder of the directive follows the conventions  
9 of the C and C++ standards for compiler directives. In particular, white space can be used before  
10 and after the #, and sometimes white space must be used to separate the words in a directive.  
11 Preprocessing tokens following the **pragma omp** are subject to macro replacement.

12 Some OpenMP directives may be composed of consecutive **pragma** preprocessing directives if  
13 specified in their syntax.

14 Directives are case-sensitive.

15 An OpenMP executable directive applies to at most one succeeding statement, which must be a  
16 structured block.

## Fortran

OpenMP directives for Fortran are specified as follows:

```
sentinel directive-name [clause[ [, ] clause]...]
```

All OpenMP compiler directives must begin with a directive *sentinel*. The format of a sentinel differs between fixed and free-form source files, as described in Section 2.1.1 on page 28 and Section 2.1.2 on page 29.

Directives are case insensitive. Directives cannot be embedded within continued statements, and statements cannot be embedded within directives.

In order to simplify the presentation, free form is used for the syntax of OpenMP directives for Fortran in the remainder of this document, except as noted.

## Fortran

Only one *directive-name* can be specified per directive (note that this includes combined directives, see Section 2.11 on page 116). The order in which clauses appear on directives is not significant. Clauses on directives may be repeated as needed, subject to the restrictions listed in the description of each clause.

Some data-sharing attribute clauses (Section 2.14.3 on page 172), data copying clauses (Section 2.14.4 on page 191), the **threadprivate** directive (Section 2.14.2 on page 166) and the **flush** directive (Section 2.12.7 on page 148) accept a *list*. A *list* consists of a comma-separated collection of one or more *list items*.

## C / C++

A *list item* is a variable or array section, subject to the restrictions specified in Section 2.4 on page 43 and in each of the sections describing clauses and directives for which a *list* appears.

## C / C++

## Fortran

A *list item* is a variable, array section or common block name (enclosed in slashes), subject to the restrictions specified in Section 2.4 on page 43 and in each of the sections describing clauses and directives for which a *list* appears.

## Fortran

## 1 2.1.1 Fixed Source Form Directives

2 The following sentinels are recognized in fixed form source files:

```
!$omp | c$omp | *$omp
```

3 Sentinels must start in column 1 and appear as a single word with no intervening characters.

4 Fortran fixed form line length, white space, continuation, and column rules apply to the directive  
5 line. Initial directive lines must have a space or zero in column 6, and continuation directive lines  
6 must have a character other than a space or a zero in column 6.

7 Comments may appear on the same line as a directive. The exclamation point initiates a comment  
8 when it appears after column 6. The comment extends to the end of the source line and is ignored.  
9 If the first non-blank character after the directive sentinel of an initial or continuation directive line  
10 is an exclamation point, the line is ignored.

11 Note – in the following example, the three formats for specifying the directive are equivalent (the  
12 first line represents the position of the first 9 columns):

```
13 c23456789
14 !$omp parallel do shared(a,b,c)
15
16 c$omp parallel do
17 c$omp+shared(a,b,c)
18
19 c$omp paralleldoshared(a,b,c)
```

## 1 2.1.2 Free Source Form Directives

2 The following sentinel is recognized in free form source files:

**!\$omp**

3 The sentinel can appear in any column as long as it is preceded only by white space (spaces and tab  
 4 characters). It must appear as a single word with no intervening character. Fortran free form line  
 5 length, white space, and continuation rules apply to the directive line. Initial directive lines must  
 6 have a space after the sentinel. Continued directive lines must have an ampersand (&) as the last  
 7 non-blank character on the line, prior to any comment placed inside the directive. Continuation  
 8 directive lines can have an ampersand after the directive sentinel with optional white space before  
 9 and after the ampersand.

10 Comments may appear on the same line as a directive. The exclamation point (!) initiates a  
 11 comment. The comment extends to the end of the source line and is ignored. If the first non-blank  
 12 character after the directive sentinel is an exclamation point, the line is ignored.

13 One or more blanks or horizontal tabs must be used to separate adjacent keywords in directives in  
 14 free source form, except in the following cases, where white space is optional between the given set  
 15 of keywords:

```

16     declare reduction
17     declare simd
18     declare target
19     distribute parallel do
20     distribute parallel do simd
21     distribute simd
22     do simd
23     end atomic
24     end critical
25     end distribute
26     end distribute parallel do
27     end distribute parallel do simd
28     end distribute simd
  
```

```

1      end do
2      end do simd
3      end master
4      end ordered
5      end parallel
6      end parallel do
7      end parallel do simd
8      end parallel sections
9      end parallel workshare
10     end sections
11     end simd
12     end single
13     end target
14     end target data
15     end target teams
16     end target teams distribute
17     end target teams distribute parallel do
18     end target teams distribute parallel do simd
19     end target teams distribute simd
20     end task
21     end task group
22     end taskloop
23     end teams
24     end teams distribute
25     end teams distribute parallel do
26     end teams distribute parallel do simd
27     end teams distribute simd

```



```

1      end workshare
2      parallel do
3      parallel do simd
4      parallel sections
5      parallel workshare
6      target data
7      target teams
8      target teams distribute
9      target teams distribute parallel do
10     target teams distribute parallel do simd
11     target teams distribute simd
12     target update
13     taskloop
14     teams distribute
15     teams distribute parallel do
16     teams distribute parallel do simd
17     teams distribute simd

```

▼

18 Note – in the following example the three formats for specifying the directive are equivalent (the  
19 first line represents the position of the first 9 columns):

```

20 !23456789
21     !$omp parallel do
22         !$omp shared(a,b,c)
23
24     !$omp parallel
25     !$omp&do shared(a,b,c)
26
27     !$omp paralleldo shared(a,b,c)

```

▲

▲

Fortran

## 1 2.1.3 Stand-Alone Directives

### 2 Summary

3 Stand-alone directives are executable directives that have no associated user code.

### 4 Description

5 Stand-alone directives do not have any associated executable user code. Instead, they represent  
6 executable statements that typically do not have succinct equivalent statements in the base  
7 languages. There are some restrictions on the placement of a stand-alone directive within a  
8 program. A stand-alone directive may be placed only at a point where a base language executable  
9 statement is allowed.

### 10 Restrictions

▼ C / C++ ▼

11 For C/C++, a stand-alone directive may not be used in place of the statement following an **if**,  
12 **while**, **do**, **switch**, or **label**. See Appendix B for the formal grammar.

▲ C / C++ ▲

▼ Fortran ▼

13 For Fortran, a stand-alone directive may not be used as the action statement in an **if** statement or  
14 as the executable statement following a label if the label is referenced in the program.

▲ Fortran ▲

## 15 2.2 Conditional Compilation

16 In implementations that support a preprocessor, the **\_OPENMP** macro name is defined to have the  
17 decimal value *yyyymm* where *yyyy* and *mm* are the year and month designations of the version of  
18 the OpenMP API that the implementation supports.

19 If this macro is the subject of a **#define** or a **#undef** preprocessing directive, the behavior is  
20 unspecified.

▼ Fortran ▼

21 The OpenMP API requires Fortran lines to be compiled conditionally, as described in the following  
22 sections.

## 1 2.2.1 Fixed Source Form Conditional Compilation Sentinels

2  
3 The following conditional compilation sentinels are recognized in fixed form source files:

**!\$ | \*\$ | c\$**

4 To enable conditional compilation, a line with a conditional compilation sentinel must satisfy the  
5 following criteria:

- 6
- 7 • The sentinel must start in column 1 and appear as a single word with no intervening white space.
  - 8 • After the sentinel is replaced with two spaces, initial lines must have a space or zero in column 6 and only white space and numbers in columns 1 through 5.
  - 9 • After the sentinel is replaced with two spaces, continuation lines must have a character other than  
10 a space or zero in column 6 and only white space in columns 1 through 5.

11 If these criteria are met, the sentinel is replaced by two spaces. If these criteria are not met, the line  
12 is left unchanged.

13 **Note** – in the following example, the two forms for specifying conditional compilation in fixed  
14 source form are equivalent (the first line represents the position of the first 9 columns):

```
15 c23456789
16 !$ 10 iam = omp_get_thread_num() +
17 !$   &           index
18
19 #ifdef _OPENMP
20     10 iam = omp_get_thread_num() +
21     &           index
22 #endif
```

## 23 2.2.2 Free Source Form Conditional Compilation Sentinel

24 The following conditional compilation sentinel is recognized in free form source files:

```
!$
```

1 To enable conditional compilation, a line with a conditional compilation sentinel must satisfy the  
2 following criteria:

- 3 • The sentinel can appear in any column but must be preceded only by white space.
- 4 • The sentinel must appear as a single word with no intervening white space.
- 5 • Initial lines must have a space after the sentinel.
- 6 • Continued lines must have an ampersand as the last non-blank character on the line, prior to any  
7 comment appearing on the conditionally compiled line. Continued lines can have an ampersand  
8 after the sentinel, with optional white space before and after the ampersand.

9 If these criteria are met, the sentinel is replaced by two spaces. If these criteria are not met, the line  
10 is left unchanged.

11 Note – in the following example, the two forms for specifying conditional compilation in free  
12 source form are equivalent (the first line represents the position of the first 9 columns):

```
13 c23456789  
14 !$ iam = omp_get_thread_num() +      &  
15 !$&   index  
16  
17 #ifdef _OPENMP  
18     iam = omp_get_thread_num() +      &  
19     index  
20 #endif
```

Fortran

## 1 2.3 Internal Control Variables

2 An OpenMP implementation must act as if there are internal control variables (ICVs) that control  
3 the behavior of an OpenMP program. These ICVs store information such as the number of threads  
4 to use for future **parallel** regions, the schedule to use for worksharing loops and whether nested  
5 parallelism is enabled or not. The ICVs are given values at various times (described below) during  
6 the execution of the program. They are initialized by the implementation itself and may be given  
7 values through OpenMP environment variables and through calls to OpenMP API routines. The  
8 program can retrieve the values of these ICVs only through OpenMP API routines.

9 For purposes of exposition, this document refers to the ICVs by certain names, but an  
10 implementation is not required to use these names or to offer any way to access the variables other  
11 than through the ways shown in Section 2.3.2 on page 36.

### 12 2.3.1 ICV Descriptions

13 The following ICVs store values that affect the operation of **parallel** regions.

- 14 • *dyn-var* - controls whether dynamic adjustment of the number of threads is enabled for  
15 encountered **parallel** regions. There is one copy of this ICV per data environment.
- 16 • *nest-var* - controls whether nested parallelism is enabled for encountered **parallel** regions.  
17 There is one copy of this ICV per data environment.
- 18 • *nthreads-var* - controls the number of threads requested for encountered **parallel** regions.  
19 There is one copy of this ICV per data environment.
- 20 • *thread-limit-var* - controls the maximum number of threads participating in the contention  
21 group. There is one copy of this ICV per data environment.
- 22 • *max-active-levels-var* - controls the maximum number of nested active **parallel** regions.  
23 There is one copy of this ICV per device.
- 24 • *place-partition-var* – controls the place partition available to the execution environment for  
25 encountered **parallel** regions. There is one copy of this ICV per implicit task.
- 26 • *active-levels-var* - the number of nested, active parallel regions enclosing the current task such  
27 that all of the **parallel** regions are enclosed by the outermost initial task region on the current  
28 device. There is one copy of this ICV per data environment.
- 29 • *levels-var* - the number of nested parallel regions enclosing the current task such that all of the  
30 **parallel** regions are enclosed by the outermost initial task region on the current device.  
31 There is one copy of this ICV per data environment.

- *bind-var* - controls the binding of OpenMP threads to places. When binding is requested, the variable indicates that the execution environment is advised not to move threads between places. The variable can also provide default thread affinity policies. There is one copy of this ICV per data environment.

The following ICVs store values that affect the operation of loop regions.

- *run-sched-var* - controls the schedule that the **runtime** schedule clause uses for loop regions. There is one copy of this ICV per data environment.
- *def-sched-var* - controls the implementation defined default scheduling of loop regions. There is one copy of this ICV per device.

The following ICVs store values that affect the program execution.

- *stacksize-var* - controls the stack size for threads that the OpenMP implementation creates. There is one copy of this ICV per device.
- *wait-policy-var* - controls the desired behavior of waiting threads. There is one copy of this ICV per device.
- *cancel-var* - controls the desired behavior of the **cancel** construct and cancellation points. There is one copy of the ICV for the whole program (the scope is global).
- *default-device-var* - controls the default target device. There is one copy of this ICV per data environment

## 2.3.2 ICV Initialization

The following table shows the ICVs, associated environment variables, and initial values:

ICV	Environment Variable	Initial value
<i>dyn-var</i>	<b>OMP_DYNAMIC</b>	See comments below
<i>nest-var</i>	<b>OMP_NESTED</b>	<i>false</i>
<i>nthreads-var</i>	<b>OMP_NUM_THREADS</b>	Implementation defined
<i>run-sched-var</i>	<b>OMP_SCHEDULE</b>	Implementation defined
<i>def-sched-var</i>	(none)	Implementation defined
<i>bind-var</i>	<b>OMP_PROC_BIND</b>	Implementation defined

*table continued on next page*

table continued from previous page

ICV	Environment Variable	Initial value
<i>stacksize-var</i>	<b>OMP_STACKSIZE</b>	Implementation defined
<i>wait-policy-var</i>	<b>OMP_WAIT_POLICY</b>	Implementation defined
<i>thread-limit-var</i>	<b>OMP_THREAD_LIMIT</b>	Implementation defined
<i>max-active-levels-var</i>	<b>OMP_MAX_ACTIVE_LEVELS</b>	See comments below
<i>active-levels-var</i>	(none)	<i>zero</i>
<i>levels-var</i>	(none)	<i>zero</i>
<i>place-partition-var</i>	<b>OMP_PLACES</b>	Implementation defined
<i>cancel-var</i>	<b>OMP_CANCELLATION</b>	<i>false</i>
<i>default-device-var</i>	<b>OMP_DEFAULT_DEVICE</b>	Implementation defined

## Comments

- Each device has its own ICVs.
- The value of the *nthreads-var* ICV is a list.
- The value of the *bind-var* ICV is a list.
- The initial value of *dyn-var* is implementation defined if the implementation supports dynamic adjustment of the number of threads; otherwise, the initial value is *false*.
- The initial value of *max-active-levels-var* is the number of levels of parallelism that the implementation supports. See the definition of *supporting n levels of parallelism* in Section 1.2.7 on page 12 for further details.

The host and target device ICVs are initialized before any OpenMP API construct or OpenMP API routine executes. After the initial values are assigned, the values of any OpenMP environment variables that were set by the user are read and the associated ICVs for the host device are modified accordingly. The method for initializing a target device's ICVs is implementation defined.

## Cross References

- **OMP\_SCHEDULE** environment variable, see Section 4.1 on page 251.
- **OMP\_NUM\_THREADS** environment variable, see Section 4.2 on page 252.
- **OMP\_DYNAMIC** environment variable, see Section 4.3 on page 253.
- **OMP\_PROC\_BIND** environment variable, see Section 4.4 on page 253.

- 1       • **OMP\_PLACES** environment variable, see Section 4.5 on page 254.
- 2       • **OMP\_NESTED** environment variable, see Section 4.6 on page 256.
- 3       • **OMP\_STACKSIZE** environment variable, see Section 4.7 on page 256.
- 4       • **OMP\_WAIT\_POLICY** environment variable, see Section 4.8 on page 257.
- 5       • **OMP\_MAX\_ACTIVE\_LEVELS** environment variable, see Section 4.9 on page 258.
- 6       • **OMP\_THREAD\_LIMIT** environment variable, see Section 4.10 on page 258.
- 7       • **OMP\_CANCELLATION** environment variable, see Section 4.11 on page 259.
- 8       • **OMP\_DEFAULT\_DEVICE** environment variable, see Section 4.13 on page 260.

### 9   2.3.3   Modifying and Retrieving ICV Values

10       The following table shows the method for modifying and retrieving the values of ICVs through  
 11       OpenMP API routines:

ICV	Ways to modify value	Ways to retrieve value
<i>dyn-var</i>	<code>omp_set_dynamic()</code>	<code>omp_get_dynamic()</code>
<i>nest-var</i>	<code>omp_set_nested()</code>	<code>omp_get_nested()</code>
<i>nthreads-var</i>	<code>omp_set_num_threads()</code>	<code>omp_get_max_threads()</code>
<i>run-sched-var</i>	<code>omp_set_schedule()</code>	<code>omp_get_schedule()</code>
<i>def-sched-var</i>	(none)	(none)
<i>bind-var</i>	(none)	<code>omp_get_proc_bind()</code>
<i>stacksize-var</i>	(none)	(none)
<i>wait-policy-var</i>	(none)	(none)
<i>thread-limit-var</i>	<b>thread_limit</b> clause	<code>omp_get_thread_limit()</code>
<i>max-active-levels-var</i>	<code>omp_set_max_active_levels()</code>	<code>omp_get_max_active_levels()</code>
<i>active-levels-var</i>	(none)	<code>omp_get_active_levels()</code>
<i>levels-var</i>	(none)	<code>omp_get_level()</code>
<i>place-partition-var</i>	(none)	(none)

*table continued on next page*



*table continued from previous page*

ICV	Ways to modify value	Ways to retrieve value
<i>cancel-var</i>	(none)	<code>omp_get_cancellation()</code>
<i>default-device-var</i>	<code>omp_set_default_device()</code>	<code>omp_get_default_device()</code>

## Comments

- The value of the *nthreads-var* ICV is a list. The runtime call `omp_set_num_threads()` sets the value of the first element of this list, and `omp_get_max_threads()` retrieves the value of the first element of this list.
- The value of the *bind-var* ICV is a list. The runtime call `omp_get_proc_bind()` retrieves the value of the first element of this list.

## Cross References

- `thread_limit` clause of the `teams` construct, see Section 2.10.5 on page 102.
- `omp_set_num_threads` routine, see Section 3.2.1 on page 208.
- `omp_get_max_threads` routine, see Section 3.2.3 on page 210.
- `omp_set_dynamic` routine, see Section 3.2.7 on page 214.
- `omp_get_dynamic` routine, see Section 3.2.8 on page 216.
- `omp_get_cancellation` routine, see Section 3.2.9 on page 217.
- `omp_set_nested` routine, see Section 3.2.10 on page 217.
- `omp_get_nested` routine, see Section 3.2.11 on page 219.
- `omp_set_schedule` routine, see Section 3.2.12 on page 220.
- `omp_get_schedule` routine, see Section 3.2.13 on page 222.
- `omp_get_thread_limit` routine, see Section 3.2.14 on page 223.
- `omp_set_max_active_levels` routine, see Section 3.2.15 on page 223.
- `omp_get_max_active_levels` routine, see Section 3.2.16 on page 225.
- `omp_get_level` routine, see Section 3.2.17 on page 226.
- `omp_get_active_level` routine, see Section 3.2.20 on page 229.
- `omp_get_proc_bind` routine, see Section 3.2.22 on page 231.
- `omp_set_default_device` routine, see Section 3.2.23 on page 233.
- `omp_get_default_device` routine, see Section 3.2.24 on page 234.

## 1 2.3.4 How ICVs are Scoped

2 The following table shows the ICVs and their scope:

ICV	Scope
<i>dyn-var</i>	data environment
<i>nest-var</i>	data environment
<i>nthreads-var</i>	data environment
<i>run-sched-var</i>	data environment
<i>def-sched-var</i>	device
<i>bind-var</i>	data environment
<i>stacksize-var</i>	device
<i>wait-policy-var</i>	device
<i>thread-limit-var</i>	data environment
<i>max-active-levels-var</i>	device
<i>active-levels-var</i>	data environment
<i>levels-var</i>	data environment
<i>place-partition-var</i>	implicit task
<i>cancel-var</i>	device
<i>default-device-var</i>	data environment

### 4 **Comments**

- 5 • There is one copy per device of each ICV with device scope
- 6 • Each data environment has its own copies of ICVs with data environment scope
- 7 • Each implicit task has its own copy of ICVs with implicit task scope

8 Calls to OpenMP API routines retrieve or modify data environment scoped ICVs in the data  
9 environment of their binding tasks.

### 10 2.3.4.1 How the Per-Data Environment ICVs Work

11 When a **task** construct or **parallel** construct is encountered, the generated task(s) inherit the  
12 values of the data environment scoped ICVs from the generating task's ICV values.

13 When a **task** construct is encountered, the generated task inherits the value of *nthreads-var* from  
14 the generating task's *nthreads-var* value. When a **parallel** construct is encountered, and the

1 generating task's *nthreads-var* list contains a single element, the generated task(s) inherit that list as  
 2 the value of *nthreads-var*. When a **parallel** construct is encountered, and the generating task's  
 3 *nthreads-var* list contains multiple elements, the generated task(s) inherit the value of *nthreads-var*  
 4 as the list obtained by deletion of the first element from the generating task's *nthreads-var* value.  
 5 The *bind-var* ICV is handled in the same way as the *nthreads-var* ICV.

6 When a **target** construct is encountered, the construct's data environment uses the values of the  
 7 data environment scoped ICVs from the device data environment ICV values of the device that will  
 8 execute the region. If a **teams** construct with a **thread\_limit** clause is encountered, the  
 9 *thread-limit-var* ICV of the construct's data environment is instead set to a value that is less than or  
 10 equal to the value specified in the clause.

11 When encountering a loop worksharing region with **schedule(runtime)**, all implicit task  
 12 regions that constitute the binding parallel region must have the same value for *run-sched-var* in  
 13 their data environments. Otherwise, the behavior is unspecified.

## 14 2.3.5 ICV Override Relationships

15 The override relationships among construct clauses and ICVs are shown in the following table:

ICV	construct clause, if used
<i>dyn-var</i>	(none)
<i>nest-var</i>	(none)
<i>nthreads-var</i>	<b>num_threads</b>
<i>run-sched-var</i>	<b>schedule</b>
<i>def-sched-var</i>	<b>schedule</b>
<i>bind-var</i>	<b>proc_bind</b>
<i>stacksize-var</i>	(none)
<i>wait-policy-var</i>	(none)
<i>thread-limit-var</i>	(none)
<i>max-active-levels-var</i>	(none)
<i>active-levels-var</i>	(none)

*table continued on next page*

table continued from previous page

ICV	construct clause, if used
<i>levels-var</i>	(none)
<i>place-partition-var</i>	(none)
<i>cancel-var</i>	(none)
<i>default-device-var</i>	(none)

## Comments

- The **num\_threads** clause overrides the value of the first element of the *nthreads-var* ICV.
- If *bind-var* is not set to *false* then the **proc\_bind** clause overrides the value of the first elements of the *bind-var* ICV; otherwise, the **proc\_bind** clause has no effect.

## Cross References

- **parallel** construct, see Section 2.5 on page 44.
- **proc\_bind** clause, Section 2.5 on page 44.
- **num\_threads** clause, see Section 2.5.1 on page 48.
- Loop construct, see Section 2.7.1 on page 55.
- **schedule** clause, see Section 2.7.1.1 on page 62.

## 1 2.4 Array Sections

2 An array section designates a subset of the elements in an array. An array section can appear only  
3 in clauses where it is explicitly allowed.

▼ C / C++ ▼

4 To specify an array section in an OpenMP construct, array subscript expressions are extended with  
5 the following syntax:

6 `[ lower-bound : length ]` or

7 `[ lower-bound : ]` or

8 `[ : length ]` or

9 `[ : ]`

10 The array section must be a subset of the original array.

11 Array sections are allowed on multidimensional arrays. Base language array subscript expressions  
12 can be used to specify length-one dimensions of multidimensional array sections.

13 The *lower-bound* and *length* are integral type expressions. When evaluated they represent a set of  
14 integer values as follows:

15 { *lower-bound*, *lower-bound* + 1, *lower-bound* + 2, ... , *lower-bound* + *length* - 1 }

16 The *lower-bound* and *length* must evaluate to non-negative integers.

17 When the size of the array dimension is not known, the *length* must be specified explicitly.

18 When the *length* is absent, it defaults to the size of the array dimension minus the *lower-bound*.

19 When the *lower-bound* is absent it defaults to 0.

20 Note – The following are examples of array sections:

21 `a[0:6]`

22 `a[:6]`


23 `a[1:10]`

24 `a[1:]`

25 `b[10][:][:0]`

26 `c[1:10][42][0:6]`

1 The first two examples are equivalent. If **a** is declared to be an eleven element array, the third and  
2 fourth examples are equivalent. The fifth example is a zero-length array section. The last example  
3 is not contiguous.



4 Fortran has built-in support for array sections but the following restrictions apply for OpenMP  
5 constructs:

- 6 • A stride expression may not be specified.
- 7 • The upper bound for the last dimension of an assumed-size dummy array must be specified.




## 8 Restrictions

9 Restrictions to array sections are as follows:

- 10 • An array section can appear only in clauses where it is explicitly allowed.



- 11 • An array section can only be specified for a base language identifier.
- 12 • The type of the variable appearing in an array section must be array, pointer, reference to array,  
13 or reference to pointer.



- 14 • An array section cannot be used in a C++ user-defined []-operator.



## 15 2.5 parallel Construct

### 16 Summary

17 This fundamental construct starts parallel execution. See Section 1.3 on page 13 for a general  
18 description of the OpenMP execution model.

## Syntax

C / C++

The syntax of the **parallel** construct is as follows:

```
#pragma omp parallel [clause[ [, ] clause] ... ] new-line
    structured-block
```

where *clause* is one of the following:

```
    if (scalar-expression)
    num_threads (integer-expression)
    default (shared | none)
    private (list)
    firstprivate (list)
    shared (list)
    copyin (list)
    reduction (reduction-identifier : list)
    proc_bind (master | close | spread)
```

C / C++

Fortran

The syntax of the **parallel** construct is as follows:

```
!$omp parallel [clause[ [, ] clause] ... ]
    structured-block
!$omp end parallel
```

1 where *clause* is one of the following:

```
2     if (scalar-logical-expression)
3     num_threads (scalar-integer-expression)
4     default (private | firstprivate | shared | none)
5     private (list)
6     firstprivate (list)
7     shared (list)
8     copyin (list)
9     reduction (reduction-identifier : list)
10    proc_bind (master | close | spread)
```

11 The **end parallel** directive denotes the end of the **parallel** construct.



Fortran

## 12 Binding

13 The binding thread set for a **parallel** region is the encountering thread. The encountering thread  
14 becomes the master thread of the new team.

## 15 Description

16 When a thread encounters a **parallel** construct, a team of threads is created to execute the  
17 **parallel** region (see Section 2.5.1 on page 48 for more information about how the number of  
18 threads in the team is determined, including the evaluation of the **if** and **num\_threads** clauses).  
19 The thread that encountered the **parallel** construct becomes the master thread of the new team,  
20 with a thread number of zero for the duration of the new **parallel** region. All threads in the new  
21 team, including the master thread, execute the region. Once the team is created, the number of  
22 threads in the team remains constant for the duration of that **parallel** region.

23 The optional **proc\_bind** clause, described in Section 2.5.2 on page 50, specifies the mapping of  
24 OpenMP threads to places within the current place partition, that is, within the places listed in the  
25 *place-partition-var* ICV for the implicit task of the encountering thread.

26 Within a **parallel** region, thread numbers uniquely identify each thread. Thread numbers are  
27 consecutive whole numbers ranging from zero for the master thread up to one less than the number  
28 of threads in the team. A thread may obtain its own thread number by a call to the  
29 **omp\_get\_thread\_num** library routine.

30 A set of implicit tasks, equal in number to the number of threads in the team, is generated by the  
31 encountering thread. The structured block of the **parallel** construct determines the code that will be



1 executed in each implicit task. Each task is assigned to a different thread in the team and becomes  
2 tied. The task region of the task being executed by the encountering thread is suspended and each  
3 thread in the team executes its implicit task. Each thread can execute a path of statements that is  
4 different from that of the other threads

5 The implementation may cause any thread to suspend execution of its implicit task at a task  
6 scheduling point, and switch to execute any explicit task generated by any of the threads in the  
7 team, before eventually resuming execution of the implicit task (for more details see Section 2.9 on  
8 page 80).

9 There is an implied barrier at the end of a **parallel** region. After the end of a **parallel**  
10 region, only the master thread of the team resumes execution of the enclosing task region.

11 If a thread in a team executing a **parallel** region encounters another **parallel** directive, it  
12 creates a new team, according to the rules in Section 2.5.1 on page 48, and it becomes the master of  
13 that new team.

14 If execution of a thread terminates while inside a **parallel** region, execution of all threads in all  
15 teams terminates. The order of termination of threads is unspecified. All work done by a team prior  
16 to any barrier that the team has passed in the program is guaranteed to be complete. The amount of  
17 work done by each thread after the last barrier that it passed and before it terminates is unspecified.

## 18 Restrictions

19 Restrictions to the **parallel** construct are as follows:

- 20 • A program that branches into or out of a **parallel** region is non-conforming.
- 21 • A program must not depend on any ordering of the evaluations of the clauses of the **parallel**  
22 directive, or on any side effects of the evaluations of the clauses.
- 23 • At most one **if** clause can appear on the directive.
- 24 • At most one **proc\_bind** clause can appear on the directive.
- 25 • At most one **num\_threads** clause can appear on the directive. The **num\_threads**  
26 expression must evaluate to a positive integer value.

▼ C / C++ ▼

27 A **throw** executed inside a **parallel** region must cause execution to resume within the same  
28 **parallel** region, and the same thread that threw the exception must catch it.

▲ C / C++ ▲

▼ Fortran ▼

29 Unsynchronized use of Fortran I/O statements by multiple threads on the same unit has unspecified  
30 behavior.

▲ Fortran ▲

## Cross References

- **default**, **shared**, **private**, **firstprivate**, and **reduction** clauses, see Section 2.14.3 on page 172.
- **copyin** clause, see Section 2.14.4 on page 191.
- **omp\_get\_thread\_num** routine, see Section 3.2.4 on page 212.

## 2.5.1 Determining the Number of Threads for a `parallel` Region

When execution encounters a **parallel** directive, the value of the **if** clause or **num\_threads** clause (if any) on the directive, the current parallel context, and the values of the *nthreads-var*, *dyn-var*, *thread-limit-var*, *max-active-levels-var*, and *nest-var* ICVs are used to determine the number of threads to use in the region.

Note that using a variable in an **if** or **num\_threads** clause expression of a **parallel** construct causes an implicit reference to the variable in all enclosing constructs. The **if** clause expression and the **num\_threads** clause expression are evaluated in the context outside of the **parallel** construct, and no ordering of those evaluations is specified. It is also unspecified whether, in what order, or how many times any side effects of the evaluation of the **num\_threads** or **if** clause expressions occur.

When a thread encounters a **parallel** construct, the number of threads is determined according to Algorithm 2.1.

---

### Algorithm 2.1

---

```
let ThreadsBusy be the number of OpenMP threads currently executing in this
contention group;

let ActiveParRegions be the number of enclosing active parallel regions;

if an if clause exists
then let IfClauseValue be the value of the if clause expression;
else let IfClauseValue = true;

if a num_threads clause exists
then let ThreadsRequested be the value of the num_threads clause expression;
else let ThreadsRequested = value of the first element of nthreads-var;
```

```

1      let ThreadsAvailable = (thread-limit-var - ThreadsBusy + 1);
2      if (IfClauseValue = false)
3      then number of threads = 1;
4      else if (ActiveParRegions >= 1) and (nest-var = false)
5      then number of threads = 1;
6      else if (ActiveParRegions = max-active-levels-var)
7      then number of threads = 1;
8      else if (dyn-var = true) and (ThreadsRequested <= ThreadsAvailable)
9      then number of threads = [ 1 : ThreadsRequested ];
10     else if (dyn-var = true) and (ThreadsRequested > ThreadsAvailable)
11     then number of threads = [ 1 : ThreadsAvailable ];
12     else if (dyn-var = false) and (ThreadsRequested <= ThreadsAvailable)
13     then number of threads = ThreadsRequested;
14     else if (dyn-var = false) and (ThreadsRequested > ThreadsAvailable)
15     then behavior is implementation defined;
17

```

---

18 **Note** – Since the initial value of the *dyn-var* ICV is implementation defined, programs that depend  
19 on a specific number of threads for correct execution should explicitly disable dynamic adjustment  
20 of the number of threads

## 21 **Cross References**

- 22 • *nthreads-var*, *dyn-var*, *thread-limit-var*, *max-active-levels-var*, and *nest-var* ICVs, see  
23 Section 2.3 on page 35.

## 1 2.5.2 Controlling OpenMP Thread Affinity

2 When a thread encounters a parallel directive without a **proc\_bind** clause, the *bind-var* ICV is  
3 used to determine the policy for assigning OpenMP threads to places within the current place  
4 partition, that is, the places listed in the *place-partition-var* ICV for the implicit task of the  
5 encountering thread. If the parallel directive has a **proc\_bind** clause then the binding policy  
6 specified by the **proc\_bind** clause overrides the policy specified by the first element of the  
7 *bind-var* ICV. Once a thread in the team is assigned to a place, the OpenMP implementation should  
8 not move it to another place.

9 The **master** thread affinity policy instructs the execution environment to assign every thread in the  
10 team to the same place as the master thread. The place partition is not changed by this policy, and  
11 each implicit task inherits the *place-partition-var* ICV of the parent implicit task.

12 The **close** thread affinity policy instructs the execution environment to assign the threads in the  
13 team to places close to the place of the parent thread. The place partition is not changed by this  
14 policy, and each implicit task inherits the *place-partition-var* ICV of the parent implicit task. If  $T$   
15 is the number of threads in the team, and  $P$  is the number of places in the parent's place partition,  
16 then the assignment of threads in the team to places is as follows:

- 17 •  $T \leq P$ . The master thread executes on the place of the parent thread. The thread with the next  
18 smallest thread number executes on the next place in the place partition, and so on, with wrap  
19 around with respect to the place partition of the master thread.
- 20 •  $T > P$  Each place  $P$  will contain  $S_p$  threads with consecutive thread numbers, where  
21  $\lfloor (T/P) \rfloor \leq S_p \leq \lceil (T/P) \rceil$ . The first  $S_0$  threads (including the master thread) are assigned to  
22 the place of the parent thread. The next  $S_1$  threads are assigned to the next place in the place  
23 partition, and so on, with wrap around with respect to the place partition of the master thread.  
24 When  $P$  does not divide  $T$  evenly, the exact number of threads in a particular place is  
25 implementation defined.

26 The purpose of the **spread** thread affinity policy is to create a sparse distribution for a team of  $T$   
27 threads among the  $P$  places of the parent's place partition. A sparse distribution is achieved by first  
28 subdividing the parent partition into  $T$  subpartitions if  $T \leq P$ , or  $P$  subpartitions if  $T > P$ . Then  
29 one thread ( $T \leq P$ ) or a set of threads ( $T > P$ ) is assigned to each subpartition. The  
30 *place-partition-var* ICV of each implicit task is set to its subpartition. The subpartitioning is not  
31 only a mechanism for achieving a sparse distribution, it also defines a subset of places for a thread  
32 to use when creating a nested parallel region. The assignment of threads to places is as follows:

- 33 •  $T \leq P$ . The parent thread's place partition is split into  $T$  subpartitions, where each subpartition  
34 contains  $\lfloor (P/T) \rfloor$  or  $\lceil (P/T) \rceil$  consecutive places. A single thread is assigned to each  
35 subpartition. The master thread executes on the place of the parent thread and is assigned to the  
36 subpartition that includes that place. The thread with the next smallest thread number is assigned  
37 to the first place in the next subpartition, and so on, with wrap around with respect to the original  
38 place partition of the master thread.

1           •  $T > P$ . The parent thread's place partition is split into  $P$  subpartitions, each consisting of a  
2           single place. Each subpartition is assigned  $S_p$  threads with consecutive thread numbers, where  
3            $\lfloor (T/P) \rfloor \leq S_p \leq \lceil (T/P) \rceil$ . The first  $S_0$  threads (including the master thread) are assigned to the  
4           subpartition containing the place of the parent thread. The next  $S_1$  threads are assigned to the  
5           next subpartition, and so on, with wrap around with respect to the original place partition of the  
6           master thread. When  $P$  does not divide  $T$  evenly, the exact number of threads in a particular  
7           subpartition is implementation defined.

8           The determination of whether the affinity request can be fulfilled is implementation defined. If the  
9           affinity request cannot be fulfilled, then the affinity of threads in the team is implementation defined.

10          Note - Wrap around is needed if the end of a place partition is reached before all thread  
11          assignments are done. For example, wrap around may be needed in the case of **close** and  $T \leq P$ ,  
12          if the master thread is assigned to a place other than the first place in the place partition. In this  
13          case, thread 1 is assigned to the place after the place of the master place, thread 2 is assigned to the  
14          place after that, and so on. The end of the place partition may be reached before all threads are  
15          assigned. In this case, assignment of threads is resumed with the first place in the place partition.

## 1 2.6 Canonical Loop Form

C / C++

2 A loop has *canonical loop form* if it conforms to the following:

3  
4  
5

---

**for** (*init-expr*; *test-expr*; *incr-expr*) *structured-block*

---

*init-expr*      One of the following:  
                  *var = lb*  
                  *integer-type var = lb*  
                  *random-access-iterator-type var = lb*  
                  *pointer-type var = lb*

*test-expr*      One of the following:  
                  *var relational-op b*  
                  *b relational-op var*

*incr-expr*      One of the following:  
                  ++*var*  
                  *var*++  
                  -- *var*  
                  *var* --  
                  *var += incr*  
                  *var -= incr*  
                  *var = var + incr*  
                  *var = incr + var*  
                  *var = var - incr*

*var*              One of the following:  
                  A variable of a signed or unsigned integer type.  
                  For C++, a variable of a random access iterator type.  
                  For C, a variable of a pointer type.  
                  If this variable would otherwise be shared, it is implicitly made private in the loop construct. This variable must not be modified during the execution of the *for-loop* other than in *incr-expr*. Unless the variable is specified **lastprivate** on the loop construct, its value after the loop is unspecified.

---

*continued on next page*

*continued from previous page*

*relational-op* One of the following:

<  
<=  
>  
>=

*lb* and *b* Loop invariant expressions of a type compatible with the type of *var*.

*incr* A loop invariant integer expression.

The canonical form allows the iteration count of all associated loops to be computed before executing the outermost loop. The computation is performed for each loop in an integer type. This type is derived from the type of *var* as follows:

- If *var* is of an integer type, then the type is the type of *var*.
- For C++, if *var* is of a random access iterator type, then the type is the type that would be used by `std::distance` applied to variables of the type of *var*.
- For C, if *var* is of a pointer type, then the type is `ptrdiff_t`.

The behavior is unspecified if any intermediate result required to compute the iteration count cannot be represented in the type determined above.

There is no implied synchronization during the evaluation of the *lb*, *b*, or *incr* expressions. It is unspecified whether, in what order, or how many times any side effects within the *lb*, *b*, or *incr* expressions occur.

**Note** – Random access iterators are required to support random access to elements in constant time. Other iterators are precluded by the restrictions since they can take linear time or offer limited functionality. It is therefore advisable to use tasks to parallelize those cases.

## 1 Restrictions

2 The following restrictions also apply:

- 3 • If *test-expr* is of the form *var relational-op b* and *relational-op* is *<* or *<=* then *incr-expr* must  
4 cause *var* to increase on each iteration of the loop. If *test-expr* is of the form *var relational-op b*  
5 and *relational-op* is *>* or *>=* then *incr-expr* must cause *var* to decrease on each iteration of the  
6 loop.
- 7 • If *test-expr* is of the form *b relational-op var* and *relational-op* is *<* or *<=* then *incr-expr* must  
8 cause *var* to decrease on each iteration of the loop. If *test-expr* is of the form *b relational-op var*  
9 and *relational-op* is *>* or *>=* then *incr-expr* must cause *var* to increase on each iteration of the  
10 loop.
- 11 • For C++, in the **simd** construct the only random access iterator types that are allowed for *var* are  
12 pointer types.

▲ C / C++ ▲

## 13 2.7 Worksharing Constructs

14 A worksharing construct distributes the execution of the associated region among the members of  
15 the team that encounters it. Threads execute portions of the region in the context of the implicit  
16 tasks each one is executing. If the team consists of only one thread then the worksharing region is  
17 not executed in parallel.

18 A worksharing region has no barrier on entry; however, an implied barrier exists at the end of the  
19 worksharing region, unless a **nowait** clause is specified. If a **nowait** clause is present, an  
20 implementation may omit the barrier at the end of the worksharing region. In this case, threads that  
21 finish early may proceed straight to the instructions following the worksharing region without  
22 waiting for the other members of the team to finish the worksharing region, and without performing  
23 a flush operation.

24 The OpenMP API defines the following worksharing constructs, and these are described in the  
25 sections that follow:

- 26 • loop construct
- 27 • **sections** construct
- 28 • **single** construct
- 29 • **workshare** construct



## 1       **Restrictions**

2       The following restrictions apply to worksharing constructs:

- 3       • Each worksharing region must be encountered by all threads in a team or by none at all, unless  
4       cancellation has been requested for the innermost enclosing parallel region.
- 5       • The sequence of worksharing regions and **barrier** regions encountered must be the same for  
6       every thread in a team

## 7   **2.7.1 Loop Construct**

### 8       **Summary**

9       The loop construct specifies that the iterations of one or more associated loops will be executed in  
10       parallel by threads in the team in the context of their implicit tasks. The iterations are distributed  
11       across threads that already exist in the team executing the **parallel** region to which the loop  
12       region binds.

### 13       **Syntax**

▼ **C / C++** ▼

14       The syntax of the loop construct is as follows:

```
#pragma omp for [clause[ [, ] clause] ... ] new-line  
          for-loops
```

15       where clause is one of the following:

```
private (list)  
firstprivate (list)  
lastprivate (list)  
linear (list)  
reduction (reduction-identifier : list)  
schedule (kind[ , chunk_size])  
collapse (n)  
ordered[ (n) ]  
nowait
```

1 The **for** directive places restrictions on the structure of all associated *for-loops*. Specifically, all  
2 associated *for-loops* must have *canonical loop form* (see Section 2.6 on page 52).



3 The syntax of the loop construct is as follows:

```
!$omp do [clause[ [, ] clause] ... ]  
    do-loops  
[$omp end do [nowait]]
```

4 where *clause* is one of the following:

- 5 **private** (*list*)
- 6 **firstprivate** (*list*)
- 7 **lastprivate** (*list*)
- 8 **linear** (*list*)
- 9 **reduction** (*reduction-identifier* : *list*)
- 10 **schedule** (*kind*[ , *chunk\_size*])
- 11 **collapse** (*n*)
- 12 **ordered**[ (*n*) ]

13 If an **end do** directive is not specified, an **end do** directive is assumed at the end of the *do-loops*.

14 All associated *do-loops* must be *do-constructs* as defined by the Fortran standard. If an **end do**  
15 directive follows a *do-construct* in which several loop statements share a **DO** termination statement,  
16 then the directive can only be specified for the outermost of these **DO** statements.

17 If any of the loop iteration variables would otherwise be shared, they are implicitly made private on  
18 the loop construct. Unless the loop iteration variables are specified **lastprivate** or **linear** on  
19 the loop construct, their values after the loop are unspecified.



## 20 Binding

21 The binding thread set for a loop region is the current team. A loop region binds to the innermost  
22 enclosing **parallel** region. Only the threads of the team executing the binding **parallel**  
23 region participate in the execution of the loop iterations and the implied barrier of the loop region if  
24 the barrier is not eliminated by a **nowait** clause.

## Description

The loop construct is associated with a loop nest consisting of one or more loops that follow the directive.

There is an implicit barrier at the end of a loop construct unless a **nowait** clause is specified.

The **collapse** clause or the **ordered** clause with the parameter may be used to specify how many loops are associated with the loop construct. An **ordered** clause without a parameter is equivalent to the **ordered** clause with a parameter value of one specified. The number of loops associated is determined by the parameters of the **collapse** clause and the **ordered** clause, which must be constant positive integer expressions.

If neither the **collapse** nor the **ordered** clause is present, the only loop that is associated with the loop construct is the one that immediately follows the loop directive. If the value of the parameter in the **collapse** or **ordered** clause is larger than the number of nested loops following the construct, the behavior is unspecified.

If both the **collapse** clause and the **ordered** clause with a parameter are specified, the **collapse** clause applies to the loops that immediately follow the directive. The **ordered** clause with the parameter then applies to the resulting loop nest (that is, the collapsed loop and the remaining loops). If the value of the parameter in the **ordered** clause is larger than the number of the loops in the resulting loop nest, the behavior is unspecified.

If more than one loop is associated with the loop construct and a **collapse** clause is specified, then the iterations of all associated loops are collapsed into one larger iteration space that is then divided according to the schedule clause.

If an **ordered** clause with the parameter is specified for the loop construct, then the associated loops form a *doacross loop nest*.

The sequential execution of the iterations in all associated loops determines the order of the iterations in the collapsed iteration space.

The iteration count for each associated loop is computed before entry to the outermost loop. If execution of any associated loop changes any of the values used to compute any of the iteration counts, then the behavior is unspecified.

The integer type (or kind, for Fortran) used to compute the iteration count for the collapsed loop is implementation defined.

A worksharing loop has logical iterations numbered 0,1,...,N-1 where N is the number of loop iterations, and the logical numbering denotes the sequence in which the iterations would be executed if the associated loop(s) were executed by a single thread. The **schedule** clause specifies how iterations of the associated loops are divided into contiguous non-empty subsets, called chunks, and how these chunks are distributed among threads of the team. Each thread executes its assigned chunk(s) in the context of its implicit task. The *chunk\_size* expression is evaluated using the original list items of any variables that are made private in the loop construct. It is unspecified whether, in what order, or how many times, any side effects of the evaluation of this

1 expression occur. The use of a variable in a **schedule** clause expression of a loop construct  
2 causes an implicit reference to the variable in all enclosing constructs.

3 Different loop regions with the same schedule and iteration count, even if they occur in the same  
4 parallel region, can distribute iterations among threads differently. The only exception is for the  
5 **static** schedule as specified in Table 2-1. Programs that depend on which thread executes a  
6 particular iteration under any other circumstances are non-conforming.

7 See Section [2.7.1.1](#) on page [62](#) for details of how the schedule for a worksharing loop is determined.  
8 The schedule *kind* can be one of those specified in Table 2-1.

---

<b>static</b>	<p>When <b>schedule (static, chunk_size)</b> is specified, iterations are divided into chunks of size <i>chunk_size</i>, and the chunks are assigned to the threads in the team in a round-robin fashion in the order of the thread number.</p> <p>When no <i>chunk_size</i> is specified, the iteration space is divided into chunks that are approximately equal in size, and at most one chunk is distributed to each thread. Note that the size of the chunks is unspecified in this case.</p> <p>A compliant implementation of the <b>static</b> schedule must ensure that the same assignment of logical iteration numbers to threads will be used in two loop regions if the following conditions are satisfied: 1) both loop regions have the same number of loop iterations, 2) both loop regions have the same value of <i>chunk_size</i> specified, or both loop regions have no <i>chunk_size</i> specified, 3) both loop regions bind to the same parallel region, and 4) neither loop is associated with a SIMD construct. A data dependence between the same logical iterations in two such loops is guaranteed to be satisfied allowing safe use of the <b>nowait</b> clause.</p>
2	<p><b>dynamic</b></p> <p>When <b>schedule (dynamic, chunk_size)</b> is specified, the iterations are distributed to threads in the team in chunks as the threads request them. Each thread executes a chunk of iterations, then requests another chunk, until no chunks remain to be distributed.</p> <p>Each chunk contains <i>chunk_size</i> iterations, except for the last chunk to be distributed, which may have fewer iterations.</p> <p>When no <i>chunk_size</i> is specified, it defaults to 1</p>
	<p><b>guided</b></p> <p>When <b>schedule (guided, chunk_size)</b> is specified, the iterations are assigned to threads in the team in chunks as the executing threads request them. Each thread executes a chunk of iterations, then requests another chunk, until no chunks remain to be assigned.</p> <p>For a <i>chunk_size</i> of 1, the size of each chunk is proportional to the number of unassigned iterations divided by the number of threads in the team, decreasing to 1. For a <i>chunk_size</i> with value <i>k</i> (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 <i>k</i> iterations (except for the last chunk to be assigned, which may have fewer than <i>k</i> iterations).</p>

---

*table continued on next page*

table continued from previous page

---

	When no <i>chunk_size</i> is specified, it defaults to 1.
<b>auto</b>	When <b>schedule (auto)</b> is specified, the decision regarding scheduling is delegated to the compiler and/or runtime system. The programmer gives the implementation the freedom to choose any possible mapping of iterations to threads in the team.
<b>runtime</b>	When <b>schedule (runtime)</b> is specified, the decision regarding scheduling is deferred until run time, and the schedule and chunk size are taken from the <i>run-sched-var</i> ICV. If the ICV is set to <b>auto</b> , the schedule is implementation defined.

---

**Note** – For a team of  $p$  threads and a loop of  $n$  iterations, let  $\lceil n/p \rceil$  be the integer  $q$  that satisfies  $n = p * q - r$ , with  $0 \leq r < p$ . One compliant implementation of the **static** schedule (with no specified *chunk\_size*) would behave as though *chunk\_size* had been specified with value  $q$ . Another compliant implementation would assign  $q$  iterations to the first  $p - r$  threads, and  $q - 1$  iterations to the remaining  $r$  threads. This illustrates why a conforming program must not rely on the details of a particular implementation.

A compliant implementation of the **guided** schedule with a *chunk\_size* value of  $k$  would assign  $q = \lceil n/p \rceil$  iterations to the first available thread and set  $n$  to the larger of  $n - q$  and  $p * k$ . It would then repeat this process until  $q$  is greater than or equal to the number of remaining iterations, at which time the remaining iterations form the final chunk. Another compliant implementation could use the same method, except with  $q = \lceil n/(2p) \rceil$ , and set  $n$  to the larger of  $n - q$  and  $2 * p * k$ .

## Restrictions

Restrictions to the loop construct are as follows:

- All loops associated with the loop construct must be perfectly nested; that is, there must be no intervening code nor any OpenMP directive between any two loops.
- The values of the loop control expressions of the loops associated with the loop construct must be the same for all the threads in the team.
- Only one **schedule** clause can appear on a loop directive.
- Only one **collapse** clause can appear on a loop directive.
- *chunk\_size* must be a loop invariant integer expression with a positive value.
- The value of the *chunk\_size* expression must be the same for all threads in the team.

- 1 • The value of the *run-sched-var* ICV must be the same for all threads in the team.
- 2 • When **schedule(runtime)** or **schedule(auto)** is specified, *chunk\_size* must not be
- 3 specified.
- 4 • Only one **ordered** clause can appear on a loop directive.
- 5 • The **ordered** clause must be present on the loop construct if any **ordered** region ever binds
- 6 to a loop region arising from the loop construct.
- 7 • The loop iteration variable may not appear in a **threadprivate** directive.

C / C++

- 8 • The associated *for-loops* must be structured blocks.
- 9 • Only an iteration of the innermost associated loop may be curtailed by a **continue** statement.
- 10 • No statement can branch to any associated **for** statement.
- 11 • Only one **nowait** clause can appear on a **for** directive.
- 12 • A throw executed inside a loop region must cause execution to resume within the same iteration
- 13 of the loop region, and the same thread that threw the exception must catch it.

C / C++

Fortran

- 14 • The associated *do-loops* must be structured blocks.
- 15 • Only an iteration of the innermost associated loop may be curtailed by a **CYCLE** statement.
- 16 • No statement in the associated loops other than the **DO** statements can cause a branch out of the
- 17 loops.
- 18 • The *do-loop* iteration variable must be of type integer.
- 19 • The *do-loop* cannot be a **DO WHILE** or a **DO** loop without loop control.

Fortran

## Cross References

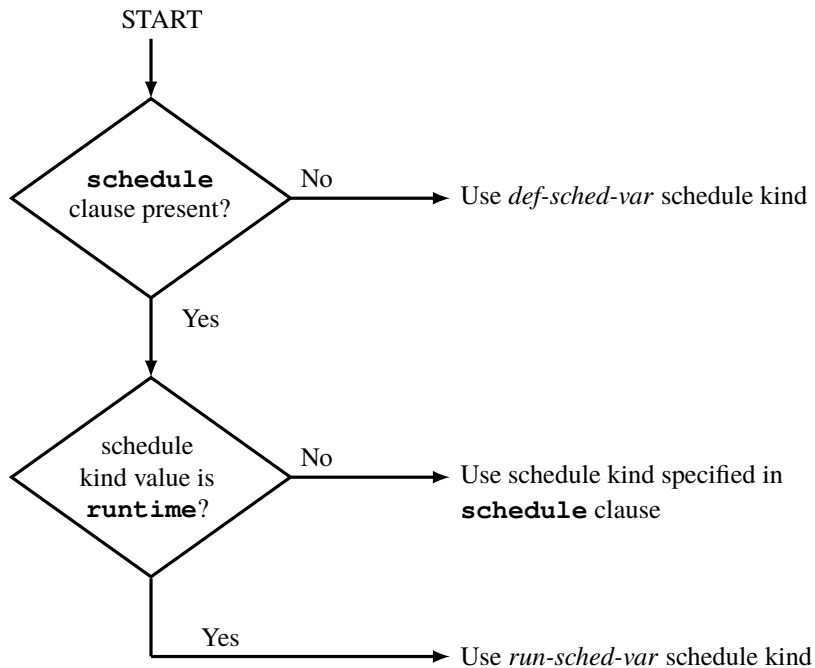
- 21 • **private**, **firstprivate**, **lastprivate**, **linear**, and **reduction** clauses, see
- 22 Section 2.14.3 on page 172.
- 23 • **OMP\_SCHEDULE** environment variable, see Section 4.1 on page 251.
- 24 • **ordered** construct, see Section 2.12.8 on page 152.
- 25 • **depend** clause, see Section 2.12.9 on page 154.

## 1 2.7.1.1 Determining the Schedule of a Worksharing Loop

2 When execution encounters a loop directive, the **schedule** clause (if any) on the directive, and  
3 the *run-sched-var* and *def-sched-var* ICVs are used to determine how loop iterations are assigned  
4 to threads. See Section 2.3 on page 35 for details of how the values of the ICVs are determined. If  
5 the loop directive does not have a **schedule** clause then the current value of the *def-sched-var*  
6 ICV determines the schedule. If the loop directive has a **schedule** clause that specifies the  
7 **runtime** schedule kind then the current value of the *run-sched-var* ICV determines the schedule.  
8 Otherwise, the value of the **schedule** clause determines the schedule. Figure 2-1 describes how  
9 the schedule for a worksharing loop is determined.

### 10 Cross References

- 11 • ICVs, see Section 2.3 on page 35



12 FIGURE 2-1 Determining the schedule for a worksharing loop.



## 1 2.7.2 sections Construct

### 2 Summary

3 The **sections** construct is a non-iterative worksharing construct that contains a set of structured  
4 blocks that are to be distributed among and executed by the threads in a team. Each structured  
5 block is executed once by one of the threads in the team in the context of its implicit task.

### 6 Syntax

C / C++

7 The syntax of the **sections** construct is as follows:

```
#pragma omp sections [clause[ [, ] clause] ... ] new-line
{
  [#pragma omp section new-line]
  structured-block
  [#pragma omp section new-line]
  structured-block]
  ...
}
```

8 where *clause* is one of the following:

```
9     private (list)
10    firstprivate (list)
11    lastprivate (list)
12    reduction (reduction-identifier : list)
13    nowait
```

C / C++

1 The syntax of the **sections** construct is as follows:

```

!$omp sections [clause[ [, ] clause] ... ]
  [!$omp section
    structured-block
  [!$omp section
    structured-block]
  ...
!$omp end sections [nowait]

```

2 where *clause* is one of the following:

```

3     private (list)
4     firstprivate (list)
5     lastprivate (list)
6     reduction (reduction-identifier : list)

```

## 7 Binding

8 The binding thread set for a **sections** region is the current team. A **sections** region binds to  
9 the innermost enclosing **parallel** region. Only the threads of the team executing the binding  
10 **parallel** region participate in the execution of the structured blocks and the implied barrier of  
11 the **sections** region if the barrier is not eliminated by a **nowait** clause.

## 12 Description

13 Each structured block in the **sections** construct is preceded by a **section** directive except  
14 possibly the first block, for which a preceding **section** directive is optional.

15 The method of scheduling the structured blocks among the threads in the team is implementation  
16 defined.

17 There is an implicit barrier at the end of a **sections** construct unless a **nowait** clause is  
18 specified.

## Restrictions

Restrictions to the **sections** construct are as follows:

- Orphaned **section** directives are prohibited. That is, the **section** directives must appear within the **sections** construct and must not be encountered elsewhere in the **sections** region.
- The code enclosed in a **sections** construct must be a structured block.
- Only a single **nowait** clause can appear on a **sections** directive.

C++

- A throw executed inside a **sections** region must cause execution to resume within the same section of the **sections** region, and the same thread that threw the exception must catch it.

C++

## Cross References

- **private**, **firstprivate**, **lastprivate**, and **reduction** clauses, see Section 2.14.3 on page 172.

## 2.7.3 single Construct

### Summary

The **single** construct specifies that the associated structured block is executed by only one of the threads in the team (not necessarily the master thread), in the context of its implicit task. The other threads in the team, which do not execute the block, wait at an implicit barrier at the end of the **single** construct unless a **nowait** clause is specified.

### Syntax

C / C++

The syntax of the single construct is as follows:

```
#pragma omp single [clause[ [, ] clause] ... ] new-line  
    structured-block
```

1 where *clause* is one of the following:

2       **private** (*list*)  
3       **firstprivate** (*list*)  
4       **copyprivate** (*list*)  
5       **nowait**



6 The syntax of the **single** construct is as follows:

```
!$omp single [clause [ , ] clause] ... ]  
    structured-block  
!$omp end single [end_clause [ , ] end_clause] ... ]
```

7 where *clause* is one of the following:

8       **private** (*list*)  
9       **firstprivate** (*list*)

10 and *end\_clause* is one of the following:

11       **copyprivate** (*list*)  
12       **nowait**



### 13 **Binding**

14 The binding thread set for a **single** region is the current team. A **single** region binds to the  
15 innermost enclosing **parallel** region. Only the threads of the team executing the binding  
16 **parallel** region participate in the execution of the structured block and the implied barrier of the  
17 **single** region if the barrier is not eliminated by a **nowait** clause.

### 18 **Description**

19 The method of choosing a thread to execute the structured block is implementation defined. There  
20 is an implicit barrier at the end of the **single** construct unless a **nowait** clause is specified.

## Restrictions

Restrictions to the **single** construct are as follows:

- The **copyprivate** clause must not be used with the **nowait** clause.
- At most one **nowait** clause can appear on a **single** construct.

C++

- A throw executed inside a **single** region must cause execution to resume within the same **single** region, and the same thread that threw the exception must catch it.

C++

## Cross References

- **private** and **firstprivate** clauses, see Section 2.14.3 on page 172.
- **copyprivate** clause, see Section 2.14.4.2 on page 193.

Fortran

## 2.7.4 workshare Construct

### Summary

The **workshare** construct divides the execution of the enclosed structured block into separate units of work, and causes the threads of the team to share the work such that each unit is executed only once by one thread, in the context of its implicit task.

### Syntax

The syntax of the **workshare** construct is as follows:

```
!$omp workshare
    structured-block
!$omp end workshare [nowait]
```

The enclosed structured block must consist of only the following:

- array assignments
- scalar assignments
- **FORALL** statements

1       • **FORALL** constructs

2       • **WHERE** statements

3       • **WHERE** constructs

4       • **atomic** constructs

5       • **critical** constructs

6       • **parallel** constructs

7       Statements contained in any enclosed **critical** construct are also subject to these restrictions.

8       Statements in any enclosed **parallel** construct are not restricted.

9       **Binding**

10       The binding thread set for a **workshare** region is the current team. A **workshare** region binds  
 11       to the innermost enclosing **parallel** region. Only the threads of the team executing the binding  
 12       **parallel** region participate in the execution of the units of work and the implied barrier of the  
 13       **workshare** region if the barrier is not eliminated by a **nowait** clause.

14       **Description**

15       There is an implicit barrier at the end of a **workshare** construct unless a **nowait** clause is  
 16       specified.

17       An implementation of the **workshare** construct must insert any synchronization that is required  
 18       to maintain standard Fortran semantics. For example, the effects of one statement within the  
 19       structured block must appear to occur before the execution of succeeding statements, and the  
 20       evaluation of the right hand side of an assignment must appear to complete prior to the effects of  
 21       assigning to the left hand side.

22       The statements in the **workshare** construct are divided into units of work as follows:

- 23       • For array expressions within each statement, including transformational array intrinsic functions  
 24       that compute scalar values from arrays:
  - 25       – Evaluation of each element of the array expression, including any references to **ELEMENTAL**  
 26       functions, is a unit of work.
  - 27       – Evaluation of transformational array intrinsic functions may be freely subdivided into any  
 28       number of units of work.
- 29       • For an array assignment statement, the assignment of each element is a unit of work.
- 30       • For a scalar assignment statement, the assignment operation is a unit of work.

- 1 • For a **WHERE** statement or construct, the evaluation of the mask expression and the masked  
2 assignments are each a unit of work.
- 3 • For a **FORALL** statement or construct, the evaluation of the mask expression, expressions  
4 occurring in the specification of the iteration space, and the masked assignments are each a unit  
5 of work
- 6 • For an **atomic** construct, the atomic operation on the storage location designated as x is a unit  
7 of work.
- 8 • For a **critical** construct, the construct is a single unit of work.
- 9 • For a **parallel** construct, the construct is a unit of work with respect to the **workshare**  
10 construct. The statements contained in the **parallel** construct are executed by a new thread  
11 team.
- 12 • If none of the rules above apply to a portion of a statement in the structured block, then that  
13 portion is a unit of work.

14 The transformational array intrinsic functions are **MATMUL**, **DOT\_PRODUCT**, **SUM**, **PRODUCT**,  
15 **MAXVAL**, **MINVAL**, **COUNT**, **ANY**, **ALL**, **SPREAD**, **PACK**, **UNPACK**, **RESHAPE**, **TRANSPOSE**,  
16 **EOSHIFT**, **CSHIFT**, **MINLOC**, and **MAXLOC**.

17 It is unspecified how the units of work are assigned to the threads executing a **workshare** region.

18 If an array expression in the block references the value, association status, or allocation status of  
19 private variables, the value of the expression is undefined, unless the same value would be  
20 computed by every thread.

21 If an array assignment, a scalar assignment, a masked array assignment, or a **FORALL** assignment  
22 assigns to a private variable in the block, the result is unspecified.

23 The **workshare** directive causes the sharing of work to occur only in the **workshare** construct,  
24 and not in the remainder of the **workshare** region.

## 25 **Restrictions**

26 The following restrictions apply to the **workshare** construct:

- 27 • All array assignments, scalar assignments, and masked array assignments must be intrinsic  
28 assignments.
- 29 • The construct must not contain any user defined function calls unless the function is  
30 **ELEMENTAL**.

## 1 2.8 SIMD Constructs

### 2 2.8.1 `simd` Construct

#### 3 Summary

4 The `simd` construct can be applied to a loop to indicate that the loop can be transformed into a  
5 SIMD loop (that is, multiple iterations of the loop can be executed concurrently using SIMD  
6 instructions).

#### 7 Syntax

8 The syntax of the `simd` construct is as follows:

▼ C / C++ ▼

```
#pragma omp simd [clause[ [, ] clause] ... ] new-line  
                  for-loops
```

9 where *clause* is one of the following:

10       **safelen** (*length*)  
11       **linear** (*list* [ : *linear-step*])  
12       **aligned** (*list* [ : *alignment*])  
13       **private** (*list*)  
14       **lastprivate** (*list*)  
15       **reduction** (*reduction-identifier* : *list*)  
16       **collapse** (*n*)

17 The `simd` directive places restrictions on the structure of the associated *for-loops*. Specifically, all  
18 associated *for-loops* must have *canonical loop form* (Section 2.6 on page 52).

▲ C / C++ ▲



```
!$omp simd [clause[ [, ] clause ... ]
    do-loops
[!$omp end simd]
```

1 where *clause* is one of the following:

- 2       **safelen** (*length*)
- 3       **linear** (*list* [ : *linear-step*])
- 4       **aligned** (*list* [ : *alignment*])
- 5       **private** (*list*)
- 6       **lastprivate** (*list*)
- 7       **reduction** (*reduction-identifier* : *list*)
- 8       **collapse** (*n*)

9 If an **end simd** directive is not specified, an **end simd** directive is assumed at the end of the  
10 *do-loops*.

11 All associated *do-loops* must be *do-constructs* as defined by the Fortran standard. If an **end simd**  
12 directive follows a *do-construct* in which several loop statements share a **DO** termination statement,  
13 then the directive can only be specified for the outermost of these **DO** statements.

## 14 **Binding**

15 A **simd** region binds to the current task region. The binding thread set of the **simd** region is the  
16 current team.

## Description

The **simd** construct enables the execution of multiple iterations of the associated loops concurrently by means of SIMD instructions.

The **collapse** clause may be used to specify how many loops are associated with the construct. The parameter of the **collapse** clause must be a constant positive integer expression. If no **collapse** clause is present, the only loop that is associated with the loop construct is the one that immediately follows the directive.

If more than one loop is associated with the **simd** construct, then the iterations of all associated loops are collapsed into one larger iteration space that is then executed with SIMD instructions. The sequential execution of the iterations in all associated loops determines the order of the iterations in the collapsed iteration space.

The iteration count for each associated loop is computed before entry to the outermost loop. If execution of any associated loop changes any of the values used to compute any of the iteration counts, then the behavior is unspecified.

The integer type (or kind, for Fortran) used to compute the iteration count for the collapsed loop is implementation defined.

A SIMD loop has logical iterations numbered 0,1,...,N-1 where N is the number of loop iterations, and the logical numbering denotes the sequence in which the iterations would be executed if the associated loop(s) were executed with no SIMD instructions. If the **safelen** clause is used then no two iterations executed concurrently with SIMD instructions can have a greater distance in the logical iteration space than its value. The parameter of the **safelen** clause must be a constant positive integer expression. The number of iterations that are executed concurrently at any given time is implementation defined. Each concurrent iteration will be executed by a different SIMD lane. Each set of concurrent iterations is a SIMD chunk. Lexical forward dependencies in the iterations of the original loop must be preserved within each SIMD chunk.

▼ C / C++ ▼

The **aligned** clause declares that the object to which each list item points is aligned to the number of bytes expressed in the optional parameter of the **aligned** clause

▲ C / C++ ▲  
▼ Fortran ▼

The **aligned** clause declares that the target of each list item is aligned to the number of bytes expressed in the optional parameter of the **aligned** clause.

▲ Fortran ▲

The optional parameter of the **aligned** clause, *alignment*, must be a constant positive integer expression. If no optional parameter is specified, implementation-defined default alignments for SIMD instructions on the target platforms are assumed.

## Restrictions

- All loops associated with the construct must be perfectly nested; that is, there must be no intervening code nor any OpenMP directive between any two loops.
- The associated loops must be structured blocks.
- A program that branches into or out of a **simd** region is non-conforming.
- Only one **collapse** clause can appear on a **simd** directive.
- A *list-item* cannot appear in more than one **aligned** clause.
- Only one **safelen** clause can appear on a **simd** directive.
- An ordered construct with the **simd** clause is the only OpenMP construct that can appear in the **simd** region.

C / C++

- The **simd** region cannot contain calls to the **longjmp** or **setjmp** functions.

C / C++

C

- The type of list items appearing in the **aligned** clause must be array or pointer.

C

C++

- The type of list items appearing in the **aligned** clause must be array, pointer, reference to array, or reference to pointer.
- No exception can be raised in the **simd** region.

C++

Fortran

- The *do-loop* iteration variable must be of type **integer**.
- The *do-loop* cannot be a **DO WHILE** or a **DO** loop without loop control.
- The type of list items appearing in the **aligned** clause must be **C\_PTR** or Cray pointer, or the list item must have the **POINTER** or **ALLOCATABLE** attribute.

Fortran

## Cross References

- **private**, **lastprivate**, **linear** and **reduction** clauses, see Section 2.14.3 on page 172.

## 1 2.8.2 declare simd Construct

### 2 Summary

3 The **declare simd** construct can be applied to a function (C, C++ and Fortran) or a subroutine  
4 (Fortran) to enable the creation of one or more versions that can process multiple arguments using  
5 SIMD instructions from a single invocation from a SIMD loop. The **declare simd** directive is a  
6 declarative directive. There may be multiple **declare simd** directives for a function (C, C++,  
7 Fortran) or subroutine (Fortran).

### 8 Syntax

9 The syntax of the **declare simd** construct is as follows:

▼ C / C++ ▼

```
#pragma omp declare simd [clause[ [, ] clause] ... ] new-line  
[#pragma omp declare simd [clause[ [, ] clause] ... ] new-line]  
[ ... ]  
    function definition or declaration
```

10 where *clause* is one of the following:

11       **simdlen**(*length*)  
12       **linear**(*argument-list* [ : *constant-linear-step*])  
13       **aligned**(*argument-list* [ : *alignment*])  
14       **uniform**(*argument-list*)  
15       **inbranch**  
16       **notinbranch**

▲ C / C++ ▲

```
!$omp declare simd(proc-name) [clause [ , ] clause] ... ]
```

where *clause* is one of the following:

```

simdlen (length)
linear (argument-list [ : constant-linear-step])
aligned (argument-list [ : alignment])
uniform (argument-list)
inbranch
notinbranch

```

## Description

The use of a **declare simd** construct on a function enables the creation of SIMD versions of the associated function that can be used to process multiple arguments from a single invocation from a SIMD loop concurrently.

The expressions appearing in the clauses of this directive are evaluated in the scope of the arguments of the function declaration or definition

The use of a **declare simd** construct enables the creation of SIMD versions of the specified subroutine or function that can be used to process multiple arguments from a single invocation from a SIMD loop concurrently.

## Fortran

1 If a **declare simd** directive contains multiple SIMD declarations, then one or more SIMD  
2 versions will be created for each declaration.

3 If a SIMD version is created, the number of concurrent arguments for the function is determined by  
4 the **simdlen** clause. If the **simdlen** clause is used its value corresponds to the number of  
5 concurrent arguments of the function. The parameter of the **simdlen** clause must be a constant  
6 positive integer expression. Otherwise, the number of concurrent arguments for the function is  
7 implementation defined.

8 The **uniform** clause declares one or more arguments to have an invariant value for all concurrent  
9 invocations of the function in the execution of a single SIMD loop.

## C / C++

10 The **aligned** clause declares that the object to which each list item points is aligned to the  
11 number of bytes expressed in the optional parameter of the **aligned** clause.

## C / C++

## Fortran

12 The **aligned** clause declares that the target of each list item is aligned to the number of bytes  
13 expressed in the optional parameter of the **aligned** clause.

## Fortran

14 The optional parameter of the **aligned** clause, *alignment*, must be a constant positive integer  
15 expression. If no optional parameter is specified, implementation-defined default alignments for  
16 SIMD instructions on the target platforms are assumed.

17 The **inbranch** clause specifies that the function will always be called from inside a conditional  
18 statement of a SIMD loop. The **notinbranch** clause specifies that the function will never be  
19 called from inside a conditional statement of a SIMD loop. If neither clause is specified, then the  
20 function may or may not be called from inside a conditional statement of a SIMD loop.

## Restrictions

- Each argument can appear in at most one **uniform** or **linear** clause.
- At most one **simdlen** clause can appear in a **declare simd** directive.
- Either **inbranch** or **notinbranch** may be specified, but not both.
- When a *constant-linear-step* expression is specified in a **linear** clause it must be a constant positive integer expression.
- The function or subroutine body must be a structured block.
- The execution of the function or subroutine, when called from a SIMD loop, cannot result in the execution of an OpenMP construct except for an **ordered** construct with the **simd** clause.
- The execution of the function or subroutine cannot have any side effects that would alter its execution for concurrent iterations of a SIMD chunk.
- A program that branches into or out of the function is non-conforming.

▼ C / C++ ▼

- If the function has any declarations, then the **declare simd** construct for any declaration that has one must be equivalent to the one specified for the definition. Otherwise, the result is unspecified.
- The function cannot contain calls to the *longjmp* or *setjmp* functions.

▲ C / C++ ▲

▼ C ▼

- The type of list items appearing in the **aligned** clause must be array or pointer.

▲ C ▲

▼ C++ ▼

- The function cannot contain any calls to **throw**.
- The type of list items appearing in the **aligned** clause must be array, pointer, reference to array, or reference to pointer.

▲ C++ ▲

- 1       • *proc-name* must not be a generic name, procedure pointer or entry name.
- 2       • Any **declare simd** directive must appear in the specification part of a subroutine subprogram,
- 3       function subprogram or interface body to which it applies.
- 4       • If a **declare simd** directive is specified in an interface block for a procedure, it must match a
- 5       **declare simd** directive in the definition of the procedure.
- 6       • If a procedure is declared via a procedure declaration statement, the procedure *proc-name* should
- 7       appear in the same specification.
- 8       • If a **declare simd** directive is specified for a procedure name with explicit interface and a
- 9       **declare simd** directive is also specified for the definition of the procedure then the two
- 10       **declare simd** directives must match. Otherwise the result is unspecified.
- 11       • Procedure pointers may not be used to access versions created by the **declare simd** directive.
- 12       • The type of list items appearing in the **aligned** clause must be **C\_PTR** or Cray pointer, or the
- 13       list item must have the **POINTER** or **ALLOCATABLE** attribute.

## Cross References

- 14       • **reduction** clause, see Section [2.14.3.6](#) on page [184](#).
- 15       • **linear** clause, see Section [2.14.3.7](#) on page [190](#).

## 2.8.3 Loop SIMD Construct

### Summary

The loop SIMD construct specifies that the iterations of one or more associated loops will be distributed across threads that already exist in the team and that the iterations executed by each thread can also be executed concurrently using SIMD instructions. The loop SIMD construct is a composite construct.



1

## Syntax

C / C++

```
#pragma omp for simd [clause[ [, ] clause] ... ] new-line
    for-loops
```

2

where *clause* can be any of the clauses accepted by the **for** or **simd** directives with identical meanings and restrictions.

3

C / C++

Fortran

```
!$omp do simd [clause[ [, ] clause] ... ]
    do-loops
[!$omp end do simd [nowait] ]
```

4

where *clause* can be any of the clauses accepted by the **simd** or **do** directives, with identical meanings and restrictions.

5

6

If an **end do simd** directive is not specified, an **end do simd** directive is assumed at the end of the *do-loops*.

7

Fortran

8

## Description

9

The loop SIMD construct will first distribute the iterations of the associated loop(s) across the implicit tasks of the parallel region in a manner consistent with any clauses that apply to the loop construct. The resulting chunks of iterations will then be converted to a SIMD loop in a manner consistent with any clauses that apply to the **simd** construct. The effect of any clause that applies to both constructs is as if it were applied to both constructs separately except the **collapse** clause, which is applied once.

10

11

12

13

14

## Restrictions

All restrictions to the loop construct and the **simd** construct apply to the loop SIMD construct. In addition, the following restriction applies:

- No **ordered** clause can be specified.
- A list item may appear in a **linear** or **firstprivate** clause but not both.

## Cross References

- loop construct, see Section 2.7.1 on page 55.
- **simd** construct, see Section 2.8.1 on page 70.
- Data attribute clauses, see Section 2.14.3 on page 172.

# 2.9 Tasking Constructs

## 2.9.1 task Construct

### Summary

The **task** construct defines an explicit task.

### Syntax

C / C++

The syntax of the **task** construct is as follows:

```
#pragma omp task [clause[ [, ] clause] ... ] new-line  
    structured-block
```

1 where *clause* is one of the following:

```
2     if (scalar-expression)  
3     final (scalar-expression)  
4     untied  
5     default (shared | none)  
6     mergeable  
7     private (list)  
8     firstprivate (list)  
9     shared (list)  
10    depend (dependence-type : list)
```



11 The syntax of the **task** construct is as follows:

```
!$omp task [clause[ [, ] clause] ... ]  
    structured-block  
!$omp end task
```

12 where *clause* is one of the following:

```
13     if (scalar-logical-expression)  
14     final (scalar-logical-expression)  
15     untied  
16     default (private | firstprivate | shared | none)  
17     mergeable  
18     private (list)  
19     firstprivate (list)  
20     shared (list)  
21     depend (dependence-type : list)
```



## 1           **Binding**

2           The binding thread set of the **task** region is the current team. A **task** region binds to the  
3           innermost enclosing **parallel** region.

## 4           **Description**

5           When a thread encounters a **task** construct, a task is generated from the code for the associated  
6           structured block. The data environment of the task is created according to the data-sharing attribute  
7           clauses on the **task** construct, per-data environment ICVs, and any defaults that apply.

8           The encountering thread may immediately execute the task, or defer its execution. In the latter case,  
9           any thread in the team may be assigned the task. Completion of the task can be guaranteed using  
10          task synchronization constructs. A **task** construct may be nested inside an outer task, but the  
11          **task** region of the inner task is not a part of the **task** region of the outer task.

12          When an **if** clause is present on a **task** construct, and the **if** clause expression evaluates to *false*,  
13          an undeferred task is generated, and the encountering thread must suspend the current task region,  
14          for which execution cannot be resumed until the generated task is completed. Note that the use of a  
15          variable in an **if** clause expression of a **task** construct causes an implicit reference to the variable  
16          in all enclosing constructs.

17          When a **final** clause is present on a **task** construct and the **final** clause expression evaluates  
18          to *true*, the generated task will be a final task. All **task** constructs encountered during execution of  
19          a final task will generate final and included tasks. Note that the use of a variable in a **final** clause  
20          expression of a **task** construct causes an implicit reference to the variable in all enclosing  
21          constructs.

22          The **if** clause expression and the **final** clause expression are evaluated in the context outside of  
23          the **task** construct, and no ordering of those evaluations is specified.

24          A thread that encounters a task scheduling point within the **task** region may temporarily suspend  
25          the **task** region. By default, a task is tied and its suspended task region can only be resumed by  
26          the thread that started its execution. If the **untied** clause is present on a **task** construct, any  
27          thread in the team can resume the **task** region after a suspension. The **untied** clause is ignored  
28          if a **final** clause is present on the same **task** construct and the **final** clause expression  
29          evaluates to *true*, or if a task is an included task.

30          The **task** construct includes a task scheduling point in the task region of its generating task,  
31          immediately following the generation of the explicit task. Each explicit **task** region includes a  
32          task scheduling point at its point of completion.

33          When a **mergeable** clause is present on a **task** construct, and the generated task is an  
34          undeferred task or an included task, the implementation may generate a merged task instead.

1 Note – When storage is shared by an explicit **task** region, it is the programmer’s responsibility to  
2 ensure, by adding proper synchronization, that the storage does not reach the end of its lifetime  
3 before the explicit **task** region completes its execution.

## 4 Restrictions

5 Restrictions to the **task** construct are as follows:

- 6 • A program that branches into or out of a **task** region is non-conforming.
- 7 • A program must not depend on any ordering of the evaluations of the clauses of the **task**  
8 directive, or on any side effects of the evaluations of the clauses.
- 9 • At most one **if** clause can appear on the directive.
- 10 • At most one **final** clause can appear on the directive.

▼ C / C++ ▼

- 11 • A throw executed inside a **task** region must cause execution to resume within the same **task**  
12 region, and the same thread that threw the exception must catch it.

▲ C / C++ ▲

▼ Fortran ▼

- 13 • Unsynchronized use of Fortran I/O statements by multiple tasks on the same unit has unspecified  
14 behavior

▲ Fortran ▲

## 15 Cross References

- 16 • Task scheduling constraints, see Section 2.9.5 on page 90.
- 17 • **depend** clause, see Section 2.12.9 on page 154.

## 18 2.9.2 **taskloop** Construct

### 19 Summary

20 The **taskloop** construct specifies that the iterations of one or more associated loops will be  
21 executed in parallel using OpenMP tasks. The iterations are distributed across tasks created by the  
22 construct and scheduled to be executed.

## Syntax

C / C++

The syntax of the **taskloop** construct is as follows:

```
#pragma omp taskloop [clause[[,] clause] ...] new-line
    for-loops
```

where *clause* is one of the following:

```
    shared (list)
    private (list)
    firstprivate (list)
    lastprivate (list)
    default (shared | none)
    grainsize (grain-size)
    num_tasks (num-tasks)
    collapse (n)
    if (scalar-expr)
    final (scalar-expr)
    untied
    mergeable
    nogroup
```

The **taskloop** directive places restrictions on the structure of all associated *for-loops*. Specifically, all associated *for-loops* must have canonical loop form (see Section 2.6 on page 52).

C / C++

1 The syntax of the **taskloop** construct is as follows:

```

!$omp taskloop [clause[[, clause] ...]
           do-loops
[!$omp end taskloop [nogroup]]

```

2 where *clause* is one of the following:

```

3     shared (list)
4     private (list)
5     firstprivate (list)
6     lastprivate (list)
7     default (private | firstprivate | shared | none)
8     grainsize (grain-size)
9     num_tasks (num-tasks)
10    collapse (n)
11    if (scalar-logical-expr)
12    final (scalar-logical-expr)
13    untied
14    mergeable

```

15 If **end taskloop** directive is not specified, an **end taskloop** directive is assumed at the end of  
16 the *do-loops*.

17 All associated *do-loops* must be *do-constructs* as defined by the Fortran standard. If an **end do**  
18 directive follows a *do-construct* in which several loop statements share a **DO** termination statement,  
19 then the directive can only be specified for the outermost of these **DO** statements.

20 If any of the loop iteration variables would otherwise be shared, they are implicitly made private for  
21 the loop-iteration tasks created by the **taskloop** construct. Unless the loop iteration variables are  
22 specified in a **lastprivate** clause on the **taskloop** construct, their values after the loop are  
23 unspecified.

## 1            **Binding**

2            The binding thread set of the **taskloop** construct is the current team. A **taskloop** region binds  
3            to the innermost enclosing **parallel** region.

## 4            **Description**

5            When a thread encounters a **taskloop** construct, the construct partitions the associated loops into  
6            tasks for parallel execution of the loops' iterations. The data environment of the created tasks is  
7            created according to the data-sharing attribute clauses on the **taskloop** construct, per-data  
8            environment ICVs, and any defaults that apply. The order of the creation of the loop tasks is  
9            unspecified.

10           If a **grainsize** clause is present on the **taskloop** construct, the number of logical loop  
11           iterations assigned to each created task is larger than or equal to the value of the *grain-size*  
12           expression, but less than two times the value of the *grain-size* expression and less than or equal to  
13           the number of logical loop iterations. If **num\_tasks** is specified, the **taskloop** construct  
14           creates as many tasks as the minimum of the *num-tasks* expression and the number of logical loop  
15           iterations. If neither a **grainsize** nor **num\_tasks** clause is present, the number of loop tasks  
16           created and the number of logical loop iterations assigned to these tasks is implementation defined.

17           The **collapse** clause may be used to specify how many loops are associated with the **taskloop**  
18           construct. The parameter of the **collapse** clause must be a constant positive integer expression.  
19           If no **collapse** clause is present, the only loop that is associated with the **taskloop** construct is  
20           the one that immediately follows the **taskloop** directive.

21           If more than one loop is associated with the **taskloop** construct, then the iterations of all  
22           associated loops are collapsed into one larger iteration space that is then divided according to the  
23           **grainsize** and **num\_tasks** clauses. The sequential execution of the iterations in all associated  
24           loops determines the order of the iterations in the collapsed iteration space.

25           The iteration count for each associated loop is computed before entry to the outermost loop. If  
26           execution of any associated loop changes any of the values used to compute any of the iteration  
27           counts, then the behavior is unspecified. The integer type (or kind, for Fortran) used to compute the  
28           iteration count for the collapsed loop is implementation defined.

29           When an **if** clause is present on a **taskloop** construct, and if the **if** clause expression evaluates  
30           to *false*, undeferred tasks are generated. The use of a variable in an **if** clause expression of a  
31           **taskloop** construct causes an implicit reference to the variable in all enclosing constructs.

32           When a **final** clause is present on a **taskloop** construct and the **final** clause expression  
33           evaluates to *true*, the generated tasks will be final tasks. The use of a variable in a **final** clause  
34           expression of a **taskloop** construct causes an implicit reference to the variable in all enclosing  
35           constructs.

36           If the **untied** clause is specified, all tasks created by the **taskloop** construct are untied tasks.



1 When a **mergeable** clause is present on a **taskloop** construct, and the generated tasks are  
2 undeferred or included tasks, the implementation may generate merged tasks instead.

3 By default, the **taskloop** construct executes as if it was enclosed in a **taskgroup** construct  
4 with no statements or directives outside of the **taskloop** construct. Thus, the **taskloop**  
5 construct creates an implicit **taskgroup** region. If the **nogroup** clause is present, no implicit  
6 **taskgroup** region is created.

---

7 **Note** – When storage is shared by a **taskloop** region, the programmer must ensure, by adding  
8 proper synchronization, that the storage does not reach the end of its lifetime before the **taskloop**  
9 region and its descendant tasks complete their execution.

---

## 10 **Restrictions**

11 The restrictions of the **taskloop** construct are as follows:

- 12 • A program that branches into or out of a **taskloop** region is non-conforming.
- 13 • All loops associated with the **taskloop** construct must be perfectly nested; that is, there must  
14 be no intervening code nor any OpenMP directive between any two loops.
- 15 • At most one **grainsize** clause can appear on a **taskloop** directive.
- 16 • At most one **num\_tasks** clause can appear on a **taskloop** directive.
- 17 • The **grainsize** clause and **num\_tasks** clause are mutually exclusive and may not appear on  
18 the same **taskloop** directive.
- 19 • At most one **collapse** clause can appear on a **taskloop** directive.
- 20 • At most one **if** clause can appear on the directive.
- 21 • At most one **final** clause can appear on the directive.

## 22 **Cross References**

- 23 • **task** construct, Section [2.9.1](#) on page [80](#).
- 24 • **taskgroup** construct, Section [2.12.5](#) on page [140](#).
- 25 • Data-sharing attribute clauses, Section [2.14.3](#) on page [172](#).

## 1 2.9.3 taskloop simd Construct

### 2 Summary

3 The **taskloop simd** construct specifies a loop that can be executed concurrently using SIMD  
4 instructions and that those iterations will also be executed in parallel using OpenMP tasks.

### 5 Syntax

C / C++

6 The syntax of the **taskloop simd** construct is as follows:

```
#pragma omp taskloop simd [clause[[,] clause] ...] new-line  
for-loops
```

7 where *clause* can be any of the clauses accepted by the **taskloop** or **simd** directives with  
8 identical meanings and restrictions.

C / C++

Fortran

9 The syntax of the **taskloop simd** construct is as follows:

```
!$omp taskloop simd [clause[[,] clause] ...]  
do-loops  
[!$omp end taskloop simd [nogroup]]
```

10 where *clause* can be any of the clauses accepted by the **taskloop** or **simd** directives with  
11 identical meanings and restrictions.

12 If **end taskloop simd** directive is not specified, an **end taskloop simd** directive is  
13 assumed at the end of the *do-loops*.

Fortran

### 14 Binding

15 The binding thread set of the **taskloop simd** construct is the current team. A **taskloop simd**  
16 region binds to the innermost enclosing parallel region.

## Description

The **taskloop simd** construct will first distribute the iterations of the associated loop(s) across tasks in a manner consistent with any clauses that apply to the **taskloop** construct. The resulting tasks will then be converted to a SIMD loop in a manner consistent with any clauses that apply to the **simd** construct. The effect of any clause that applies to both constructs is as if it were applied to both constructs separately except the **collapse** clause, which is applied once.

## Restrictions

- The restrictions for the **taskloop** and **simd** constructs apply.
- No **reduction** clause can be specified.

## Cross References


- **taskloop** construct, see Section 2.9.2 on page 83.
- **simd** construct, see Section 2.8.1 on page 70.
- Data-sharing attribute clauses, see Section 2.14.3 on page 172.

## 2.9.4 taskyield Construct

### Summary

The **taskyield** construct specifies that the current task can be suspended in favor of execution of a different task. The **taskyield** construct is a stand-alone directive.

### Syntax



The syntax of the **taskyield** construct is as follows:

```
!$omp taskyield
```

Fortran

## 1 Binding

2 A **taskyield** region binds to the current task region. The binding thread set of the **taskyield**  
3 region is the current team.

## 4 Description

5 The **taskyield** region includes an explicit task scheduling point in the current task region.

## 6 Cross References

- 7
- Task scheduling, see Section [2.9.5](#) on page [90](#).

## 8 2.9.5 Task Scheduling

9 Whenever a thread reaches a task scheduling point, the implementation may cause it to perform a  
10 task switch, beginning or resuming execution of a different task bound to the current team. Task  
11 scheduling points are implied at the following locations:

- 12
- the point immediately following the generation of an explicit task

13

  - after the point of completion of a **task** region

14

  - in a **taskyield** region

15

  - in a **taskwait** region

16

  - at the end of a **taskgroup** region

17

  - in an implicit and explicit **barrier** region

18

  - the point immediately following the generation of a **target** region

19

  - at the beginning and end of a **target data** region

20

  - in a **target update** region

21 When a thread encounters a task scheduling point it may do one of the following, subject to the  
22 *Task Scheduling Constraints* (below):

- 23
- begin execution of a tied task bound to the current team

- 1           ● resume any suspended task region, bound to the current team, to which it is tied
- 2           ● begin execution of an untied task bound to the current team
- 3           ● resume any suspended untied task region bound to the current team.

4           If more than one of the above choices is available, it is unspecified as to which will be chosen.

5           *Task Scheduling Constraints* are as follows:

- 6           1. An included task is executed immediately after generation of the task.
- 7           2. Scheduling of new tied tasks is constrained by the set of task regions that are currently tied to the  
8           thread, and that are not suspended in a **barrier** region. If this set is empty, any new tied task  
9           may be scheduled. Otherwise, a new tied task may be scheduled only if it is a descendent task of  
10          every task in the set.
- 11          3. A dependent task shall not be scheduled until its task dependences are fulfilled.
- 12          4. When an explicit task is generated by a construct containing an **if** clause for which the  
13          expression evaluated to *false*, and the previous constraints are already met, the task is executed  
14          immediately after generation of the task.

15          A program relying on any other assumption about task scheduling is non-conforming.

---

16          **Note** – Task scheduling points dynamically divide task regions into parts. Each part is executed  
17          uninterrupted from start to end. Different parts of the same task region are executed in the order in  
18          which they are encountered. In the absence of task synchronization constructs, the order in which a  
19          thread executes parts of different schedulable tasks is unspecified.

20          A correct program must behave correctly and consistently with all conceivable scheduling  
21          sequences that are compatible with the rules above.

22          For example, if **threadprivate** storage is accessed (explicitly in the source code or implicitly  
23          in calls to library routines) in one part of a task region, its value cannot be assumed to be preserved  
24          into the next part of the same task region if another schedulable task exists that modifies it.

25          As another example, if a lock acquire and release happen in different parts of a task region, no  
26          attempt should be made to acquire the same lock in any part of another task that the executing  
27          thread may schedule. Otherwise, a deadlock is possible. A similar situation can occur when a  
28          critical region spans multiple parts of a task and another schedulable task contains a critical region  
29          with the same name.

30          The use of threadprivate variables and the use of locks or critical sections in an explicit task with an  
31          **if** clause must take into account that when the **if** clause evaluates to *false*, the task is executed  
32          immediately, without regard to *Task Scheduling Constraint 2*.

---

## 1 2.10 Device Constructs

### 2 2.10.1 target data Construct

#### 3 Summary

4 Map variables to a device data environment for the extent of the region.

#### 5 Syntax

▼ C / C++ ▼

6 The syntax of the **target data** construct is as follows:

```
#pragma omp target data [clause[ [, ] clause] ... ] new-line  
structured-block
```

7 where *clause* is one of the following:

```
device (integer-expression)  
map ([[map-type-modifier[,]] map-type : ] list)  
if (scalar-expression)
```

▲ C / C++ ▲

▼ Fortran ▼

11 The syntax of the **target data** construct is as follows:

```
!$omp target data [clause[ [, ] clause] ... ]  
structured-block  
!$omp end target data
```

12 where *clause* is one of the following:

```
device (scalar-integer-expression)  
map ([[map-type-modifier[,]] map-type : ] list)  
if (scalar-logical-expression)
```

16 The **end target data** directive denotes the end of the **target data** construct.

▲ Fortran ▲

## 1        **Binding**

2        The binding task region for a **target data** construct is the encountering task. The target region  
3        binds to the enclosing parallel or task region.

## 4        **Description**

5        When a **target data** construct is encountered, the encountering task executes the region. If  
6        there is no **device** clause, the default device is determined by the *default-device-var* ICV.  
7        Variables are mapped for the extent of the region, according to any data-mapping clauses, from the  
8        data environment of the encountering task to the device data environment. When an **if** clause is  
9        present and the **if** clause expression evaluates to *false*, the device is the host.

## 10       **Restrictions**

- 11       • A program must not depend on any ordering of the evaluations of the clauses of the  
12       **target data** directive, or on any side effects of the evaluations of the clauses.
- 13       • At most one **device** clause can appear on the directive. The **device** expression must evaluate  
14       to a non-negative integer value.
- 15       • At most one **if** clause can appear on the directive.
- 16       • A *map-type* in a **map** clause must be **to**, **from**, **tofrom** or **alloc**.
- 17       • At least one **map** clause must appear on the directive.

## 18       **Cross References**

- 19       • **map** clause, see Section [2.14.5](#) on page [195](#).
- 20       • *default-device-var*, see Section [2.3](#) on page [35](#).

## 21    **2.10.2 target Construct**

### 22       **Summary**

23       Map variables to a device data environment and execute the construct on that device.

### 24       **Syntax**



25       The syntax of the **target** construct is as follows:

```
#pragma omp target [clause[ [, ] clause] ... ] new-line
structured-block
```

1 where *clause* is one of the following:

```
2     device (integer-expression)
3     map ([[map-type-modifier[,]] map-type : ] list)
4     nowait
5     depend (dependence-type : list)
6     if (scalar-expression)
```



7 The syntax of the **target** construct is as follows:

```
!$omp target [clause[ [, ] clause] ... ]
structured-block
!$omp end target
```

8 where *clause* is one of the following:

```
9     device (scalar-integer-expression)
10    map ([[map-type-modifier[,]] map-type : ] list)
11    nowait
12    depend (dependence-type : list)
13    if (scalar-logical-expression)
```

14 The **end target** directive denotes the end of the **target** construct



## 15 Binding

16 The binding task for a **target** construct is the encountering task. The target region binds to the  
17 enclosing parallel or task region.



## Description

The **target** construct provides a superset of the functionality and restrictions provided by the **target data** directive.

The functionality added to the **target** directive is the inclusion of an executable region to be executed by a device. That is, the **target** directive is an executable directive.

The **target** construct executes as if it was enclosed in a **task** construct with no statements or directives outside of the **target** construct. The generated task is a *target task*.

A *target task* is executed immediately and waits at a task scheduling point for the device to complete the **target** region. The encountering thread becomes available to execute other tasks at that task scheduling point. If the *target task* is undeferred then the encountering thread must suspend the current task region, for which execution cannot be resumed until the generated *target task* is complete.

By default the generated task is undeferred. When a **nowait** clause is present, the current task may resume execution before the generated task completes its execution.

If a **depend** clause is present, then it is treated as if it had appeared on the implicit **task** construct that encloses the **target** construct.

When an **if** clause is present and the **if** clause expression evaluates to *false*, the **target** region is executed by the host device in the host data environment.

## Restrictions

- If a **target**, **target update**, **target data**, **target enter data**, or **target exit data** construct appears within a **target** region the behavior is unspecified.
- The result of an **omp\_set\_default\_device**, **omp\_get\_default\_device**, or **omp\_get\_num\_devices** routine called within a target region is unspecified.
- The effect of an access to a **threadprivate** variable in a target region is unspecified.
- A variable referenced in a **target** construct that is not declared in the construct is implicitly treated as if it had appeared in a **map** clause with a *map-type* of **tofrom**.
- A variable referenced in a **target** region but not the target construct that is not declared in the target region must appear in a **declare target** directive.
- The restrictions for the **task** construct apply.
- A *map-type* in a **map** clause must be **to**, **from**, **tofrom** or **alloc**.

C++

- 1
- A throw executed inside a **target** region must cause execution to resume within the same **target** region, and the same thread that threw the exception must catch it.

C++

### Cross References

- 3
- **target data** construct, see Section 2.10.1 on page 92.
  - **task** construct, see Section 2.9.1 on page 80.
  - **task** scheduling constraints, see Section 2.9.5 on page 90
  - *default-device-var*, see Section 2.3 on page 35.
  - **map** clause, see Section 2.14.5 on page 195.

## 9 2.10.3 target update Construct

### Summary

10 The **target update** directive makes the corresponding list items in the device data environment  
11 consistent with their original list items, according to the specified motion clauses. The  
12 **target update** construct is a stand-alone directive.  
13

### Syntax

C / C++

14 The syntax of the **target update** construct is as follows:  
15

```
#pragma omp target update clause[ [, ] clause] ... ] new-line
```

16 where *clause* is either *motion-clause* or one of the following:

17 **device** (*integer-expression*)  
18 **if** (*scalar-expression*)  
19 **nowait**  
20 **depend** (*dependence-type* : *list*)

1 and *motion-clause* is one of the following:

2 **to** (*list*)

3 **from** (*list*)



4 The syntax of the **target update** construct is as follows:

```
!$omp target update clause[ [, ] clause] ... ]
```

5 where *clause* is either *motion-clause* or one of the following:

6 **device** (*scalar-integer-expression*)

7 **if** (*scalar-logical-expression*)

8 **nowait**

9 **depend** (*dependence-type* : *list*)

10 and *motion-clause* is one of the following:

11 **to** (*list*)

12 **from** (*list*)



### 13 **Binding**

14 The binding task for a **target update** construct is the encountering task. The **target update**  
15 directive is a stand-alone directive.

## Description

For each list item in a **to** or **from** clause there is a corresponding list item and an original list item. If the corresponding list item is not present in the device data environment then no assignment occurs to or from the original list item. Otherwise, each corresponding list item in the device data environment has an original list item in the current task's data environment.

For each list item in a **from** clause the value of the corresponding list item is assigned to the original list item.

For each list item in a **to** clause the value of the original list item is assigned to the corresponding list item.

The list items that appear in the **to** or **from** clauses may include array sections.

The **target update** construct executes as if it was enclosed in a **task** construct with no statements or directives outside of the **target update** construct. The generated task is a *target task*.

A *target task* is executed immediately and waits at a task scheduling point for the device to complete the **target update** region. The encountering thread becomes available to execute other tasks at that task scheduling point. If the *target task* is undeferred then the encountering thread must suspend the current task region, for which execution cannot be resumed until the generated *target task* is complete.

By default the generated task is undeferred. When a **nowait** clause is present, the current task may resume execution before the generated task completes its execution.

If a **depend** clause is present, then it is treated as if it had appeared on the implicit **task** construct that encloses the **target update** construct.

The device is specified in the **device** clause. If there is no **device** clause, the device is determined by the *default-device-var* ICV. When an **if** clause is present and the **if** clause expression evaluates to *false* then no assignments occur.

## Restrictions

- A program must not depend on any ordering of the evaluations of the clauses of the **target update** directive, or on any side effects of the evaluations of the clauses.
- At least one *motion-clause* must be specified.
- If a list item is an array section it must specify contiguous storage.
- A variable that is part of another variable (such as a field of a structure) but is not an array element or an array section cannot appear as a list item in a clause of a **target update** construct.
- A list item can only appear in a **to** or **from** clause, but not both.

- 1 • A list item in a **to** or **from** clause must have a mappable type
- 2 • At most one **device** clause can appear on the directive. The **device** expression must evaluate
- 3 to a non-negative integer value.
- 4 • At most one **if** clause can appear on the directive.
- 5 • The restrictions for the **task** construct apply.

## 6 **Cross References**

- 7 • *default-device-var*, see Section 2.3 on page 35.
- 8 • **target data**, see Section 2.10.1 on page 92.
- 9 • Array sections, Section 2.4 on page 43
- 10 • **task** construct, see Section 2.9.1 on page 80.
- 11 • **task** scheduling constraints, see Section 2.9.5 on page 90

## 12 **2.10.4 declare target Directive**

### 13 **Summary**

14 The **declare target** directive specifies that variables, functions (C, C++ and Fortran), and  
15 subroutines (Fortran) are mapped to a device. The **declare target** directive is a declarative  
16 directive.

### 17 **Syntax**

▼  C / C++ 

18 The syntax of the **declare target** directive is as follows:

```
#pragma omp declare target new-line  
declaration-definition-seq  
#pragma omp end declare target new-line
```

▲  C / C++   
▼  Fortran 

19 The syntax of the **declare target** directive is as follows:

20 For variables, functions and subroutines:

```
!$omp declare target (list)
```

1 where *list* is a comma-separated list of named variables, procedure names and named common  
2 blocks. Common block names must appear between slashes.

3 For functions and subroutines:

```
!$omp declare target
```

▲────────────────── Fortran ───────────────────▲

4 **Description**

▼────────────────── C / C++ ───────────────────▼

5 Variable declarations at file or namespace scope that appear between the **declare target** and  
6 **end declare target** directives form an implicit list where each list item is the variable name.

7 Function declarations at file, namespace or class scope that appear between the **declare target**  
8 and **end declare target** form an implicit list where each list item is the function name.

▲────────────────── C / C++ ───────────────────▲

▼────────────────── Fortran ───────────────────▼

9 If a **declare target** does not have an explicit list, then an implicit list of one item is formed  
10 from the name of the enclosing subroutine subprogram, function subprogram or interface body to  
11 which it applies.

▲────────────────── Fortran ───────────────────▲

12 If a list item is a function (C, C++, Fortran) or subroutine (Fortran) then a device-specific version of  
13 the routine is created that can be called from a target region.

14 If a list item is a variable then the original variable is mapped to a corresponding variable in the  
15 device data environment of all devices as if it had appeared in a **map** clause with a *map-type* to on  
16 the implicit **target data** construct for each device. The list item is never removed from those  
17 device data environments as if its reference count is initialized to positive infinity.

## Restrictions

- A threadprivate variable cannot appear in a **declare target** directive.
- A variable declared in a **declare target** directive must have a mappable type

C / C++

- All declarations and definitions for a function must have a **declare target** directive if one is specified for any of them. Otherwise, the result is unspecified

C / C++

Fortran

- If a list item is a procedure name, it must not be a generic name, procedure pointer or entry name.
- Any **declare target** directive with a list can only appear in a specification part of a subroutine subprogram, function subprogram, program or module.
- Any **declare target** directive without a list can only appear in a specification part of a subroutine subprogram, function subprogram or interface body to which it applies.
- If a **declare target** directive is specified in an interface block for a procedure, it must match a **declare target** directive in the definition of the procedure.
- If an external procedure is a type-bound procedure of a derived type and a **declare target** directive is specified in the definition of the external procedure, such a directive must appear in the interface block that is accessible to the derived type definition.
- If any procedure is declared via a procedure declaration statement that is not in the type-bound procedure part of a derived-type definition, any **declare target** with the procedure name must appear in the same specification part.
- A variable that is part of another variable (as an array or structure element) cannot appear in a **declare target** directive.
- The **declare target** directive must appear in the declaration section of a scoping unit in which the common block or variable is declared. Although variables in common blocks can be accessed by use association or host association, common block names cannot. This means that a common block name specified in a **declare target** directive must be declared to be a common block in the same scoping unit in which the **declare target** directive appears.
- If a **declare target** directive specifying a common block name appears in one program unit, then such a directive must also appear in every other program unit that contains a **COMMON** statement specifying the same name. It must appear after the last such **COMMON** statement in the program unit.
- If a **declare target** variable or a **declare target** common block is declared with the **BIND** attribute, the corresponding C entities must also be specified in a **declare target** directive in the C program.

- 1 • A blank common block cannot appear in a **declare target** directive.
- 2 • A variable can only appear in a **declare target** directive in the scope in which it is declared.
- 3 It must not be an element of a common block or appear in an **EQUIVALENCE** statement.
- 4 • A variable that appears in a **declare target** directive must be declared in the Fortran scope
- 5 of a module or have the **SAVE** attribute, either explicitly or implicitly.

Fortran

## 6 2.10.5 teams Construct

### 7 Summary

8 The **teams** construct creates a league of thread teams and the master thread of each team executes  
9 the region.

### 10 Syntax

C / C++

11 The syntax of the **teams** construct is as follows:

```
#pragma omp teams [clause[ [, ] clause] ... ] new-line
structured-block
```

12 where *clause* is one of the following:

13 **num\_teams** (*integer-expression*)

14 **thread\_limit** (*integer-expression*)

15 **default** (**shared** | **none**)

16 **private** (*list*)

17 **firstprivate** (*list*)

18 **shared** (*list*)

19 **reduction** (*reduction-identifier* : *list*)

C / C++

Fortran

20 The syntax of the **teams** construct is as follows:



```
!$omp teams [clause[ [, ] clause] ... ]
structured-block
!$omp end teams
```

1 where *clause* is one of the following:

2 **num\_teams** (*scalar-integer-expression*)

3 **thread\_limit** (*scalar-integer-expression*)

4 **default** (**shared** | **firstprivate** | **private** | **none**)

5 **private** (*list*)

6 **firstprivate** (*list*)

7 **shared** (*list*)

8 **reduction** (*reduction-identifier* : *list*)

9 The **end teams** directive denotes the end of the **teams** construct.

Fortran

## 10 Binding

11 The binding thread set for a **teams** region is the encountering thread.

## 12 Description

13 When a thread encounters a **teams** construct, a league of thread teams is created and the master  
14 thread of each thread team executes the **teams** region.

15 The number of teams created is implementation defined, but is less than or equal to the value  
16 specified in the **num\_teams** clause.

17 The maximum number of threads participating in the contention group that each team initiates is  
18 implementation defined, but is less than or equal to the value specified in the **thread\_limit**  
19 clause.

20 Once the teams are created, the number of teams remains constant for the duration of the **teams**  
21 region.

22 Within a **teams** region, team numbers uniquely identify each team. Team numbers are consecutive  
23 whole numbers ranging from zero to one less than the number of teams. A thread may obtain its  
24 own team number by a call to the **omp\_get\_team\_num** library routine.

25 The threads other than the master thread do not begin execution until the master thread encounters a  
26 **parallel** region.

1 After the teams have completed execution of the **teams** region, the encountering thread resumes  
2 execution of the enclosing **target** region.

3 There is no implicit barrier at the end of a **teams** construct.

## 4 **Restrictions**

5 Restrictions to the **teams** construct are as follows:

- 6 • A program that branches into or out of a **teams** region is non-conforming.
- 7 • A program must not depend on any ordering of the evaluations of the clauses of the **teams**  
8 directive, or on any side effects of the evaluation of the clauses.
- 9 • At most one **thread\_limit** clause can appear on the directive. The **thread\_limit**  
10 expression must evaluate to a positive integer value.
- 11 • At most one **num\_teams** clause can appear on the directive. The **num\_teams** expression must  
12 evaluate to a positive integer value.
- 13 • If specified, a **teams** construct must be contained within a **target** construct. That **target**  
14 construct must contain no statements, declarations or directives outside of the **teams** construct.
- 15 • **distribute**, **parallel**, **parallel sections**, **parallel workshare**, and the  
16 parallel loop and parallel loop SIMD constructs are the only OpenMP constructs that can be  
17 closely nested in the **teams** region.

## 18 **Cross References**

- 19 • **default**, **shared**, **private**, **firstprivate**, and **reduction** clauses, see  
20 Section [2.14.3](#) on page [172](#).
- 21 • **omp\_get\_num\_teams** routine, see Section [3.2.26](#) on page [235](#).
- 22 • **omp\_get\_team\_num** routine, see Section [3.2.27](#) on page [237](#).

## 23 **2.10.6 distribute Construct**

### 24 **Summary**

25 The **distribute** construct specifies that the iterations of one or more loops will be executed by  
26 the thread teams in the context of their implicit tasks. The iterations are distributed across the  
27 master threads of all teams that execute the **teams** region to which the **distribute** region binds.

## Syntax

C / C++

The syntax of the **distribute** construct is as follows:

```
#pragma omp distribute [clause[ [, ] clause] ... ] new-line  
for-loops
```

Where *clause* is one of the following:

```
private (list)  
firstprivate (list)  
collapse (n)  
dist_schedule (kind[ , chunk_size])
```

All associated *for-loops* must have the canonical form described in Section 2.6 on page 52.

C / C++

Fortran

The syntax of the **distribute** construct is as follows:

```
!$omp distribute [clause[ [, ] clause] ... ]  
do-loops  
[!$omp end distribute]
```

Where *clause* is one of the following:

```
private (list)  
firstprivate (list)  
collapse (n)  
dist_schedule (kind[ , chunk_size])
```

If an **end distribute** directive is not specified, an **end distribute** directive is assumed at the end of the *do-loops*.

All associated *do-loops* must be *do-constructs* as defined by the Fortran standard. If an **end do** directive follows a *do-construct* in which several loop statements share a **DO** termination statement, then the directive can only be specified for the outermost of these **DO** statements.

Fortran

## 1 Binding

2 The binding thread set for a **distribute** region is the set of master threads created by a **teams**  
3 construct. A **distribute** region binds to the innermost enclosing **teams** region. Only the  
4 threads executing the binding **teams** region participate in the execution of the loop iterations.

## 5 Description

6 The **distribute** construct is associated with a loop nest consisting of one or more loops that  
7 follow the directive.

8 There is no implicit barrier at the end of a **distribute** construct.

9 The **collapse** clause may be used to specify how many loops are associated with the  
10 **distribute** construct. The parameter of the **collapse** clause must be a constant positive  
11 integer expression. If no **collapse** clause is present, the only loop that is associated with the  
12 **distribute** construct is the one that immediately follows the **distribute** construct.

13 If more than one loop is associated with the **distribute** construct, then the iteration of all  
14 associated loops are collapsed into one larger iteration space. The sequential execution of the  
15 iterations in all associated loops determines the order of the iterations in the collapsed iteration  
16 space.

17 If **dist\_schedule** is specified, *kind* must be **static**. If specified, iterations are divided into  
18 chunks of size *chunk\_size*, chunks are assigned to the teams of the league in a round-robin fashion  
19 in the order of the team number. When no *chunk\_size* is specified, the iteration space is divided into  
20 chunks that are approximately equal in size, and at most one chunk is distributed to each team of  
21 the league. Note that the size of the chunks is unspecified in this case.

22 When no **dist\_schedule** clause is specified, the schedule is implementation defined.

## 23 Restrictions

24 Restrictions to the **distribute** construct are as follows:

- 25 • The **distribute** construct inherits the restrictions of the loop construct.
- 26 • A **distribute** construct must be closely nested in a **teams** region.

## 27 Cross References

- 28 • loop construct, see Section [2.7.1](#) on page [55](#).
- 29 • **teams** construct, see Section [2.10.5](#) on page [102](#)

## 1 2.10.7 **distribute simd Construct**

### 2 **Summary**

3 The **distribute simd** construct specifies a loop that will be distributed across the master  
4 threads of the **teams** region and executed concurrently using SIMD instructions. The  
5 **distribute simd** construct is a composite construct.

### 6 **Syntax**

7 The syntax of the **distribute simd** construct is as follows:

▼ C / C++ ▼

```
#pragma omp distribute simd [clause[ [, ] clause] ... ]  
    for-loops
```

8 where *clause* can be any of the clauses accepted by the **distribute** or **simd** directives with  
9 identical meanings and restrictions.

▲ C / C++ ▲

```
!$omp distribute simd [clause[ [, ] clause]... ]
    do-loops
[!$omp end distribute simd]
```

1 where *clause* can be any of the clauses accepted by the **distribute** or **simd** directives with  
 2 identical meanings and restrictions.

3 If an **end distribute simd** directive is not specified, an **end distribute simd** directive is  
 4 assumed at the end of the *do-loops*

## 5 Description

6 The **distribute simd** construct will first distribute the iterations of the associated loop(s)  
 7 according to the semantics of the **distribute** construct and any clauses that apply to the  
 8 distribute construct. The resulting chunks of iterations will then be converted to a SIMD loop in a  
 9 manner consistent with any clauses that apply to the **simd** construct. The effect of any clause that  
 10 applies to both constructs is as if it were applied to both constructs separately except the  
 11 **collapse** clause, which is applied once.

## 12 Restrictions

- 13 • The restrictions for the **distribute** and **simd** constructs apply.
- 14 • A list item may appear in a **linear** or **firstprivate** clause but not both.
- 15 • A list item may appear in a **linear** or **lastprivate** clause but not both.

## 16 Cross References

- 17 • **simd** construct, see Section [2.8.1](#) on page [70](#).
- 18 • **distribute** construct, see Section [2.10.6](#) on page [104](#).
- 19 • Data attribute clauses, see Section [2.14.3](#) on page [172](#).

## 1 2.10.8 Distribute Parallel Loop Construct

### 2 Summary

3 The distribute parallel loop construct specifies a loop that can be executed in parallel by multiple  
4 threads that are members of multiple teams. The distribute parallel loop construct is a composite  
5 construct.

### 6 Syntax

7 The syntax of the distribute parallel loop construct is as follows:

▼ C / C++ ▼

```
#pragma omp distribute parallel for [clause[ [, ] clause] ... ]  
for-loops
```

8 where *clause* can be any of the clauses accepted by the **distribute** or parallel loop directives  
9 with identical meanings and restrictions.

▲ C / C++ ▲

▼ Fortran ▼

```
!$omp distribute parallel do [clause[ [, ] clause] ... ]  
do-loops  
[!$omp end distribute parallel do]
```

10 where *clause* can be any of the clauses accepted by the **distribute** or parallel loop directives  
11 with identical meanings and restrictions.

12 If an **end distribute parallel do** directive is not specified, an  
13 **end distribute parallel do** directive is assumed at the end of the *do-loops*.

▲ Fortran ▲

### 14 Description

15 The distribute parallel loop construct will first distribute the iterations of the associated loop(s) into  
16 chunks according to the semantics of the **distribute** construct and any clauses that apply to the  
17 **distribute** construct. Each of these chunks will form a loop. Each resulting loop will then be  
18 distributed across the threads within the teams region to which the **distribute** construct binds  
19 in a manner consistent with any clauses that apply to the parallel loop construct. The effect of any  
20 clause that applies to both constructs is as if it were applied to both constructs separately except the  
21 **collapse** clause, which is applied once.

## Restrictions

- The restrictions for the **distribute** and parallel loop constructs apply.
- A list item may appear in a **linear** or **firstprivate** clause but not both.
- A list item may appear in a **linear** or **lastprivate** clause but not both.

## Cross References

- **distribute** construct, see Section 2.10.6 on page 104.
- Parallel loop construct, see Section 2.11.1 on page 117.
- Data attribute clauses, see Section 2.14.3 on page 172.

## 2.10.9 Distribute Parallel Loop SIMD Construct

### Summary

The distribute parallel loop SIMD construct specifies a loop that can be executed concurrently using SIMD instructions in parallel by multiple threads that are members of multiple teams. The distribute parallel loop SIMD construct is a composite construct.

### Syntax

C / C++

The syntax of the distribute parallel loop SIMD construct is as follows:

```
#pragma omp distribute parallel for simd [clause[ [, ] clause] ... ]  
for-loops
```

where *clause* can be any of the clauses accepted by the **distribute** or parallel loop SIMD directives with identical meanings and restrictions

C / C++



1 The syntax of the distribute parallel loop SIMD construct is as follows:

```

2      !$omp distribute parallel do simd [clause[ [, ] clause] ... ]
3          do-loops
4      [!$omp end distribute parallel do simd]

```

2 where *clause* can be any of the clauses accepted by the **distribute** or parallel loop SIMD  
3 directives with identical meanings and restrictions.

4 If an **end distribute parallel do simd** directive is not specified, an  
5 **end distribute parallel do simd** directive is assumed at the end of the *do-loops*.

## 6 Description

7 The distribute parallel loop SIMD construct will first distribute the iterations of the associated  
8 loop(s) according to the semantics of the **distribute** construct and any clauses that apply to the  
9 **distribute** construct. The resulting loops will then be distributed across the threads contained  
10 within the **teams** region to which the **distribute** construct binds in a manner consistent with  
11 any clauses that apply to the parallel loop construct. The resulting chunks of iterations will then be  
12 converted to a SIMD loop in a manner consistent with any clauses that apply to the **simd** construct.  
13 The effect of any clause that applies to both constructs is as if it were applied to both constructs  
14 separately except the **collapse** clause, which is applied once.

## 15 Restrictions

- 16 • The restrictions for the **distribute** and parallel loop SIMD constructs apply.
- 17 • A list item may appear in a **linear** or **firstprivate** clause but not both.
- 18 • A list item may appear in a **linear** or **lastprivate** clause but not both.

## 19 Cross References

- 20 • **distribute** construct, see Section [2.10.6](#) on page [104](#).
- 21 • Parallel loop SIMD construct, see Section [2.11.4](#) on page [121](#).
- 22 • Data attribute clauses, see Section [2.14.3](#) on page [172](#).

## 1 2.10.10 target enter data Construct

### 2 Summary

3 The **target enter data** directive specifies that variables are mapped to a device data  
4 environment. The **target enter data** directive is a stand-alone directive.

### 5 Syntax

C / C++

6 The syntax of the **target enter data** construct is as follows:

```
#pragma omp target enter data [clause[[,] clause]...] new-line
```

7 where *clause* is one of the following:

```
8     device (integer-expression)  
9     map ([ [map-type-modifier[,]] map-type : ] list)  
10    if (scalar-expression)  
11    depend (dependence-type : list)  
12    nowait
```

C / C++

Fortran

13 The syntax of the **target enter data** is as follows:

```
!$omp target enter data [clause[[,] clause]...]
```

14 where *clause* is one of the following:

```
15     device (scalar-integer-expression)  
16     map ([ [map-type-modifier[,]] map-type : ] list)  
17     if (scalar-logical-expression)  
18     depend (dependence-type : list)  
19     nowait
```

Fortran

## 1            **Binding**

2            The binding task for a **target enter data** construct is the encountering task.

## 3            **Description**

4            When a **target enter data** construct is encountered, the list items are mapped to the device  
5            data environment according to the **map** clause semantics.

6            The **target enter data** construct executes as if it was enclosed in a **task** construct with no  
7            statements or directives outside of the **target enter data** construct. The generated task is a  
8            *target task*.

9            A *target task* is executed immediately and waits at a task scheduling point for the device to  
10           complete the **target enter data** region. The encountering thread becomes available to execute  
11           other tasks at that task scheduling point. If the *target task* is undeferred then the encountering  
12           thread must suspend the current task region, for which execution cannot be resumed until the  
13           generated target task is complete.

14           By default the generated task is undeferred. When a **nowait** clause is present, the current task  
15           may resume execution before the generated task completes its execution.

16           If a **depend** clause is present, then it is treated as if it had appeared on the implicit **task** construct  
17           that encloses the **target enter data** construct.

18           If there is no **device** clause, the default device is determined by the *default-device-var* ICV.

19           When an **if** clause is present and the **if** clause expression evaluates to *false*, the device is the host.

## 20           **Restrictions**

- 21           • A program must not depend on any ordering of the evaluations of the clauses of the  
22           **target enter data** directive, or on any side effects of the evaluations of the clauses.
- 23           • At least one **map** clause must appear on the directive.
- 24           • At most one **device** clause can appear on the directive. The **device** expression must evaluate  
25           to a non-negative integer value.
- 26           • At most one **if** clause can appear on the directive.
- 27           • A *map-type* must be specified in all **map** clauses and must be either **to** or **alloc**.
- 28           • The restrictions for the **task** construct apply.

## Cross References

- **target data**, see Section 2.10.1 on page 92.
- **target exit data**, see Section 2.10.11 on page 114.
- **map** clause, see Section 2.14.5 on page 195.
- *default-device-var*, see Section 2.3.1 on page 35.
- **task**, see Section 2.9.1 on page 80.
- **task scheduling constraints**, see Section 2.9.5 on page 90.

## 2.10.11 target exit data Construct

### Summary

The **target exit data** directive specifies that list items are unmapped from a device data environment. The **target exit data** directive is a stand-alone directive.

### Syntax

C / C++

The syntax of the **target exit data** construct is as follows:

```
#pragma omp target exit data [clause[[,] clause]...] new-line
```

where *clause* is one of the following:

```
device (integer-expression)
map ([ [map-type-modifier[,] map-type : ] list)
if (scalar-expression)
depend (dependence-type : list)
nowait
```

C / C++

Fortran

The syntax of the **target exit data** is as follows:

```
!$omp target exit data [clause[[,] clause]...]
```

1 where clause is one of the following:

```
2     device (scalar-integer-expression)  
3     map ([ [map-type-modifier[,]] map-type : ] list)  
4     if (scalar-logical-expression)  
5     depend (dependence-type : list)  
6     nowait
```

Fortran

## 7 Binding

8 The binding task for a **target exit data** construct is the encountering task.

## 9 Description

10 When a **target exit data** construct is encountered, the list items in the **map** clauses are  
11 unmapped from the device data environment according to the **map** clause semantics.

12 The **target exit data** construct executes as if it was enclosed in a **task** construct with no  
13 statements or directives outside of the **target exit data** construct. The generated task is a  
14 *target task*.

15 A *target task* is executed immediately and waits at a task scheduling point for the device to  
16 complete the **target exit data** region. The encountering thread becomes available to execute  
17 other tasks at that task scheduling point. If the *target task* is undeferred then the encountering  
18 thread must suspend the current task region, for which execution cannot be resumed until the  
19 generated *target task* is complete.

20 By default the generated task is undeferred. When a **nowait** clause is present, the current task  
21 may resume execution before the generated task completes its execution.

22 If a **depend** clause is present, then it is treated as if it had appeared on the implicit **task** construct  
23 that encloses the **target exit data** construct.

24 If there is no **device** clause, the default device is determined by the *default-device-var* ICV.

25 When an **if** clause is present and the **if** clause expression evaluates to *false*, the device is the host.

## Restrictions

- A program must not depend on any ordering of the evaluations of the clauses of the **target exit data** directive, or on any side effects of the evaluations of the clauses.
- At least one **map** clause must appear on the directive.
- At most one **device** clause can appear on the directive. The **device** expression must evaluate to a non-negative integer value.
- At most one **if** clause can appear on the directive.
- A *map-type* must be specified in all **map** clauses and must be either **from**, **release**, or **delete**.
- The restrictions for the **task** construct apply.

## Cross References

- **target data**, see Section 2.10.1 on page 92.
- **target enter data**, see Section 2.10.10 on page 112.
- **map** clause, see Section 2.14.5 on page 195.
- *default-device-var*, see Section 2.3.1 on page 35.
- **task**, see Section 2.9.1 on page 80.
- **task scheduling constraints**, see Section 2.9.5 on page 90.

## 2.11 Combined Constructs

Combined constructs are shortcuts for specifying one construct immediately nested inside another construct. The semantics of the combined constructs are identical to that of explicitly specifying the first construct containing one instance of the second construct and no other statements.

Some combined constructs have clauses that are permitted on both constructs that were combined. Where specified, the effect is as if applying the clauses to one or both constructs. If not specified and applying the clause to one construct would result in different program behavior than applying the clause to the other construct then the program's behavior is unspecified.

## 1 2.11.1 Parallel Loop Construct

### 2 Summary

3 The parallel loop construct is a shortcut for specifying a **parallel** construct containing one or  
4 more associated loops and no other statements.

### 5 Syntax

C / C++

6 The syntax of the parallel loop construct is as follows:

```
#pragma omp parallel for [clause[ [, ] clause] ... ] new-line  
for-loop
```

7 where *clause* can be any of the clauses accepted by the **parallel** or **for** directives, except the  
8 **nowait** clause, with identical meanings and restrictions.

C / C++

Fortran

9 The syntax of the parallel loop construct is as follows:

```
!$omp parallel do [clause[ [, ] clause] ... ]  
do-loops  
[!$omp end parallel do]
```

10 where *clause* can be any of the clauses accepted by the **parallel** or **do** directives, with identical  
11 meanings and restrictions.

12 If an **end parallel do** directive is not specified, an **end parallel do** directive is assumed at  
13 the end of the *do-loops*. **nowait** may not be specified on an **end parallel do** directive.

Fortran

## Description

C / C++

The semantics are identical to explicitly specifying a **parallel** directive immediately followed by a **for** directive.

C / C++

Fortran

The semantics are identical to explicitly specifying a **parallel** directive immediately followed by a **do** directive, and an **end do** directive immediately followed by an **end parallel** directive.

Fortran

## Restrictions

- The restrictions for the **parallel** construct and the loop construct apply

## Cross References

- **parallel** construct, see Section 2.5 on page 44.
- loop construct, see Section 2.7.1 on page 55.
- Data attribute clauses, see Section 2.14.3 on page 172.

## 2.11.2 parallel sections Construct

### Summary

The **parallel sections** construct is a shortcut for specifying a **parallel** construct containing one **sections** construct and no other statements.

### Syntax

C / C++

The syntax of the **parallel sections** construct is as follows:

```
#pragma omp parallel sections [clause [ , ] clause] ... ] new-line
{
  [#pragma omp section new-line]
    structured-block
  [#pragma omp section new-line]
    structured-block]
  ...
}
```





1 where *clause* can be any of the clauses accepted by the **parallel** or **sections** directives,  
2 except the **nowait** clause, with identical meanings and restrictions.



3 The syntax of the **parallel sections** construct is as follows:

```
!$omp parallel sections [clause[ [, ] clause] ... ]  
    [!$omp section  
      structured-block  
    [!$omp section  
      structured-block]  
    ...  
!$omp end parallel sections
```

4 where *clause* can be any of the clauses accepted by the **parallel** or **sections** directives, with  
5 identical meanings and restrictions.

6 The last section ends at the **end parallel sections** directive. **nowait** cannot be specified  
7 on an **end parallel sections** directive.



1

## Description

▼ C / C++ ▼

2

The semantics are identical to explicitly specifying a **parallel** directive immediately followed by a **sections** directive.

3

▲ C / C++ ▲

▼ Fortran ▼

4

The semantics are identical to explicitly specifying a **parallel** directive immediately followed by a **sections** directive, and an **end sections** directive immediately followed by an **end parallel** directive.

5

6

▲ Fortran ▲

7

## Restrictions

8

The restrictions for the **parallel** construct and the **sections** construct apply.

9

## Cross References

10

• **parallel** construct, see Section 2.5 on page 44.

11

• **sections** construct, see Section 2.7.2 on page 63.

12

• Data attribute clauses, see Section 2.14.3 on page 172.

▼ Fortran ▼

13

## 2.11.3 parallel workshare Construct

14

### Summary

15

The **parallel workshare** construct is a shortcut for specifying a **parallel** construct containing one **workshare** construct and no other statements.

16

## Syntax

The syntax of the **parallel workshare** construct is as follows:

```
!$omp parallel workshare [clause[ [, ] clause] ... ]  
    structured-block  
!$omp end parallel workshare
```

where *clause* can be any of the clauses accepted by the **parallel** directive, with identical meanings and restrictions. **nowait** may not be specified on an **end parallel workshare** directive.

## Description

The semantics are identical to explicitly specifying a **parallel** directive immediately followed by a **workshare** directive, and an **end workshare** directive immediately followed by an **end parallel** directive.

## Restrictions

The restrictions for the **parallel** construct and the **workshare** construct apply.

## Cross References

- **parallel** construct, see Section 2.5 on page 44.
- **workshare** construct, see Section 2.7.4 on page 67.
- Data attribute clauses, see Section 2.14.3 on page 172.

Fortran

## 2.11.4 Parallel Loop SIMD Construct

### Summary

The parallel loop SIMD construct is a shortcut for specifying a **parallel** construct containing one loop SIMD construct and no other statement.

### Syntax

C / C++

```
#pragma omp parallel for simd [clause[ [, ] clause]... ] new-line
for-loops
```

1 where *clause* can be any of the clauses accepted by the **parallel**, **for** or **simd** directives,  
2 except the **nowait** clause, with identical meanings and restrictions.



```
!$omp parallel do simd [clause[ [, ] clause]... ]  
do-loops  
!$omp end parallel do simd
```

3 where *clause* can be any of the clauses accepted by the **parallel**, **do** or **simd** directives, with  
4 identical meanings and restrictions.

5 If an **end parallel do simd** directive is not specified, an **end parallel do simd** directive  
6 is assumed at the end of the *do-loops*. **nowait** may not be specified on an  
7 **end parallel do simd** directive.



## 8 Description

9 The semantics of the parallel loop SIMD construct are identical to explicitly specifying a  
10 **parallel** directive immediately followed by a loop SIMD directive. The effect of any clause that  
11 applies to both constructs is as if it were applied to the loop SIMD construct and not to the  
12 **parallel** construct.

## 13 Restrictions

14 The restrictions for the **parallel** construct and the loop SIMD construct apply.

## 15 Cross References

- 16 • **parallel** construct, see Section 2.5 on page 44.
- 17 • loop SIMD construct, see Section 2.8.3 on page 78.
- 18 • Data attribute clauses, see Section 2.14.3 on page 172.

## 1 2.11.5 target teams construct

### 2 Summary

3 The **target teams** construct is a shortcut for specifying a **target** construct containing a  
4 **teams** construct.

### 5 Syntax

6 The syntax of the **target teams** construct is as follows:

▼ C / C++ ▼

```
#pragma omp target teams [clause[ [, ] clause] ... ]  
    structured-block
```

7 where *clause* can be any of the clauses accepted by the **target** or **teams** directives with identical  
8 meanings and restrictions.

▲ C / C++ ▲

## Fortran

```
!$omp target teams [clause[ [, ] clause] ... ]  
    structured-block  
!$omp end target teams
```

1 where *clause* can be any of the clauses accepted by the **target** or **teams** directives with identical  
2 meanings and restrictions.

## Fortran

### Description

#### C / C++

4 The semantics are identical to explicitly specifying a **target** directive immediately followed by a  
5 **teams** directive.

#### C / C++

#### Fortran

6 The semantics are identical to explicitly specifying a **target** directive immediately followed by a  
7 **teams** directive, and an **end teams** directive immediately followed by an **end target**  
8 directive.

## Fortran

### Restrictions

9 The restrictions for the **target** and **teams** constructs apply.

### Cross References

- 12 • **target** construct, see Section 2.10.2 on page 93.
- 13 • **teams** construct, see Section 2.10.5 on page 102.
- 14 • Data attribute clauses, see Section 2.14.3 on page 172.

## 1 2.11.6 teams distribute Construct

### 2 Summary

3 The **teams distribute** construct is a shortcut for specifying a **teams** construct containing a  
4 **distribute** construct.

### 5 Syntax

6 The syntax of the **teams distribute** construct is as follows:

▼ C / C++ ▲

```
#pragma omp teams distribute [clause[ [, ] clause] ... ]  
    for-loops
```

7 where *clause* can be any of the clauses accepted by the **teams** or **distribute** directives with  
8 identical meanings and restrictions.

▲ C / C++ ▲

▼ Fortran ▼

```
!$omp teams distribute [clause[ [, ] clause] ... ]  
    do-loops  
[!$omp end teams distribute]
```

9 where *clause* can be any of the clauses accepted by the **teams** or **distribute** directives with  
10 identical meanings and restrictions.

11 If an **end teams distribute** directive is not specified, an **end teams distribute**  
12 directive is assumed at the end of the *do-loops*.

▲ Fortran ▲

### 13 Description

14 The semantics are identical to explicitly specifying a **teams** directive immediately followed by a  
15 **distribute** directive. Some clauses are permitted on both constructs.

### 16 Restrictions

17 The restrictions for the **teams** and **distribute** constructs apply.

## Cross References

- **teams** construct, see Section 2.10.5 on page 102.
- **distribute** construct, see Section 2.10.6 on page 104.
- Data attribute clauses, see Section 2.14.3 on page 172.

## 2.11.7 teams distribute simd Construct

### Summary

The **teams distribute simd** construct is a shortcut for specifying a **teams** construct containing a **distribute simd** construct.

### Syntax

The syntax of the **teams distribute simd** construct is as follows:

C / C++

```
#pragma omp teams distribute simd [clause[ [, ] clause]... ]  
    for-loops
```

where *clause* can be any of the clauses accepted by the **teams** or **distribute simd** directives with identical meanings and restrictions.

C / C++

Fortran

```
!$omp teams distribute simd [clause[ [, ] clause]... ]  
    do-loops  
[!$omp end teams distribute simd]
```

where *clause* can be any of the clauses accepted by the **teams** or **distribute simd** directives with identical meanings and restrictions.

If an **end teams distribute** directive is not specified, an **end teams distribute** directive is assumed at the end of the *do-loops*.

Fortran



## Description

The semantics are identical to explicitly specifying a **teams** directive immediately followed by a **distribute simd** directive. Some clauses are permitted on both constructs.

## Restrictions

The restrictions for the **teams** and **distribute simd** constructs apply.

## Cross References

- **teams** construct, see Section 2.10.5 on page 102.
- **distribute simd** construct, see Section 2.10.7 on page 107.
- Data attribute clauses, see Section 2.14.3 on page 172.

## 2.11.8 target teams distribute construct

### Summary

The **target teams distribute** construct is a shortcut for specifying a **target** construct containing a **teams distribute** construct.

### Syntax

The syntax of the **target teams distribute** construct is as follows:

C / C++

```
#pragma omp target teams distribute [clause[ [, ] clause] ... ]  
    for-loops
```

where *clause* can be any of the clauses accepted by the **target** or **teams distribute** directives with identical meanings and restrictions.

C / C++

```

1  !$omp target teams distribute [clause[ [, ] clause]... ]
2  do-loops
3  [!$omp end target teams distribute]

```

1 where *clause* can be any of the clauses accepted by the **target** or **teams distribute**  
2 directives with identical meanings and restrictions.

3 If an **end target teams distribute** directive is not specified, an  
4 **end target teams distribute** directive is assumed at the end of the *do-loops*.

## 5 Description

6 The semantics are identical to explicitly specifying a **target** directive immediately followed by a  
7 **teams distribute** directive.

## 8 Restrictions

9 The restrictions for the **target** and **teams distribute** constructs apply.

## 10 Cross References

- 11 • **target** construct, see Section [2.10.1](#) on page [92](#).
- 12 • **teams distribute** construct, see Section [2.11.6](#) on page [125](#).
- 13 • Data attribute clauses, see Section [2.14.3](#) on page [172](#).

## 14 2.11.9 target teams distribute simd Construct

### 15 Summary

16 The **target teams distribute simd** construct is a shortcut for specifying a **target**  
17 construct containing a **teams distribute simd** construct.

### 18 Syntax

19 The syntax of the **target teams distribute simd** construct is as follows:

## C / C++

```
#pragma omp target teams distribute simd [clause[ [, ] clause]... ]  
    for-loops
```

1 where *clause* can be any of the clauses accepted by the **target** or **teams distribute simd**  
2 directives with identical meanings and restrictions.

## C / C++

## Fortran

```
!$omp target teams distribute simd [clause[ [, ] clause]... ]  
    do-loops  
[!$omp end target teams distribute simd]
```

3 where *clause* can be any of the clauses accepted by the **target** or **teams distribute simd**  
4 directives with identical meanings and restrictions.

5 If an **end target teams distribute simd** directive is not specified, an  
6 **end target teams distribute simd** directive is assumed at the end of the *do-loops*.

## Fortran

### Description

8 The semantics are identical to explicitly specifying a **target** directive immediately followed by a  
9 **teams distribute simd** directive.

### Restrictions

11 The restrictions for the **target** and **teams distribute simd** constructs apply.

### Cross References

- 13 • **target** construct, see Section [2.10.1](#) on page [92](#).
- 14 • **teams distribute simd** construct, see Section [2.11.7](#) on page [126](#).
- 15 • Data attribute clauses, see Section [2.14.3](#) on page [172](#).

## 1 2.11.10 Teams Distribute Parallel Loop Construct

### 2 Summary

3 The teams distribute parallel loop construct is a shortcut for specifying a **teams** construct  
4 containing a distribute parallel loop construct.

### 5 Syntax

6 The syntax of the teams distribute parallel loop construct is as follows:

▼ C / C++ ▲

```
#pragma omp teams distribute parallel for [clause[ [, ] clause] ... ]  
for-loops
```

7 where *clause* can be any of the clauses accepted by the **teams** or **distribute parallel for**  
8 directives with identical meanings and restrictions.

▲ C / C++ ▲  
▼ Fortran ▼

```
!$omp teams distribute parallel do [clause[ [, ] clause] ... ]  
do-loops  
[ !$omp end teams distribute parallel do ]
```

9 where *clause* can be any of the clauses accepted by the **teams** or **distribute parallel do**  
10 directives with identical meanings and restrictions.

11 If an **end teams distribute parallel do** directive is not specified, an  
12 **end teams distribute parallel do** directive is assumed at the end of the *do-loops*.

▲ Fortran ▲

### 13 Description

14 The semantics are identical to explicitly specifying a **teams** directive immediately followed by a  
15 distribute parallel loop directive. The effect of any clause that applies to both constructs is as if it  
16 were applied to both constructs separately.

### 17 Restrictions

18 The restrictions for the **teams** and distribute parallel loop constructs apply.

## Cross References

- **teams** construct, see Section 2.10.5 on page 102.
- Distribute parallel loop construct, see Section 2.10.8 on page 109.
- Data attribute clauses, see Section 2.14.3 on page 172.

## 2.11.11 Target Teams Distribute Parallel Loop Construct

### Summary

The target teams distribute parallel loop construct is a shortcut for specifying a **target** construct containing a teams distribute parallel loop construct.

### Syntax

The syntax of the target teams distribute parallel loop construct is as follows:

C / C++

```
#pragma omp target teams distribute parallel for [clause[ [, ] clause] ... ]  
for-loops
```

where *clause* can be any of the clauses accepted by the **target** or **teams distribute parallel for** directives with identical meanings and restrictions.

C / C++

Fortran

```
!$omp target teams distribute parallel do [clause[ [, ] clause] ... ]  
do-loops  
[!$omp end target teams distribute parallel do]
```

where *clause* can be any of the clauses accepted by the **target** or **teams distribute parallel do** directives with identical meanings and restrictions.

If an **end target teams distribute parallel do** directive is not specified, an **end target teams distribute parallel do** directive is assumed at the end of the *do-loops*.

Fortran

1           **Description**

2           The semantics are identical to explicitly specifying a **target** directive immediately followed by a  
3           teams distribute parallel loop directive.

4           **Restrictions**

5           The restrictions for the **target** and teams distribute parallel loop constructs apply.

6           **Cross References**

- 7           • **target** construct, see Section 2.10.2 on page 93.  
8           • Distribute parallel loop construct, see Section 2.11.10 on page 130.  
9           • Data attribute clauses, see Section 2.14.3 on page 172.

10   **2.11.12 Teams Distribute Parallel Loop SIMD Construct**

11           **Summary**

12           The teams distribute parallel loop SIMD construct is a shortcut for specifying a **teams** construct  
13           containing a distribute parallel loop SIMD construct.

14           **Syntax**

15           The syntax of the teams distribute parallel loop construct is as follows:

▼────────────────────────────────── C / C++ ───────────────────────────────────▼

```
#pragma omp teams distribute parallel for simd [clause[ [, ] clause]... ]  
    for-loops
```

16           where *clause* can be any of the clauses accepted by the **teams** or  
17           **distribute parallel for simd** directives with identical meanings and restrictions.

▲────────────────────────────────── C / C++ ───────────────────────────────────▲

```
!$omp teams distribute parallel do simd [clause[ [, ] clause]... ]
      do-loops
[!$omp end teams distribute parallel do simd]
```

1 where *clause* can be any of the clauses accepted by the **teams** or  
2 **distribute parallel do simd** directives with identical meanings and restrictions.

3 If an **end teams distribute parallel do simd** directive is not specified, an  
4 **end teams distribute parallel do simd** directive is assumed at the end of the *do-loops*.

## 5 Description

6 The semantics are identical to explicitly specifying a **teams** directive immediately followed by a  
7 distribute parallel loop SIMD directive. The effect of any clause that applies to both constructs is as  
8 if it were applied to both constructs separately.

## 9 Restrictions

10 The restrictions for the teams and distribute parallel loop SIMD constructs apply.

## 11 Cross References

- 12 • **teams** construct, see Section [2.10.5](#) on page [102](#).
- 13 • Distribute parallel loop SIMD construct, see Section [2.10.9](#) on page [110](#).
- 14 • Data attribute clauses, see Section [2.14.3](#) on page [172](#).

## 15 2.11.13 Target Teams Distribute Parallel Loop SIMD 16 Construct

### 17 Summary

18 The target teams distribute parallel loop SIMD construct is a shortcut for specifying a **target**  
19 construct containing a teams distribute parallel loop SIMD construct.

## Syntax

The syntax of the target teams distribute parallel loop SIMD construct is as follows:

C / C++

```
#pragma omp target teams distribute parallel for simd [clause[ [, ] clause]... ]  
for-loops
```

where *clause* can be any of the clauses accepted by the **target** or **teams distribute parallel for simd** directives with identical meanings and restrictions.

C / C++

Fortran

```
!$omp target teams distribute parallel do simd [clause[ [, ] clause]... ]  
do-loops  
[!$omp end target teams distribute parallel do simd]
```

where *clause* can be any of the clauses accepted by the **target** or **teams distribute parallel do simd** directives with identical meanings and restrictions.

If an **end target teams distribute parallel do simd** directive is not specified, an **end target teams distribute parallel do simd** directive is assumed at the end of the *do-loops*.

Fortran

## Description

The semantics are identical to explicitly specifying a **target** directive immediately followed by a teams distribute parallel loop SIMD directive.

## Restrictions

The restrictions for the **target** and teams distribute parallel loop SIMD constructs apply.

## Cross References

- **target** construct, see Section [2.10.2](#) on page [93](#).
- Teams distribute parallel loop SIMD construct, see Section [2.11.12](#) on page [132](#).
- Data attribute clauses, see Section [2.14.3](#) on page [172](#).



## 1 2.12 Master and Synchronization Constructs 2 and Clauses

3 OpenMP provides the following synchronization constructs:

- 4 • the **master** construct.
- 5 • the **critical** construct.
- 6 • the **barrier** construct.
- 7 • the **taskwait** construct.
- 8 • the **taskgroup** construct.
- 9 • the **atomic** construct.
- 10 • the **flush** construct.
- 11 • the **ordered** construct.

### 12 2.12.1 master Construct

#### 13 Summary

14 The **master** construct specifies a structured block that is executed by the master thread of the team.

#### 15 Syntax

16  C / C++

The syntax of the **master** construct is as follows:

```
#pragma omp master new-line  
    structured-block
```

 C / C++  
 Fortran

17 The syntax of the **master** construct is as follows:

```
!$omp master  
    structured-block  
!$omp end master
```

 Fortran

1       **Binding**

2       The binding thread set for a **master** region is the current team. A **master** region binds to the  
3       innermost enclosing **parallel** region. Only the master thread of the team executing the binding  
4       **parallel** region participates in the execution of the structured block of the **master** region.

5       **Description**

6       Other threads in the team do not execute the associated structured block. There is no implied  
7       barrier either on entry to, or exit from, the **master** construct.

8       **Restrictions**

9       ▼── C++ ───▼

- A throw executed inside a **master** region must cause execution to resume within the same **master** region, and the same thread that threw the exception must catch it

10      ▲── C++ ───▲

11   **2.12.2 critical Construct**

12       **Summary**

13       The **critical** construct restricts execution of the associated structured block to a single thread at  
14       a time.

15       **Syntax**

16       ▼── C / C++ ───▼

16       The syntax of the **critical** construct is as follows:

```
#pragma omp critical [(name) ] new-line
    structured-block
```

17       ▲── C / C++ ───▲

17       ▼── Fortran ───▼

17       The syntax of the **critical** construct is as follows:

```
!$omp critical [ (name) ]
    structured-block
!$omp end critical [ (name) ]
```

Fortran

## 1 Binding

2 The binding thread set for a **critical** region is all threads in the contention group. Region  
3 execution is restricted to a single thread at a time among all threads in the contention group,  
4 without regard to the team(s) to which the threads belong.

## 5 Description

6 An optional *name* may be used to identify the **critical** construct. All **critical** constructs  
7 without a name are considered to have the same unspecified name. A thread waits at the beginning  
8 of a **critical** region until no thread in the contention group is executing a **critical** region  
9 with the same name. The **critical** construct enforces exclusive access with respect to all  
10 **critical** constructs with the same name in all threads in the contention group, not just those  
11 threads in the current team.

C / C++

12 Identifiers used to identify a **critical** construct have external linkage and are in a name space  
13 that is separate from the name spaces used by labels, tags, members, and ordinary identifiers.

C / C++

Fortran

14 The names of **critical** constructs are global entities of the program. If a name conflicts with  
15 any other entity, the behavior of the program is unspecified.

Fortran

## Restrictions

C++

- A throw executed inside a **critical** region must cause execution to resume within the same **critical** region, and the same thread that threw the exception must catch it.

C++

Fortran

The following restrictions apply to the critical construct:

- If a *name* is specified on a **critical** directive, the same *name* must also be specified on the **end critical** directive.
- If no *name* appears on the **critical** directive, no *name* can appear on the **end critical** directive.

Fortran

## 2.12.3 barrier Construct

### Summary

The **barrier** construct specifies an explicit barrier at the point at which the construct appears. The **barrier** construct is a stand-alone directive.

### Syntax

C / C++

The syntax of the **barrier** construct is as follows:

```
#pragma omp barrier new-line
```

C / C++

Fortran

The syntax of the **barrier** construct is as follows:

```
!$omp barrier
```

Fortran

1       **Binding**

2       The binding thread set for a **barrier** region is the current team. A **barrier** region binds to the  
3       innermost enclosing **parallel** region.

4       **Description**

5       All threads of the team executing the binding **parallel** region must execute the **barrier**  
6       region and complete execution of all explicit tasks bound to this **parallel** region before any are  
7       allowed to continue execution beyond the barrier.

8       The **barrier** region includes an implicit task scheduling point in the current task region.

9       **Restrictions**

10      The following restrictions apply to the **barrier** construct:

- 11      • Each **barrier** region must be encountered by all threads in a team or by none at all, unless  
12      cancellation has been requested for the innermost enclosing parallel region.
- 13      • The sequence of worksharing regions and **barrier** regions encountered must be the same for  
14      every thread in a team.

15   **2.12.4 taskwait Construct**

16      **Summary**

17      The **taskwait** construct specifies a wait on the completion of child tasks of the current task. The  
18      **taskwait** construct is a stand-alone directive.

19      **Syntax**

▼──────────────────────────────── C / C++ ─────────────────────────────────▼

20      The syntax of the **taskwait** construct is as follows:

```
#pragma omp taskwait newline
```

▲──────────────────────────────── C / C++ ─────────────────────────────────▲

▼──────────────────────────────── Fortran ─────────────────────────────────▼

21      The syntax of the **taskwait** construct is as follows:

```
!$omp taskwait
```

Fortran

## 1 Binding

2 The **taskwait** region binds to the current task region. The binding thread set of the **taskwait**  
3 region is the current team.

## 4 Description

5 The **taskwait** region includes an implicit task scheduling point in the current task region. The  
6 current task region is suspended at the task scheduling point until all child tasks that it generated  
7 before the **taskwait** region complete execution.

## 8 2.12.5 taskgroup Construct

### 9 Summary

10 The **taskgroup** construct specifies a wait on completion of child tasks of the current task and  
11 their descendent tasks.

### 12 Syntax

C / C++

13 The syntax of the **taskgroup** construct is as follows:

```
#pragma omp taskgroup new-line  
                  structured-block
```

C / C++

Fortran

14 The syntax of the **taskgroup** construct is as follows:

```
!$omp taskgroup  
      structured-block  
!$omp end taskgroup
```

Fortran

## 1        **Binding**

2        A **taskgroup** region binds to the current task region. The binding thread set of the **taskgroup**  
3        region is the current team.

## 4        **Description**

5        When a thread encounters a **taskgroup** construct, it starts executing the region. There is an  
6        implicit task scheduling point at the end of the **taskgroup** region. The current task is suspended  
7        at the task scheduling point until all child tasks that it generated in the **taskgroup** region and all  
8        of their descendent tasks complete execution.

## 9        **Cross References**

- 10        • Task scheduling, see Section [2.9.5](#) on page [90](#).

## 11    **2.12.6 atomic Construct**

### 12        **Summary**

13        The **atomic** construct ensures that a specific storage location is accessed atomically, rather than  
14        exposing it to the possibility of multiple, simultaneous reading and writing threads that may result  
15        in indeterminate values

### 16        **Syntax**

▼ C / C++ ▼

17        The syntax of the **atomic** construct takes either of the following forms:

```
#pragma omp atomic [read | write | update |  
capture] [seq_cst] new-line  
expression-stmt
```

18        or

```
#pragma omp atomic capture [seq_cst] new-line  
structured-block
```

where *expression-stmt* is an expression statement with one of the following forms:

- If clause is **read**:  
 $v = x;$
- If clause is **write**:  
 $x = \textit{expr};$
- If clause is **update** or not present:  
 $x++;$   
 $x--;$   
 $++x;$   
 $--x;$   
 $x \textit{ binop} = \textit{expr};$   
 $x = x \textit{ binop} \textit{ expr};$   
 $x = \textit{expr} \textit{ binop} x;$
- If clause is **capture**:  
 $v = x++;$   
 $v = x--;$   
 $v = ++x;$   
 $v = --x;$   
 $v = x \textit{ binop} = \textit{expr};$   
 $v = x = x \textit{ binop} \textit{ expr};$   
 $v = x = \textit{expr} \textit{ binop} x;$

and where *structured-block* is a structured block with one of the following forms:

- { $v = x; x \textit{ binop} = \textit{expr};$ }
- { $x \textit{ binop} = \textit{expr}; v = x;$ }
- { $v = x; x = x \textit{ binop} \textit{ expr};$ }
- { $v = x; x = \textit{expr} \textit{ binop} x;$ }
- { $x = x \textit{ binop} \textit{ expr}; v = x;$ }
- { $x = \textit{expr} \textit{ binop} x; v = x;$ }
- { $v = x; x = \textit{expr};$ }
- { $v = x; x++;$ }
- { $v = x; ++x;$ }
- { $++x; v = x;$ }
- { $x++; v = x;$ }
- { $v = x; x--;$ }
- { $v = x; --x;$ }
- { $--x; v = x;$ }
- { $x--; v = x;$ }



- 1 In the preceding expressions:
- 2 •  $x$  and  $v$  (as applicable) are both *l-value* expressions with scalar type.
- 3 • During the execution of an atomic region, multiple syntactic occurrences of  $x$  must designate the
- 4 same storage location.
- 5 • Neither of  $v$  and  $expr$  (as applicable) may access the storage location designated by  $x$ .
- 6 • Neither of  $x$  and  $expr$  (as applicable) may access the storage location designated by  $v$ .
- 7 •  $expr$  is an expression with scalar type.
- 8 •  $binop$  is one of  $+$ ,  $*$ ,  $-$ ,  $/$ ,  $\&$ ,  $\wedge$ ,  $|$ ,  $\ll$ , or  $\gg$ .
- 9 •  $binop$ ,  $binop=$ ,  $++$ , and  $--$  are not overloaded operators.
- 10 • The expression  $x\ binop\ expr$  must be numerically equivalent to  $x\ binop\ (expr)$ . This requirement
- 11 is satisfied if the operators in  $expr$  have precedence greater than  $binop$ , or by using parentheses
- 12 around  $expr$  or subexpressions of  $expr$ .
- 13 • The expression  $expr\ binop\ x$  must be numerically equivalent to  $(expr)\ binop\ x$ . This requirement
- 14 is satisfied if the operators in  $expr$  have precedence equal to or greater than  $binop$ , or by using
- 15 parentheses around  $expr$  or subexpressions of  $expr$ .
- 16 • For forms that allow multiple occurrences of  $x$ , the number of times that  $x$  is evaluated is
- 17 unspecified.



18 The syntax of the **atomic** construct takes any of the following forms:

```
!$omp atomic read [seq_cst]
    capture-statement
[!$omp end atomic]
```

19 or

```
!$omp atomic write [seq_cst]
    write-statement
[!$omp end atomic]
```

20 or

```
!$omp atomic [update] [seq_cst]
    update-statement
[!$omp end atomic]
```

1 or

```
!$omp atomic capture [seq_cst]
    update-statement
    capture-statement
!$omp end atomic
```

2 or

```
!$omp atomic capture [seq_cst]
    capture-statement
    update-statement
!$omp end atomic
```

3 or

```
!$omp atomic capture [seq_cst]
    capture-statement
    write-statement
!$omp end atomic
```

4 where *write-statement* has the following form (if clause is **write**):

5  $x = \text{expr}$

6 where *capture-statement* has the following form (if clause is **capture** or **read**):

7  $v = x$

8 and where *update-statement* has one of the following forms (if clause is **update**, **capture**, or  
9 not present):

10  $x = x \text{ operator } \text{expr}$

11  $x = \text{expr operator } x$

12  $x = \text{intrinsic\_procedure\_name } (x, \text{expr\_list})$

13  $x = \text{intrinsic\_procedure\_name } (\text{expr\_list}, x)$

14 In the preceding statements:

- 15 •  $x$  and  $v$  (as applicable) are both scalar variables of intrinsic type.
- 16 •  $x$  must not have the **ALLOCATABLE** attribute. .

- 1           • During the execution of an atomic region, multiple syntactic occurrences of  $x$  must designate the
- 2           same storage location.
- 3           • None of  $v$ ,  $expr$  and  $expr\_list$  (as applicable) may access the same storage location as  $x$ .
- 4           • None of  $x$ ,  $expr$  and  $expr\_list$  (as applicable) may access the same storage location as  $v$ .
- 5           •  $expr$  is a scalar expression.
- 6           •  $expr\_list$  is a comma-separated, non-empty list of scalar expressions. If
- 7           *intrinsic\_procedure\_name* refers to **IAND**, **IOR**, or **IEOR**, exactly one expression must appear in
- 8           *expr\_list*.
- 9           • *intrinsic\_procedure\_name* is one of **MAX**, **MIN**, **IAND**, **IOR**, or **IEOR**.
- 10          • *operator* is one of **+**, **\***, **-**, **/**, **.AND.**, **.OR.**, **.EQV.**, or **.NEQV.**
- 11          • The expression  $x \text{ operator } expr$  must be numerically equivalent to  $x \text{ operator } (expr)$ . This
- 12          requirement is satisfied if the operators in  $expr$  have precedence greater than *operator*, or by
- 13          using parentheses around  $expr$  or subexpressions of  $expr$ .
- 14          • The expression  $expr \text{ operator } x$  must be numerically equivalent to  $(expr) \text{ operator } x$ . This
- 15          requirement is satisfied if the operators in  $expr$  have precedence equal to or greater than
- 16          *operator*, or by using parentheses around  $expr$  or subexpressions of  $expr$ .
- 17          • *intrinsic\_procedure\_name* must refer to the intrinsic procedure name and not to other program
- 18          entities.
- 19          • *operator* must refer to the intrinsic operator and not to a user-defined operator.
- 20          • All assignments must be intrinsic assignments.
- 21          • For forms that allow multiple occurrences of  $x$ , the number of times that  $x$  is evaluated is
- 22          unspecified.

## Fortran

- 23          • In all **atomic** construct forms, the **seq\_cst** clause and the clause that denotes the type of the
- 24          atomic construct can appear in any order. In addition, an optional comma may be used to
- 25          separate the clauses

### Binding

27          The binding thread set for an atomic region is all threads in the contention group. **atomic** regions

28          enforce exclusive access with respect to other **atomic** regions that access the same storage

29          location  $x$  among all threads in the contention group without regard to the teams to which the

30          threads belong.

## Description

The **atomic** construct with the **read** clause forces an atomic read of the location designated by *x* regardless of the native machine word size.

The **atomic** construct with the **write** clause forces an atomic write of the location designated by *x* regardless of the native machine word size.

The **atomic** construct with the **update** clause forces an atomic update of the location designated by *x* using the designated operator or intrinsic. Note that when no clause is present, the semantics are equivalent to atomic update. Only the read and write of the location designated by *x* are performed mutually atomically. The evaluation of *expr* or *expr\_list* need not be atomic with respect to the read or write of the location designated by *x*. No task scheduling points are allowed between the read and the write of the location designated by *x*.

The **atomic** construct with the **capture** clause forces an atomic update of the location designated by *x* using the designated operator or intrinsic while also capturing the original or final value of the location designated by *x* with respect to the atomic update. The original or final value of the location designated by *x* is written in the location designated by *v* depending on the form of the **atomic** construct structured block or statements following the usual language semantics. Only the read and write of the location designated by *x* are performed mutually atomically. Neither the evaluation of *expr* or *expr\_list*, nor the write to the location designated by *v* need be atomic with respect to the read or write of the location designated by *x*. No task scheduling points are allowed between the read and the write of the location designated by *x*.

Any **atomic** construct with a **seq\_cst** clause forces the atomically performed operation to include an implicit flush operation without a list.

---

**Note** – As with other implicit flush regions, Section 1.4.4 on page 20 reduces the ordering that must be enforced. The intent is that, when the analogous operation exists in C++11 or C11, a sequentially consistent **atomic** construct has the same semantics as a **memory\_order\_seq\_cst** atomic operation in C++11/C11. Similarly, a non-sequentially consistent **atomic** construct has the same semantics as a **memory\_order\_relaxed** atomic operation in C++11/C11.

Unlike non-sequentially consistent **atomic** constructs, sequentially consistent **atomic** constructs preserve the interleaving (sequentially consistent) behavior of correct, data-race-free programs. However, they are not designed to replace the **flush** directive as a mechanism to enforce ordering for non-sequentially consistent **atomic** constructs, and attempts to do so require extreme caution. For example, a sequentially consistent **atomic write** construct may appear to be reordered with a subsequent non-sequentially consistent **atomic write** construct, since such reordering would not be observable by a correct program if the second write were outside an **atomic** directive.

1 For all forms of the **atomic** construct, any combination of two or more of these **atomic**  
2 constructs enforces mutually exclusive access to the locations designated by  $x$ . To avoid race  
3 conditions, all accesses of the locations designated by  $x$  that could potentially occur in parallel must  
4 be protected with an **atomic** construct.

5 **atomic** regions do not guarantee exclusive access with respect to any accesses outside of  
6 **atomic** regions to the same storage location  $x$  even if those accesses occur during a **critical**  
7 or **ordered** region, while an OpenMP lock is owned by the executing task, or during the  
8 execution of a **reduction** clause.

9 However, other OpenMP synchronization can ensure the desired exclusive access. For example, a  
10 barrier following a series of atomic updates to  $x$  guarantees that subsequent accesses do not form a  
11 race with the atomic accesses.

12 A compliant implementation may enforce exclusive access between **atomic** regions that update  
13 different storage locations. The circumstances under which this occurs are implementation defined.

14 If the storage location designated by  $x$  is not size-aligned (that is, if the byte alignment of  $x$  is not a  
15 multiple of the size of  $x$ ), then the behavior of the **atomic** region is implementation defined.

## 16 Restrictions

▼ C / C++ ▼

17 The following restriction applies to the **atomic** construct:

- 18 • All atomic accesses to the storage locations designated by  $x$  throughout the program are required  
19 to have a compatible type.

▲ C / C++ ▲

▼ Fortran ▼

20 The following restriction applies to the **atomic** construct:

- 21 • All atomic accesses to the storage locations designated by  $x$  throughout the program are required  
22 to have the same type and type parameters.

▲ Fortran ▲

## Cross References

- **critical** construct, see Section 2.12.2 on page 136.
- **barrier** construct, see Section 2.12.3 on page 138.
- **flush** construct, see Section 2.12.7 on page 148.
- **ordered** construct, see Section 2.12.8 on page 152.
- **reduction** clause, see Section 2.14.3.6 on page 184.
- lock routines, see Section 3.3 on page 239.

## 2.12.7 flush Construct

### Summary

The **flush** construct executes the OpenMP flush operation. This operation makes a thread's temporary view of memory consistent with memory, and enforces an order on the memory operations of the variables explicitly specified or implied. See the memory model description in Section 1.4 on page 17 for more details. The **flush** construct is a stand-alone directive.

### Syntax

C / C++

The syntax of the **flush** construct is as follows:

```
#pragma omp flush [(list)] new-line
```

C / C++

Fortran

The syntax of the **flush** construct is as follows:

```
!$omp flush [(list)]
```

Fortran

1           **Binding**

2           The binding thread set for a **flush** region is the encountering thread. Execution of a **flush**  
3           region affects the memory and the temporary view of memory of only the thread that executes the  
4           region. It does not affect the temporary view of other threads. Other threads must themselves  
5           execute a flush operation in order to be guaranteed to observe the effects of the encountering  
6           thread's flush operation

7           **Description**

8           A **flush** construct without a list, executed on a given thread, operates as if the whole  
9           thread-visible data state of the program, as defined by the base language, is flushed. A **flush**  
10          construct with a list applies the flush operation to the items in the list, and does not return until the  
11          operation is complete for all specified list items. An implementation may implement a **flush** with  
12          a list by ignoring the list, and treating it the same as a **flush** without a list.



13          If a pointer is present in the list, the pointer itself is flushed, not the memory block to which the  
14          pointer refers.



15          If the list item or a subobject of the list item has the **POINTER** attribute, the allocation or  
16          association status of the **POINTER** item is flushed, but the pointer target is not. If the list item is a  
17          Cray pointer, the pointer is flushed, but the object to which it points is not. If the list item is of type  
18          **C\_PTR**, the variable is flushed, but the storage that corresponds to that address is not flushed. If the  
19          list item or the subobject of the list item has the **ALLOCATABLE** attribute and has an allocation  
20          status of currently allocated, the allocated variable is flushed; otherwise the allocation status is  
21          flushed.



1 Note – Use of a **flush** construct with a list is extremely error prone and users are strongly  
2 discouraged from attempting it. The following examples illustrate the ordering properties of the  
3 flush operation. In the following incorrect pseudocode example, the programmer intends to prevent  
4 simultaneous execution of the protected section by the two threads, but the program does not work  
5 properly because it does not enforce the proper ordering of the operations on variables **a** and **b**.  
6 Any shared data accessed in the protected section is not guaranteed to be current or consistent  
7 during or after the protected section. The atomic notation in the pseudocode in the following two  
8 examples indicates that the accesses to **a** and **b** are **ATOMIC** writes and captures. Otherwise both  
9 examples would contain data races and automatically result in unspecified behavior.

*Incorrect example:*

**a = b = 0**

*thread 1*

**atomic(b = 1)**

*flush(b)*

*flush(a)*

**atomic(tmp = a)**

**if (tmp == 0) then**

*protected section*

**end if**

*thread 2*

**atomic(a = 1)**

*flush(a)*

*flush(b)*

**atomic(tmp = b)**

**if (tmp == 0) then**

*protected section*

**end if**

11 The problem with this example is that operations on variables **a** and **b** are not ordered with respect  
12 to each other. For instance, nothing prevents the compiler from moving the flush of **b** on thread 1 or  
13 the flush of **a** on thread 2 to a position completely after the protected section (assuming that the  
14 protected section on thread 1 does not reference **b** and the protected section on thread 2 does not  
15 reference **a**). If either re-ordering happens, both threads can simultaneously execute the protected  
16 section.

17 The following pseudocode example correctly ensures that the protected section is executed by not  
18 more than one of the two threads at any one time. Notice that execution of the protected section by  
19 neither thread is considered correct in this example. This occurs if both flushes complete prior to  
20 either thread executing its **if** statement.



Correct example:

**a = b = 0**

*thread 1*

*thread 2*

**atomic(b = 1)**

**atomic(a = 1)**

*flush(a, b)*

*flush(a, b)*

**atomic(tmp = a)**

**atomic(tmp = b)**

**if (tmp == 0) then**

**if (tmp == 0) then**

*protected section*

*protected section*

**end if**

**end if**

The compiler is prohibited from moving the flush at all for either thread, ensuring that the respective assignment is complete and the data is flushed before the **if** statement is executed.



A **flush** region without a list is implied at the following locations:

- During a barrier region.
- At entry to a **target update** region whose corresponding construct has a **to** clause.
- At exit from a **target update** region whose corresponding construct has a **from** clause.
- At entry to and exit from **parallel**, **critical**, **ordered**, **target** and **target data** regions.
- At exit from worksharing regions unless a **nowait** is present.
- At entry to and exit from the **atomic** operation (read, write, update, or capture) performed in a sequentially consistent atomic region.
- During **omp\_set\_lock** and **omp\_unset\_lock** regions.
- During **omp\_test\_lock**, **omp\_set\_nest\_lock**, **omp\_unset\_nest\_lock** and **omp\_test\_nest\_lock** regions, if the region causes the lock to be set or unset.
- Immediately before and immediately after every task scheduling point.

A **flush** region with a list is implied at the following locations:

- 1 • At entry to and exit from the **atomic** operation (read, write, update, or capture) performed in a  
2 non-sequentially consistent **atomic** region, where the list contains only the storage location  
3 designated as x according to the description of the syntax of the **atomic** construct in  
4 Section 2.12.6 on page 141.

5 Note – A **flush** region is not implied at the following locations:

- 6 • At entry to worksharing regions.  
7 • At entry to or exit from a **master** region.

## 8 2.12.8 ordered Construct

### 9 Summary

10 The **ordered** construct specifies a structured block in a loop region that will be executed in the  
11 order of the loop iterations. This sequentializes and orders the code within an **ordered** region  
12 while allowing code outside the region to run in parallel.

### 13 Syntax

14 C / C++

The syntax of the **ordered** construct is as follows:

```
#pragma omp ordered [clause[ [, ] clause] ... ] new-line  
    structured-block
```

15 where *clause* is one of the following:

- 16 • **threads**  
17 • **simd**

18 or

```
#pragma omp ordered clause [[[ , ] clause] ... ] new-line
```

19 where *clause* is:

- 20 • **depend** (*dependence-type* [ : *vec*])

1 The syntax of the **ordered** construct is as follows:

```
!$omp ordered [clause[ [, ] clause] ... ]
    structured-block
!$omp end ordered
```

2 where *clause* is one of the following:

- 3 • **threads**
- 4 • **simd**

5 or

```
!$omp ordered clause [[[, ] clause] ... ]
```

6 where *clause* is:

- 7 • **depend** (*dependence-type* [: *vec*])

8 If the **depend** clause is specified, the **ordered** construct is a stand-alone directive.

## 9 Binding

10 The binding thread set for an **ordered** region is the current team. An **ordered** region binds to  
 11 the innermost enclosing loop region. **ordered** regions that bind to different loop regions execute  
 12 independently of each other.

## 13 Description

14 If no clause is specified, the **ordered** construct behaves as if the **threads** clause had been  
 15 specified. If the **threads** clause is specified, the threads in the team executing the loop region  
 16 execute **ordered** regions sequentially in the order of the loop iterations. If any **depend** clauses  
 17 are specified then those clauses specify the order in which the threads in the team execute  
 18 **ordered** regions. When the thread executing the first iteration of the loop encounters an  
 19 **ordered** construct, it can enter the **ordered** region without waiting. When a thread executing  
 20 any subsequent iteration encounters an **ordered** region, it waits at the beginning of that ordered  
 21 region until execution of all **ordered** regions belonging to all previous iterations or iterations  
 22 specified by the **depend** clauses have completed. If the **simd** clause is specified, the **ordered**  
 23 regions encountered by any thread will use only a single SIMD lane to execute the **ordered**  
 24 regions in the order of the loop iterations.

## Restrictions

Restrictions to the **ordered** construct are as follows:

- The loop region to which an **ordered** region without any clause or with a **threads** clause binds must have an **ordered** clause without the parameter specified on the corresponding loop directive.
- The loop region to which an **ordered** region with any **depend** clauses binds must have an **ordered** clause with the parameter specified on the corresponding loop directive.
- An **ordered** construct with the **depend** clause specified must be closely nested inside a loop (or parallel loop) construct.
- An **ordered** construct with the **depend** clause specified must not be closely nested in a **simd** construct.
- During execution of an iteration of a loop or a loop nest within a loop region, a thread must not execute more than one ordered region that arises from an **ordered** clause with no parameter and binds to the same loop region.

C++

- A throw executed inside a **ordered** region must cause execution to resume within the same **ordered** region, and the same thread that threw the exception must catch it.

C++

## Cross References

- loop construct, see Section [2.7.1](#) on page [55](#).
- parallel loop construct, see Section [2.11.1](#) on page [117](#).
- **depend** Clause, see Section [2.12.9](#) on page [154](#)

### 2.12.9 depend Clause

#### Summary

The **depend** clause enforces additional constraints on the scheduling of tasks or loop iterations. These constraints establish dependences only between sibling tasks or between loop iterations. The clause consists of a *dependence-type* with one or more list items.

## Syntax

The syntax of the **depend** clause is as follows:

```
depend (dependence-type : list)
```

or

```
depend (dependence-type [: vec])
```

where *vec* is the iteration vector.

## Description

Task dependences are derived from the *dependence-type* of a **depend** clause and its list items, where *dependence-type* is one of the following:

The **in** *dependence-type*. The generated task will be a dependent task of all previously generated sibling tasks that reference at least one of the list items in an **out** or **inout** *dependence-type* list.

The **out** and **inout** *dependence-types*. The generated task will be a dependent task of all previously generated sibling tasks that reference at least one of the list items in an **in**, **out**, or **inout** *dependence-type* list.

The list items that appear in the **depend** clause may include array sections.

---

**Note** – The enforced task dependence establishes a synchronization of memory accesses performed by a dependent task with respect to accesses performed by the predecessor tasks. However, it is the responsibility of the programmer to synchronize properly with respect to other concurrent accesses that occur outside of those tasks.

---

Loop dependences are derived from the *dependence-type* of a **depend** clause and the iteration vector *vec*.

The **source** *dependence-type* specifies the completion of cross-iteration dependences that arise from the current iteration.

The **sink** *dependence-type* specifies a sink of cross-iteration dependences. The current iteration is blocked until all source iterations specified complete execution.

The iteration vector *vec* must have the form of  $(x_1 \pm d_1, x_2 \pm d_2, \dots, x_n \pm d_n)$ , where *n* is the value specified by the **ordered** clause in the loop directive,  $x_i$  denotes the loop iteration variable of the *i*-th nested loop associated with the loop directive, and  $d_i$  is a constant non-negative integer. If the iteration vector *vec* indicates a lexicographically later iteration, it can cause a deadlock. If the iteration vector *vec* indicates an invalid iteration, the **ordered** construct with the **depend** clause is ignored.

## Restrictions

Restrictions to the **depend** clause are as follows:

- List items used in **depend** clauses of the same task or sibling tasks must indicate identical storage or disjoint storage.
- List items used in **depend** clauses cannot be zero-length array sections.
- A variable that is part of another variable (such as a field of a structure) but is not an array element or an array section cannot appear in a **depend** clause.
- If the *dependence-type* is **source**, the iteration vector *vec* must not be specified.

## Cross References

- Array sections, see Section 2.4 on page 43.
- **task** construct, see Section 2.9.1 on page 80.
- Task scheduling constraints, see Section 2.9.5 on page 90.
- **ordered** construct, see Section 2.12.8 on page 152.

## 2.13 Cancellation Constructs

### 2.13.1 **cancel** Construct

#### Summary

The **cancel** construct activates cancellation of the innermost enclosing region of the type specified. The **cancel** construct is a stand-alone directive.

#### Syntax

C / C++

The syntax of the **cancel** construct is as follows:

```
#pragma omp cancel construct-type-clause [ , ] if-clause new-line
```

1 where *construct-type-clause* is one of the following:

2       **parallel**  
3       **sections**  
4       **for**  
5       **taskgroup**

6 and *if-clause* is

7       **if** (*scalar-expression*)



8 The syntax of the **cancel** construct is as follows:

```
!$omp cancel construct-type-clause [ [, ] if-clause] new-line
```

9 where *construct-type-clause* is one of the following:

10       **parallel**  
11       **sections**  
12       **do**  
13       **taskgroup**

14 and *if-clause* is

15       **if** (*scalar-logical-expression*)



## 16 **Binding**

17 The binding thread set of the **cancel** region is the current team. The **cancel** region binds to the  
18 innermost enclosing construct of the type corresponding to the *type-clause* specified in the directive  
19 (that is, the innermost **parallel**, **sections**, loop, or **taskgroup** construct).

## Description

The **cancel** construct activates cancellation of the binding construct only if *cancel-var* is **true**, in which case the construct causes the encountering task to continue execution at the end of the canceled construct. If *cancel-var* is **false**, the **cancel** construct is ignored.

Threads check for active cancellation only at cancellation points. Cancellation points are implied at the following locations:

- implicit barriers;
- **barrier** regions;
- **cancel** regions;
- **cancellation point** regions;

When a thread reaches one of the above cancellation points and if *cancel-var* is *true*, the thread immediately checks for active cancellation (that is, if cancellation has been activated by a **cancel** construct). If cancellation is active, the encountering thread continues execution at the end of the canceled construct.

---

**Note** – If one thread activates cancellation and another thread encounters a cancellation point, the absolute order of execution between the two threads is non-deterministic. Whether the thread that encounters a cancellation point detects the activated cancellation depends on the underlying hardware and operating system.

---

When cancellation of tasks is activated through the **cancel taskgroup** construct, the innermost enclosing **taskgroup** will be canceled. The task that encountered the **cancel taskgroup** construct continues execution at the end of its **task** region, which implies completion of that task. Any task that belongs to the innermost enclosing **taskgroup** and has already begun execution must run to completion or until a cancellation point is reached. Upon reaching a cancellation point and if cancellation is active, the task continues execution at the end of its **task** region, which implies the task's completion. Any task that belongs to the innermost enclosing **taskgroup** and that has not begun execution may be discarded, which implies its completion.

When cancellation is active for a **parallel**, **sections**, **for**, or **do** region, each thread of the binding thread set resumes execution at the end of the canceled region if a cancellation point is encountered. If the canceled region is a **parallel** region, any tasks that have been created by a **task** construct and their descendent tasks are canceled according to the above **taskgroup** cancellation semantics. If the canceled region is a **sections**, **for**, or **do** region, no task cancellation occurs.



## C++

1 The usual C++ rules for object destruction are followed when cancellation is performed.

## C++

## Fortran

2 All private objects or subobjects with **ALLOCATABLE** attribute that are allocated inside the  
3 canceled construct are deallocated.

## Fortran

4 **Note** – The programmer is responsible for releasing locks and other synchronization data structures  
5 that might cause a deadlock when a **cancel** construct is encountered and blocked threads cannot  
6 be canceled. The programmer is also responsible for ensuring proper synchronization to avoid  
7 deadlocks that might arise from cancellation of OpenMP regions that contain OpenMP  
8 synchronization constructs.

9 If the canceled construct contains a **reduction** or **lastprivate** clause, the final value of the  
10 **reduction** or **lastprivate** variable is undefined.

11 When an **if** clause is present on a **cancel** construct and the **if** expression evaluates to *false*, the  
12 **cancel** construct does not activate cancellation. The cancellation point associated with the  
13 **cancel** construct is always encountered regardless of the value of the **if** expression.

## Restrictions

14 The restrictions to the **cancel** construct are as follows:

- 15 • The behavior for concurrent cancellation of a region and a region nested within it is unspecified.
- 16 • If *construct-type-clause* is **taskgroup**, the **cancel** construct must be closely nested inside a  
17 **task** construct. Otherwise, the **cancel** construct must be closely nested inside an OpenMP  
18 construct that matches the type specified in *construct-type-clause* of the **cancel** construct.
- 19 • If *construct-type-clause* is **taskgroup** and the **cancel** construct is not nested inside a  
20 **taskgroup** region, then the behavior is unspecified.
- 21 • A worksharing construct that is canceled must not have a **nowait** clause.
- 22 • A loop construct that is canceled must not have an **ordered** clause.
- 23 • A construct that may be subject to cancellation must not encounter an orphaned cancellation  
24 point. That is, a cancellation point must only be encountered within that construct and must not  
25 be encountered elsewhere in its region.
- 26

## Cross References

- *cancel-var*, see Section 2.3.1 on page 35.
- **cancellation point** construct, see Section 2.13.2 on page 160.
- **omp\_get\_cancellation** routine, see Section 3.2.9 on page 217.

## 2.13.2 cancellation point Construct

### Summary

The **cancellation point** construct introduces a user-defined cancellation point at which implicit or explicit tasks check if cancellation of the innermost enclosing region of the type specified has been activated. The **cancellation point** construct is a stand-alone directive.

### Syntax

C / C++

The syntax of the **cancellation point** construct is as follows:

```
#pragma omp cancellation point construct-type-clause new-line
```

where *construct-type-clause* is one of the following:

```
parallel  
sections  
for  
taskgroup
```

C / C++

1 The syntax of the **cancellation point** construct is as follows:

```
!$omp cancellation point construct-type-clause
```

2 where *construct-type-clause* is one of the following:

3       **parallel**

4       **sections**

5       **do**

6       **taskgroup**

## 7 **Binding**

8 A **cancellation point** region binds to the current task region.

## 9 **Description**

10 This directive introduces a user-defined cancellation point at which an implicit or explicit task must  
 11 check if cancellation of the innermost enclosing region of the type specified in the clause has been  
 12 requested. This construct does not implement a synchronization between threads or tasks.

13 When an implicit or explicit task reaches a user-defined cancellation point and if *cancel-var* is  
 14 **true** the task immediately checks whether cancellation of the region specified in the clause has  
 15 been activated. If so, the encountering task continues execution at the end of the canceled construct.

## 16 **Restrictions**

- 17 • A **cancellation point** construct for which *construct-type-clause* is **taskgroup** must be  
 18 closely nested inside a **task** construct. A **cancellation point** construct for which  
 19 *construct-type-clause* is not **taskgroup** must be closely nested inside an OpenMP construct  
 20 that matches the type specified in *construct-type-clause*.
- 21 • An OpenMP program with orphaned **cancellation point** constructs is non-conforming.

## 22 **Cross References**

- 23 • *cancel-var*, see Section 2.3.1 on page 35.
- 24 • **cancel** construct, see Section 2.13.1 on page 156.
- 25 • **omp\_get\_cancellation** routine, see Section 3.2.9 on page 217.

## 1 2.14 Data Environment

2 This section presents a directive and several clauses for controlling the data environment during the  
3 execution of **parallel**, **task**, **simd**, and worksharing regions.

- 4 • Section 2.14.1 on page 162 describes how the data-sharing attributes of variables referenced in  
5 **parallel**, **task**, **simd**, and worksharing regions are determined.
- 6 • The **threadprivate** directive, which is provided to create threadprivate memory, is  
7 described in Section 2.14.2 on page 166.
- 8 • Clauses that may be specified on directives to control the data-sharing attributes of variables  
9 referenced in **parallel**, **task**, **simd** or worksharing constructs are described in  
10 Section 2.14.3 on page 172
- 11 • Clauses that may be specified on directives to copy data values from private or threadprivate  
12 variables on one thread to the corresponding variables on other threads in the team are described  
13 in Section 2.14.4 on page 191.
- 14 • Clauses that may be specified on directives to map variables to devices are described in  
15 Section 2.14.5 on page 195.

### 16 2.14.1 Data-sharing Attribute Rules

17 This section describes how the data-sharing attributes of variables referenced in **parallel**, **task**,  
18 **simd**, and worksharing regions are determined. The following two cases are described separately:

- 19 • Section 2.14.1.1 on page 162 describes the data-sharing attribute rules for variables referenced in  
20 a construct.
- 21 • Section 2.14.1.2 on page 166 describes the data-sharing attribute rules for variables referenced in  
22 a region, but outside any construct.

#### 23 2.14.1.1 Data-sharing Attribute Rules for Variables Referenced 24 in a Construct

25 The data-sharing attributes of variables that are referenced in a construct can be *predetermined*,  
26 *explicitly determined*, or *implicitly determined*, according to the rules outlined in this section.

27 Specifying a variable on a **firstprivate**, **lastprivate**, **linear**, **reduction**, or  
28 **copyprivate** clause of an enclosed construct causes an implicit reference to the variable in the  
29 enclosing construct. Specifying a variable on a **map** clause of an enclosed construct may cause an

1 implicit reference to the variable in the enclosing construct. Such implicit references are also  
2 subject to the data-sharing attribute rules outlined in this section.

3 Certain variables and objects have predetermined data-sharing attributes as follows:

▼ **C / C++** ▼

- 4 ● Variables appearing in **threadprivate** directives are threadprivate.
- 5 ● Variables with automatic storage duration that are declared in a scope inside the construct are  
6 private.
- 7 ● Objects with dynamic storage duration are shared.
- 8 ● Static data members are shared.
- 9 ● The loop iteration variable(s) in the associated *for-loop(s)* of a **for** or **parallel for**  
10 construct is (are) private.
- 11 ● The loop iteration variable in the associated *for-loop* of a **simd** construct with just one associated  
12 *for-loop* is linear with a *constant-linear-step* that is the increment of the associated *for-loop*.
- 13 ● The loop iteration variables in the associated *for-loops* of a **simd** construct with multiple  
14 associated *for-loops* are lastprivate.
- 15 ● Variables with static storage duration that are declared in a scope inside the construct are shared.

▲ **C / C++** ▲

▼ **Fortran** ▼

- 16 ● Variables and common blocks appearing in **threadprivate** directives are threadprivate.
- 17 ● The loop iteration variable(s) in the associated *do-loop(s)* of a **do** or **parallel do** construct is  
18 (are) private.
- 19 ● The loop iteration variable in the associated *do-loop* of a **simd** construct with just one associated  
20 *do-loop* is linear with a *constant-linear-step* that is the increment of the associated *do-loop*.
- 21 ● The loop iteration variables in the associated *do-loops* of a **simd** construct with multiple  
22 associated *do-loops* are lastprivate.
- 23 ● A loop iteration variable for a sequential loop in a **parallel** or **task** construct is private in  
24 the innermost such construct that encloses the loop.
- 25 ● Implied-do indices and **forall** indices are private.
- 26 ● Cray pointees have the same the data-sharing attribute as the storage with which their Cray  
27 pointers are associated.
- 28 ● Assumed-size arrays are shared.
- 29 ● An associate name preserves the association with the selector established at the **ASSOCIATE**  
30 statement.

## Fortran

1 Variables with predetermined data-sharing attributes may not be listed in data-sharing attribute  
2 clauses, except for the cases listed below. For these exceptions only, listing a predetermined  
3 variable in a data-sharing attribute clause is allowed and overrides the variable's predetermined  
4 data-sharing attributes.

## C / C++

- 5 • The loop iteration variable(s) in the associated *for-loop(s)* of a **for** or **parallel for**  
6 construct may be listed in a **private** or **lastprivate** clause.
- 7 • The loop iteration variable in the associated *for-loop* of a **simd** construct with just one  
8 associated *for-loop* may be listed in a **linear** clause with a *constant-linear-step* that is the  
9 increment of the associated *for-loop*.
- 10 • The loop iteration variables in the associated *for-loops* of a **simd** construct with multiple  
11 associated *for-loops* may be listed in a **lastprivate** clause.
- 12 • Variables with **const**-qualified type having no mutable member may be listed in a  
13 **firstprivate** clause, even if they are static data members.

## C / C++

## Fortran

- 14 • The loop iteration variable(s) in the associated *do-loop(s)* of a **do** or **parallel do** construct  
15 may be listed in a **private** or **lastprivate** clause.
- 16 • The loop iteration variable in the associated *do-loop* of a **simd** construct with just one  
17 associated *do-loop* may be listed in a **linear** clause with a *constant-linear-step* that is the  
18 increment of the associated loop.
- 19 • The loop iteration variables in the associated *do-loops* of a **simd** construct with multiple  
20 associated *do-loops* may be listed in a **lastprivate** clause.
- 21 • Variables used as loop iteration variables in sequential loops in a **parallel** or **task** construct  
22 may be listed in data-sharing clauses on the construct itself, and on enclosed constructs, subject  
23 to other restrictions.
- 24 • Assumed-size arrays may be listed in a **shared** clause.

## Fortran

1 Additional restrictions on the variables that may appear in individual clauses are described with  
2 each clause in Section 2.14.3 on page 172.

3 Variables with *explicitly determined* data-sharing attributes are those that are referenced in a given  
4 construct and are listed in a data-sharing attribute clause on the construct.

5 Variables with *implicitly determined* data-sharing attributes are those that are referenced in a given  
6 construct, do not have predetermined data-sharing attributes, and are not listed in a data-sharing  
7 attribute clause on the construct.

8 Rules for variables with *implicitly determined* data-sharing attributes are as follows:

- 9 • In a **parallel** or **task** construct, the data-sharing attributes of these variables are determined  
10 by the **default** clause, if present (see Section 2.14.3.1 on page 173).
- 11 • In a **parallel** construct, if no **default** clause is present, these variables are shared.
- 12 • For constructs other than **task**, if no **default** clause is present, these variables reference the  
13 variables with the same names that exist in the enclosing context.
- 14 • In a **task** construct, if no **default** clause is present, a variable that in the enclosing context is  
15 determined to be shared by all implicit tasks bound to the current team is shared.

## Fortran

- 16 • In an orphaned **task** construct, if no **default** clause is present, dummy arguments are  
17 firstprivate.

## Fortran

- 18 • In a **task** construct, if no **default** clause is present, a variable whose data-sharing attribute is  
19 not determined by the rules above is firstprivate.

20 Additional restrictions on the variables for which data-sharing attributes cannot be implicitly  
21 determined in a **task** construct are described in Section 2.14.3.4 on page 179.

## 1 2.14.1.2 Data-sharing Attribute Rules for Variables Referenced 2 in a Region but not in a Construct

3 The data-sharing attributes of variables that are referenced in a region, but not in a construct, are  
4 determined as follows:

▼ C / C++ ▼

- 5 • Variables with static storage duration that are declared in called routines in the region are shared.
- 6 • File-scope or namespace-scope variables referenced in called routines in the region are shared  
7 unless they appear in a **threadprivate** directive.
- 8 • Objects with dynamic storage duration are shared.
- 9 • Static data members are shared unless they appear in a **threadprivate** directive.
- 10 • In C++, formal arguments of called routines in the region that are passed by reference have the  
11 same data-sharing attributes as the associated actual arguments.
- 12 • Other variables declared in called routines in the region are private.

▲ C / C++ ▲  
▼ Fortran ▼

- 13 • Local variables declared in called routines in the region and that have the **save** attribute, or that  
14 are data initialized, are shared unless they appear in a **threadprivate** directive.
- 15 • Variables belonging to common blocks, or accessed by host or use association, and referenced in  
16 called routines in the region are shared unless they appear in a **threadprivate** directive.
- 17 • Dummy arguments of called routines in the region that are passed by reference have the same  
18 data-sharing attributes as the associated actual arguments.
- 19 • Cray pointees have the same data-sharing attribute as the storage with which their Cray pointers  
20 are associated.
- 21 • Implied-do indices, **forall** indices, and other local variables declared in called routines in the  
22 region are private.

▲ Fortran ▲

## 23 2.14.2 **threadprivate** Directive

### 24 Summary

25 The **threadprivate** directive specifies that variables are replicated, with each thread having its  
26 own copy. The **threadprivate** directive is a declarative directive.



## Syntax

C / C++

The syntax of the **threadprivate** directive is as follows:

```
#pragma omp threadprivate (list) new-line
```

where *list* is a comma-separated list of file-scope, namespace-scope, or static block-scope variables that do not have incomplete types.

C / C++

Fortran

The syntax of the **threadprivate** directive is as follows:

```
!$omp threadprivate (list)
```

where *list* is a comma-separated list of named variables and named common blocks. Common block names must appear between slashes.

Fortran

## Description

Each copy of a threadprivate variable is initialized once, in the manner specified by the program, but at an unspecified point in the program prior to the first reference to that copy. The storage of all copies of a threadprivate variable is freed according to how static variables are handled in the base language, but at an unspecified point in the program.

A program in which a thread references another thread's copy of a threadprivate variable is non-conforming.

The content of a threadprivate variable can change across a task scheduling point if the executing thread switches to another task that modifies the variable. For more details on task scheduling, see Section 1.3 on page 13 and Section 2.9 on page 80.

In **parallel** regions, references by the master thread will be to the copy of the variable in the thread that encountered the **parallel** region.

During a sequential part references will be to the initial thread's copy of the variable. The values of data in the initial thread's copy of a threadprivate variable are guaranteed to persist between any two consecutive references to the variable in the program.

The values of data in the threadprivate variables of non-initial threads are guaranteed to persist between two consecutive active **parallel** regions only if all the following conditions hold:

- Neither **parallel** region is nested inside another explicit **parallel** region.
- The number of threads used to execute both **parallel** regions is the same.
- The thread affinity policies used to execute both **parallel** regions are the same.
- The value of the *dyn-var* internal control variable in the enclosing task region is *false* at entry to both **parallel** regions.

If these conditions all hold, and if a threadprivate variable is referenced in both regions, then threads with the same thread number in their respective regions will reference the same copy of that variable.

▼ C / C++ ▼

If the above conditions hold, the storage duration, lifetime, and value of a thread's copy of a threadprivate variable that does not appear in any **copyin** clause on the second region will be retained. Otherwise, the storage duration, lifetime, and value of a thread's copy of the variable in the second region is unspecified.

If the value of a variable referenced in an explicit initializer of a threadprivate variable is modified prior to the first reference to any instance of the threadprivate variable, then the behavior is unspecified.

▲ C / C++ ▲

## C++

1 The order in which any constructors for different threadprivate variables of class type are called is  
2 unspecified. The order in which any destructors for different threadprivate variables of class type  
3 are called is unspecified.

## C++

## Fortran

4 A variable is affected by a **copyin** clause if the variable appears in the **copyin** clause or it is in a  
5 common block that appears in the **copyin** clause.

6 If the above conditions hold, the definition, association, or allocation status of a thread's copy of a  
7 **threadprivate** variable or a variable in a **threadprivate** common block, that is not  
8 affected by any **copyin** clause that appears on the second region, will be retained. Otherwise, the  
9 definition and association status of a thread's copy of the variable in the second region is undefined,  
10 and the allocation status of an allocatable variable will be implementation defined.

11 If a **threadprivate** variable or a variable in a **threadprivate** common block is not  
12 affected by any **copyin** clause that appears on the first **parallel** region in which it is  
13 referenced, the variable or any subobject of the variable is initially defined or undefined according  
14 to the following rules:

- 15 ● If it has the **ALLOCATABLE** attribute, each copy created will have an initial allocation status of  
16 not currently allocated.
- 17 ● If it has the **POINTER** attribute:
  - 18 – if it has an initial association status of disassociated, either through explicit initialization or  
19 default initialization, each copy created will have an association status of disassociated;
  - 20 – otherwise, each copy created will have an association status of undefined.
- 21 ● If it does not have either the **POINTER** or the **ALLOCATABLE** attribute:
  - 22 – if it is initially defined, either through explicit initialization or default initialization, each copy  
23 created is so defined;
  - 24 – otherwise, each copy created is undefined.

## Fortran

## Restrictions

The restrictions to the **threadprivate** directive are as follows:

- A **threadprivate** variable must not appear in any clause except the **copyin**, **copyprivate**, **schedule**, **num\_threads**, **thread\_limit**, and **if** clauses.
- A program in which an untied task accesses **threadprivate** storage is non-conforming.

C / C++

- A variable that is part of another variable (as an array or structure element) cannot appear in a **threadprivate** clause unless it is a static data member of a C++ class.
- A **threadprivate** directive for file-scope variables must appear outside any definition or declaration, and must lexically precede all references to any of the variables in its list.
- A **threadprivate** directive for namespace-scope variables must appear outside any definition or declaration other than the namespace definition itself, and must lexically precede all references to any of the variables in its list.
- Each variable in the list of a **threadprivate** directive at file, namespace, or class scope must refer to a variable declaration at file, namespace, or class scope that lexically precedes the directive.
- A **threadprivate** directive for static block-scope variables must appear in the scope of the variable and not in a nested scope. The directive must lexically precede all references to any of the variables in its list.
- Each variable in the list of a **threadprivate** directive in block scope must refer to a variable declaration in the same scope that lexically precedes the directive. The variable declaration must use the static storage-class specifier.
- If a variable is specified in a **threadprivate** directive in one translation unit, it must be specified in a **threadprivate** directive in every translation unit in which it is declared.
- The address of a **threadprivate** variable is not an address constant.

C / C++

## C++

- 1 • A **threadprivate** directive for static class member variables must appear in the class  
2 definition, in the same scope in which the member variables are declared, and must lexically  
3 precede all references to any of the variables in its list.
- 4 • A threadprivate variable must not have an incomplete type or a reference type.
- 5 • A threadprivate variable with class type must have:
  - 6 – an accessible, unambiguous default constructor in case of default initialization without a given  
7 initializer;
  - 8 – an accessible, unambiguous constructor accepting the given argument in case of direct  
9 initialization;
  - 10 – an accessible, unambiguous copy constructor in case of copy initialization with an explicit  
11 initializer

## C++

## Fortran

- 12 • A variable that is part of another variable (as an array or structure element) cannot appear in a  
13 **threadprivate** clause.
- 14 • The **threadprivate** directive must appear in the declaration section of a scoping unit in  
15 which the common block or variable is declared. Although variables in common blocks can be  
16 accessed by use association or host association, common block names cannot. This means that a  
17 common block name specified in a **threadprivate** directive must be declared to be a  
18 common block in the same scoping unit in which the **threadprivate** directive appears.
- 19 • If a **threadprivate** directive specifying a common block name appears in one program unit,  
20 then such a directive must also appear in every other program unit that contains a **COMMON**  
21 statement specifying the same name. It must appear after the last such **COMMON** statement in the  
22 program unit.
- 23 • If a **threadprivate** variable or a **threadprivate** common block is declared with the  
24 **BIND** attribute, the corresponding C entities must also be specified in a **threadprivate**  
25 directive in the C program.
- 26 • A blank common block cannot appear in a **threadprivate** directive.
- 27 • A variable can only appear in a **threadprivate** directive in the scope in which it is declared.  
28 It must not be an element of a common block or appear in an **EQUIVALENCE** statement.
- 29 • A variable that appears in a **threadprivate** directive must be declared in the scope of a  
30 module or have the **SAVE** attribute, either explicitly or implicitly.

## Fortran

## Cross References

- *dyn-var* ICV, see Section 2.3 on page 35.
- number of threads used to execute a **parallel** region, see Section 2.5.1 on page 48.
- **copyin** clause, see Section 2.14.4.1 on page 192.

### 2.14.3 Data-Sharing Attribute Clauses

Several constructs accept clauses that allow a user to control the data-sharing attributes of variables referenced in the construct. Data-sharing attribute clauses apply only to variables for which the names are visible in the construct on which the clause appears.

Not all of the clauses listed in this section are valid on all directives. The set of clauses that is valid on a particular directive is described with the directive.

Most of the clauses accept a comma-separated list of list items (see Section 2.1 on page 26). All list items appearing in a clause must be visible, according to the scoping rules of the base language. With the exception of the **default** clause, clauses may be repeated as needed. A list item that specifies a given variable may not appear in more than one clause on the same directive, except that a variable may be specified in both **firstprivate** and **lastprivate** clauses.

C++

If a variable referenced in a data-sharing attribute clause has a type derived from a template, and there are no other references to that variable in the program, then any behavior related to that variable is unspecified.

C++

Fortran

A named common block may be specified in a list by enclosing the name in slashes. When a named common block appears in a list, it has the same meaning as if every explicit member of the common block appeared in the list. An explicit member of a common block is a variable that is named in a **COMMON** statement that specifies the common block name and is declared in the same scoping unit in which the clause appears.

Although variables in common blocks can be accessed by use association or host association, common block names cannot. As a result, a common block name specified in a data-sharing attribute clause must be declared to be a common block in the same scoping unit in which the data-sharing attribute clause appears.

When a named common block appears in a **private**, **firstprivate**, **lastprivate**, or **shared** clause of a directive, none of its members may be declared in another data-sharing

1 attribute clause in that directive. When individual members of a common block appear in a  
2 **private**, **firstprivate**, **lastprivate**, **reduction**, or **linear** clause of a directive,  
3 the storage of the specified variables is no longer Fortran associated with the storage of the common  
4 block itself.



### 5 **2.14.3.1 default clause**

#### 6 **Summary**

7 The **default** clause explicitly determines the data-sharing attributes of variables that are  
8 referenced in a **parallel**, **task** or **teams** construct and would otherwise be implicitly  
9 determined (see Section 2.14.1.1 on page 162).

#### 10 **Syntax**



11 The syntax of the **default** clause is as follows:

```
default (shared | none)
```



12 The syntax of the **default** clause is as follows:

```
default (private | firstprivate | shared | none)
```



## Description

The **default (shared)** clause causes all variables referenced in the construct that have implicitly determined data-sharing attributes to be shared.

Fortran

The **default (firstprivate)** clause causes all variables in the construct that have implicitly determined data-sharing attributes to be firstprivate.

The **default (private)** clause causes all variables referenced in the construct that have implicitly determined data-sharing attributes to be private.

Fortran

The **default (none)** clause requires that each variable that is referenced in the construct, and that does not have a predetermined data-sharing attribute, must have its data-sharing attribute explicitly determined by being listed in a data-sharing attribute clause.

## Restrictions

The restrictions to the **default** clause are as follows:

- Only a single default clause may be specified on a **parallel**, **task**, or **teams** directive.

### 2.14.3.2 shared clause

#### Summary

The **shared** clause declares one or more list items to be shared by tasks generated by a **parallel**, **task** or **teams** construct.

#### Syntax

The syntax of the **shared** clause is as follows:

```
shared (list)
```



## Description

All references to a list item within a task refer to the storage area of the original variable at the point the directive was encountered.

It is the programmer's responsibility to ensure, by adding proper synchronization, that storage shared by an explicit **task** region does not reach the end of its lifetime before the explicit **task** region completes its execution.

### Fortran

The association status of a shared pointer becomes undefined upon entry to and on exit from the **parallel**, **task** or **teams** construct if it is associated with a target or a subobject of a target that is in a **private**, **firstprivate**, **lastprivate**, or **reduction** clause inside the construct.

Under certain conditions, passing a shared variable to a non-intrinsic procedure may result in the value of the shared variable being copied into temporary storage before the procedure reference, and back out of the temporary storage into the actual argument storage after the procedure reference. It is implementation defined when this situation occurs

**Note** – Use of intervening temporary storage may occur when the following three conditions hold regarding an actual argument in a reference to a non-intrinsic procedure:

a The actual argument is one of the following:

- A shared variable.
- A subobject of a shared variable.
- An object associated with a shared variable.
- An object associated with a subobject of a shared variable.

b The actual argument is also one of the following:

- An array section.
- An array section with a vector subscript.
- An assumed-shape array.
- A pointer array.

c The associated dummy argument for this actual argument is an explicit-shape array or an assumed-size array.

1 These conditions effectively result in references to, and definitions of, the temporary storage during  
2 the procedure reference. Any references to (or definitions of) the shared storage that is associated  
3 with the dummy argument by any other task must be synchronized with the procedure reference to  
4 avoid possible race conditions.



## 5 **Restrictions**

6 The restrictions for the **shared** clause are as follows:

- 7 • A variable that is part of another variable (as an array or structure element) cannot appear in a  
8 shared clause.

### 9 **2.14.3.3 private clause**

#### 10 **Summary**

11 The **private** clause declares one or more list items to be private to a task or to a SIMD lane.

#### 12 **Syntax**

13 The syntax of the private clause is as follows:

```
private (list)
```

## Description

Each task that references a list item that appears in a **private** clause in any statement in the construct receives a new list item. Each SIMD lane used in a **simd** construct that references a list item that appears in a private clause in any statement in the construct receives a new list item. Language-specific attributes for new list items are derived from the corresponding original list item. Inside the construct, all references to the original list item are replaced by references to the new list item. In the rest of the region, it is unspecified whether references are to the new list item or the original list item. Therefore, if an attempt is made to reference the original item, its value after the region is also unspecified. If a SIMD construct or a task does not reference a list item that appears in a **private** clause, it is unspecified whether SIMD lanes or the task receive a new list item.

The value and/or allocation status of the original list item will change only:

- if accessed and modified via pointer,
- if possibly accessed in the region but outside of the construct,
- as a side effect of directives or clauses, or

▼ Fortran ▼

- if accessed and modified via construct association.

▲ Fortran ▲

List items that appear in a **private**, **firstprivate**, or **reduction** clause in a **parallel** construct may also appear in a **private** clause in an enclosed **parallel**, **task**, or worksharing, or **simd** construct.

List items that appear in a **private** or **firstprivate** clause in a **task** construct may also appear in a **private** clause in an enclosed **parallel** or **task** construct.

List items that appear in a **private**, **firstprivate**, **lastprivate**, or **reduction** clause in a worksharing construct may also appear in a **private** clause in an enclosed **parallel** or **task** construct.

▼ C / C++ ▼

If the type of a list item is a reference to a type  $T$  then the type will be considered to be  $T$  for all purposes of this clause.

A new list item of the same type, with automatic storage duration, is allocated for the construct. The storage and thus lifetime of these list items lasts until the block in which they are created exits. The size and alignment of the new list item are determined by the type of the variable. This allocation occurs once for each task generated by the construct and/or once for each SIMD lane used by the construct.

The new list item is initialized, or has an undefined initial value, as if it had been locally declared without an initializer.

▲ C / C++ ▲

## C++

1 The order in which any default constructors for different private variables of class type are called is  
2 unspecified. The order in which any destructors for different private variables of class type are  
3 called is unspecified.

## C++

## Fortran

4 If any statement of the construct references a list item, a new list item of the same type and type  
5 parameters is allocated: once for each implicit task in the **parallel** construct; once for each task  
6 generated by a **task** construct; and once for each SIMD lane used by a **simd** construct. The initial  
7 value of the new list item is undefined. Within a **parallel**, **worksharing**, **task**, **teams**, or  
8 **simd** region, the initial status of a private pointer is undefined.

9 For a list item or the subobject of a list item with the **ALLOCATABLE** attribute:

- 10 • if the allocation status is “not currently allocated”, the new list item or the subobject of the new  
11 list item will have an initial allocation status of "not currently allocated";
- 12 • if the allocation status is “currently allocated”, the new list item or the subobject of the new list  
13 item will have an initial allocation status of "currently allocated". If the new list item or the  
14 subobject of the new list item is an array, its bounds will be the same as those of the original list  
15 item or the subobject of the original list item.

16 A list item that appears in a **private** clause may be storage-associated with other variables when  
17 the **private** clause is encountered. Storage association may exist because of constructs such as  
18 **EQUIVALENCE** or **COMMON**. If *A* is a variable appearing in a **private** clause and *B* is a variable  
19 that is storage-associated with *A*, then:

- 20 • The contents, allocation, and association status of *B* are undefined on entry to the **parallel**,  
21 **task**, **simd**, or **teams** region.
- 22 • Any definition of *A*, or of its allocation or association status, causes the contents, allocation, and  
23 association status of *B* to become undefined.
- 24 • Any definition of *B*, or of its allocation or association status, causes the contents, allocation, and  
25 association status of *A* to become undefined.

26 A list item that appears in a **private** clause may be a selector of an **ASSOCIATE** construct. If the  
27 construct association is established prior to a **parallel** region, the association between the  
28 associate name and the original list item will be retained in the region.

29 Finalization of a list item of a finalizable type or subobjects of a list item of a finalizable type occurs  
30 at the end of the region. The order in which any final subroutines for different variables of a  
31 finalizable type are called is unspecified.

## Fortran

## Restrictions

The restrictions to the **private** clause are as follows:

- A variable that is part of another variable (as an array or structure element) cannot appear in a **private** clause.

C / C++

- A variable of class type (or array thereof) that appears in a **private** clause requires an accessible, unambiguous default constructor for the class type.
- A variable that appears in a **private** clause must not have a **const**-qualified type unless it is of class type with a **mutable** member. This restriction does not apply to the **firstprivate** clause.
- A variable that appears in a **private** clause must not have an incomplete type or be a reference to an incomplete type.
- If a list item is a reference type then it must bind to the same object for all threads of the team.

C / C++

Fortran

- A variable that appears in a **private** clause must either be definable, or an allocatable variable. This restriction does not apply to the **firstprivate** clause.
- Variables that appear in namelist statements, in variable format expressions, and in expressions for statement function definitions, may not appear in a **private** clause.
- Pointers with the **INTENT (IN)** attribute may not appear in a **private** clause. This restriction does not apply to the **firstprivate** clause.

Fortran

### 2.14.3.4 **firstprivate** clause

#### Summary

The **firstprivate** clause declares one or more list items to be private to a task, and initializes each of them with the value that the corresponding original item has when the construct is encountered.

#### Syntax

The syntax of the **firstprivate** clause is as follows:

## `firstprivate (list)`

### Description

The **firstprivate** clause provides a superset of the functionality provided by the **private** clause.

A list item that appears in a **firstprivate** clause is subject to the **private** clause semantics described in Section 2.14.3.3 on page 176, except as noted. In addition, the new list item is initialized from the original list item existing before the construct. The initialization of the new list item is done once for each task that references the list item in any statement in the construct. The initialization is done prior to the execution of the construct.

For a **firstprivate** clause on a **parallel**, **task**, or **teams** construct, the initial value of the new list item is the value of the original list item that exists immediately prior to the construct in the task region where the construct is encountered. For a **firstprivate** clause on a worksharing construct, the initial value of the new list item for each implicit task of the threads that execute the worksharing construct is the value of the original list item that exists in the implicit task immediately prior to the point in time that the worksharing construct is encountered.

To avoid race conditions, concurrent updates of the original list item must be synchronized with the read of the original list item that occurs as a result of the **firstprivate** clause.

If a list item appears in both **firstprivate** and **lastprivate** clauses, the update required for **lastprivate** occurs after all the initializations for **firstprivate**.

▼ C / C++ ▼

For variables of non-array type, the initialization occurs by copy assignment. For an array of elements of non-array type, each element is initialized as if by assignment from an element of the original array to the corresponding element of the new array.

▲ C / C++ ▲

▼ C++ ▼

For variables of class type, a copy constructor is invoked to perform the initialization. The order in which copy constructors for different variables of class type are called is unspecified.

▲ C++ ▲

▼ Fortran ▼

If the original list item does not have the **POINTER** attribute, initialization of the new allocation status of not currently allocated, in which case the new list items will have the same status.

If the original list item has the **POINTER** attribute, the new list items receive the same association status of the original list item as if by pointer assignment.

▲ Fortran ▲

## Restrictions

The restrictions to the **firstprivate** clause are as follows:

- A variable that is part of another variable (as an array or structure element) cannot appear in a **firstprivate** clause.
- A list item that is private within a **parallel** region must not appear in a **firstprivate** clause on a worksharing construct if any of the worksharing regions arising from the worksharing construct ever bind to any of the **parallel** regions arising from the **parallel** construct.
- A list item that is private within a **teams** region must not appear in a **firstprivate** clause on a **distribute** construct if any of the **distribute** regions arising from the **distribute** construct ever bind to any of the **teams** regions arising from the **teams** construct.
- A list item that appears in a **reduction** clause of a **parallel** construct must not appear in a **firstprivate** clause on a worksharing or **task** construct if any of the worksharing or **task** regions arising from the worksharing or **task** construct ever bind to any of the **parallel** regions arising from the **parallel** construct.
- A list item that appears in a **reduction** clause of a **teams** construct must not appear in a **firstprivate** clause on a **distribute** construct if any of the **distribute** regions arising from the **distribute** construct ever bind to any of the **teams** regions arising from the **teams** construct.
- A list item that appears in a **reduction** clause in a worksharing construct must not appear in a **firstprivate** clause in a task construct encountered during execution of any of the worksharing regions arising from the worksharing construct.

C++

- A variable of class type (or array thereof) that appears in a **firstprivate** clause requires an accessible, unambiguous copy constructor for the class type.

C++

C / C++

- A variable that appears in a **firstprivate** clause must not have an incomplete C/C++ type or be a reference to an incomplete type.
- If a list item is a reference type then it must bind to the same object for all threads of the team.

C / C++

Fortran

- Variables that appear in namelist statements, in variable format expressions, and in expressions for statement function definitions, may not appear in a **firstprivate** clause.

Fortran

## 1 2.14.3.5 `lastprivate` clause

### 2 Summary

3 The `lastprivate` clause declares one or more list items to be private to an implicit task or to a  
4 SIMD lane, and causes the corresponding original list item to be updated after the end of the region.

### 5 Syntax

6 The syntax of the `lastprivate` clause is as follows:

```
lastprivate (list)
```

### 7 Description

8 The `lastprivate` clause provides a superset of the functionality provided by the `private`  
9 clause.

10 A list item that appears in a `lastprivate` clause is subject to the `private` clause semantics  
11 described in Section 2.14.3.3 on page 176. In addition, when a `lastprivate` clause appears on  
12 the directive that identifies a worksharing construct or a SIMD construct, the value of each new list  
13 item from the sequentially last iteration of the associated loops, or the lexically last `section`  
14 construct, is assigned to the original list item.

▼ C / C++ ▼

15 For an array of elements of non-array type, each element is assigned to the corresponding element  
16 of the original array.

▲ C / C++ ▲

▼ Fortran ▼

17 If the original list item does not have the `POINTER` attribute, its update occurs as if by intrinsic  
18 assignment.

19 If the original list item has the `POINTER` attribute, its update occurs as if by pointer assignment.



1 List items that are not assigned a value by the sequentially last iteration of the loops, or by the  
 2 lexically last **section** construct, have unspecified values after the construct. Unassigned  
 3 subcomponents also have unspecified values after the construct.

4 The original list item becomes defined at the end of the construct if there is an implicit barrier at  
 5 that point. To avoid race conditions, concurrent reads or updates of the original list item must be  
 6 synchronized with the update of the original list item that occurs as a result of the **lastprivate**  
 7 clause.

8 If the **lastprivate** clause is used on a construct to which **nowait** is applied, accesses to the  
 9 original list item may create a data race. To avoid this, synchronization must be inserted to ensure  
 10 that the sequentially last iteration or lexically last section construct has stored and flushed that list  
 11 item.

12 If the **lastprivate** clause is used on a **distribute simd**, distribute parallel loop, or  
 13 distribute parallel loop SIMD, accesses to the original list item may create a data race. To avoid  
 14 this, synchronization must be inserted to ensure that the sequentially last iteration has stored and  
 15 flushed that list item.

16 If a list item appears in both **firstprivate** and **lastprivate** clauses, the update required  
 17 for **lastprivate** occurs after all initializations for **firstprivate**.

## 18 Restrictions

19 The restrictions to the **lastprivate** clause are as follows:

- 20 • A variable that is part of another variable (as an array or structure element) cannot appear in a  
 21 **lastprivate** clause.
- 22 • A list item that is private within a **parallel** region, or that appears in the **reduction** clause  
 23 of a **parallel** construct, must not appear in a **lastprivate** clause on a worksharing  
 24 construct if any of the corresponding worksharing regions ever binds to any of the corresponding  
 25 **parallel** regions.

## C++

26 • A variable of class type (or array thereof) that appears in a **lastprivate** clause requires an  
 27 accessible, unambiguous default constructor for the class type, unless the list item is also  
 28 specified in a **firstprivate** clause.

29 • A variable of class type (or array thereof) that appears in a **lastprivate** clause requires an  
 30 accessible, unambiguous copy assignment operator for the class type. The order in which copy  
 31 assignment operators for different variables of class type are called is unspecified.

## C++

## C / C++

- 1 • A variable that appears in a **lastprivate** clause must not have a **const**-qualified type unless
- 2 it is of class type with a **mutable** member.
- 3 • A variable that appears in a **lastprivate** clause must not have an incomplete C/C++ type or
- 4 be a reference to an incomplete type.
- 5 • If a list item is a reference type then it must bind to the same object for all threads of the team.

## C / C++

## Fortran

- 6 • A variable that appears in a **lastprivate** clause must be definable.
- 7 • If the original list item has the **ALLOCATABLE** attribute, the corresponding list item in the
- 8 sequentially last iteration or lexically last section must have an allocation status of allocated upon
- 9 exit from that iteration or section.
- 10 • Variables that appear in namelist statements, in variable format expressions, and in expressions
- 11 for statement function definitions, may not appear in a **lastprivate** clause.

## Fortran

### 12 2.14.3.6 reduction clause

#### 13 Summary

14 The **reduction** clause specifies a *reduction-identifier* and one or more list items. For each list  
15 item, a private copy is created in each implicit task or SIMD lane, and is initialized with the  
16 initializer value of the *reduction-identifier*. After the end of the region, the original list item is  
17 updated with the values of the private copies using the combiner associated with the  
18 *reduction-identifier*.

## Syntax

C / C++

The syntax of the **reduction** clause is as follows:

```
reduction (reduction-identifier : list)
```

where:

C

*reduction-identifier* is either an *identifier* or one of the following operators: +, -, \*, &, |, ^, && and ||

C

C++

*reduction-identifier* is either an *id-expression* or one of the following operators: +, -, \*, &, |, ^, && and ||

C++

The following table lists each *reduction-identifier* that is implicitly declared at every scope for arithmetic types and its semantic initializer value. The actual initializer value is that value as expressed in the data type of the reduction list item.


Identifier	Initializer	Combiner
+	<code>omp_priv = 0</code>	<code>omp_out += omp_in</code>
*	<code>omp_priv = 1</code>	<code>omp_out *= omp_in</code>
-	<code>omp_priv = 0</code>	<code>omp_out -= omp_in</code>
&	<code>omp_priv = 0</code>	<code>omp_out &amp;= omp_in</code>
	<code>omp_priv = 0</code>	<code>omp_out  = omp_in</code>
^	<code>omp_priv = 0</code>	<code>omp_out ^= omp_in</code>
&&	<code>omp_priv = 1</code>	<code>omp_out = omp_in &amp;&amp; omp_out</code>
	<code>omp_priv = 0</code>	<code>omp_out = omp_in    omp_out</code>
max	<code>omp_priv = Least representable number in the reduction list item type</code>	<code>omp_out = omp_in &gt; omp_out ? omp_in : omp_out</code>

*table continued on next page*

table continued from previous page

Identifier	Initializer	Combiner
<b>min</b>	<b>omp_priv</b> = <i>Largest representable number in the reduction list item type</i>	<b>omp_out</b> = <b>omp_in</b> < <b>omp_out</b> ? <b>omp_in</b> : <b>omp_out</b>

1 where **omp\_in** and **omp\_out** correspond to two identifiers that refer to storage of the type of the  
2 list item. **omp\_out** holds the final value of the combiner operation.



3 The syntax of the **reduction** clause is as follows:

```
reduction(reduction-identifier : list)
```

4 where *reduction-identifier* is either a base language identifier, or a user-defined operator, or one of  
5 the following operators: **+**, **-**, **\***, **.and.**, **.or.**, **.eqv.**, **.neqv.**, or one of the following intrinsic  
6 procedure names: **max**, **min**, **iand**, **ior**, **ieor**.

7 The following table lists each *reduction-identifier* that is implicitly declared for numeric and logical  
8 types and its semantic initializer value. The actual initializer value is that value as expressed in the  
9 data type of the reduction list item.

Identifier	Initializer	Combiner
<b>+</b>	<b>omp_priv</b> = 0	<b>omp_out</b> = <b>omp_in</b> + <b>omp_out</b>
<b>*</b>	<b>omp_priv</b> = 1	<b>omp_out</b> = <b>omp_in</b> * <b>omp_out</b>
<b>-</b>	<b>omp_priv</b> = 0	<b>omp_out</b> = <b>omp_in</b> + <b>omp_out</b>
<b>.and.</b>	<b>omp_priv</b> = <b>.true.</b>	<b>omp_out</b> = <b>omp_in</b> <b>.and.</b> <b>omp_out</b>
<b>.or.</b>	<b>omp_priv</b> = <b>.false.</b>	<b>omp_out</b> = <b>omp_in</b> <b>.or.</b> <b>omp_out</b>
<b>.eqv.</b>	<b>omp_priv</b> = <b>.true.</b>	<b>omp_out</b> = <b>omp_in</b> <b>.eqv.</b> <b>omp_out</b>
<b>.neqv.</b>	<b>omp_priv</b> = <b>.false.</b>	<b>omp_out</b> = <b>omp_in</b> <b>.neqv.</b> <b>omp_out</b>

table continued on next page

table continued from previous page

Identifier	Initializer	Combiner
<code>max</code>	<code>omp_priv = Least representable number in the reduction list item type</code>	<code>omp_out = max(omp_in, omp_out)</code>
<code>min</code>	<code>omp_priv = Largest representable number in the reduction list item type</code>	<code>omp_out = min(omp_in, omp_out)</code>
<code>iand</code>	<code>omp_priv = All bits on</code>	<code>omp_out = iand(omp_in, omp_out)</code>
<code>ior</code>	<code>omp_priv = 0</code>	<code>omp_out = ior(omp_in, omp_out)</code>
<code>ieor</code>	<code>omp_priv = 0</code>	<code>omp_out = ieor(omp_in, omp_out)</code>

## Fortran

Any *reduction-identifier* that is defined with the **declare reduction** directive is also valid. In that case, the initializer and combiner of the *reduction-identifier* are specified by the *initializer-clause* and the combiner in the **declare reduction** directive.

### Description

The reduction clause can be used to perform some forms of recurrence calculations (involving mathematically associative and commutative operators) in parallel.

For **parallel** and worksharing constructs, a private copy of each list item is created, one for each implicit task, as if the **private** clause had been used. For the **simd** construct, a private copy of each list item is created, one for each SIMD lane as if the **private** clause had been used. For the **teams** construct, a private copy of each list item is created, one for each team in the league as if the **private** clause had been used. The private copy is then initialized as specified above. At the end of the region for which the **reduction** clause was specified, the original list item is updated by combining its original value with the final value of each of the private copies, using the combiner of the specified *reduction-identifier*.

The *reduction-identifier* specified in the **reduction** clause must match a previously declared *reduction-identifier* of the same name and type for each of the list items. This match is done by means of a name lookup in the base language.

- 1 The list items that appear in the **reduction** clause may include array sections.
- 2 If the list item is an array or an array section it will be treated as if a **reduction** clause would be  
3 applied to each separate element of the array section. The elements of each private array section  
4 will be allocated contiguously.
- 5 If the type is a derived class, then any *reduction-identifier* that matches its base classes are also a  
6 match, if there is no specific match for the type.
- 7 If the *reduction-identifier* is not an *id-expression* then it is implicitly converted to one by  
8 prepending the keyword operator (for example, **+** becomes *operator+*).
- 9 If the *reduction-identifier* is qualified then a qualified name lookup is used to find the declaration.
- 10 If the *reduction-identifier* is unqualified then an *argument-dependent name lookup* must be  
11 performed using the type of each list item.

- 12 If **nowait** is not used, the reduction computation will be complete at the end of the construct;  
13 however, if the reduction clause is used on a construct to which **nowait** is also applied, accesses to  
14 the original list item will create a race and, thus, have unspecified effect unless synchronization  
15 ensures that they occur after all threads have executed all of their iterations or **section** constructs,  
16 and the reduction computation has completed and stored the computed value of that list item. This  
17 can most simply be ensured through a barrier synchronization.
- 18 The location in the OpenMP program at which the values are combined and the order in which the  
19 values are combined are unspecified. Therefore, when comparing sequential and parallel runs, or  
20 when comparing one parallel run to another (even if the number of threads used is the same), there  
21 is no guarantee that bit-identical results will be obtained or that side effects (such as floating-point  
22 exceptions) will be identical or take place at the same location in the OpenMP program.
- 23 To avoid race conditions, concurrent reads or updates of the original list item must be synchronized  
24 with the update of the original list item that occurs as a result of the **reduction** computation.

## 25 Restrictions

- 26 The restrictions to the **reduction** clause are as follows:
- 27 • A list item that appears in a **reduction** clause of a worksharing construct must be shared in  
28 the **parallel** regions to which any of the worksharing regions arising from the worksharing  
29 construct bind.
  - 30 • A list item that appears in a **reduction** clause of the innermost enclosing worksharing or  
31 **parallel** construct may not be accessed in an explicit task.

- 1 • Any number of **reduction** clauses can be specified on the directive, but a list item can appear  
2 only once in the **reduction** clauses for that directive.
- 3 • For a *reduction-identifier* declared with the **declare reduction** construct, the directive  
4 must appear before its use in a **reduction** clause.
- 5 • If a list item is an array section, it must specify contiguous storage and it cannot be a zero-length  
6 array section.
- 7 • If a list item is an array section, its lower-bound must be zero.
- 8 • If a list item is an array section, accesses to the elements of the array outside the specified array  
9 section result in unspecified behavior.

C / C++

- 10 • The type of a list item that appears in a **reduction** clause must be valid for the  
11 *reduction-identifier*. For a **max** or **min** reduction in C, the type of the list item must be an  
12 allowed arithmetic data type: **char**, **int**, **float**, **double**, or **\_Bool**, possibly modified with  
13 **long**, **short**, **signed**, or **unsigned**. For a **max** or **min** reduction in C++, the type of the  
14 list item must be an allowed arithmetic data type: **char**, **wchar\_t**, **int**, **float**, **double**, or  
15 **bool**, possibly modified with **long**, **short**, **signed**, or **unsigned**.
- 16 • A list item that appears in a **reduction** clause must not be **const**-qualified.
- 17 • If a list item is a reference type then it must bind to the same object for all threads of the team.
- 18 • The *reduction-identifier* for any list item must be unambiguous and accessible.

C / C++

Fortran

- 19 • The type and the rank of a list item that appears in a **reduction** clause must be valid for the  
20 *combiner* and *initializer*.
- 21 • A list item that appears in a **reduction** clause must be definable.
- 22 • A procedure pointer may not appear in a **reduction** clause.
- 23 • A pointer with the **INTENT (IN)** attribute may not appear in the **reduction** clause.
- 24 • A pointer must be associated upon entry and exit to the region.
- 25 • A pointer must not have its association status changed within the region.
- 26 • An original list item with the **POINTER** attribute must be associated at entry to the construct  
27 containing the **reduction** clause. Additionally, the list item must not be deallocated, allocated,  
28 or pointer assigned within the region.

- 1 • An original list item with the **ALLOCATABLE** attribute must be in the allocated state at entry to  
2 the construct containing the **reduction** clause. Additionally, the list item must not be  
3 deallocated and/or allocated within the region.
- 4 • If the *reduction-identifier* is defined in a **declare reduction** directive, the  
5 **declare reduction** directive must be in the same subprogram, or accessible by host or use  
6 association.
- 7 • If the *reduction-identifier* is a user-defined operator, the same explicit interface for that operator  
8 must be accessible as at the **declare reduction** directive.
- 9 • If the *reduction-identifier* is defined in a **declare reduction** directive, any subroutine or  
10 function referenced in the initializer clause or combiner expression must be an intrinsic function,  
11 or must have an explicit interface where the same explicit interface is accessible as at the  
12 **declare reduction** directive.

Fortran

### 13 2.14.3.7 linear clause

#### 14 Summary

15 The **linear** clause declares one or more list items to be private to a SIMD lane and to have a  
16 linear relationship with respect to the iteration space of a loop.

#### 17 Syntax

18 The syntax of the **linear** clause is as follows:

```
linear (list [ : linear-step ])
```

#### 19 Description

20 The **linear** clause provides a superset of the functionality provided by the **private** clause.

21 A list item that appears in a **linear** clause is subject to the **private** clause semantics described  
22 in Section 2.14.3.3 on page 176 except as noted. In addition, the value of the new list item on each  
23 iteration of the associated loop(s) corresponds to the value of the original list item before entering  
24 the construct plus the logical number of the iteration times *linear-step*. If *linear-step* is not  
25 specified it is assumed to be 1. The value corresponding to the sequentially last iteration of the  
26 associated loops is assigned to the original list item.



## Restrictions

- The *linear-step* expression must be invariant during the execution of the region associated with the construct. Otherwise, the execution results in unspecified behavior.
- A *list-item* cannot appear in more than one **linear** clause.
- A *list-item* that appears in a **linear** clause cannot appear in any other data-sharing attribute clause.

C / C++

- A *list-item* that appears in a **linear** clause must be of integral or pointer type, or must be a reference to an integral or pointer type.

C / C++

Fortran

- A *list-item* that appears in a **linear** clause must be of type **integer**.
- Variables that have the **POINTER** attribute and Cray pointers may not appear in a linear clause.
- The list item with the **ALLOCATABLE** attribute in the sequentially last iteration must have an allocation status of allocated upon exit from that iteration.

Fortran

## 2.14.4 Data Copying Clauses

This section describes the **copyin** clause (allowed on the **parallel** directive and combined parallel worksharing directives) and the **copyprivate** clause (allowed on the **single** directive).

These clauses support the copying of data values from private or threadprivate variables on one implicit task or thread to the corresponding variables on other implicit tasks or threads in the team.

The clauses accept a comma-separated list of list items (see Section 2.1 on page 26). All list items appearing in a clause must be visible, according to the scoping rules of the base language. Clauses may be repeated as needed, but a list item that specifies a given variable may not appear in more than one clause on the same directive

Fortran

An associate name preserves the association with the selector established at the **ASSOCIATE** statement. A list item that appears in a data copying clause may be a selector of an **ASSOCIATE** construct. If the construct association is established prior to a parallel region, the association between the associate name and the original list item will be retained in the region.

Fortran

## 1 2.14.4.1 `copyin` clause

### 2 Summary

3 The **copyin** clause provides a mechanism to copy the value of the master thread's threadprivate  
4 variable to the threadprivate variable of each other member of the team executing the **parallel**  
5 region.

### 6 Syntax

7 The syntax of the **copyin** clause is as follows:

```
copyin (list)
```

### 8 Description

▼ C / C++ ▼

9 The copy is done after the team is formed and prior to the start of execution of the associated  
10 structured block. For variables of non-array type, the copy occurs by copy assignment. For an array  
11 of elements of non-array type, each element is copied as if by assignment from an element of the  
12 master thread's array to the corresponding element of the other thread's array.

▲ C / C++ ▲

▼ C++ ▼

13 For class types, the copy assignment operator is invoked. The order in which copy assignment  
14 operators for different variables of class type are called is unspecified.

▲ C++ ▲

▼ Fortran ▼

15 The copy is done, as if by assignment, after the team is formed and prior to the start of execution of  
16 the associated structured block.

17 On entry to any **parallel** region, each thread's copy of a variable that is affected by a **copyin**  
18 clause for the **parallel** region will acquire the allocation, association, and definition status of the  
19 master thread's copy, according to the following rules:

- 20 • If the original list item has the **POINTER** attribute, each copy receives the same association  
21 status of the master thread's copy as if by pointer assignment.
- 22 • If the original list item does not have the **POINTER** attribute, each copy becomes defined with  
23 the value of the master thread's copy as if by intrinsic assignment, unless it has the allocation  
24 status of not currently allocated, in which case each copy will have the same status.

▲ Fortran ▲

## Restrictions

The restrictions to the **copyin** clause are as follows:

C / C++

- A list item that appears in a **copyin** clause must be **threadprivate**.
- A variable of class type (or array thereof) that appears in a **copyin** clause requires an accessible, unambiguous copy assignment operator for the class type.

C / C++

Fortran

- A list item that appears in a **copyin** clause must be **threadprivate**. Named variables appearing in a **threadprivate** common block may be specified: it is not necessary to specify the whole common block.
- A common block name that appears in a **copyin** clause must be declared to be a common block in the same scoping unit in which the **copyin** clause appears.

Fortran

### 2.14.4.2 **copyprivate** clause

#### Summary

The **copyprivate** clause provides a mechanism to use a private variable to broadcast a value from the data environment of one implicit task to the data environments of the other implicit tasks belonging to the **parallel** region.

To avoid race conditions, concurrent reads or updates of the list item must be synchronized with the update of the list item that occurs as a result of the **copyprivate** clause.

#### Syntax

The syntax of the **copyprivate** clause is as follows:

```
copyprivate (list)
```


1  
2  
3  
4  
  
5  
6  
7  
8  
9  
10  
11  
  
12  
13  
  
14  
15  
16  
17  
  
18  
19  
20  
21  
  
22  
23  
  
24  
25  
26

## Description


The effect of the **copyprivate** clause on the specified list items occurs after the execution of the structured block associated with the **single** construct (see Section 2.7.3 on page 65), and before any of the threads in the team have left the barrier at the end of the construct.



In all other implicit tasks belonging to the **parallel** region, each specified list item becomes defined with the value of the corresponding list item in the implicit task whose thread executed the structured block. For variables of non-array type, the definition occurs by copy assignment. For an array of elements of non-array type, each element is copied by copy assignment from an element of the array in the data environment of the implicit task associated with the thread that executed the structured block to the corresponding element of the array in the data environment of the other implicit tasks




For class types, a copy assignment operator is invoked. The order in which copy assignment operators for different variables of class type are called is unspecified.



If a list item does not have the **POINTER** attribute, then in all other implicit tasks belonging to the **parallel** region, the list item becomes defined as if by intrinsic assignment with the value of the corresponding list item in the implicit task associated with the thread that executed the structured block.

If the list item has the **POINTER** attribute, then, in all other implicit tasks belonging to the **parallel** region, the list item receives, as if by pointer assignment, the same association status of the corresponding list item in the implicit task associated with the thread that executed the structured block.

The order in which any final subroutines for different variables of a finalizable type are called is unspecified.



**Note** – The **copyprivate** clause is an alternative to using a shared variable for the value when providing such a shared variable would be difficult (for example, in a recursion requiring a different variable at each level).



## Restrictions

The restrictions to the **copyprivate** clause are as follows:

- All list items that appear in the **copyprivate** clause must be either **threadprivate** or **private** in the enclosing context.
- A list item that appears in a **copyprivate** clause may not appear in a **private** or **firstprivate** clause on the **single** construct.

C++

- A variable of class type (or array thereof) that appears in a **copyprivate** clause requires an accessible unambiguous copy assignment operator for the class type.

C++

Fortran

- A common block that appears in a **copyprivate** clause must be **threadprivate**.
- Pointers with the **INTENT (IN)** attribute may not appear in the **copyprivate** clause.
- The list item with the **ALLOCATABLE** attribute must have the allocation status of **allocated** when the intrinsic assignment is performed.

Fortran

### 2.14.5 map Clause

#### Summary

The **map** clause specifies how an original list item is mapped from the current task's data environment to a corresponding list item in the device data environment of the device identified by the construct.

#### Syntax

The syntax of the map clause is as follows:

```
map ([map-type-modifier[,]] map-type : ] list)
```

## Description

The list items that appear in a **map** clause may include array sections.

The *map-type* and *map-type-modifier* specify the effect of the **map** clause, as described below.

The original and corresponding list items may share storage such that writes to either item by one task followed by a read or write of the other item by another task without intervening synchronization can result in data races.

If the **map** clause appears on a **target**, **target data**, or **target enter data** construct then on entry to the region the following sequence of steps occurs:

1. If a corresponding list item of the original list item is not present in the device data environment, then:
  - a) A new list item with language-specific attributes is derived from the original list item and created in the device data environment.
  - b) The new list item becomes the corresponding list item to the original list item in the device data environment.
  - c) The corresponding list item has a reference count that is initialized to zero.
2. The corresponding list item's reference count is incremented by one.
3. If the corresponding list item's reference count is one or the **always map-type-modifier** is present, then:
  - a) If the *map-type* is **to** or **tofrom**, then the corresponding list item is assigned the value of the original list item.
4. If the corresponding list item's reference count is one, then:
  - a) If the *map-type* is **from** or **alloc** the value of the corresponding list item is undefined.

If the **map** clause appears on a **target**, **target data**, or **target exit data** construct then on exit from the region the following sequence of steps occurs:

1. If a corresponding list item of the original list item is not present in the device data environment, then the list item is ignored.
2. If a corresponding list item of the original list item is present in the device data environment, then:
  - a) If the corresponding list item's reference count is greater than zero, then:
    - i. the corresponding list item's reference count is decremented by one.
    - ii. If the *map-type* is **delete**, then the corresponding list item's reference count is set to zero.

- 1           b) If the corresponding list item's reference count is zero or the **always** *map-type-modifier* is  
 2           present, then:
- 3                 i. If the *map-type* is **from** or **tofrom**, then the original list item is assigned the value of  
 4                 the corresponding list item.
- 5           c) If the corresponding list item's reference count is zero, then the corresponding list item is  
 6           removed from the device data environment

▼ C / C++ ▲

7           If a new list item is created then a new list item of the same type, with automatic storage duration, is  
 8           allocated for the construct. The size and alignment of the new list item are determined by the type  
 9           of the variable. This allocation occurs if the region references the list item in any statement.

10          If a new list item is created for an array section and the type of the variable appearing in that array  
 11          section is pointer, reference to array, or reference to pointer then the variable is implicitly treated as  
 12          if it had appeared in a map clause with a *map-type* of **alloc**. The corresponding variable is  
 13          assigned the address of the storage location of the corresponding array section in the device data  
 14          environment. If the variable appears in a **to** or **from** clause in a **target update** region during  
 15          the lifetime of the new list item but not as part of the specification of an array section, the behavior  
 16          is unspecified.

▲ C / C++ ▲

▼ Fortran ▼

17          If a new list item is created then a new list item of the same type, type parameter, and rank is  
 18          allocated.

▲ Fortran ▲

19          The *map-type* determines how the new list item is initialized.

20          If a *map-type* is not specified, the *map-type* defaults to **tofrom**.

## Restrictions

- If a list item is an array section, it must specify contiguous storage.
- At most one list item can be an array item derived from a given variable in **map** clauses of the same construct.
- List items of **map** clauses in the same construct must not share original storage.
- If any part of the original storage of a list item has corresponding storage in the device data environment, all of the original storage must have corresponding storage in the device data environment.
- A variable that is part of another variable (such as a field of a structure) but is not an array element or an array section cannot appear in a **map** clause.
- If variables that share storage are mapped, the behavior is unspecified.
- A list item must have a mappable type.
- **threadprivate** variables cannot appear in a **map** clause.

C / C++

- Initialization and assignment are through bitwise copy.
- A variable for which the type is pointer, reference to array, or reference to pointer and an array section derived from that variable must not appear as list items of **map** clauses of the same construct.
- A variable for which the type is pointer, reference to array, or reference to pointer must not appear as a list item if the device data environment already contains an array section derived from that variable.
- An array section derived from a variable for which the type is pointer, reference to array, or reference to pointer must not appear as a list item if the device data environment already contains that variable.

C / C++

Fortran

- The value of the new list item becomes that of the original list item in the map initialization and assignment.

Fortran



## 1 2.15 declare reduction Directive

### 2 Summary

3 The following section describes the directive for declaring user-defined reductions. The  
4 **declare reduction** directive declares a *reduction-identifier* that can be used in a  
5 **reduction** clause. The **declare reduction** directive is a declarative directive.

### 6 Syntax

C

```
#pragma omp declare reduction(reduction-identifier : typename-list :  
combiner ) [initializer-clause] new-line
```

7 where:

- 8 • *reduction-identifier* is either a base language identifier or one of the following operators: +, -, \*,  
9 &, |, ^, && and ||
- 10 • *typename-list* is list of type names
- 11 • *combiner* is an expression
- 12 • *initializer-clause* is **initializer**(*initializer-expr*) where *initializer-expr* is  
13 **omp\_priv** = *initializer* or *function-name* (*argument-list*)

C

C++

```
#pragma omp declare reduction(reduction-identifier : typename-list :  
combiner) [initializer-clause] new-line
```

14 where:

- 15 • *reduction-identifier* is either a base language identifier or one of the following operators: +, -, \*,  
16 &, |, ^, && and ||
- 17 • *typename-list* is list of type names
- 18 • *combiner* is an expression
- 19 • *initializer-clause* is **initializer**(*initializer-expr*) where *initializer-expr* is  
20 **omp\_priv** *initializer* or *function-name* (*argument-list*)

C++

```
!$omp declare reduction(reduction-identifier : type-list : combiner)
[initializer-clause]
```

1 where:

- 2 • *reduction-identifier* is either a base language identifier, or a user-defined operator, or one of the
- 3 following operators: +, -, \*, .and., .or., .eqv., .neqv., or one of the following intrinsic
- 4 procedure names: max, min, iand, ior, ieor.
- 5 • *type-list* is a list of type specifiers
- 6 • *combiner* is either an assignment statement or a subroutine name followed by an argument list
- 7 • *initializer-clause* is **initializer** (*initializer-expr*), where *initializer-expr* is
- 8 **omp\_priv** = *expression* or *subroutine-name* (*argument-list*)

## 9 Description

10 Custom reductions can be defined using the **declare reduction** directive; the

11 *reduction-identifier* and the type identify the **declare reduction** directive. The

12 *reduction-identifier* can later be used in a **reduction** clause using variables of the type or types

13 specified in the **declare reduction** directive. If the directive applies to several types then it is

14 considered as if there were multiple **declare reduction** directives, one for each type.

15 If a type with deferred or assumed length type parameter is specified in a **declare reduction**

16 directive, the *reduction-identifier* of that directive can be used in a **reduction** clause with any

17 variable of the same type and the same kind parameter, regardless of the length type Fortran

18 parameters with which the variable is declared.

19 The visibility and accessibility of this declaration are the same as those of a variable declared at the

20 same point in the program. The enclosing context of the *combiner* and of the *initializer-expr* will be

21 that of the **declare reduction** directive. The *combiner* and the *initializer-expr* must be correct

22 in the base language as if they were the body of a function defined at the same point in the program.

## Fortran

1 If the *reduction-identifier* is the same as the name of a user-defined operator or an extended  
2 operator, or the same as a generic name that is one of the allowed intrinsic procedures, and if the  
3 operator or procedure name appears in an accessibility statement in the same module, the  
4 accessibility of the corresponding **declare reduction** directive is determined by the  
5 accessibility attribute of the statement.

6 If the *reduction-identifier* is the same as a generic name that is one of the allowed intrinsic  
7 procedures and is accessible, and if it has the same name as a derived type in the same module, the  
8 accessibility of the corresponding **declare reduction** directive is determined by the  
9 accessibility of the generic name according to the base language.

## Fortran

### C++

10 The **declare reduction** directive can also appear at points in the program at which a static  
11 data member could be declared. In this case, the visibility and accessibility of the declaration are  
12 the same as those of a static data member declared at the same point in the program.

### C++

13 The *combiner* specifies how partial results can be combined into a single value. The *combiner* can  
14 use the special variable identifiers **omp\_in** and **omp\_out** that are of the type of the variables  
15 being reduced with this *reduction-identifier*. Each of them will denote one of the values to be  
16 combined before executing the *combiner*. It is assumed that the special **omp\_out** identifier will  
17 refer to the storage that holds the resulting combined value after executing the *combiner*.

18 The number of times the *combiner* is executed, and the order of these executions, for any  
19 **reduction** clause is unspecified.

## Fortran

20 If the *combiner* is a subroutine name with an argument list, the *combiner* is evaluated by calling the  
21 subroutine with the specified argument list.

22 If the *combiner* is an assignment statement, the *combiner* is evaluated by executing the assignment  
23 statement.

## Fortran

24 As the *initializer-expr* value of a user-defined reduction is not known *a priori* the *initializer-clause*  
25 can be used to specify one. Then the contents of the *initializer-clause* will be used as the initializer  
26 for private copies of reduction list items where the **omp\_priv** identifier will refer to the storage to  
27 be initialized. The special identifier **omp\_orig** can also appear in the *initializer-clause* and it will  
28 refer to the storage of the original variable to be reduced.

29 The number of times that the *initializer-expr* is evaluated, and the order of these evaluations, is  
30 unspecified.

C / C++

1 If the *initializer-expr* is a function name with an argument list, the *initializer-expr* is evaluated by  
2 calling the function with the specified argument list. Otherwise, the *initializer-expr* specifies how  
3 **omp\_priv** is declared and initialized.

C / C++

4 If no *initializer-clause* is specified, the private variables will be initialized following the rules for  
5 initialization of objects with static storage duration.

C

6 If no *initializer-expr* is specified, the private variables will be initialized following the rules for  
7 *default-initialization*.

C++

C++

Fortran

8 If the *initializer-expr* is a subroutine name with an argument list, the *initializer-expr* is evaluated by  
9 calling the subroutine with the specified argument list.

10 If the *initializer-expr* is an assignment statement, the *initializer-expr* is evaluated by executing the  
11 assignment statement.

12 If no *initializer-clause* is specified, the private variables will be initialized as follows:

- 13 • For **complex**, **real**, or **integer** types, the value 0 will be used.
- 14 • For **logical** types, the value **.false.** will be used.
- 15 • For derived types for which default initialization is specified, default initialization will be used.
- 16 • Otherwise, not specifying an *initializer-clause* results in unspecified behavior.

Fortran

C / C++

17 If *reduction-identifier* is used in a **target** region then a **declare target** construct must be  
18 specified for any function that can be accessed through *combiner* and *initializer-expr*.

C / C++

Fortran

1 If *reduction-identifier* is used in a **target** region then a **declare target** construct must be  
2 specified for any function or subroutine that can be accessed through *combiner* and *initializer-expr*.

Fortran

3 **Restrictions**

- 4 • Only the variables **omp\_in** and **omp\_out** are allowed in the *combiner*.
- 5 • Only the variables **omp\_priv** and **omp\_orig** are allowed in the *initializer-clause*.
- 6 • If the variable **omp\_orig** is modified in the *initializer-clause*, the behavior is unspecified.
- 7 • If execution of the *combiner* or the *initializer-expr* results in the execution of an OpenMP  
8 construct or an OpenMP API call, then the behavior is unspecified.
- 9 • A *reduction-identifier* may not be re-declared in the current scope for the same type or for a type  
10 that is compatible according to the base language rules.
- 11 • At most one *initializer-clause* can be specified.

C / C++

- 12 • A type name in a **declare reduction** directive cannot be a function type, an array type, a  
13 reference type, or a type qualified with **const**, **volatile** or **restrict**.

C / C++

C

- 14 • If the *initializer-expr* is a function name with an argument list, then one of the arguments must be  
15 the address of **omp\_priv**.

C

C++

- 16 • If the *initializer-expr* is a function name with an argument list, then one of the arguments must be  
17 **omp\_priv** or the address of **omp\_priv**.

C++

- 1       • If the *initializer-expr* is a subroutine name with an argument list, then one of the arguments must  
2       be `omp_priv`.
- 3       • If the `declare reduction` directive appears in the specification part of a module and the  
4       corresponding reduction clause does not appear in the same module, the *reduction-identifier* must  
5       be the same as the name of a user-defined operator, one of the allowed operators that is extended  
6       or a generic name that is the same as the name of one of the allowed intrinsic procedures.
- 7       • If the `declare reduction` directive appears in the specification of a module, if the  
8       corresponding `reduction` clause does not appear in the same module, and if the  
9       *reduction-identifier* is the same as the name of a user-defined operator or an extended operator, or  
10      the same as a generic name that is the same as one of the allowed intrinsic procedures, the  
11      interface for that operator or the generic name must be defined in the specification of the same  
12      module, or must be accessible by use association.
- 13      • Any subroutine, or function used in the `initializer` clause or *combiner* expression must be  
14      an intrinsic function, or must have an accessible interface.
- 15      • Any user-defined operator, or extended operator used in the `initializer` clause or *combiner*  
16      expression must have an accessible interface.
- 17      • If any subroutine, function, user-defined operator or extended operator is used in the  
18      `initializer` clause or *combiner* expression, it must be accessible to the subprogram in  
19      which the corresponding `reduction` clause is specified.
- 20      • If the length type parameter is specified for a character type, it must be a constant, a colon or an `*`.
- 21      • If a character type with deferred or assumed length parameter is specified in a  
22      `declare reduction` directive, no other `declare reduction` directives with Fortran  
23      character type of the same kind parameter and the same *reduction-identifier* are allowed in the  
24      same scope.
- 25      • Any subroutine used in the `initializer` clause or *combiner* expression must not have any  
26      alternate returns appear in the argument list.

## Cross References

- `reduction` clause, Section [2.14.3.6](#) on page 184.

## 1 2.16 Nesting of Regions

2 This section describes a set of restrictions on the nesting of regions. The restrictions on nesting are  
3 as follows:

- 4 • A worksharing region may not be closely nested inside a worksharing, explicit **task**,  
5 **critical**, **ordered**, **atomic**, or **master** region.
- 6 • A **barrier** region may not be closely nested inside a worksharing, explicit **task**, **critical**,  
7 **ordered**, **atomic**, or **master** region.
- 8 • A **master** region may not be closely nested inside a worksharing, **atomic**, or explicit **task**  
9 region.
- 10 • An **ordered** region may not be closely nested inside a **critical**, **atomic**, or explicit **task**  
11 region.
- 12 • An **ordered** region must be closely nested inside a loop region (or parallel loop region) with an  
13 **ordered** clause.
- 14 • A **critical** region may not be nested (closely or otherwise) inside a **critical** region with  
15 the same name. Note that this restriction is not sufficient to prevent deadlock.
- 16 • OpenMP constructs may not be nested inside an **atomic** region.
- 17 • OpenMP constructs may not be nested inside a **simd** region.
- 18 • If a **target**, **target update**, or **target data** construct appears within a **target** region  
19 then the behavior is unspecified.
- 20 • If specified, a **teams** construct must be contained within a **target** construct. That **target**  
21 construct must contain no statements or directives outside of the **teams** construct.
- 22 • **distribute**, **parallel**, **parallel sections**, **parallel workshare**, and the  
23 parallel loop and parallel loop SIMD constructs are the only OpenMP constructs that can be  
24 closely nested in the **teams** region.
- 25 • A **distribute** construct must be closely nested in a **teams** region.
- 26 • If *construct-type-clause* is **taskgroup**, the **cancel** construct must be closely nested inside a  
27 **task** construct and the **cancel** construct must be nested inside a **taskgroup** region.  
28 Otherwise, the **cancel** construct must be closely nested inside an OpenMP construct that  
29 matches the type specified in *construct-type-clause* of the **cancel** construct.
- 30 • A **cancellation point** construct for which *construct-type-clause* is **taskgroup** must be  
31 nested inside a **task** construct. A **cancellation point** construct for which  
32 *construct-type-clause* is not **taskgroup** must be closely nested inside an OpenMP construct  
33 that matches the type specified in *construct-type-clause*.

# Runtime Library Routines

---

This chapter describes the OpenMP API runtime library routines and is divided into the following sections:

- Runtime library definitions (Section 3.1 on page 207).

- Execution environment routines that can be used to control and to query the parallel execution environment (Section 3.2 on page 208).

- Lock routines that can be used to synchronize access to data (Section 3.3 on page 239).

- Portable timer routines (Section 3.4 on page 245).

Throughout this chapter, *true* and *false* are used as generic terms to simplify the description of the routines.

▼ C / C++ ▼

*true* means a nonzero integer value and *false* means an integer value of zero.

▲ C / C++ ▲

▼ Fortran ▼

*true* means a logical value of `.TRUE.` and *false* means a logical value of `.FALSE.`

▲ Fortran ▲

▼ Fortran ▼

## Restrictions

The following restriction applies to all OpenMP runtime library routines:

- OpenMP runtime library routines may not be called from **PURE** or **ELEMENTAL** procedures.

▲ Fortran ▲



# 1 3.1 Runtime Library Definitions

2 For each base language, a compliant implementation must supply a set of definitions for the  
3 OpenMP API runtime library routines and the special data types of their parameters. The set of  
4 definitions must contain a declaration for each OpenMP API runtime library routine and a  
5 declaration for the *simple lock*, *nestable lock*, *schedule*, and *thread affinity policy* data types. In  
6 addition, each set of definitions may specify other implementation specific values.



## C / C++

7 The library routines are external functions with “C” linkage.

8 Prototypes for the C/C++ runtime library routines described in this chapter shall be provided in a  
9 header file named **omp.h**. This file defines the following:

- 10 • The prototypes of all the routines in the chapter.
- 11 • The type **omp\_lock\_t**.
- 12 • The type **omp\_nest\_lock\_t**.
- 13 • The type **omp\_sched\_t**.
- 14 • The type **omp\_proc\_bind\_t**.

15 See Section Section C.1 on page 297 for an example of this file.



## C / C++



## Fortran

16 The OpenMP Fortran API runtime library routines are external procedures. The return values of  
17 these routines are of default kind, unless otherwise specified.

18 Interface declarations for the OpenMP Fortran runtime library routines described in this chapter  
19 shall be provided in the form of a Fortran **include** file named **omp\_lib.h** or a Fortran 90  
20 **module** named **omp\_lib**. It is implementation defined whether the **include** file or the  
21 **module** file (or both) is provided.

22 These files define the following:

- 23 • The interfaces of all of the routines in this chapter.
- 24 • The **integer parameter** **omp\_lock\_kind**.
- 25 • The **integer parameter** **omp\_nest\_lock\_kind**.
- 26 • The **integer parameter** **omp\_sched\_kind**.
- 27 • The **integer parameter** **omp\_proc\_bind\_kind**.

- The **integer parameter `openmp_version`** with a value `yyyymm` where `yyyy` and `mm` are the year and month designations of the version of the OpenMP Fortran API that the implementation supports. This value matches that of the C preprocessor macro `_OPENMP`, when a macro preprocessor is supported (see Section 2.2 on page 32).

See Section C.1 on page 299 and Section C.3 on page 302 for examples of these files.

It is implementation defined whether any of the OpenMP runtime library routines that take an argument are extended with a generic interface so arguments of different **KIND** type can be accommodated. See Appendix C.4 for an example of such an extension.

Fortran

## 3.2 Execution Environment Routines

This section describes routines that affect and monitor threads, processors, and the parallel environment.

### 3.2.1 `omp_set_num_threads`

#### Summary

The `omp_set_num_threads` routine affects the number of threads to be used for subsequent parallel regions that do not specify a `num_threads` clause, by setting the value of the first element of the `nthreads-var` ICV of the current task.

#### Format

C / C++

```
void omp_set_num_threads(int num_threads);
```

C / C++

Fortran

```
subroutine omp_set_num_threads(num_threads)  
integer num_threads
```

Fortran

## Constraints on Arguments

The value of the argument passed to this routine must evaluate to a positive integer, or else the behavior of this routine is implementation defined.

## Binding

The binding task set for an `omp_set_num_threads` region is the generating task.

## Effect

The effect of this routine is to set the value of the first element of the *nthreads-var* ICV of the current task to the value specified in the argument.

See Section 2.5.1 on page 48 for the rules governing the number of threads used to execute a `parallel` region.

## Cross References

- *nthreads-var* ICV, see Section 2.3 on page 35.
- `OMP_NUM_THREADS` environment variable, see Section 4.2 on page 252.
- `omp_get_max_threads` routine, see Section 3.2.3 on page 210.
- `parallel` construct, see Section 2.5 on page 44.
- `num_threads` clause, see Section 2.5 on page 44.

## 3.2.2 `omp_get_num_threads`

### Summary

The `omp_get_num_threads` routine returns the number of threads in the current team.

### Format

C / C++

```
int omp_get_num_threads(void);
```

C / C++

```
integer function omp_get_num_threads ()
```

## 1 Binding

2 The binding region for an `omp_get_num_threads` region is the innermost enclosing  
3 `parallel` region.

## 4 Effect

5 The `omp_get_num_threads` routine returns the number of threads in the team executing the  
6 `parallel` region to which the routine region binds. If called from the sequential part of a  
7 program, this routine returns 1.

8 See Section 2.5.1 on page 48 for the rules governing the number of threads used to execute a  
9 `parallel` region.

## 10 Cross References

- 11 • `parallel` construct, see Section 2.5 on page 44.
- 12 • `omp_set_num_threads` routine, see Section 3.2.1 on page 208.
- 13 • `OMP_NUM_THREADS` environment variable, see Section 4.2 on page 252.

## 14 3.2.3 omp\_get\_max\_threads

### 15 Summary

16 The `omp_get_max_threads` routine returns an upper bound on the number of threads that  
17 could be used to form a new team if a `parallel` construct without a `num_threads` clause were  
18 encountered after execution returns from this routine.

1

## Format

C / C++

```
int omp_get_max_threads(void);
```

C / C++

Fortran

```
integer function omp_get_max_threads()
```

Fortran

2

## Binding

3

The binding task set for an `omp_get_max_threads` region is the generating task.

4

## Effect

5

The value returned by `omp_get_max_threads` is the value of the first element of the *nthreads-var* ICV of the current task. This value is also an upper bound on the number of threads that could be used to form a new team if a parallel region without a `num_threads` clause were encountered after execution returns from this routine.

6

7

8

9

10

See Section 2.5.1 on page 48 for the rules governing the number of threads used to execute a `parallel` region.

11

12

13

Note – The return value of the `omp_get_max_threads` routine can be used to dynamically allocate sufficient storage for all threads in the team formed at the subsequent active `parallel` region.

14

## Cross References

15

- *nthreads-var* ICV, see Section 2.3 on page 35.

16

- `parallel` construct, see Section 2.5 on page 44.

17

- `num_threads` clause, see Section 2.5 on page 44.

18

- `omp_set_num_threads` routine, see Section 3.2.1 on page 208.

19

- `OMP_NUM_THREADS` environment variable, see Section 4.2 on page 252.

## 1 3.2.4 `omp_get_thread_num`

### 2 Summary

3 The `omp_get_thread_num` routine returns the thread number, within the current team, of the  
4 calling thread.

### 5 Format

▼	C / C++	▼
<pre>int omp_get_thread_num(void);</pre>		
▲	C / C++	▲
▼	Fortran	▼
<pre>integer function omp_get_thread_num()</pre>		
▲	Fortran	▲

### 6 Binding

7 The binding thread set for an `omp_get_thread_num` region is the current team. The binding  
8 region for an `omp_get_thread_num` region is the innermost enclosing `parallel` region.

### 9 Effect

10 The `omp_get_thread_num` routine returns the thread number of the calling thread, within the  
11 team executing the `parallel` region to which the routine region binds. The thread number is an  
12 integer between 0 and one less than the value returned by `omp_get_num_threads`, inclusive.  
13 The thread number of the master thread of the team is 0. The routine returns 0 if it is called from  
14 the sequential part of a program.

15 Note – The thread number may change during the execution of an untied task. The value returned  
16 by `omp_get_thread_num` is not generally useful during the execution of such a task region.

### 17 Cross References

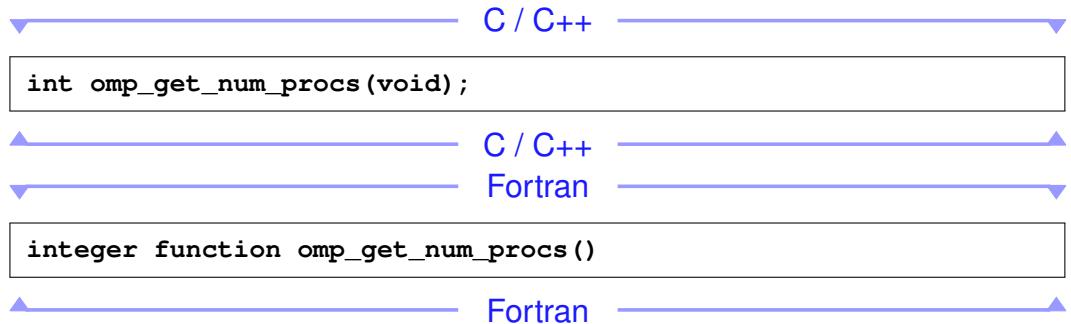
- 18 • `omp_get_num_threads` routine, see Section [3.2.2](#) on page [209](#).

## 1 3.2.5 `omp_get_num_procs`

### 2 Summary

3 The `omp_get_num_procs` routine returns the number of processors available to the device.

### 4 Format



### 5 Binding

6 The binding thread set for an `omp_get_num_procs` region is all threads on a device. The effect  
7 of executing this routine is not related to any specific region corresponding to any construct or API  
8 routine.

### 9 Effect

10 The `omp_get_num_procs` routine returns the number of processors that are available to the  
11 device at the time the routine is called. Note that this value may change between the time that it is  
12 determined by the `omp_get_num_procs` routine and the time that it is read in the calling  
13 context due to system actions outside the control of the OpenMP implementation.


## 14 3.2.6 `omp_in_parallel`

### 15 Summary

16 The `omp_in_parallel` routine returns *true* if the *active-levels-var* ICV is greater than zero;  
17 otherwise, it returns *false*.

1

## Format



2

## Binding

3

The binding task set for an `omp_in_parallel` region is the generating task.

4

## Effect

5

The effect of the `omp_in_parallel` routine is to return *true* if the current task is enclosed by an active `parallel` region, and the `parallel` region is enclosed by the outermost initial task region on the device; otherwise it returns *false*.

6

7

8

## Cross References

9

- *active-levels-var*, see Section [2.3](#) on page [35](#).

10

- `omp_get_active_level` routine, see Section [3.2.20](#) on page [229](#).

## 11 3.2.7 `omp_set_dynamic`

12

### Summary

13

The `omp_set_dynamic` routine enables or disables dynamic adjustment of the number of threads available for the execution of subsequent `parallel` regions by setting the value of the *dyn-var* ICV

14

15



1

## Format

C / C++

```
void omp_set_dynamic(int dynamic_threads);
```

C / C++

Fortran

```
subroutine omp_set_dynamic(dynamic_threads)
  logical dynamic_threads
```

Fortran

2

## Binding

3

The binding task set for an `omp_set_dynamic` region is the generating task.

4

## Effect

5

For implementations that support dynamic adjustment of the number of threads, if the argument to `omp_set_dynamic` evaluates to *true*, dynamic adjustment is enabled for the current task; otherwise, dynamic adjustment is disabled for the current task. For implementations that do not support dynamic adjustment of the number of threads this routine has no effect: the value of *dyn-var* remains *false*.

6

7

8

9

10

See Section 2.5.1 on page 48 for the rules governing the number of threads used to execute a `parallel` region.

11

12

## Cross References

13

- *dyn-var* ICV, see Section 2.3 on page 35.

14

- `omp_get_num_threads` routine, see Section 3.2.2 on page 209.

15

- `omp_get_dynamic` routine, see Section 3.2.8 on page 216.

16

- `OMP_DYNAMIC` environment variable, see Section 4.3 on page 253.

## 1 3.2.8 `omp_get_dynamic`

### 2 Summary

3 The `omp_get_dynamic` routine returns the value of the *dyn-var* ICV, which determines whether  
4 dynamic adjustment of the number of threads is enabled or disabled.

### 5 Format

▼	C / C++	▼
<pre>int omp_get_dynamic(void);</pre>		
▲	C / C++	▲
▼	Fortran	▼
<pre>logical function omp_get_dynamic()</pre>		
▲	Fortran	▲

### 6 Binding

7 The binding task set for an `omp_get_dynamic` region is the generating task.

### 8 Effect

9 This routine returns *true* if dynamic adjustment of the number of threads is enabled for the current  
10 task; it returns *false*, otherwise. If an implementation does not support dynamic adjustment of the  
11 number of threads, then this routine always returns *false*.

12 See Section 2.5.1 on page 48 for the rules governing the number of threads used to execute a  
13 `parallel` region.

### 14 Cross References

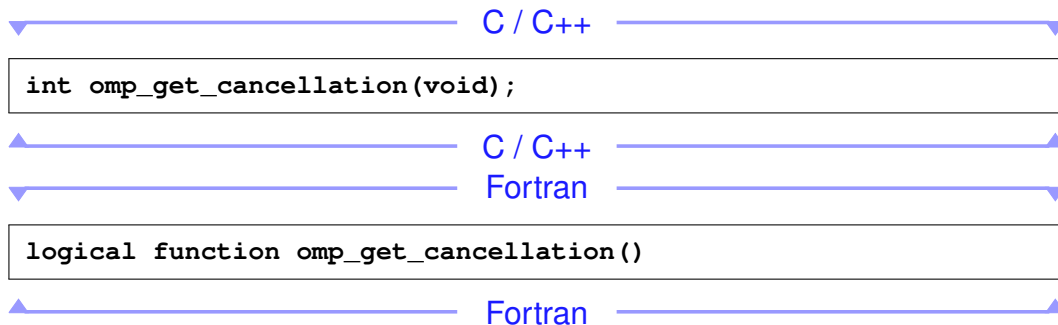
- 15 • *dyn-var* ICV, see Section 2.3 on page 35.
- 16 • `omp_set_dynamic` routine, see Section 3.2.7 on page 214.
- 17 • `OMP_DYNAMIC` environment variable, see Section 4.3 on page 253.

## 1 3.2.9 `omp_get_cancellation`

### 2 Summary

3 The `omp_get_cancellation` routine returns the value of the *cancel-var* ICV, which controls  
4 the behavior of the `cancel` construct and cancellation points.

### 5 Format



### 6 Binding

7 The binding task set for an `omp_get_cancellation` region is the whole program.

### 8 Effect

9 This routine returns *true* if cancellation is activated. It returns *false* otherwise.

### 10 Cross References

- 11 • *cancel-var* ICV, see Section 2.3.1 on page 35.
- 12 • `OMP_CANCELLATION` environment variable, see Section 4.11 on page 259


## 13 3.2.10 `omp_set_nested`

### 14 Summary

15 The `omp_set_nested` routine enables or disables nested parallelism, by setting the *nest-var*  
16 ICV.

1

## Format



2

## Binding

3

The binding task set for an **omp\_set\_nested** region is the generating task.

4

## Effect

5

For implementations that support nested parallelism, if the argument to **omp\_set\_nested** evaluates to *true*, nested parallelism is enabled for the current task; otherwise, nested parallelism is disabled for the current task. For implementations that do not support nested parallelism, this routine has no effect: the value of *nest-var* remains *false*.

6

7

8

9

See Section 2.5.1 on page 48 for the rules governing the number of threads used to execute a **parallel** region.

10

11

## Cross References

12

- *nest-var* ICV, see Section 2.3 on page 35.

13

- **omp\_set\_max\_active\_levels** routine, see Section 3.2.15 on page 223.

14

- **omp\_get\_max\_active\_levels** routine, see Section 3.2.16 on page 225.

15

- **omp\_get\_nested** routine, see Section 3.2.11 on page 219.

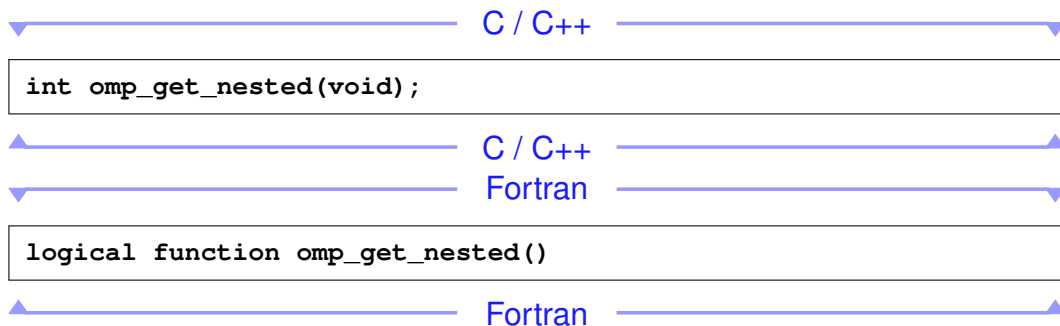
16

- **OMP\_NESTED** environment variable, see Section 4.6 on page 256.

## 1 3.2.11 `omp_get_nested`

### 2 Summary

3 The `omp_get_nested` routine returns the value of the *nest-var* ICV, which determines if nested  
4 parallelism is enabled or disabled.



### 5 Binding

6 The binding task set for an `omp_get_nested` region is the generating task.

### 7 Effect

8 This routine returns *true* if nested parallelism is enabled for the current task; it returns *false*,  
9 otherwise. If an implementation does not support nested parallelism, this routine always returns  
10 *false*.

11 See Section 2.5.1 on page 48 for the rules governing the number of threads used to execute a  
12 `parallel` region.

### 13 Cross References

- 14 • *nest-var* ICV, see Section 2.3 on page 35.
- 15 • `omp_set_nested` routine, see Section 3.2.10 on page 217.
- 16 • `OMP_NESTED` environment variable, see Section 4.6 on page 256.

## 1 3.2.12 `omp_set_schedule`

### 2 Summary

3 The `omp_set_schedule` routine affects the schedule that is applied when `runtime` is used as  
4 schedule kind, by setting the value of the *run-sched-var* ICV.

### 5 Format

▼ C / C++ ▼

```
void omp_set_schedule(omp_sched_t kind, int modifier);
```

▲ C / C++ ▲

▼ Fortran ▼

```
subroutine omp_set_schedule(kind, modifier)
integer (kind=omp_sched_kind) kind
integer modifier
```

▲ Fortran ▲

### 6 Constraints on Arguments

7 The first argument passed to this routine can be one of the valid OpenMP schedule kinds (except for  
8 `runtime`) or any implementation specific schedule. The C/C++ header file (`omp.h`) and the  
9 Fortran include file (`omp_lib.h`) and/or Fortran 90 module file (`omp_lib`) define the valid  
10 constants. The valid constants must include the following, which can be extended with  
11 implementation specific values:

```
typedef enum omp_sched_t
    omp_sched_static = 1,
    omp_sched_dynamic = 2,
    omp_sched_guided = 3,
    omp_sched_auto = 4
    omp_sched_t;
```

```
integer(kind=omp_sched_kind), parameter :: omp_sched_static = 1
integer(kind=omp_sched_kind), parameter :: omp_sched_dynamic = 2
integer(kind=omp_sched_kind), parameter :: omp_sched_guided = 3
integer(kind=omp_sched_kind), parameter :: omp_sched_auto = 4
```

## 1 Binding

2 The binding task set for an **omp\_set\_schedule** region is the generating task.

## 3 Effect

4 The effect of this routine is to set the value of the *run-sched-var* ICV of the current task to the  
 5 values specified in the two arguments. The schedule is set to the schedule type specified by the first  
 6 argument *kind*. It can be any of the standard schedule types or any other implementation specific  
 7 one. For the schedule types **static**, **dynamic**, and **guided** the *chunk\_size* is set to the value of  
 8 the second argument, or to the default *chunk\_size* if the value of the second argument is less than 1;  
 9 for the schedule type **auto** the second argument has no meaning; for implementation specific  
 10 schedule types, the values and associated meanings of the second argument are implementation  
 11 defined.

## 12 Cross References

- 13 • *run-sched-var* ICV, see Section 2.3 on page 35.
- 14 • **omp\_get\_schedule** routine, see Section 3.2.13 on page 222.
- 15 • **OMP\_SCHEDULE** environment variable, see Section 4.1 on page 251.
- 16 • Determining the schedule of a worksharing loop, see Section 2.7.1.1 on page 62.

## 1 3.2.13 `omp_get_schedule`

### 2 Summary

3 The `omp_get_schedule` routine returns the schedule that is applied when the runtime schedule  
4 is used.

### 5 Format

C / C++

```
void omp_get_schedule(omp_sched_t * kind, int * modifier);
```

C / C++

Fortran

```
subroutine omp_get_schedule(kind, modifier)
integer (kind=omp_sched_kind) kind
integer modifier
```

Fortran

### 6 Binding

7 The binding task set for an `omp_get_schedule` region is the generating task.

### 8 Effect

9 This routine returns the *run-sched-var* ICV in the task to which the routine binds. The first  
10 argument *kind* returns the schedule to be used. It can be any of the standard schedule types as  
11 defined in Section 3.2.12 on page 220, or any implementation specific schedule type. The second  
12 argument is interpreted as in the `omp_set_schedule` call, defined in Section 3.2.12 on  
13 page 220.

### 14 Cross References

- 15 • *run-sched-var* ICV, see Section 2.3 on page 35.
- 16 • `omp_set_schedule` routine, see Section 3.2.12 on page 220.
- 17 • `OMP_SCHEDULE` environment variable, see Section 4.1 on page 251.
- 18 • Determining the schedule of a worksharing loop, see Section 2.7.1.1 on page 62.



## 1 3.2.14 `omp_get_thread_limit`

### 2 Summary

3 The `omp_get_thread_limit` routine returns the maximum number of OpenMP threads  
4 available to participate in the current contention group.

### 5 Format

C / C++

```
int omp_get_thread_limit(void);
```

C / C++

Fortran

```
integer function omp_get_thread_limit()
```

Fortran

### 6 Binding

7 The binding thread set for an `omp_get_thread_limit` region is all threads on the device. The  
8 effect of executing this routine is not related to any specific region corresponding to any construct  
9 or API routine.

### 10 Effect

11 The `omp_get_thread_limit` routine returns the value of the *thread-limit-var* ICV.

### 12 Cross References

- 13 • *thread-limit-var* ICV, see Section 2.3 on page 35.
- 14 • `OMP_THREAD_LIMIT` environment variable, see Section 4.10 on page 258.


## 15 3.2.15 `omp_set_max_active_levels`

### 16 Summary

17 The `omp_set_max_active_levels` routine limits the number of nested active parallel  
18 regions on the device, by setting the *max-active-levels-var* ICV

1

## Format



2

## Constraints on Arguments

3

The value of the argument passed to this routine must evaluate to a non-negative integer, otherwise the behavior of this routine is implementation defined.

4

5

## Binding

6

When called from a sequential part of the program, the binding thread set for an **omp\_set\_max\_active\_levels** region is the encountering thread. When called from within any explicit parallel region, the binding thread set (and binding region, if required) for the **omp\_set\_max\_active\_levels** region is implementation defined.

7

8

9

10

## Effect

11

The effect of this routine is to set the value of the *max-active-levels-var* ICV to the value specified in the argument.

12

13

If the number of parallel levels requested exceeds the number of levels of parallelism supported by the implementation, the value of the *max-active-levels-var* ICV will be set to the number of parallel levels supported by the implementation.

14

15

16

This routine has the described effect only when called from a sequential part of the program. When called from within an explicit **parallel** region, the effect of this routine is implementation defined.

17

18

19

## Cross References

20

- *max-active-levels-var* ICV, see Section 2.3 on page 35.

21

- **omp\_get\_max\_active\_levels** routine, see Section 3.2.16 on page 225.

22

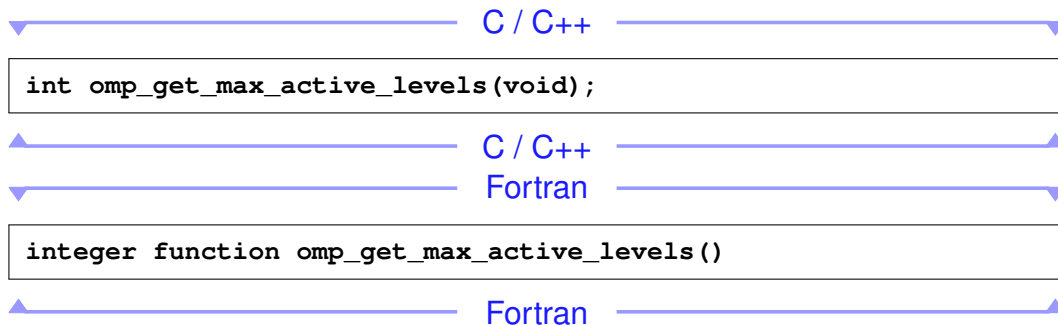
- **OMP\_MAX\_ACTIVE\_LEVELS** environment variable, see Section 4.9 on page 258.

## 1 3.2.16 `omp_get_max_active_levels`

### 2 Summary

3 The `omp_get_max_active_levels` routine returns the value of the *max-active-levels-var*  
4 ICV, which determines the maximum number of nested active parallel regions on the device.

### 5 Format



### 6 Binding

7 When called from a sequential part of the program, the binding thread set for an  
8 `omp_get_max_active_levels` region is the encountering thread. When called from within  
9 any explicit parallel region, the binding thread set (and binding region, if required) for the  
10 `omp_get_max_active_levels` region is implementation defined.

### 11 Effect

12 The `omp_get_max_active_levels` routine returns the value of the *max-active-levels-var*  
13 ICV, which determines the maximum number of nested active parallel regions on the device.

### 14 Cross References

- 15 • *max-active-levels-var* ICV, see Section 2.3 on page 35.
- 16 • `omp_set_max_active_levels` routine, see Section 3.2.15 on page 223.
- 17 • `OMP_MAX_ACTIVE_LEVELS` environment variable, see Section 4.9 on page 258.

## 1 3.2.17 `omp_get_level`

### 2 Summary

3 The `omp_get_level` routine returns the value of the *levels-var* ICV.

### 4 Format

C / C++

```
int omp_get_level(void);
```

C / C++

Fortran

```
integer function omp_get_level()
```

Fortran

### 5 Binding

6 The binding task set for an `omp_get_level` region is the generating task.

### 7 Effect

8 The effect of the `omp_get_level` routine is to return the number of nested `parallel` regions  
9 (whether active or inactive) enclosing the current task such that all of the `parallel` regions are  
10 enclosed by the outermost initial task region on the current device.

### 11 Cross References

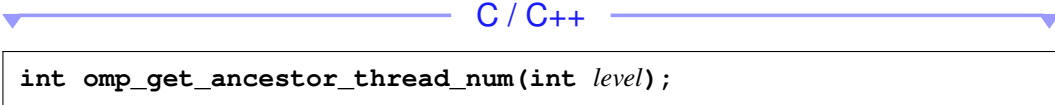

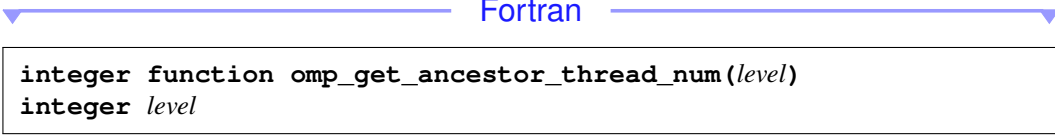

- 12 • *levels-var* ICV, see Section 2.3 on page 35.
- 13 • `omp_get_active_level` routine, see Section 3.2.20 on page 229.
- 14 • `OMP_MAX_ACTIVE_LEVELS` environment variable, see Section 4.9 on page 258.

## 1 3.2.18 `omp_get_ancestor_thread_num`

### 2 Summary

3 The `omp_get_ancestor_thread_num` routine returns, for a given nested level of the current  
4 thread, the thread number of the ancestor of the current thread.

### Format

5  C / C++  
`int omp_get_ancestor_thread_num(int level);`  
 C / C++  
 Fortran  
`integer function omp_get_ancestor_thread_num(level)  
integer level`  
 Fortran

### 6 Binding

7 The binding thread set for an `omp_get_ancestor_thread_num` region is the encountering  
8 thread. The binding region for an `omp_get_ancestor_thread_num` region is the innermost  
9 enclosing `parallel` region.

### 10 Effect

11 The `omp_get_ancestor_thread_num` routine returns the thread number of the ancestor at a  
12 given nest level of the current thread or the thread number of the current thread. If the requested  
13 nest level is outside the range of 0 and the nest level of the current thread, as returned by the  
14 `omp_get_level` routine, the routine returns -1.

15 Note – When the `omp_get_ancestor_thread_num` routine is called with a value of  
16 `level=0`, the routine always returns 0. If `level=omp_get_level()`, the routine has the  
17 same effect as the `omp_get_thread_num` routine.

## Cross References

- `omp_get_level` routine, see Section 3.2.17 on page 226.
- `omp_get_thread_num` routine, see Section 3.2.4 on page 212.
- `omp_get_team_size` routine, see Section 3.2.19 on page 228.

## 3.2.19 `omp_get_team_size`

### Summary

The `omp_get_team_size` routine returns, for a given nested level of the current thread, the size of the thread team to which the ancestor or the current thread belongs.

### Format

C / C++

```
int omp_get_team_size(int level);
```

C / C++

Fortran

```
integer function omp_get_team_size(level)  
integer level
```

Fortran

### Binding

The binding thread set for an `omp_get_team_size` region is the encountering thread. The binding region for an `omp_get_team_size` region is the innermost enclosing `parallel` region.

## Effect

The `omp_get_team_size` routine returns the size of the thread team to which the ancestor or the current thread belongs. If the requested nested level is outside the range of 0 and the nested level of the current thread, as returned by the `omp_get_level` routine, the routine returns -1. Inactive parallel regions are regarded like active parallel regions executed with one thread.

Note – When the `omp_get_team_size` routine is called with a value of `level=0`, the routine always returns 1. If `level=omp_get_level()`, the routine has the same effect as the `omp_get_num_threads` routine.

## Cross References

- `omp_get_num_threads` routine, see Section 3.2.2 on page 209.
- `omp_get_level` routine, see Section 3.2.17 on page 226.
- `omp_get_ancestor_thread_num` routine, see Section 3.2.18 on page 227.

## 3.2.20 `omp_get_active_level`

### Summary

The `omp_get_active_level` routine returns the value of the *active-level-vars* ICV..

### Format

C / C++

```
int omp_get_active_level(void);
```

C / C++

## Fortran

```
integer function omp_get_active_level()
```

## Fortran

### 1 Binding

2 The binding task set for the an `omp_get_active_level` region is the generating task.

### 3 Effect

4 The effect of the `omp_get_active_level` routine is to return the number of nested, active  
5 **parallel** regions enclosing the current task such that all of the **parallel** regions are enclosed  
6 by the outermost initial task region on the current device.

### 7 Cross References

- 8 • *active-levels-var* ICV, see Section [2.3](#) on page [35](#).
- 9 • `omp_get_level` routine, see Section [3.2.17](#) on page [226](#).

## 10 3.2.21 `omp_in_final`

### 11 Summary

12 The `omp_in_final` routine returns *true* if the routine is executed in a final task region;  
13 otherwise, it returns *false*.

### 14 Format

#### C / C++

```
int omp_in_final(void);
```

#### C / C++

#### Fortran

```
logical function omp_in_final()
```

#### Fortran



1        **Binding**

2        The binding task set for an `omp_in_final` region is the generating task.

3        **Effect**

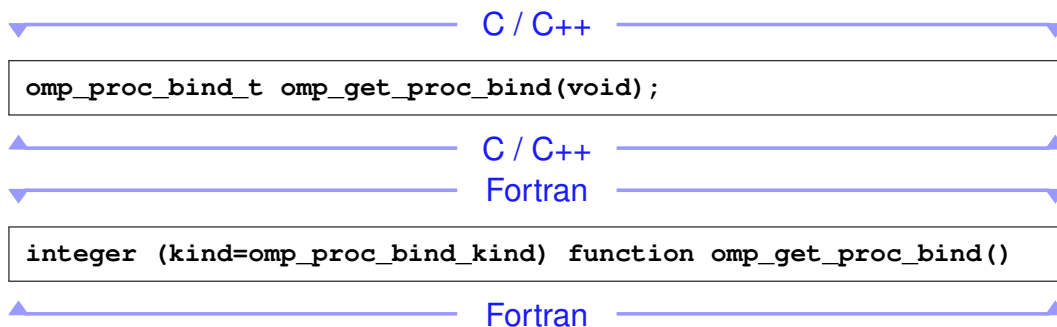
4        `omp_in_final` returns *true* if the enclosing task region is final. Otherwise, it returns *false*.

5    **3.2.22   `omp_get_proc_bind`**

6        **Summary**

7        The `omp_get_proc_bind` routine returns the thread affinity policy to be used for the  
8        subsequent nested `parallel` regions that do not specify a `proc_bind` clause.

9        **Format**



## Constraints on Arguments

The value returned by this routine must be one of the valid affinity policy kinds. The C/ C++ header file (`omp.h`) and the Fortran include file (`omp_lib.h`) and/or Fortran 90 module file (`omp_lib`) define the valid constants. The valid constants must include the following:

C / C++

```
typedef enum omp_proc_bind_t {
    omp_proc_bind_false = 0,
    omp_proc_bind_true = 1,
    omp_proc_bind_master = 2,
    omp_proc_bind_close = 3,
    omp_proc_bind_spread = 4
} omp_proc_bind_t;
```

C / C++

Fortran

```
integer (kind=omp_proc_bind_kind), &
    parameter :: omp_proc_bind_false = 0
integer (kind=omp_proc_bind_kind), &
    parameter :: omp_proc_bind_true = 1
integer (kind=omp_proc_bind_kind), &
    parameter :: omp_proc_bind_master = 2
integer (kind=omp_proc_bind_kind), &
    parameter :: omp_proc_bind_close = 3
integer (kind=omp_proc_bind_kind), &
    parameter :: omp_proc_bind_spread = 4
```

Fortran

## Binding

The binding task set for an `omp_get_proc_bind` region is the generating task

## Effect

The effect of this routine is to return the value of the first element of the *bind-var* ICV of the current task. See Section 2.5.2 on page 50 for the rules governing the thread affinity policy.

## Cross References

- *bind-var* ICV, see Section 2.3 on page 35.
- `OMP_PROC_BIND` environment variable, see Section 4.4 on page 253.
- Controlling OpenMP thread affinity, see Section 2.5.2 on page 50.

### 3.2.23 `omp_set_default_device`

#### Summary

The `omp_set_default_device` routine controls the default target device by assigning the value of the *default-device-var* ICV.

#### Format

C / C++

```
void omp_set_default_device(int device_num);
```

C / C++

Fortran

```
subroutine omp_set_default_device(device_num)  
integer device_num
```

Fortran

#### Binding

The binding task set for an `omp_set_default_device` region is the generating task.

#### Effect

The effect of this routine is to set the value of the *default-device-var* ICV of the current task to the value specified in the argument. When called from within a **target** region the effect of this routine is unspecified.

## Cross References

- *default-device-var*, see Section 2.3 on page 35.
- `omp_get_default_device`, see Section 3.2.24 on page 234.
- `OMP_DEFAULT_DEVICE` environment variable, see Section 4.13 on page 260

## 3.2.24 `omp_get_default_device`

### Summary

The `omp_get_default_device` routine returns the default target device.

### Format

C / C++

```
int omp_get_default_device(void);
```

C / C++

Fortran

```
integer function omp_get_default_device()
```

Fortran

### Binding

The binding task set for an `omp_get_default_device` region is the generating task.

### Effect

The `omp_get_default_device` routine returns the value of the *default-device-var* ICV of the current task. When called from within a **target** region the effect of this routine is unspecified.

### Cross References

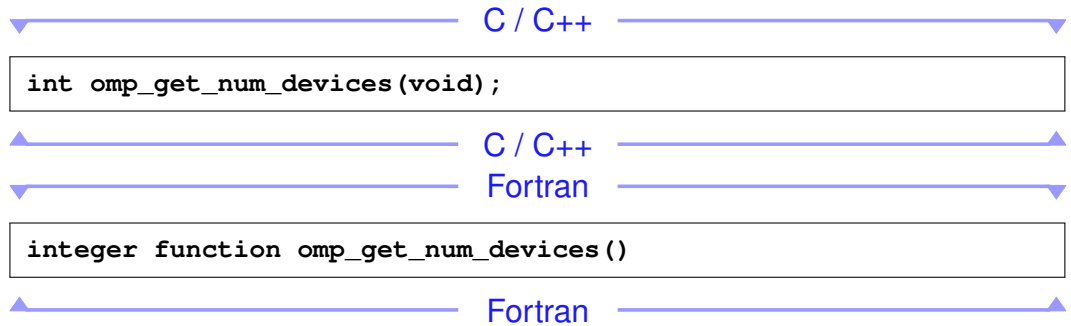
- *default-device-var*, see Section 2.3 on page 35.
- `omp_set_default_device`, see Section 3.2.23 on page 233.
- `OMP_DEFAULT_DEVICE` environment variable, see Section 4.13 on page 260.

## 1 3.2.25 `omp_get_num_devices`

### 2 Summary

3 The `omp_get_num_devices` routine returns the number of target devices.

### 4 Format



### 5 Binding

6 The binding task set for an `omp_get_num_devices` region is the generating task.

### 7 Effect

8 The `omp_get_num_devices` routine returns the number of available target devices. When  
9 called from within a `target` region the effect of this routine is unspecified.

### 10 Cross References

11 None.

## 12 3.2.26 `omp_get_num_teams`

### 13 Summary

14 The `omp_get_num_teams` routine returns the number of teams in the current `teams` region.

1

## Format

▼ C / C++ ▼

```
int omp_get_num_teams(void);
```

▲ C / C++ ▲

▼ Fortran ▼

```
integer function omp_get_num_teams()
```

▲ Fortran ▲

2

## Binding

3

The binding task set for an `omp_get_num_teams` region is the generating task

4

## Effect

5

The effect of this routine is to return the number of teams in the current `teams` region. The routine returns 1 if it is called from outside of a `teams` region.

6

7

## Cross References

8

- `teams` construct, see Section [2.10.5](#) on page [102](#).

## 1 3.2.27 `omp_get_team_num`

### 2 Summary

3 The `omp_get_team_num` routine returns the team number of the calling thread.

### 4 Format

▼ C / C++ ▼

```
int omp_get_team_num(void);
```

▲ C / C++ ▲

▼ Fortran ▼

```
integer function omp_get_team_num()
```

▲ Fortran ▲

### 5 Binding

6 The binding task set for an `omp_get_team_num` region is the generating task.

### 7 Effect

8 The `omp_get_team_num` routine returns the team number of the calling thread. The team  
9 number is an integer between 0 and one less than the value returned by `omp_get_num_teams`,  
10 inclusive. The routine returns 0 if it is called outside of a `teams` region.

### 11 Cross References

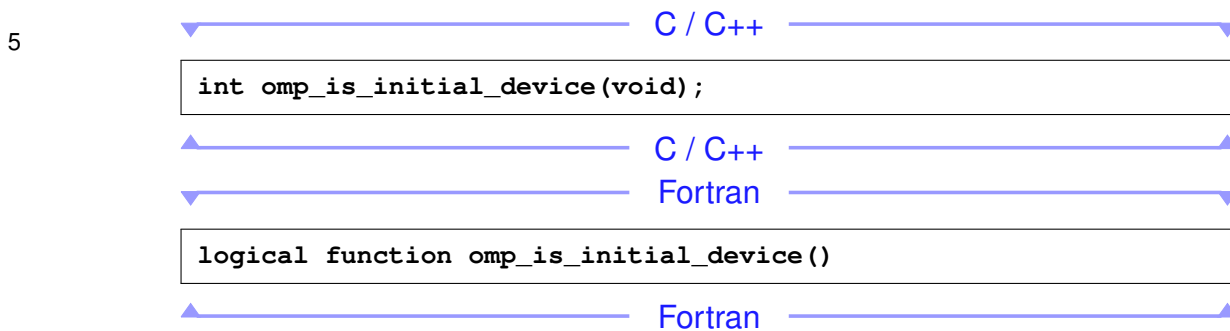
- 12 • `teams` construct, see Section [2.10.5](#) on page [102](#).
- 13 • `omp_get_num_teams` routine, see Section [3.2.26](#) on page [235](#).

## 1 3.2.28 `omp_is_initial_device`

### 2 Summary

3 The `omp_is_initial_device` routine returns *true* if the current task is executing on the host  
4 device; otherwise, it returns *false*.

### Format



### 6 Binding

7 The binding task set for an `omp_is_initial_device` region is the generating task.

### 8 Effect

9 The effect of this routine is to return *true* if the current task is executing on the host device;  
10 otherwise, it returns *false*.

### 11 Cross References

- 12 • **target** construct, see Section [2.10.2](#) on page [93](#)



## 1 3.3 Lock Routines

2 The OpenMP runtime library includes a set of general-purpose lock routines that can be used for  
3 synchronization. These general-purpose lock routines operate on OpenMP locks that are  
4 represented by OpenMP lock variables. OpenMP lock variables must be accessed only through the  
5 routines described in this section; programs that otherwise access OpenMP lock variables are  
6 non-conforming.

7 An OpenMP lock can be in one of the following states: *uninitialized*, *unlocked*, or *locked*. If a lock  
8 is in the unlocked state, a task can *set* the lock, which changes its state to *locked*. The task that sets  
9 the lock is then said to *own* the lock. A task that owns a lock can *unset* that lock, returning it to the  
10 *unlocked* state. A program in which a task unsets a lock that is owned by another task is  
11 non-conforming.

12 Two types of locks are supported: *simple locks* and *nestable locks*. A *nestable lock* can be set  
13 multiple times by the same task before being unset; a *simple lock* cannot be set if it is already  
14 owned by the task trying to set it. *Simple lock* variables are associated with *simple locks* and can  
15 only be passed to *simple lock* routines. *Nestable lock* variables are associated with *nestable locks*  
16 and can only be passed to *nestable lock* routines.

17 Constraints on the state and ownership of the lock accessed by each of the lock routines are  
18 described with the routine. If these constraints are not met, the behavior of the routine is  
19 unspecified.

20 The OpenMP lock routines access a lock variable in such a way that they always read and update  
21 the most current value of the lock variable. It is not necessary for an OpenMP program to include  
22 explicit **flush** directives to ensure that the lock variable's value is consistent among different  
23 tasks.

### 24 Binding

25 The binding thread set for all lock routine regions is all threads in the contention group. As a  
26 consequence, for each OpenMP lock, the lock routine effects relate to all tasks that call the routines,  
27 without regard to which teams the threads in the contention group executing the tasks belong.

### 28 Simple Lock Routines

▼ C / C++ ▼

29 The type **omp\_lock\_t** is a data type capable of representing a simple lock. For the following  
30 routines, a simple lock variable must be of **omp\_lock\_t** type. All simple lock routines require an  
31 argument that is a pointer to a variable of type **omp\_lock\_t**.

▲ C / C++ ▲

## Fortran

1 For the following routines, a simple lock variable must be an integer variable of  
2 **kind=omp\_lock\_kind**.

## Fortran

3 The simple lock routines are as follows:

- 4 • The **omp\_init\_lock** routine initializes a simple lock.
- 5 • The **omp\_destroy\_lock** routine uninitialized a simple lock.
- 6 • The **omp\_set\_lock** routine waits until a simple lock is available, and then sets it.
- 7 • The **omp\_unset\_lock** routine unsets a simple lock.
- 8 • The **omp\_test\_lock** routine tests a simple lock, and sets it if it is available.

## 9 Nestable Lock Routines

### C / C++

10 The type **omp\_nest\_lock\_t** is a data type capable of representing a nestable lock. For the  
11 following routines, a nested lock variable must be of **omp\_nest\_lock\_t** type. All nestable lock  
12 routines require an argument that is a pointer to a variable of type **omp\_nest\_lock\_t**.

### C / C++

### Fortran

13 For the following routines, a nested lock variable must be an integer variable of  
14 **kind=omp\_nest\_lock\_kind**.

### Fortran

15 The nestable lock routines are as follows:

- 16 • The **omp\_init\_nest\_lock** routine initializes a nestable lock.
- 17 • The **omp\_destroy\_nest\_lock** routine uninitialized a nestable lock.
- 18 • The **omp\_set\_nest\_lock** routine waits until a nestable lock is available, and then sets it.
- 19 • The **omp\_unset\_nest\_lock** routine unsets a nestable lock.
- 20 • The **omp\_test\_nest\_lock** routine tests a nestable lock, and sets it if it is available

## 21 Restrictions

22 OpenMP lock routines have the following restrictions:

- 23 • The use of the same OpenMP lock in different contention groups results in unspecified behavior.

### 1 3.3.1 `omp_init_lock` and `omp_init_nest_lock`

#### 2 Summary

3 These routines provide the only means of initializing an OpenMP lock.

#### 4 Format

C / C++

```
void omp_init_lock(omp_lock_t *lock);  
void omp_init_nest_lock(omp_nest_lock_t *lock);
```

C / C++

Fortran

```
subroutine omp_init_lock(svar)  
integer (kind=omp_lock_kind) svar  
  
subroutine omp_init_nest_lock(nvar)  
integer (kind=omp_nest_lock_kind) nvar
```

Fortran

#### 5 Constraints on Arguments

6 A program that accesses a lock that is not in the uninitialized state through either routine is  
7 non-conforming.

#### 8 Effect

9 The effect of these routines is to initialize the lock to the unlocked state; that is, no task owns the  
10 lock. In addition, the nesting count for a nestable lock is set to zero.

### 11 3.3.2 `omp_destroy_lock` and 12 `omp_destroy_nest_lock`

#### 13 Summary

14 These routines ensure that the OpenMP lock is uninitialized

1

## Format

C / C++

```
void omp_destroy_lock(omp_lock_t *lock);
void omp_destroy_nest_lock(omp_nest_lock_t *lock);
```

C / C++

Fortran

```
subroutine omp_destroy_lock(svar)
integer (kind=omp_lock_kind) svar

subroutine omp_destroy_nest_lock(nvar)
integer (kind=omp_nest_lock_kind) nvar
```

Fortran

2

## Constraints on Arguments

3

A program that accesses a lock that is not in the unlocked state through either routine is non-conforming.

4

5

## Effect

6

The effect of these routines is to change the state of the lock to uninitialized.

## 7 3.3.3 omp\_set\_lock and omp\_set\_nest\_lock

8

### Summary

9

These routines provide a means of setting an OpenMP lock. The calling task region is suspended until the lock is set.

10

1

## Format

C / C++

```
void omp_set_lock(omp_lock_t *lock);
void omp_set_nest_lock(omp_nest_lock_t *lock);
```

C / C++

Fortran

```
subroutine omp_set_lock(svar)
integer (kind=omp_lock_kind) svar

subroutine omp_set_nest_lock(nvar)
integer (kind=omp_nest_lock_kind) nvar
```

Fortran

2

## Constraints on Arguments

3

A program that accesses a lock that is in the uninitialized state through either routine is

4

non-conforming. A simple lock accessed by `omp_set_lock` that is in the locked state must not

5

be owned by the task that contains the call or deadlock will result.

6

## Effect

7

Each of these routines causes suspension of the task executing the routine until the specified lock is available and then sets the lock.

8

9

A simple lock is available if it is unlocked. Ownership of the lock is granted to the task executing the routine.

10

11

A nestable lock is available if it is unlocked or if it is already owned by the task executing the

12

routine. The task executing the routine is granted, or retains, ownership of the lock, and the nesting

13

count for the lock is incremented.

### 14 3.3.4 `omp_unset_lock` and `omp_unset_nest_lock`

15

## Summary

16

These routines provide the means of unsetting an OpenMP lock.

1

## Format

C / C++

```
void omp_unset_lock(omp_lock_t *lock);
void omp_unset_nest_lock(omp_nest_lock_t *lock);
```

C / C++

Fortran

```
subroutine omp_unset_lock(svar)
integer (kind=omp_lock_kind) svar

subroutine omp_unset_nest_lock(nvar)
integer (kind=omp_nest_lock_kind) nvar
```

Fortran

2

## Constraints on Arguments

3

A program that accesses a lock that is not in the locked state or that is not owned by the task that contains the call through either routine is non-conforming.

4

5

## Effect

6

For a simple lock, the **omp\_unset\_lock** routine causes the lock to become unlocked.

7

For a nestable lock, the **omp\_unset\_nest\_lock** routine decrements the nesting count, and causes the lock to become unlocked if the resulting nesting count is zero.

8

9

For either routine, if the lock becomes unlocked, and if one or more task regions were suspended because the lock was unavailable, the effect is that one task is chosen and given ownership of the lock.

10

11

## 12 3.3.5 omp\_test\_lock and omp\_test\_nest\_lock

13

### Summary

14

These routines attempt to set an OpenMP lock but do not suspend execution of the task executing the routine.

15

1

## Format

C / C++

```
int omp_test_lock(omp_lock_t *lock);
int omp_test_nest_lock(omp_nest_lock_t *lock);
```

C / C++

Fortran

```
logical function omp_test_lock(svar)
integer (kind=omp_lock_kind) svar
integer function omp_test_nest_lock(nvar)
integer (kind=omp_nest_lock_kind) nvar
```

Fortran

2

## Constraints on Arguments

3

A program that accesses a lock that is in the uninitialized state through either routine is non-conforming. The behavior is unspecified if a simple lock accessed by **omp\_test\_lock** is in the locked state and is owned by the task that contains the call.

4

6

## Effect

7

These routines attempt to set a lock in the same manner as **omp\_set\_lock** and **omp\_set\_nest\_lock**, except that they do not suspend execution of the task executing the routine.

8

9

For a simple lock, the **omp\_test\_lock** routine returns *true* if the lock is successfully set; otherwise, it returns *false*.

10

11

For a nestable lock, the **omp\_test\_nest\_lock** routine returns the new nesting count if the lock is successfully set; otherwise, it returns zero.

12

13

## 14 3.4 Timing Routines

15

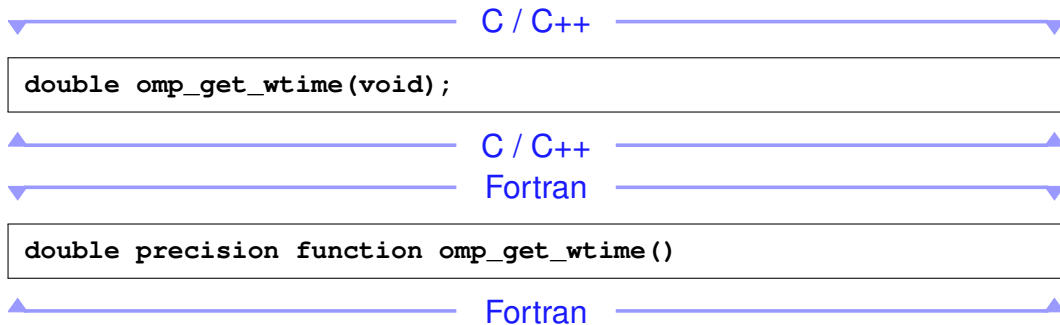
This section describes routines that support a portable wall clock timer.

## 1 3.4.1 `omp_get_wtime`

### 2 Summary

3 The `omp_get_wtime` routine returns elapsed wall clock time in seconds.

### 4 Format



### 5 Binding

6 The binding thread set for an `omp_get_wtime` region is the encountering thread. The routine's  
7 return value is not guaranteed to be consistent across any set of threads.



1       **Effect**

2       The `omp_get_wtime` routine returns a value equal to the elapsed wall clock time in seconds  
3       since some “time in the past”. The actual “time in the past” is arbitrary, but it is guaranteed not to  
4       change during the execution of the application program. The time returned is a “per-thread time”,  
5       so it is not required to be globally consistent across all the threads participating in an application.

6       **Note** – It is anticipated that the routine will be used to measure elapsed times as shown in the  
7       following example:

▼────────────────── C / C++ ───────────────────▶

```
double start;  
double end;  
start = omp_get_wtime();  
... work to be timed ...  
end = omp_get_wtime();  
printf("Work took %f seconds\n", end - start);
```

▲────────────────── C / C++ ───────────────────▶

▼────────────────── Fortran ───────────────────▶

```
DOUBLE PRECISION START, END  
START = omp_get_wtime()  
... work to be timed ...  
END = omp_get_wtime()  
PRINT *, "Work took", END - START, "seconds"
```

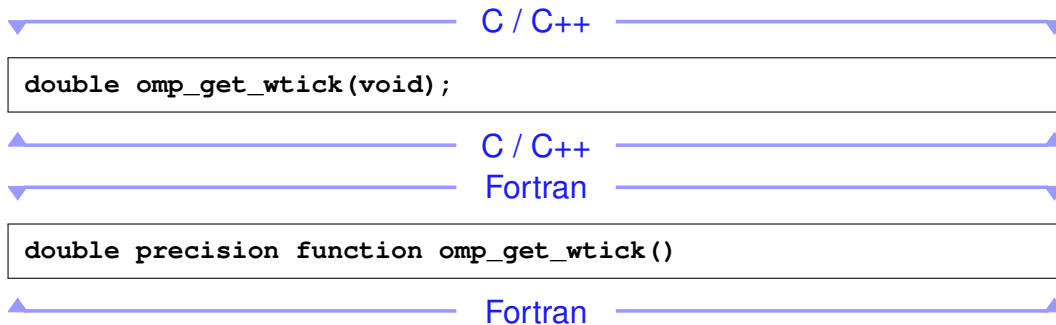
▲────────────────── Fortran ───────────────────▶

## 1 3.4.2 `omp_get_wtick`

### 2 Summary

3 The `omp_get_wtick` routine returns the precision of the timer used by `omp_get_wtime`.

### 4 Format



### 5 Binding

6 The binding thread set for an `omp_get_wtick` region is the encountering thread. The routine's  
7 return value is not guaranteed to be consistent across any set of threads.

### 8 Effect

9 The `omp_get_wtick` routine returns a value equal to the number of seconds between successive  
10 clock ticks of the timer used by `omp_get_wtime`.

2 

## Environment Variables

---

3 This chapter describes the OpenMP environment variables that specify the settings of the ICVs that  
4 affect the execution of OpenMP programs (see Section 2.3 on page 35). The names of the  
5 environment variables must be upper case. The values assigned to the environment variables are  
6 case insensitive and may have leading and trailing white space. Modifications to the environment  
7 variables after the program has started, even if modified by the program itself, are ignored by the  
8 OpenMP implementation. However, the settings of some of the ICVs can be modified during the  
9 execution of the OpenMP program by the use of the appropriate directive clauses or OpenMP API  
10 routines.

11 The environment variables are as follows:

- 12 ● **OMP\_SCHEDULE** sets the *run-sched-var* ICV that specifies the runtime schedule type and chunk  
13 size. It can be set to any of the valid OpenMP schedule types.
- 14 ● **OMP\_NUM\_THREADS** sets the *nthreads-var* ICV that specifies the number of threads to use for  
15 parallel regions.
- 16 ● **OMP\_DYNAMIC** sets the *dyn-var* ICV that specifies the dynamic adjustment of threads to use for  
17 **parallel** regions.
- 18 ● **OMP\_PROC\_BIND** sets the *bind-var* ICV that controls the OpenMP thread affinity policy.
- 19 ● **OMP\_PLACES** sets the *place-partition-var* ICV that defines the OpenMP places that are  
20 available to the execution environment.
- 21 ● **OMP\_NESTED** sets the *nest-var* ICV that enables or disables nested parallelism.
- 22 ● **OMP\_STACKSIZE** sets the *stacksize-var* ICV that specifies the size of the stack for threads  
23 created by the OpenMP implementation.
- 24 ● **OMP\_WAIT\_POLICY** sets the *wait-policy-var* ICV that controls the desired behavior of waiting  
25 threads.
- 26 ● **OMP\_MAX\_ACTIVE\_LEVELS** sets the *max-active-levels-var* ICV that controls the maximum  
27 number of nested active **parallel** regions.

- 1       • **OMP\_THREAD\_LIMIT** sets the *thread-limit-var* ICV that controls the maximum number of  
2       threads participating in a contention group.
- 3       • **OMP\_CANCELLATION** sets the *cancel-var* ICV that enables or disables cancellation.
- 4       • **OMP\_DISPLAY\_ENV** instructs the runtime to display the OpenMP version number and the  
5       initial values of the ICVs, once, during initialization of the runtime.
- 6       • **OMP\_DEFAULT\_DEVICE** sets the *default-device-var* ICV that controls the default device  
7       number.

8       The examples in this chapter only demonstrate how these variables might be set in Unix C shell  
9       (csh) environments. In Korn shell (ksh) and DOS environments the actions are similar, as follows:

- 10      • csh:

```
setenv OMP_SCHEDULE "dynamic"
```

- 11      • ksh:

```
export OMP_SCHEDULE="dynamic"
```

- 12      • DOS:

```
set OMP_SCHEDULE=dynamic
```

## 1 4.1 OMP\_SCHEDULE

2 The **OMP\_SCHEDULE** environment variable controls the schedule type and chunk size of all loop  
3 directives that have the schedule type **runtime**, by setting the value of the *run-sched-var* ICV.

4 The value of this environment variable takes the form:

5 *type*[, *chunk*]

6 where

- 7 • *type* is one of **static**, **dynamic**, **guided**, or **auto**
- 8 • *chunk* is an optional positive integer that specifies the chunk size

9 If *chunk* is present, there may be white space on either side of the “,”. See Section 2.7.1 on  
10 page 55 for a detailed description of the schedule types.

11 The behavior of the program is implementation defined if the value of **OMP\_SCHEDULE** does not  
12 conform to the above format.

13 Implementation specific schedules cannot be specified in **OMP\_SCHEDULE**. They can only be  
14 specified by calling **omp\_set\_schedule**, described in Section 3.2.12 on page 220.

15 Example:

```
setenv OMP_SCHEDULE "guided, 4"  
setenv OMP_SCHEDULE "dynamic"
```

### 16 Cross References

- 17 • *run-sched-var* ICV, see Section 2.3 on page 35.
- 18 • Loop construct, see Section 2.7.1 on page 55.
- 19 • Parallel loop construct, see Section 2.11.1 on page 117.
- 20 • **omp\_set\_schedule** routine, see Section 3.2.12 on page 220.
- 21 • **omp\_get\_schedule** routine, see Section 3.2.13 on page 222.

## 1 4.2 OMP\_NUM\_THREADS

2 The **OMP\_NUM\_THREADS** environment variable sets the number of threads to use for **parallel**  
3 regions by setting the initial value of the *nthreads-var* ICV. See Section 2.3 on page 35 for a  
4 comprehensive set of rules about the interaction between the **OMP\_NUM\_THREADS** environment  
5 variable, the **num\_threads** clause, the **omp\_set\_num\_threads** library routine and dynamic  
6 adjustment of threads, and Section 2.5.1 on page 48 for a complete algorithm that describes how the  
7 number of threads for a **parallel** region is determined.

8 The value of this environment variable must be a list of positive integer values. The values of the  
9 list set the number of threads to use for **parallel** regions at the corresponding nested levels.

10 The behavior of the program is implementation defined if any value of the list specified in the  
11 **OMP\_NUM\_THREADS** environment variable leads to a number of threads which is greater than an  
12 implementation can support, or if any value is not a positive integer.

13 Example:

```
setenv OMP_NUM_THREADS 4,3,2
```

### 14 Cross References

- 15 • *nthreads-var* ICV, see Section 2.3 on page 35.
- 16 • **num\_threads** clause, Section 2.5 on page 44.
- 17 • **omp\_set\_num\_threads** routine, see Section 3.2.1 on page 208.
- 18 • **omp\_get\_num\_threads** routine, see Section 3.2.2 on page 209.
- 19 • **omp\_get\_max\_threads** routine, see Section 3.2.3 on page 210.
- 20 • **omp\_get\_team\_size** routine, see Section 3.2.19 on page 228.

## 1 4.3 OMP\_DYNAMIC

2 The **OMP\_DYNAMIC** environment variable controls dynamic adjustment of the number of threads  
3 to use for executing **parallel** regions by setting the initial value of the *dyn-var* ICV. The value of  
4 this environment variable must be **true** or **false**. If the environment variable is set to **true**, the  
5 OpenMP implementation may adjust the number of threads to use for executing **parallel**  
6 regions in order to optimize the use of system resources. If the environment variable is set to  
7 **false**, the dynamic adjustment of the number of threads is disabled. The behavior of the program  
8 is implementation defined if the value of **OMP\_DYNAMIC** is neither **true** nor **false**.

9 Example:

```
setenv OMP_DYNAMIC true
```

### 10 Cross References

- 11 • *dyn-var* ICV, see Section 2.3 on page 35.
- 12 • **omp\_set\_dynamic** routine, see Section 3.2.7 on page 214.
- 13 • **omp\_get\_dynamic** routine, see Section 3.2.8 on page 216.

## 14 4.4 OMP\_PROC\_BIND

15 The **OMP\_PROC\_BIND** environment variable sets the initial value of the *bind-var* ICV. The value  
16 of this environment variable is either **true**, **false**, or a comma separated list of **master**,  
17 **close**, or **spread**. The values of the list set the thread affinity policy to be used for parallel  
18 regions at the corresponding nested level.

19 If the environment variable is set to **false**, the execution environment may move OpenMP threads  
20 between OpenMP places, thread affinity is disabled, and **proc\_bind** clauses on **parallel**  
21 constructs are ignored.

22 Otherwise, the execution environment should not move OpenMP threads between OpenMP places,  
23 thread affinity is enabled, and the initial thread is bound to the first place in the OpenMP place list.

24 The behavior of the program is implementation defined if any of the values in the  
25 **OMP\_PROC\_BIND** environment variable is not **true**, **false**, or a comma separated list of  
26 **master**, **close**, or **spread**. The behavior is also implementation defined if an initial thread  
27 cannot be bound to the first place in the OpenMP place list.

28 Example:

```
setenv OMP_PROC_BIND false
setenv OMP_PROC_BIND "spread, spread, close"
```

## Cross References

- *bind-var* ICV, see Section 2.3 on page 35.
- `proc_bind` clause, see Section 2.5.2 on page 50.
- `omp_get_proc_bind` routine, see Section 3.2.22 on page 231.

## 4.5 OMP\_PLACES

A list of places can be specified in the **OMP\_PLACES** environment variable. The *place-partition-var* ICV obtains its initial value from the **OMP\_PLACES** value, and makes the list available to the execution environment. The value of **OMP\_PLACES** can be one of two types of values: either an abstract name describing a set of places or an explicit list of places described by non-negative numbers.

The **OMP\_PLACES** environment variable can be defined using an explicit ordered list of comma-separated places. A place is defined by an unordered set of comma-separated non-negative numbers enclosed by braces. The meaning of the numbers and how the numbering is done are implementation defined. Generally, the numbers represent the smallest unit of execution exposed by the execution environment, typically a hardware thread.

Intervals may also be used to define places. Intervals can be specified using the *<lower-bound> : <length> : <stride>* notation to represent the following list of numbers: “*<lower-bound>*, *<lower-bound> + <stride>*, ..., *<lower-bound> + (<length>- 1)\*<stride>*.” When *<stride>* is omitted, a unit stride is assumed. Intervals can specify numbers within a place as well as sequences of places.

An exclusion operator “!” can also be used to exclude the number or place immediately following the operator.

Alternatively, the abstract names listed in TABLE 4-1 should be understood by the execution and runtime environment. The precise definitions of the abstract names are implementation defined. An implementation may also add abstract names as appropriate for the target platform.

The abstract name may be appended by a positive number in parentheses to denote the length of the place list to be created, that is *abstract\_name(num\_places)*. When requesting fewer places than available on the system, the determination of which resources of type *abstract\_name* are to be



1 included in the place list is implementation defined. When requesting more resources than  
2 available, the length of the place list is implementation defined.

3 **TABLE 4-1** List of defined abstract names for **OMP\_PLACES**

<b>Abstract Name</b>	<b>Meaning</b>
<b>threads</b>	Each place corresponds to a single hardware thread on the target machine.
<b>cores</b>	Each place corresponds to a single core (having one or more hardware threads) on the target machine.
<b>sockets</b>	Each place corresponds to a single socket (consisting of one or more cores) on the target machine.

5 The behavior of the program is implementation defined when the execution environment cannot  
6 map a numerical value (either explicitly defined or implicitly derived from an interval) within the  
7 **OMP\_PLACES** list to a processor on the target platform, or if it maps to an unavailable processor.  
8 The behavior is also implementation defined when the **OMP\_PLACES** environment variable is  
9 defined using an abstract name.

10 Example:  
11

```
setenv OMP_PLACES threads
setenv OMP_PLACES "threads(4)"
setenv OMP_PLACES "{0,1,2,3},{4,5,6,7},{8,9,10,11},{12,13,14,15}"
setenv OMP_PLACES "{0:4},{4:4},{8:4},{12:4}"
setenv OMP_PLACES "{0:4}:4:4"
```

12 where each of the last three definitions corresponds to the same 4 places including the smallest  
13 units of execution exposed by the execution environment numbered, in turn, 0 to 3, 4 to 7, 8 to 11,  
14 and 12 to 15.

## 15 **Cross References**

- 16 • *place-partition-var*, Section 2.3 on page 35.
- 17 • Controlling OpenMP thread affinity, Section 2.5.2 on page 50.

## 1 4.6 OMP\_NESTED

2 The **OMP\_NESTED** environment variable controls nested parallelism by setting the initial value of  
3 the *nest-var* ICV. The value of this environment variable must be **true** or **false**. If the  
4 environment variable is set to **true**, nested parallelism is enabled; if set to **false**, nested  
5 parallelism is disabled. The behavior of the program is implementation defined if the value of  
6 **OMP\_NESTED** is neither **true** nor **false**.

7 Example:

```
setenv OMP_NESTED false
```

### 8 Cross References

- 9 • *nest-var* ICV, see Section 2.3 on page 35.
- 10 • **omp\_set\_nested** routine, see Section 3.2.10 on page 217.
- 11 • **omp\_get\_team\_size** routine, see Section 3.2.19 on page 228.

## 12 4.7 OMP\_STACKSIZE

13 The **OMP\_STACKSIZE** environment variable controls the size of the stack for threads created by  
14 the OpenMP implementation, by setting the value of the *stacksize-var* ICV. The environment  
15 variable does not control the size of the stack for an initial thread.

16 The value of this environment variable takes the form:

17 *size* | *size***B** | *size***K** | *size***M** | *size***G**

18 where:

- 19 • *size* is a positive integer that specifies the size of the stack for threads that are created by the  
20 OpenMP implementation.
- 21 • **B**, **K**, **M**, and **G** are letters that specify whether the given size is in Bytes, Kilobytes (1024 Bytes),  
22 Megabytes (1024 Kilobytes), or Gigabytes (1024 Megabytes), respectively. If one of these letters  
23 is present, there may be white space between *size* and the letter.

1 If only *size* is specified and none of **B**, **K**, **M**, or **G** is specified, then *size* is assumed to be in Kilobytes.  
2 The behavior of the program is implementation defined if **OMP\_STACKSIZE** does not conform to  
3 the above format, or if the implementation cannot provide a stack with the requested size.

4 Examples:

```
setenv OMP_STACKSIZE 2000500B
setenv OMP_STACKSIZE "3000 k "
setenv OMP_STACKSIZE 10M
setenv OMP_STACKSIZE " 10 M "
setenv OMP_STACKSIZE "20 m "
setenv OMP_STACKSIZE " 1G"
setenv OMP_STACKSIZE 20000
```

## 5 **Cross References**

- 6 • *stacksize-var* ICV, see Section [2.3](#) on page [35](#).

## 7 **4.8 OMP\_WAIT\_POLICY**

8 The **OMP\_WAIT\_POLICY** environment variable provides a hint to an OpenMP  
9 implementation about the desired behavior of waiting threads by setting the *wait-policy-var* ICV. A  
10 compliant OpenMP implementation may or may not abide by the setting of the environment  
11 variable.

12 The value of this environment variable takes the form:

13 **ACTIVE** | **PASSIVE**

14 The **ACTIVE** value specifies that waiting threads should mostly be active, consuming processor  
15 cycles, while waiting. An OpenMP implementation may, for example, make waiting threads spin.

16 The **PASSIVE** value specifies that waiting threads should mostly be passive, not consuming  
17 processor cycles, while waiting. For example, an OpenMP implementation may make waiting  
18 threads yield the processor to other threads or go to sleep.

19 The details of the **ACTIVE** and **PASSIVE** behaviors are implementation defined.

20 Examples:

```
setenv OMP_WAIT_POLICY ACTIVE
setenv OMP_WAIT_POLICY active
setenv OMP_WAIT_POLICY PASSIVE
setenv OMP_WAIT_POLICY passive
```

## Cross References

- *wait-policy-var* ICV, see Section 2.3 on page 35.

## 4.9 OMP\_MAX\_ACTIVE\_LEVELS

The **OMP\_MAX\_ACTIVE\_LEVELS** environment variable controls the maximum number of nested active **parallel** regions by setting the initial value of the *max-active-levels-var* ICV.

The value of this environment variable must be a non-negative integer. The behavior of the program is implementation defined if the requested value of **OMP\_MAX\_ACTIVE\_LEVELS** is greater than the maximum number of nested active parallel levels an implementation can support, or if the value is not a non-negative integer.

## Cross References

- *max-active-levels-var* ICV, see Section 2.3 on page 35.
- **omp\_set\_max\_active\_levels** routine, see Section 3.2.15 on page 223.
- **omp\_get\_max\_active\_levels** routine, see Section 3.2.16 on page 225.

## 4.10 OMP\_THREAD\_LIMIT

The **OMP\_THREAD\_LIMIT** environment variable sets the maximum number of OpenMP threads to use in a contention group by setting the *thread-limit-var* ICV.

The value of this environment variable must be a positive integer. The behavior of the program is implementation defined if the requested value of **OMP\_THREAD\_LIMIT** is greater than the number of threads an implementation can support, or if the value is not a positive integer.

## Cross References

- *thread-limit-var* ICV, see Section 2.3 on page 35.
- `omp_get_thread_limit` routine, see Section 3.2.14 on page 223.

## 4.11 OMP\_CANCELLATION

The `OMP_CANCELLATION` environment variable sets the initial value of the *cancel-var* ICV.

The value of this environment variable must be `true` or `false`. If set to `true`, the effects of the `cancel` construct and of cancellation points are enabled and cancellation is activated. If set to `false`, cancellation is disabled and the `cancel` construct and cancellation points are effectively ignored.

## Cross References

- *cancel-var*, see Section 2.3.1 on page 35.
- `cancel` construct, see Section 2.13.1 on page 156.
- `cancellation point` construct, see Section 2.13.2 on page 160.
- `omp_get_cancellation` routine, see Section 3.2.9 on page 217.

## 4.12 OMP\_DISPLAY\_ENV

The `OMP_DISPLAY_ENV` environment variable instructs the runtime to display the OpenMP version number and the value of the ICVs associated with the environment variables described in Chapter 4, as *name = value* pairs. The runtime displays this information once, after processing the environment variables and before any user calls to change the ICV values by runtime routines defined in Chapter 3.

The value of the `OMP_DISPLAY_ENV` environment variable may be set to one of these values:

**TRUE | FALSE | VERBOSE**

The `TRUE` value instructs the runtime to display the OpenMP version number defined by the `_OPENMP` version macro (or the `openmp_version` Fortran parameter) value and the initial ICV

1 values for the environment variables listed in Chapter 4. The **VERBOSE** value indicates that the  
2 runtime may also display the values of runtime variables that may be modified by vendor-specific  
3 environment variables. The runtime does not display any information when the  
4 **OMP\_DISPLAY\_ENV** environment variable is **FALSE**, undefined, or any other value than **TRUE** or  
5 **VERBOSE**.

6 The display begins with "OPENMP DISPLAY ENVIRONMENT BEGIN", followed by the  
7 **\_OPENMP** version macro (or the **openmp\_version** Fortran parameter) value and ICV values, in  
8 the format *NAME* '=' *VALUE*. *NAME* corresponds to the macro or environment variable name,  
9 optionally prepended by a bracketed *device-type*. *VALUE* corresponds to the value of the macro or  
10 ICV associated with this environment variable. Values should be enclosed in single quotes. The  
11 display is terminated with "OPENMP DISPLAY ENVIRONMENT END".

12 Example:

```
% setenv OMP_DISPLAY_ENV TRUE
```

13 The above example causes an OpenMP implementation to generate output of the following form:

```
OPENMP DISPLAY ENVIRONMENT BEGIN
_OPENMP='201307'
[host] OMP_SCHEDULE='GUIDED,4'
[host] OMP_NUM_THREADS='4,3,2'
[device] OMP_NUM_THREADS='2'
[host,device] OMP_DYNAMIC='TRUE'
[host] OMP_PLACES='0:4,4:4,8:4,12:4'
...
OPENMP DISPLAY ENVIRONMENT END
```

## 14 4.13 OMP\_DEFAULT\_DEVICE

15 The **OMP\_DEFAULT\_DEVICE** environment variable sets the device number to use in device  
16 constructs by setting the initial value of the *default-device-var* ICV.

17 The value of this environment variable must be a non-negative integer value.

### 18 Cross References

- 19 • *default-device-var* ICV, see Section 2.3 on page 35.
- 20 • device constructs, Section 2.10 on page 92.

## 2 Stubs for Runtime Library Routines

---

3 This section provides stubs for the runtime library routines defined in the OpenMP API. The stubs  
4 are provided to enable portability to platforms that do not support the OpenMP API. On these  
5 platforms, OpenMP programs must be linked with a library containing these stub routines. The stub  
6 routines assume that the directives in the OpenMP program are ignored. As such, they emulate  
7 serial semantics.

8 Note that the lock variable that appears in the lock routines must be accessed exclusively through  
9 these routines. It should not be initialized or otherwise modified in the user program.

10 In an actual implementation the lock variable might be used to hold the address of an allocated  
11 memory block, but here it is used to hold an integer value. Users should not make assumptions  
12 about mechanisms used by OpenMP implementations to implement locks based on the scheme  
13 used by the stub procedures.

▼ Fortran ▼

14 **Note** – In order to be able to compile the Fortran stubs file, the include file `omp_lib.h` was split  
15 into two files: `omp_lib_kinds.h` and `omp_lib.h` and the `omp_lib_kinds.h` file included  
16 where needed. There is no requirement for the implementation to provide separate files.

▲ Fortran ▲

## 1 A.1 C/C++ Stub Routines

```
2     #include <stdio.h>
3     #include <stdlib.h>
4     #include "omp.h"
5
6     void omp_set_num_threads(int num_threads)
7     {
8     }
9
10    int omp_get_num_threads(void)
11    {
12        return 1;
13    }
14
15    int omp_get_max_threads(void)
16    {
17        return 1;
18    }
19
20    int omp_get_thread_num(void)
21    {
22        return 0;
23    }
24
25    int omp_get_num_procs(void)
26    {
27        return 1;
28    }
29
30    int omp_in_parallel(void)
31    {
32        return 0;
33    }
34
35    void omp_set_dynamic(int dynamic_threads)
36    {
37    }
38
39    int omp_get_dynamic(void)
40    {
41        return 0;
42    }
43
44    int omp_get_cancellation(void)
45    {
46        return 0;
```



```

1      }
2
3      void omp_set_nested(int nested)
4      {
5      }
6
7      int omp_get_nested(void)
8      {
9          return 0;
10     }
11
12     void omp_set_schedule(omp_sched_t kind, int modifier)
13     {
14     }
15
16     void omp_get_schedule(omp_sched_t *kind, int *modifier)
17     {
18         *kind = omp_sched_static;
19         *modifier = 0;
20     }
21
22     int omp_get_thread_limit(void)
23     {
24         return 1;
25     }
26
27     void omp_set_max_active_levels(int max_active_levels)
28     {
29     }
30
31     int omp_get_max_active_levels(void)
32     {
33         return 0;
34     }
35
36     int omp_get_level(void)
37     {
38         return 0;
39     }
40
41     int omp_get_ancestor_thread_num(int level)
42     {
43         if (level == 0)
44         {
45             return 0;
46         }
47         else

```

```

1      {
2          return -1;
3      }
4  }
5
6  int omp_get_team_size(int level)
7  {
8      if (level == 0)
9      {
10         return 1;
11     }
12     else
13     {
14         return -1;
15     }
16 }
17
18 int omp_get_active_level(void)
19 {
20     return 0;
21 }
22
23 int omp_in_final(void)
24 {
25     return 1;
26 }
27
28 omp_proc_bind_t omp_get_proc_bind(void)
29 {
30     return omp_proc_bind_false;
31 }
32
33 void omp_set_default_device(int device_num)
34 {
35 }
36
37 int omp_get_default_device(void)
38 {
39     return 0;
40 }
41
42 int omp_get_num_devices(void)
43 {
44     return 0;
45 }
46
47 int omp_get_num_teams(void)

```

```

1      {
2          return 1;
3      }
4
5      int omp_get_team_num(void)
6      {
7          return 0;
8      }
9
10     int omp_is_initial_device(void)
11     {
12         return 1;
13     }
14
15     struct __omp_lock
16     {
17         int lock;
18     };
19
20     enum { UNLOCKED = -1, INIT, LOCKED };
21
22     void omp_init_lock(omp_lock_t *arg)
23     {
24         struct __omp_lock *lock = (struct __omp_lock *)arg;
25         lock->lock = UNLOCKED;
26     }
27
28     void omp_destroy_lock(omp_lock_t *arg)
29     {
30         struct __omp_lock *lock = (struct __omp_lock *)arg;
31         lock->lock = INIT;
32     }
33
34     void omp_set_lock(omp_lock_t *arg)
35     {
36         struct __omp_lock *lock = (struct __omp_lock *)arg;
37         if (lock->lock == UNLOCKED)
38         {
39             lock->lock = LOCKED;
40         }
41         else if (lock->lock == LOCKED)
42         {
43             fprintf(stderr, "error: deadlock in using lock variable\n");
44             exit(1);
45         }
46         else
47

```

```

1      {
2          fprintf(stderr, "error: lock not initialized\n");
3          exit(1);
4      }
5  }
6
7  void omp_unset_lock(omp_lock_t *arg)
8  {
9      struct __omp_lock *lock = (struct __omp_lock *)arg;
10     if (lock->lock == LOCKED)
11     {
12         lock->lock = UNLOCKED;
13     }
14     else if (lock->lock == UNLOCKED)
15     {
16         fprintf(stderr, "error: lock not set\n");
17         exit(1);
18     }
19     else
20     {
21         fprintf(stderr, "error: lock not initialized\n");
22         exit(1);
23     }
24 }
25
26 int omp_test_lock(omp_lock_t *arg)
27 {
28     struct __omp_lock *lock = (struct __omp_lock *)arg;
29     if (lock->lock == UNLOCKED)
30     {
31         lock->lock = LOCKED;
32         return 1;
33     }
34     else if (lock->lock == LOCKED)
35     {
36         return 0;
37     }
38     else
39     {
40         fprintf(stderr, "error: lock not initialized\ n");
41         exit(1);
42     }
43 }
44
45 struct __omp_nest_lock
46 {
47     short owner;

```

```

1      short count;
2  };
3
4  enum { NOOWNER = -1, MASTER = 0 };
5
6  void omp_init_nest_lock(omp_nest_lock_t *arg)
7  {
8      struct __omp_nest_lock *nlock=(struct __omp_nest_lock *)arg;
9      nlock->owner = NOOWNER;
10     nlock->count = 0;
11 }
12
13 void omp_destroy_nest_lock(omp_nest_lock_t *arg)
14 {
15     struct __omp_nest_lock *nlock=(struct __omp_nest_lock *)arg;
16     nlock->owner = NOOWNER;
17     nlock->count = UNLOCKED;
18 }
19
20 void omp_set_nest_lock(omp_nest_lock_t *arg)
21 {
22     struct __omp_nest_lock *nlock=(struct __omp_nest_lock *)arg;
23     if (nlock->owner == MASTER && nlock->count >= 1)
24     {
25         nlock->count++;
26     }
27     else if (nlock->owner == NOOWNER && nlock->count == 0)
28     {
29         nlock->owner = MASTER;
30         nlock->count = 1;
31     }
32     else
33     {
34         fprintf(stderr, "error: lock corrupted or not initialized\n");
35         exit(1);
36     }
37 }
38
39 void omp_unset_nest_lock(omp_nest_lock_t *arg)
40 {
41     struct __omp_nest_lock *nlock=(struct __omp_nest_lock *)arg;
42     if (nlock->owner == MASTER && nlock->count >= 1)
43     {
44         nlock->count--;
45         if (nlock->count == 0)
46         {
47             nlock->owner = NOOWNER;

```

```

1         }
2     }
3     else if (nlock->owner == NOOWNER && nlock->count == 0)
4     {
5         fprintf(stderr, "error: lock not set\n");
6         exit(1);
7     }
8     else
9     {
10        fprintf(stderr, "error: lock corrupted or not initialized\n");
11        exit(1);
12    }
13 }
14
15 int omp_test_nest_lock(omp_nest_lock_t *arg)
16 {
17     struct __omp_nest_lock *nlock=(struct __omp_nest_lock *)arg;
18     omp_set_nest_lock(arg);
19     return nlock->count;
20 }
21
22 double omp_get_wtime(void)
23 {
24     /* This function does not provide a working
25      * wallclock timer. Replace it with a version
26      * customized for the target machine.
27      */
28     return 0.0;
29 }
30
31 double omp_get_wtick(void)
32 {
33     /* This function does not provide a working
34      * clock tick function. Replace it with
35      * a version customized for the target machine.
36      */
37     return 365. * 86400.;
38 }
39

```

## 1 A.2 Fortran Stub Routines

```
2      subroutine omp_set_num_threads(num_threads)
3          integer num_threads
4          return
5      end subroutine
6
7      integer function omp_get_num_threads()
8          omp_get_num_threads = 1
9          return
10     end function
11
12     integer function omp_get_max_threads()
13         omp_get_max_threads = 1
14         return
15     end function
16
17     integer function omp_get_thread_num()
18         omp_get_thread_num = 0
19         return
20     end function
21
22     integer function omp_get_num_procs()
23         omp_get_num_procs = 1
24         return
25     end function
26
27     logical function omp_in_parallel()
28         omp_in_parallel = .false.
29         return
30     end function
31
32     subroutine omp_set_dynamic(dynamic_threads)
33         logical dynamic_threads
34         return
35     end subroutine
36
37     logical function omp_get_dynamic()
38         omp_get_dynamic = .false.
39         return
40     end function
41
42     logical function omp_get_cancellation()
43         omp_get_cancellation = .false.
44         return
45     end function
46
```

```

1      subroutine omp_set_nested(nested)
2          logical nested
3          return
4      end subroutine
5
6      logical function omp_get_nested()
7          omp_get_nested = .false.
8          return
9      end function
10
11     subroutine omp_set_schedule(kind, modifier)
12         include 'omp_lib_kinds.h'
13         integer (kind=omp_sched_kind) kind
14         integer modifier
15         return
16     end subroutine
17
18     subroutine omp_get_schedule(kind, modifier)
19         include 'omp_lib_kinds.h'
20         integer (kind=omp_sched_kind) kind
21         integer modifier
22         kind = omp_sched_static
23         modifier = 0
24         return
25     end subroutine
26
27     integer function omp_get_thread_limit()
28         omp_get_thread_limit = 1
29         return
30     end function
31
32     subroutine omp_set_max_active_levels( level )
33         integer level
34     end subroutine
35     integer function omp_get_max_active_levels()
36         omp_get_max_active_levels = 0
37         return
38     end function
39
40     integer function omp_get_level()
41         omp_get_level = 0
42         return
43     end function
44
45     integer function omp_get_ancestor_thread_num( level )
46         integer level
47         if ( level .eq. 0 ) then

```



```

1         omp_get_ancestor_thread_num = 0
2     else
3         omp_get_ancestor_thread_num = -1
4     end if
5     return
6 end function
7
8 integer function omp_get_team_size( level )
9     integer level
10    if ( level .eq. 0 ) then
11        omp_get_team_size = 1
12    else
13        omp_get_team_size = -1
14    end if
15    return
16 end function
17
18 integer function omp_get_active_level()
19     omp_get_active_level = 0
20     return
21 end function
22
23 logical function omp_in_final()
24     omp_in_final = .true.
25     return
26 end function
27
28 function omp_get_proc_bind()
29     include 'omp_lib_kinds.h'
30     integer (kind=omp_proc_bind_kind) omp_get_proc_bind
31     omp_get_proc_bind = omp_proc_bind_false
32 end function omp_get_proc_bind
33
34 subroutine omp_set_default_device(device_num)
35     integer device_num
36     return
37 end subroutine
38
39 integer function omp_get_default_device()
40     omp_get_default_device = 0
41     return
42 end function
43
44 integer function omp_get_num_devices()
45     omp_get_num_devices = 0
46     return
47 end function

```

```

1
2 integer function omp_get_num_teams()
3     omp_get_num_teams = 1
4     return
5 end function
6
7 integer function omp_get_team_num()
8     omp_get_team_num = 0
9     return
10 end function
11
12 logical function omp_is_initial_device()
13     omp_is_initial_device = .true.
14     return
15 end function
16
17 subroutine omp_init_lock(lock)
18     ! lock is 0 if the simple lock is not initialized
19     !         -1 if the simple lock is initialized but not set
20     !         1 if the simple lock is set
21     include 'omp_lib_kinds.h'
22     integer(kind=omp_lock_kind) lock
23
24     lock = -1
25     return
26 end subroutine
27
28 subroutine omp_destroy_lock(lock)
29     include 'omp_lib_kinds.h'
30     integer(kind=omp_lock_kind) lock
31
32     lock = 0
33     return
34 end subroutine
35
36 subroutine omp_set_lock(lock)
37     include 'omp_lib_kinds.h'
38     integer(kind=omp_lock_kind) lock
39
40     if (lock .eq. -1) then
41         lock = 1
42     elseif (lock .eq. 1) then
43         print *, 'error: deadlock in using lock variable'
44         stop
45     else
46         print *, 'error: lock not initialized'
47         stop

```

```

1      endif
2      return
3  end subroutine
4
5  subroutine omp_unset_lock(lock)
6      include 'omp_lib_kinds.h'
7      integer(kind=omp_lock_kind) lock
8
9      if (lock .eq. 1) then
10         lock = -1
11     elseif (lock .eq. -1) then
12         print *, 'error: lock not set'
13         stop
14     else
15         print *, 'error: lock not initialized'
16         stop
17     endif
18     return
19 end subroutine
20
21 logical function omp_test_lock(lock)
22     include 'omp_lib_kinds.h'
23     integer(kind=omp_lock_kind) lock
24
25     if (lock .eq. -1) then
26         lock = 1
27         omp_test_lock = .true.
28     elseif (lock .eq. 1) then
29         omp_test_lock = .false.
30     else
31         print *, 'error: lock not initialized'
32         stop
33     endif
34
35     return
36 end function
37
38 subroutine omp_init_nest_lock(nlock)
39     ! nlock is
40     ! 0 if the nestable lock is not initialized
41     ! -1 if the nestable lock is initialized but not set
42     ! 1 if the nestable lock is set
43     ! no use count is maintained
44     include 'omp_lib_kinds.h'
45     integer(kind=omp_nest_lock_kind) nlock
46
47     nlock = -1

```

```

1
2     return
3 end subroutine
4
5 subroutine omp_destroy_nest_lock(nlock)
6     include 'omp_lib_kinds.h'
7     integer(kind=omp_nest_lock_kind) nlock
8
9     nlock = 0
10
11    return
12 end subroutine
13
14 subroutine omp_set_nest_lock(nlock)
15     include 'omp_lib_kinds.h'
16     integer(kind=omp_nest_lock_kind) nlock
17
18     if (nlock .eq. -1) then
19         nlock = 1
20     elseif (nlock .eq. 0) then
21         print *, 'error: nested lock not initialized'
22         stop
23     else
24         print *, 'error: deadlock using nested lock variable'
25         stop
26     endif
27
28    return
29 end subroutine
30
31 subroutine omp_unset_nest_lock(nlock)
32     include 'omp_lib_kinds.h'
33     integer(kind=omp_nest_lock_kind) nlock
34
35     if (nlock .eq. 1) then
36         nlock = -1
37     elseif (nlock .eq. 0) then
38         print *, 'error: nested lock not initialized'
39         stop
40     else
41         print *, 'error: nested lock not set'
42         stop
43     endif
44
45    return
46 end subroutine
47

```

```

1      integer function omp_test_nest_lock(nlock)
2          include 'omp_lib_kinds.h'
3          integer(kind=omp_nest_lock_kind) nlock
4
5          if (nlock .eq. -1) then
6              nlock = 1
7              omp_test_nest_lock = 1
8          elseif (nlock .eq. 1) then
9              omp_test_nest_lock = 0
10         else
11             print *, 'error: nested lock not initialized'
12             stop
13         endif
14
15         return
16     end function
17
18     double precision function omp_get_wtime()
19         ! this function does not provide a working
20         ! wall clock timer. replace it with a version
21         ! customized for the target machine.
22
23         omp_get_wtime = 0.0d0
24
25         return
26     end function
27
28     double precision function omp_get_wtick()
29         ! this function does not provide a working
30         ! clock tick function. replace it with
31         ! a version customized for the target machine.
32         double precision one_year
33         parameter (one_year=365.d0*86400.d0)
34
35         omp_get_wtick = one_year
36
37         return
38     end function

```

## 1 APPENDIX B

# 2 OpenMP C and C++ Grammar

---

## 3 B.1 Notation

4 The grammar rules consist of the name for a non-terminal, followed by a colon, followed by  
5 replacement alternatives on separate lines.

6 The syntactic expression  $term_{opt}$  indicates that the term is optional within the replacement.

7 The syntactic expression  $term_{optseq}$  is equivalent to  $term-seq_{opt}$  with the following additional rules:

8  $term-seq :$

9  $term$

10  $term-seq term$

11  $term-seq , term$

## 1 B.2 Rules

2 The notation is described in Section 6.1 of the C standard. This grammar appendix shows the  
3 extensions to the base language grammar for the OpenMP C and C++ directives.

4		C++	
5	<i>statement-seq:</i>		
6	<i>statement</i>		
7	<i>openmp-directive</i>		
8	<i>statement-seq statement</i>		
9	<i>statement-seq openmp-directive</i>		
		C++	
		C90	
10	<i>statement-list:</i>		
11	<i>statement</i>		
12	<i>openmp-directive</i>		
13	<i>statement-list statement</i>		
14	<i>statement-list openmp-directive</i>		
		C90	
		C99	
15	<i>block-item:</i>		
16	<i>declaration</i>		
17	<i>statement</i>		
18	<i>openmp-directive</i>		
		C99	
19	<i>statement:</i>		
20	<i>/* standard statements */</i>		
21	<i>openmp-construct</i>		
22	<i>declaration-definition:</i>		
23	<i>/* Any C or C++ declaration or definition statement */</i>		
24	<i>function-statement:</i>		

1                    **/\* C or C++ function definition or declaration \*/**

2                    *declaration-definition-seq:*

3                    *declaration-definition*

4                    *declaration-definition-seq declaration-definition*

5                    *openmp-construct:*

6                    *parallel-construct*

7                    *for-construct*

8                    *sections-construct*

9                    *single-construct*

10                   *simd-construct*

11                   *for-simd-construct*

12                   *parallel-for-simd-construct*

13                   *target-data-construct*

14                   *target-construct*

15                   *target-update-construct*

16                   *teams-construct*

17                   *distribute-construct*

18                   *distribute-simd-construct*

19                   *distribute-parallel-for-construct*

20                   *distribute-parallel-for-simd-construct*

21                   *target-teams-construct*

22                   *teams-distribute-construct*

23                   *teams-distribute-simd-construct*

24                   *target-teams-distribute-construct*

25                   *target-teams-distribute-simd-construct*

26                   *teams-distribute-parallel-for-construct*

27                   *target-teams-distribute-parallel-for-construct*

28                   *teams-distribute-parallel-for-simd-construct*

29                   *target-teams-distribute-parallel-for-simd-construct*

30                   *parallel-for-construct*

31                   *parallel-sections-construct*



1           *task-construct*

2           *taskloop-construct*

3           *master-construct*

4           *critical-construct*

5           *atomic-construct*

6           *ordered-construct*

7       *openmp-directive:*

8           *barrier-directive*

9           *taskwait-directive*

10          *taskyield-directive*

11          *flush-directive*

12       *structured-block:*

13           *statement*

14       *parallel-construct:*

15           *parallel-directive structured-block*

16       *parallel-directive:*

17           **# pragma omp parallel** *parallel-clause<sub>optseq</sub> new-line*

18       *parallel-clause:*

19           *unique-parallel-clause*

20           *data-default-clause*

21           *data-privatization-clause*

22           *data-privatization-in-clause*

23           *data-sharing-clause*

24           *data-reduction-clause*

25           *collapse-clause*

26       *unique-parallel-clause:*

27           *if-clause*

28           **num\_threads** ( *expression* )

29           **copyin** ( *variable-list* )

30       *for-construct:*

31           *for-directive iteration-statement*

```

1      for-directive:
2          # pragma omp for for-clauseoptseq new-line
3      for-clause:
4          unique-for-clause
5          data-privatization-clause
6          data-privatization-in-clause
7          data-privatization-out-clause
8          data-reduction-clause
9          collapse-clause
10         nowait
11     unique-for-clause:
12         ordered
13         schedule ( schedule-kind )
14         schedule ( schedule-kind , expression )
15     collapse-clause:
16         collapse ( expression )
17     schedule-kind:
18         static
19         dynamic
20         guided
21         auto
22         runtime
23     sections-construct:
24         sections-directive section-scope
25     sections-directive:
26         # pragma omp sections sections-clauseoptseq new-line
27     sections-clause:
28         data-privatization-clause
29         data-privatization-in-clause
30         data-privatization-out-clause
31         data-reduction-clause

```

```

1      nowait
2  section-scope:
3      { section-sequence }
4  section-sequence:
5      section-directiveopt structured-block
6      section-sequence section-directive structured-block
7  section-directive:
8      # pragma omp section new-line
9  single-construct:
10     single-directive structured-block
11  single-directive:
12     # pragma omp single single-clauseoptseq new-line
13  single-clause:
14     unique-single-clause
15     data-privatization-clause
16     data-privatization-in-clause
17     nowait
18  unique-single-clause:
19     copyprivate ( variable-list )
20  simd-construct:
21     simd-directive iteration-statement
22  simd-directive:
23     # pragma omp simd simd-clauseoptseq new-line
24  simd-clause:
25     collapse-clause
26     aligned-clause
27     linear-clause
28     uniform-clause
29     data-reduction-clause
30     inbranch-clause
31  inbranch-clause:

```

```

1      inbranch
2      notinbranch
3  uniform-clause:
4      uniform ( variable-list )
5  linear-clause:
6      linear ( variable-list )
7      linear ( variable-list :expression )
8  aligned-clause:
9      aligned ( variable-list )
10     aligned ( variable-list :expression )
11  declare-simd-construct:
12     declare-simd-directive-seq function-statement
13  declare-simd-directive-seq:
14     declare-simd-directive
15     declare-simd-directive-seq declare-simd-directive
16  declare-simd-directive:
17     # pragma omp declare simd declare-simd-clauseoptseq new-line
18  declare-simd-clause:
19     simdlen ( expression )
20     aligned-clause
21     linear-clause
22     uniform-clause
23     data-reduction-clause
24     inbranch-clause
25  for-simd-construct:
26     for-simd-directive iteration-statement
27  for-simd-directive:
28     # pragma omp for simd for-simd-clauseoptseq new-line
29  for-simd-clause:
30     for-clause
31     simd-clause

```

```

1      parallel-for-simd-construct:
2          parallel-for-simd-directive iteration-statement
3      parallel-for-simd-directive:
4          # pragma omp parallel for simd parallel-for-simd-clauseoptseq new-line
5      parallel-for-simd-clause:
6          parallel-for-clause
7          simd-clause
8      target-data-construct:
9          target-data-directive structured-block
10     target-data-directive:
11         # pragma omp target data target-data-clauseoptseq new-line
12     target-data-clause:
13         device-clause
14         map-clause
15         if-clause
16     device-clause:
17         device ( expression )
18     map-clause:
19         map ( map-typeopt variable-array-section-list )
20     map-type:
21         alloc:
22         to:
23         from:
24         tofrom:
25     target-construct:
26         target-directive structured-block
27     target-directive:
28         # pragma omp target target-clauseoptseq new-line
29     target-clause:
30         device-clause
31         map-clause

```

```

1         if-clause
2     target-update-construct:
3         target-update-directive structured-block
4     target-update-directive:
5         # pragma omp target update target-update-clauseoptseq new-line
6     target-update-clause:
7         motion-clause
8         device-clause
9         if-clause
10    motion-clause:
11        to ( variable-array-section-list )
12        from ( variable-array-section-list )
13    declare-target-construct:
14        declare-target-directive declaration-definition-seq end-declare-target-directive
15    declare-target-directive:
16        # pragma omp declare target new-line
17    end-declare-target-directive:
18        # pragma omp end declare target new-line
19    teams-construct:
20        teams-directive structured-block
21    teams-directive:
22        # pragma omp teams teams-clauseoptseq new-line
23    teams-clause:
24        num_teams ( expression )
25        thread_limit ( expression )
26        data-default-clause
27        data-privatization-clause
28        data-privatization-in-clause
29        data-sharing-clause
30        data-reduction-clause
31    distribute-construct:

```

```

1      distribute-directive iteration-statement
2  distribute-directive:
3      # pragma omp distribute distribute-clauseoptseq new-line
4  distribute-clause:
5      data-privatization-clause
6      data-privatization-in-clause
7      collapse-clause
8      dist_schedule ( static )
9      dist_schedule ( static , expression )
10 distribute-simd-construct:
11     distribute-simd-directive iteration-statement
12 distribute-simd-directive:
13     #pragma omp distribute simd distribute-simd-clauseoptseq new-line
14 distribute-simd-clause:
15     distribute-clause
16     simd-clause
17 distribute-parallel-for-construct:
18     distribute-parallel-for-directive iteration-statement
19 distribute-parallel-for-directive:
20     #pragma omp distribute parallel for distribute-parallel-for-clauseoptseq new-line
21 distribute-parallel-for-clause:
22     distribute-clause
23     parallel-for-clause
24 distribute-parallel-for-simd-construct:
25     distribute-parallel-for-simd-directive iteration-statement
26 distribute-parallel-for-simd-directive:
27     #pragma omp distribute parallel for distribute-parallel-for-simd-clauseoptseq
28 new-line
29     distribute-parallel-for-simd-clause:
30     distribute-clause
31     parallel-for-simd-clause

```

1           *target-teams-construct:*

2                 *target-teams-directive iteration-statement*

3    *target-teams-directive:*

4                 **#pragma omp target teams** *target-teams-clause<sub>optseq</sub> new-line*

5    *target-teams-clause:*

6                 *target-clause*

7                 *teams-clause*

8    *teams-distribute-construct:*

9                 *teams-distribute-directive iteration-statement*

10   *teams-distribute-directive:*

11                **#pragma omp teams distribute** *teams-distribute-clause<sub>optseq</sub> new-line*

12   *teams-distribute-clause:*

13                *teams-clause*

14                *distribute-clause*

15   *teams-distribute-simd-construct:*

16                *teams-distribute-simd-directive iteration-statement*

17   *teams-distribute-simd-directive:*

18                **#pragma omp teams distribute simd** *teams-distribute-simd-clause<sub>optseq</sub> new-line*

19   *teams-distribute-simd-clause:*

20                *teams-clause*

21                *distribute-simd-clause*

22   *target-teams-distribute-construct:*

23                *target-teams-distribute-directive iteration-statement*

24   *target-teams-distribute-directive:*

25                **#pragma omp target teams distribute** *target-teams-distribute-clause<sub>optseq</sub> new-line*

26   *target-teams-distribute-clause:*

27                *target-clause*

28                *teams-distribute-clause*

29   *target-teams-distribute-simd-construct:*

30                *target-teams-distribute-simd-directive iteration-statement*

31   *target-teams-distribute-simd-directive:*



```

1      #pragma omp target teams distribute simd
2      target-teams-distribute-simd-clauseoptseq new-line
3
4      target-teams-distribute-simd-clause:
5
6      target-clause
7
8      teams-distribute-simd-clause
9
10     teams-distribute-parallel-for-construct:
11
12     teams-distribute-parallel-for-directive iteration-statement
13
14     teams-distribute-parallel-for-directive:
15
16     #pragma omp teams distribute parallel for
17     teams-distribute-parallel-for-clauseoptseq new-line
18
19     teams-distribute-parallel-for-clause:
20
21     teams-clause
22
23     distribute-parallel-for-clause
24
25     target-teams-distribute-parallel-for-construct:
26
27     target-teams-distribute-parallel-for-directive iteration-statement
28
29     target-teams-distribute-parallel-for-directive:
30
31     #pragma omp teams distribute parallel for
32     target-teams-distribute-parallel-for-clauseoptseq new-line
33
34     target-teams-distribute-parallel-for-clause:
35
36     target-clause
37
38     teams-distribute-parallel-for-clause
39
40     teams-distribute-parallel-for-simd-construct:
41
42     teams-distribute-parallel-for-simd-directive iteration-statement
43
44     teams-distribute-parallel-for-simd-directive:
45
46     #pragma omp teams distribute parallel for simd
47     teams-distribute-parallel-for-simd-clauseoptseq new-line
48
49     teams-distribute-parallel-for-simd-clause:
50
51     teams-clause
52
53     distribute-parallel-for-simd-clause
54
55     target-teams-distribute-parallel-for-simd-construct:
56
57     target-teams-distribute-parallel-for-simd-directive iteration-statement
58
59     target-teams-distribute-parallel-for-simd-directive:

```

```

1      #pragma omp target teams distribute parallel for simd
2      target-teams-distribute-parallel-for-simd-clauseoptseq new-line
3
4      target-teams-distribute-parallel-for-simd-clause:
5
6      target-clause
7
8      teams-distribute-parallel-for-simd-clause
9
10     task-construct:
11
12     task-directive structured-block
13
14     task-directive:
15
16     # pragma omp task task-clauseoptseq new-line
17
18     task-clause:
19
20     unique-task-clause
21
22     data-default-clause
23
24     data-privatization-clause
25
26     data-privatization-in-clause
27
28     data-sharing-clause
29
30     unique-task-clause:
31
32     if-clause
33
34     final-clause
35
36     untied-clause
37
38     mergeable-clause
39
40     depend ( dependence-type : variable-array-section-list )
41
42     final-clause:
43
44     final ( scalar-expression )
45
46     untied-clause:
47
48     untied
49
50     mergeable-clause:
51
52     mergeable
53
54     dependence-type:
55
56     in
57
58     out
59
60     inout

```

```

1      taskloop-construct:
2          taskloop-directive iteration-statement
3      taskloop-directive:
4          #pragma omp taskloop taskloop-clauseoptseq new-line
5      taskloop-clause:
6          unique-taskloop-clause
7          data-sharing-clause
8          data-privatization-clause
9          data-privatization-in-clause
10         data-privatization-out-clause
11         data-default-clause
12     unique-taskloop-clause:
13         grainsize ( expression )
14         num_tasks ( expression )
15         collapse-clause
16         if-clause
17         final-clause
18         untied-clause
19         mergeable-clause
20         nogroup
21     taskloop-simd-construct:
22         taskloop-simd-directive iteration-statement
23     taskloop-simd-directive:
24         #pragma omp taskloop simd taskloop-simd-clauseoptseq new-line
25     taskloop-simd-clause:
26         taskloop-clause
27         simd-clause
28     parallel-for-construct:
29         parallel-for-directive iteration-statement
30     parallel-for-directive:
31         # pragma omp parallel for parallel-for-clauseoptseq new-line

```

1           *parallel-for-clause:*

2                 *unique-parallel-clause*

3                 *unique-for-clause*

4                 *data-default-clause*

5                 *data-privatization-clause*

6                 *data-privatization-in-clause*

7                 *data-privatization-out-clause*

8                 *data-sharing-clause*

9                 *data-reduction-clause*

10           *parallel-sections-construct:*

11                 *parallel-sections-directive section-scope*

12           *parallel-sections-directive:*

13                 **# pragma omp parallel sections** *parallel-sections-clause<sub>optseq</sub> new-line*

14           *parallel-sections-clause:*

15                 *unique-parallel-clause*

16                 *data-default-clause*

17                 *data-privatization-clause*

18                 *data-privatization-in-clause*

19                 *data-privatization-out-clause*

20                 *data-sharing-clause*

21                 *data-reduction-clause*

22           *master-construct:*

23                 *master-directive structured-block*

24           *master-directive:*

25                 **# pragma omp master** *new-line*

26           *critical-construct:*

27                 *critical-directive structured-block*

28           *critical-directive:*

29                 **# pragma omp critical** *region-phrase<sub>opt</sub> new-line*

30           *region-phrase:*

31                 ( *identifier* )

```

1      barrier-directive:
2          # pragma omp barrier new-line
3      taskwait-directive:
4          # pragma omp taskwait new-line
5      taskgroup-construct:
6          taskgroup-directive structured-block
7      taskgroup-directive:
8          # pragma omp taskgroup new-line
9      taskyield-directive:
10         # pragma omp taskyield new-line
11     atomic-construct:
12         atomic-directive expression-statement
13         atomic-directive structured block
14     atomic-directive:
15         # pragma omp atomic atomic-clauseopt seq_cst-clauseopt new-line
16     atomic-clause:
17         read
18         write
19         update
20         capture
21     seq_cst-clause:
22         seq_cst
23     flush-directive:
24         # pragma omp flush flush-varsopt new-line
25     flush-vars:
26         ( variable-list )
27     ordered-construct:
28         ordered-directive structured-block
29     ordered-directive:
30         # pragma omp ordered new-line
31     cancel-directive:

```

```

1      # pragma omp cancel construct-type-clause if-clauseopt new-line
2  construct-type-clause:
3      parallel
4      sections
5      for
6      taskgroup
7  cancellation-point-directive:
8      # pragma omp cancellation point construct-type-clause new-line
9  declaration:
10     /* standard declarations */
11     threadprivate-directive
12     declare-simd-directive
13     declare-target-construct
14     declare-reduction-directive
15  threadprivate-directive:
16     # pragma omp threadprivate ( variable-list ) new-line
17  declare-reduction-directive:
18     # pragma omp declare reduction ( reduction-identifier : reduction-type-list :
19  expression ) initializer-clauseopt new-line
20  reduction-identifier:
21  ─────────────────────────── C ───────────────────────────
22  identifier
23  ─────────────────────────── C ───────────────────────────
24  ─────────────────────────── C++ ───────────────────────────
25  id-expression
26  ─────────────────────────── C++ ───────────────────────────
27  ─────────────────────────── C / C++ ───────────────────────────
28  one of: + * - & ^ | && || min max

```

C / C++

- 1 *reduction-type-list:*
- 2     *type-id*
- 3     *reduction-type-list, type-id*
- 4 *initializer-clause:*
- 5     **initializer** ( *identifier = initializer* )
- 6     **initializer** ( *identifier* ( *argument-expression-list* ) )
- 7     **initializer** ( *identifier* *initializer* )
- 8     **initializer** ( *id-expression* ( *expression-list* ) )
- 9 *data-default-clause:*
- 10     **default** ( *shared* )
- 11     **default** ( *none* )
- 12 *data-privatization-clause:*
- 13     **private** ( *variable-list* )
- 14 *data-privatization-in-clause:*
- 15     **firstprivate** ( *variable-list* )
- 16 *data-privatization-out-clause:*
- 17     **lastprivate** ( *variable-list* )
- 18 *data-sharing-clause:*
- 19     **shared** ( *variable-list* )
- 20 *data-reduction-clause:*
- 21     **reduction** ( *reduction-identifier* : *variable-list* )
- 22 *if-clause:*
- 23     **if** ( *scalar-expression* )

C

1 *array-section:*  
 2     *identifier array-section-subscript*  
 3 *variable-list:*  
 4     *identifier*  
 5     *variable-list , identifier*  
 6 *variable-array-section-list:*  
 7     *identifier*  
 8     *array-section*  
 9     *variable-array-section-list , identifier*  
 10    *variable-array-section-list , array-section*

C

C++

11 *array-section:*  
 12     *id-expression array-section-subscript*  
 13 *variable-list:*  
 14     *id-expression*  
 15     *variable-list , id-expression*  
 16 *variable-array-section-list:*  
 17     *id-expression*  
 18     *array-section*  
 19     *variable-array-section-list , id-expression*  
 20    *variable-array-section-list , array-section*

C++

21 *array-section-subscript:*  
 22     *array-section-subscript [ expression<sub>opt</sub> : expression<sub>opt</sub> ]*  
 23     *array-section-subscript [ expression ]*  
 24     *[ expression<sub>opt</sub> : expression<sub>opt</sub> ]*  
 25     *[ expression ]*





## 1 APPENDIX C

2

# Interface Declarations

---

3

This appendix gives examples of the C/C++ header file, the Fortran **include** file and Fortran **module** that shall be provided by implementations as specified in Chapter 3. It also includes an example of a Fortran 90 generic interface for a library routine. This is a non-normative section, implementation files may differ.

4

5

6

## 1 C.1 Example of the omp.h Header File

```
2     #ifndef _OMP_H_DEF
3     #define _OMP_H_DEF
4
5     /*
6     * define the lock data types
7     */
8     typedef void *omp_lock_t;
9
10    typedef void *omp_nest_lock_t;
11
12    /*
13    * define the schedule kinds
14    */
15    typedef enum omp_sched_t
16    {
17        omp_sched_static = 1,
18        omp_sched_dynamic = 2,
19        omp_sched_guided = 3,
20        omp_sched_auto = 4
21    /* , Add vendor specific schedule constants here */
22    } omp_sched_t;
23
24    /*
25    * define the proc bind values
26    */
27    typedef enum omp_proc_bind_t
28    {
29        omp_proc_bind_false = 0,
30        omp_proc_bind_true = 1,
31        omp_proc_bind_master = 2,
32        omp_proc_bind_close = 3,
33        omp_proc_bind_spread = 4
34    } omp_proc_bind_t;
35
36    /*
37    * exported OpenMP functions
38    */
39    #ifdef __cplusplus
40    extern "C"
41    {
42    #endif
43
44    extern void omp_set_num_threads(int num_threads);
45    extern int omp_get_num_threads(void);
46    extern int omp_get_max_threads(void);
```

```

1     extern int omp_get_thread_num(void);
2     extern int omp_get_num_procs(void);
3     extern int omp_in_parallel(void);
4     extern void omp_set_dynamic(int dynamic_threads);
5     extern int omp_get_dynamic(void);
6     extern void omp_set_nested(int nested);
7     extern int omp_get_cancellation(void);
8     extern int omp_get_nested(void);
9     extern void omp_set_schedule(omp_sched_t kind, int modifier);
10    extern void omp_get_schedule(omp_sched_t *kind, int *modifier);
11    extern int omp_get_thread_limit(void);
12    extern void omp_set_max_active_levels(int max_active_levels);
13    extern int omp_get_max_active_levels(void);
14    extern int omp_get_level(void);
15    extern int omp_get_ancestor_thread_num(int level);
16    extern int omp_get_team_size(int level);
17    extern int omp_get_active_level(void);
18    extern int omp_in_final(void);
19    extern omp_proc_bind_t omp_get_proc_bind(void);
20    extern void omp_set_default_device(int device_num);
21    extern int omp_get_default_device(void);
22    extern int omp_get_num_devices(void);
23    extern int omp_get_num_teams(void);
24    extern int omp_get_team_num(void);
25    extern int omp_is_initial_device(void);
26
27    extern void omp_init_lock(omp_lock_t *lock);
28    extern void omp_destroy_lock(omp_lock_t *lock);
29    extern void omp_set_lock(omp_lock_t *lock);
30    extern void omp_unset_lock(omp_lock_t *lock);
31    extern int omp_test_lock(omp_lock_t *lock);
32
33    extern void omp_init_nest_lock(omp_nest_lock_t *lock);
34    extern void omp_destroy_nest_lock(omp_nest_lock_t *lock);
35    extern void omp_set_nest_lock(omp_nest_lock_t *lock);
36    extern void omp_unset_nest_lock(omp_nest_lock_t *lock);
37    extern int omp_test_nest_lock(omp_nest_lock_t *lock);
38
39    extern double omp_get_wtime(void);
40    extern double omp_get_wtick(void);
41
42    #ifdef __cplusplus
43    }
44    #endif
45
46    #endif

```

## 1 C.2 Example of an Interface Declaration include 2 File

```
3      omp_lib_kinds.h:
4      integer omp_lock_kind
5          integer omp_nest_lock_kind
6      ! this selects an integer that is large enough to hold a 64 bit integer
7          parameter ( omp_lock_kind = selected_int_kind( 10 ) )
8          parameter ( omp_nest_lock_kind = selected_int_kind( 10 ) )
9      integer omp_sched_kind
10     ! this selects an integer that is large enough to hold a 32 bit integer
11     parameter ( omp_sched_kind = selected_int_kind( 8 ) )
12     integer ( omp_sched_kind ) omp_sched_static
13     parameter ( omp_sched_static = 1 )
14     integer ( omp_sched_kind ) omp_sched_dynamic
15     parameter ( omp_sched_dynamic = 2 )
16     integer ( omp_sched_kind ) omp_sched_guided
17     parameter ( omp_sched_guided = 3 )
18     integer ( omp_sched_kind ) omp_sched_auto
19     parameter ( omp_sched_auto = 4 )
20     integer omp_proc_bind_kind
21     parameter ( omp_proc_bind_kind = selected_int_kind( 8 ) )
22     integer ( omp_proc_bind_kind ) omp_proc_bind_false
23     parameter ( omp_proc_bind_false = 0 )
24     integer ( omp_proc_bind_kind ) omp_proc_bind_true
25     parameter ( omp_proc_bind_true = 1 )
26     integer ( omp_proc_bind_kind ) omp_proc_bind_master
27     parameter ( omp_proc_bind_master = 2 )
28     integer ( omp_proc_bind_kind ) omp_proc_bind_close
29     parameter ( omp_proc_bind_close = 3 )
30     integer ( omp_proc_bind_kind ) omp_proc_bind_spread
31     parameter ( omp_proc_bind_spread = 4 )
32
33     omp_lib.h:
34     ! default integer type assumed below
35     ! default logical type assumed below
36     ! OpenMP API v4.0
37
38     include 'omp_lib_kinds.h'
39     integer openmp_version
40     parameter ( openmp_version = 201307 )
41
42     external omp_set_num_threads
43     external omp_get_num_threads
44     integer omp_get_num_threads
```

```

1      external omp_get_max_threads
2      integer omp_get_max_threads
3      external omp_get_thread_num
4      integer omp_get_thread_num
5      external omp_get_num_procs
6      integer omp_get_num_procs
7      external omp_in_parallel
8      logical omp_in_parallel
9      external omp_set_dynamic
10     external omp_get_dynamic
11     logical omp_get_dynamic
12     external omp_get_cancellation
13     integer omp_get_cancellation
14     external omp_set_nested
15     external omp_get_nested
16     logical omp_get_nested
17     external omp_set_schedule
18     external omp_get_schedule
19     external omp_get_thread_limit
20     integer omp_get_thread_limit
21     external omp_set_max_active_levels
22     external omp_get_max_active_levels
23     integer omp_get_max_active_levels
24     external omp_get_level
25     integer omp_get_level
26     external omp_get_ancestor_thread_num
27     integer omp_get_ancestor_thread_num
28     external omp_get_team_size
29     integer omp_get_team_size
30     external omp_get_active_level
31     integer omp_get_active_level
32     external omp_set_default_device
33     external omp_get_default_device
34     integer omp_get_default_device
35     external omp_get_num_devices
36     integer omp_get_num_devices
37     external omp_get_num_teams
38     integer omp_get_num_teams
39     external omp_get_team_num
40     integer omp_get_team_num
41     external omp_is_initial_device
42     logical omp_is_initial_device
43
44     external omp_in_final
45     logical omp_in_final
46
47     integer ( omp_proc_bind_kind ) omp_get_proc_bind

```

```
1      external omp_get_proc_bind
2
3      external omp_init_lock
4      external omp_destroy_lock
5      external omp_set_lock
6      external omp_unset_lock
7      external omp_test_lock
8      logical omp_test_lock
9
10     external omp_init_nest_lock
11     external omp_destroy_nest_lock
12     external omp_set_nest_lock
13     external omp_unset_nest_lock
14     external omp_test_nest_lock
15     integer omp_test_nest_lock
16
17     external omp_get_wtick
18     double precision omp_get_wtick
19     external omp_get_wtime
20     double precision omp_get_wtime
```

## 1 C.3 Example of a Fortran Interface Declaration 2 module

```
3         !      the "!" of this comment starts in column 1
4         !23456
5
6         module omp_lib_kinds
7             integer, parameter :: omp_lock_kind = selected_int_kind( 10 )
8             integer, parameter :: omp_nest_lock_kind = selected_int_kind( 10 )
9             integer, parameter :: omp_sched_kind = selected_int_kind( 8 )
10            integer(kind=omp_sched_kind), parameter ::
11            &    omp_sched_static = 1
12            integer(kind=omp_sched_kind), parameter ::
13            &    omp_sched_dynamic = 2
14            integer(kind=omp_sched_kind), parameter ::
15            &    omp_sched_guided = 3
16            integer(kind=omp_sched_kind), parameter ::
17            &    omp_sched_auto = 4
18            integer, parameter :: omp_proc_bind_kind = selected_int_kind( 8 )
19            integer (kind=omp_proc_bind_kind), parameter ::
20            &    omp_proc_bind_false = 0
21            integer (kind=omp_proc_bind_kind), parameter ::
22            &    omp_proc_bind_true = 1
23            integer (kind=omp_proc_bind_kind), parameter ::
24            &    omp_proc_bind_master = 2
25            integer (kind=omp_proc_bind_kind), parameter ::
26            &    omp_proc_bind_close = 3
27            integer (kind=omp_proc_bind_kind), parameter ::
28            &    omp_proc_bind_spread = 4
29            end module omp_lib_kinds
30
31            module omp_lib
32
33                use omp_lib_kinds
34
35                !
36                OpenMP API v4.0
37                integer, parameter :: openmp_version = 201307
38
39                interface
40
41                    subroutine omp_set_num_threads (number_of_threads_expr)
42                        integer, intent(in) :: number_of_threads_expr
43                    end subroutine omp_set_num_threads
44
45                    function omp_get_num_threads ()
46                        integer :: omp_get_num_threads
```



```

1      end function omp_get_num_threads
2
3      function omp_get_max_threads ()
4          integer :: omp_get_max_threads
5      end function omp_get_max_threads
6
7      function omp_get_thread_num ()
8          integer :: omp_get_thread_num
9      end function omp_get_thread_num
10
11     function omp_get_num_procs ()
12         integer :: omp_get_num_procs
13     end function omp_get_num_procs
14
15     function omp_in_parallel ()
16         logical :: omp_in_parallel
17     end function omp_in_parallel
18
19     subroutine omp_set_dynamic (enable_expr)
20         logical, intent(in) :: enable_expr
21     end subroutine omp_set_dynamic
22
23     function omp_get_dynamic ()
24         logical :: omp_get_dynamic
25     end function omp_get_dynamic
26
27     function omp_get_cancellation ()
28         integer :: omp_get_cancellation
29     end function omp_get_cancellation
30
31     subroutine omp_set_nested (enable_expr)
32         logical, intent(in) :: enable_expr
33     end subroutine omp_set_nested
34
35     function omp_get_nested ()
36         logical :: omp_get_nested
37     end function omp_get_nested
38
39     subroutine omp_set_schedule (kind, modifier)
40         use omp_lib_kinds
41         integer(kind=omp_sched_kind), intent(in) :: kind
42         integer, intent(in) :: modifier
43     end subroutine omp_set_schedule
44
45     subroutine omp_get_schedule (kind, modifier)
46         use omp_lib_kinds
47         integer(kind=omp_sched_kind), intent(out) :: kind

```

```

1         integer, intent(out)::modifier
2     end subroutine omp_get_schedule
3
4     function omp_get_thread_limit()
5         integer :: omp_get_thread_limit
6     end function omp_get_thread_limit
7
8     subroutine omp_set_max_active_levels(var)
9         integer, intent(in) :: var
10    end subroutine omp_set_max_active_levels
11
12    function omp_get_max_active_levels()
13        integer :: omp_get_max_active_levels
14    end function omp_get_max_active_levels
15
16    function omp_get_level()
17        integer :: omp_get_level
18    end function omp_get_level
19
20    function omp_get_ancestor_thread_num(level)
21        integer, intent(in) :: level
22        integer :: omp_get_ancestor_thread_num
23    end function omp_get_ancestor_thread_num
24
25    function omp_get_team_size(level)
26        integer, intent(in) :: level
27        integer :: omp_get_team_size
28    end function omp_get_team_size
29
30    function omp_get_active_level()
31        integer :: omp_get_active_level
32    end function omp_get_active_level
33
34    function omp_in_final()
35        logical omp_in_final
36    end function omp_in_final
37
38    function omp_get_proc_bind( )
39        include 'omp_lib_kinds.h'
40        integer (kind=omp_proc_bind_kind) omp_get_proc_bind
41        omp_get_proc_bind = omp_proc_bind_false
42    end function omp_get_proc_bind
43
44    subroutine omp_set_default_device (device_num)
45        integer :: device_num
46    end subroutine omp_set_default_device
47

```

```

1      function omp_get_default_device ()
2          integer :: omp_get_default_device
3      end function omp_get_default_device
4
5      function omp_get_num_devices ()
6          integer :: omp_get_num_devices
7      end function omp_get_num_devices
8
9      function omp_get_num_teams ()
10         integer :: omp_get_num_teams
11     end function omp_get_num_teams
12
13     function omp_get_team_num ()
14         integer :: omp_get_team_num
15     end function omp_get_team_num
16
17     function omp_is_initial_device ()
18         logical :: omp_is_initial_device
19     end function omp_is_initial_device
20
21     subroutine omp_init_lock (var)
22         use omp_lib_kinds
23         integer (kind=omp_lock_kind), intent(out) :: var
24     end subroutine omp_init_lock
25
26     subroutine omp_destroy_lock (var)
27         use omp_lib_kinds
28         integer (kind=omp_lock_kind), intent(inout) :: var
29     end subroutine omp_destroy_lock
30
31     subroutine omp_set_lock (var)
32         use omp_lib_kinds
33         integer (kind=omp_lock_kind), intent(inout) :: var
34     end subroutine omp_set_lock
35
36     subroutine omp_unset_lock (var)
37         use omp_lib_kinds
38         integer (kind=omp_lock_kind), intent(inout) :: var
39     end subroutine omp_unset_lock
40
41     function omp_test_lock (var)
42         use omp_lib_kinds
43         logical :: omp_test_lock
44         integer (kind=omp_lock_kind), intent(inout) :: var
45     end function omp_test_lock
46
47     subroutine omp_init_nest_lock (var)

```

```

1         use omp_lib_kinds
2         integer (kind=omp_nest_lock_kind), intent(out) :: var
3     end subroutine omp_init_nest_lock
4
5     subroutine omp_destroy_nest_lock (var)
6         use omp_lib_kinds
7         integer (kind=omp_nest_lock_kind), intent(inout) :: var
8     end subroutine omp_destroy_nest_lock
9
10    subroutine omp_set_nest_lock (var)
11        use omp_lib_kinds
12        integer (kind=omp_nest_lock_kind), intent(inout) :: var
13    end subroutine omp_set_nest_lock
14
15    subroutine omp_unset_nest_lock (var)
16        use omp_lib_kinds
17        integer (kind=omp_nest_lock_kind), intent(inout) :: var
18    end subroutine omp_unset_nest_lock
19
20    function omp_test_nest_lock (var)
21        use omp_lib_kinds
22        integer :: omp_test_nest_lock
23        integer (kind=omp_nest_lock_kind), intent(inout) :: var
24    end function omp_test_nest_lock
25
26    function omp_get_wtick ()
27        double precision :: omp_get_wtick
28    end function omp_get_wtick
29
30    function omp_get_wtime ()
31        double precision :: omp_get_wtime
32    end function omp_get_wtime
33
34    end interface
35
36 end module omp_lib

```

## 1 C.4 Example of a Generic Interface for a Library 2 Routine

3 Any of the OpenMP runtime library routines that take an argument may be extended with a generic  
4 interface so arguments of different **KIND** type can be accommodated.

5 The **OMP\_SET\_NUM\_THREADS** interface could be specified in the **omp\_lib** module as follows:

```
interface omp_set_num_threads

    subroutine omp_set_num_threads_4(number_of_threads_expr)
        use omp_lib_kinds
        integer(4), intent(in) :: number_of_threads_expr
    end subroutine omp_set_num_threads_4

    subroutine omp_set_num_threads_8(number_of_threads_expr)
        use omp_lib_kinds
        integer(8), intent(in) :: number_of_threads_expr
    end subroutine omp_set_num_threads_8

end interface omp_set_num_threads
```

## OpenMP Implementation-Defined Behaviors

---

4 This appendix summarizes the behaviors that are described as implementation defined in this API.  
5 Each behavior is cross-referenced back to its description in the main specification. An  
6 implementation is required to define and document its behavior in these cases.

- 7 • **Processor:** a hardware unit that is implementation defined (see Section 1.2.1 on page 2).
- 8 • **Device:** an implementation defined logical execution engine (see Section 1.2.1 on page 2).
- 9 • **Memory model:** the minimum size at which a memory update may also read and write back  
10 adjacent variables that are part of another variable (as array or structure elements) is  
11 implementation defined but is no larger than required by the base language (see Section 1.4.1 on  
12 page 17).
- 13 • **Memory model:** Implementations are allowed to relax the ordering imposed by implicit flush  
14 operations when the result is only visible to programs using non-sequentially consistent atomic  
15 directives (see Section 1.4.4 on page 20).
- 16 • **Internal control variables:** the initial values of *dyn-var*, *nthreads-var*, *run-sched-var*,  
17 *def-sched-var*, *bind-var*, *stacksize-var*, *wait-policy-var*, *thread-limit-var*, *max-active-levels-var*,  
18 *place-partition-var*, and *default-device-var* are implementation defined (see Section 2.3.2 on  
19 page 36).
- 20 • **Dynamic adjustment of threads:** providing the ability to dynamically adjust the number of  
21 threads is implementation defined. Implementations are allowed to deliver fewer threads (but at  
22 least one) than indicated in Algorithm 2-1 even if dynamic adjustment is disabled (see  
23 Section 2.5.1 on page 48).
- 24 • **Thread affinity:** With  $T \leq P$ , when  $T$  does not divide  $P$  evenly, the assignment of the  
25 remaining  $P - T * S$  places into subpartitions is implementation defined. With  $T > P$ , when  $P$   
26 does not divide  $T$  evenly, the assignment of the remaining  $T - P * S$  threads into places is  
27 implementation defined. The determination of whether the affinity request can be fulfilled is

- 1 implementation defined. If not, the number of threads in the team and their mapping to places  
2 become implementation defined (see Section 2.5.2 on page 50).
- 3 • **Loop directive:** the integer type (or kind, for Fortran) used to compute the iteration count of a  
4 collapsed loop is implementation defined. The effect of the **schedule(runtime)** clause  
5 when the *run-sched-var* ICV is set to **auto** is implementation defined. See Section 2.7.1 on  
6 page 55.
  - 7 • **sections construct:** the method of scheduling the structured blocks among threads in the  
8 team is implementation defined (see Section 2.7.2 on page 63).
  - 9 • **single construct:** the method of choosing a thread to execute the structured block is  
10 implementation defined (see Section 2.7.3 on page 65)
  - 11 • **simd construct:** the integer type (or kind, for Fortran) used to compute the iteration count for the  
12 collapsed loop is implementation defined. The number of iterations that are executed  
13 concurrently at any given time is implementation defined. If the **aligned** clause is not  
14 specified, the assumed alignment is implementation defined (see Section 2.8.1 on page 70).
  - 15 • **declare simd construct:** if the **simdlen** clause is not specified, the number of concurrent  
16 arguments for the function is implementation defined. If the **aligned** clause is not specified,  
17 the assumed alignment is implementation defined (see Section 2.8.2 on page 74).
  - 18 • The number of loop iterations assigned to a task created from a **taskloop** construct is  
19 implementation defined, unless the **grainsize** or **num\_tasks** clauses are specified (see  
20 Section 2.9.2 on page 83).
  - 21 • **teams construct:** the number of teams that are created is implementation defined but less than or  
22 equal to the value of the **num\_teams** clause if specified. The maximum number of threads  
23 participating in the contention group that each team initiates is implementation defined but less  
24 than or equal to the value of the **thread\_limit** clause if specified (see Section 2.10.5 on  
25 page 102).
  - 26 • If no **dist\_schedule** clause is specified then the schedule for the **distribute** construct is  
27 implementation defined (see Section 2.10.6 on page 104).
  - 28 • **atomic construct:** a compliant implementation may enforce exclusive access between  
29 **atomic** regions that update different storage locations. The circumstances under which this  
30 occurs are implementation defined. If the storage location designated by *x* is not size-aligned  
31 (that is, if the byte alignment of *x* is not a multiple of the size of *x*), then the behavior of the  
32 atomic region is implementation defined (see Section 2.12.6 on page 141).
  - 33 • **omp\_set\_num\_threads routine:** if the argument is not a positive integer the behavior is  
34 implementation defined (see Section 3.2.1 on page 208).
  - 35 • **omp\_set\_schedule routine:** for implementation specific schedule types, the values and  
36 associated meanings of the second argument are implementation defined. (see Section 3.2.12 on  
37 page 220).

- 1       • **omp\_set\_max\_active\_levels routine:** when called from within any explicit parallel  
2       region the binding thread set (and binding region, if required) for the  
3       **omp\_set\_max\_active\_levels** region is implementation defined and the behavior is  
4       implementation defined. If the argument is not a non-negative integer then the behavior is  
5       implementation defined (see Section 3.2.15 on page 223).
- 6       • **omp\_get\_max\_active\_levels routine:** when called from within any explicit parallel  
7       region the binding thread set (and binding region, if required) for the  
8       **omp\_get\_max\_active\_levels** region is implementation defined (see Section 3.2.16 on  
9       page 225).
- 10      • **OMP\_SCHEDULE environment variable:** if the value of the variable does not conform to the  
11      specified format then the result is implementation defined (see Section 4.1 on page 251).
- 12      • **OMP\_NUM\_THREADS environment variable:** if any value of the list specified in the  
13      **OMP\_NUM\_THREADS** environment variable leads to a number of threads that is greater than the  
14      implementation can support, or if any value is not a positive integer, then the result is  
15      implementation defined (see Section 4.2 on page 252).
- 16      • **OMP\_PROC\_BIND environment variable:** if the value is not **true**, **false**, or a comma  
17      separated list of **master**, **close**, or **spread**, the behavior is implementation defined. The  
18      behavior is also implementation defined if an initial thread cannot be bound to the first place in  
19      the OpenMP place list (see Section 4.4 on page 253).
- 20      • **OMP\_DYNAMIC environment variable:** if the value is neither **true** nor **false** the behavior is  
21      implementation defined (see Section 4.3 on page 253).
- 22      • **OMP\_NESTED environment variable:** if the value is neither **true** nor **false** the behavior is  
23      implementation defined (see Section 4.6 on page 256).
- 24      • **OMP\_STACKSIZE environment variable:** if the value does not conform to the specified format  
25      or the implementation cannot provide a stack of the specified size then the behavior is  
26      implementation defined (see Section 4.7 on page 256).
- 27      • **OMP\_WAIT\_POLICY environment variable:** the details of the **ACTIVE** and **PASSIVE**  
28      behaviors are implementation defined (see Section 4.8 on page 257).
- 29      • **OMP\_MAX\_ACTIVE\_LEVELS environment variable:** if the value is not a non-negative integer  
30      or is greater than the number of parallel levels an implementation can support then the behavior  
31      is implementation defined (see Section 4.9 on page 258).
- 32      • **OMP\_THREAD\_LIMIT environment variable:** if the requested value is greater than the number  
33      of threads an implementation can support, or if the value is not a positive integer, the behavior of  
34      the program is implementation defined (see Section 4.10 on page 258).
- 35      • **OMP\_PLACES environment variable:** the meaning of the numbers specified in the environment  
36      variable and how the numbering is done are implementation defined. The precise definitions of  
37      the abstract names are implementation defined. An implementation may add  
38      implementation-defined abstract names as appropriate for the target platform. When creating a



1 place list of  $n$  elements by appending the number  $n$  to an abstract name, the determination of  
2 which resources to include in the place list is implementation defined. When requesting more  
3 resources than available, the length of the place list is also implementation defined. The behavior  
4 of the program is implementation defined when the execution environment cannot map a  
5 numerical value (either explicitly defined or implicitly derived from an interval) within the  
6 **OMP\_PLACES** list to a processor on the target platform, or if it maps to an unavailable processor.  
7 The behavior is also implementation defined when the **OMP\_PLACES** environment variable is  
8 defined using an abstract name (see Section 4.5 on page 254).

- 9 • **Thread affinity policy:** if the affinity request for a **parallel** construct cannot be fulfilled, the  
10 behavior of the thread affinity policy is implementation defined for that **parallel** construct.

---

### Fortran

- 11 • **threadprivate directive:** if the conditions for values of data in the threadprivate objects of  
12 threads (other than an initial thread) to persist between two consecutive active parallel regions do  
13 not all hold, the allocation status of an allocatable variable in the second region is  
14 implementation defined (see Section 2.14.2 on page 166).
- 15 • **shared clause:** passing a shared variable to a non-intrinsic procedure may result in the value of  
16 the shared variable being copied into temporary storage before the procedure reference, and back  
17 out of the temporary storage into the actual argument storage after the procedure reference.  
18 Situations where this occurs other than those specified are implementation defined (see  
19 Section 2.14.3.2 on page 174).
- 20 • **Runtime library definitions:** it is implementation defined whether the include file **omp\_lib.h**  
21 or the module **omp\_lib** (or both) is provided. It is implementation defined whether any of the  
22 OpenMP runtime library routines that take an argument are extended with a generic interface so  
23 arguments of different **KIND** type can be accommodated (see Section 3.1 on page 207).

---

### Fortran

## 1 APPENDIX E

# 2 Features History

---

3 This appendix summarizes the major changes between recent versions of the OpenMP API since  
4 version 2.5.

## 5 E.1 Version 4.0 to 4.1 Differences

- 6 • Taskloop constructs (see Section 2.9.2 on page 83 and Section 2.9.3 on page 88) were added to  
7 support nestable parallel loops that create OpenMP tasks.

## 8 E.2 Version 3.1 to 4.0 Differences

- 9 • Various changes throughout the specification were made to provide initial support of Fortran  
10 2003 (see Section 1.6 on page 21).
- 11 • C/C++ array syntax was extended to support array sections (see Section 2.4 on page 43).
- 12 • The `proc_bind` clause (see Section 2.5.2 on page 50), the `OMP_PLACES` environment  
13 variable (see Section 4.5 on page 254), and the `omp_get_proc_bind` runtime routine (see  
14 Section 3.2.22 on page 231) were added to support thread affinity policies.
- 15 • SIMD constructs were added to support SIMD parallelism (see Section 2.8 on page 70).

- 1       • Device constructs (see Section 2.10 on page 92), the `OMP_DEFAULT_DEVICE` environment  
2       variable (see Section 4.13 on page 260), the `omp_set_default_device`,  
3       `omp_get_default_device`, `omp_get_num_devices`, `omp_get_num_teams`,  
4       `omp_get_team_num`, and `omp_is_initial_device` routines were added to support  
5       execution on devices.
- 6       • Implementation defined task scheduling points for untied tasks were removed (see Section 2.9.5  
7       on page 90).
- 8       • The `depend` clause (see Section 2.12.9 on page 154) was added to support task dependences.
- 9       • The `taskgroup` construct (see Section 2.12.5 on page 140) was added to support more flexible  
10      deep task synchronization.
- 11      • The `reduction` clause (see Section 2.14.3.6 on page 184) was extended and the  
12      `declare reduction` construct (see Section 2.15 on page 199) was added to support user  
13      defined reductions.
- 14      • The `atomic` construct (see Section 2.12.6 on page 141) was extended to support atomic swap  
15      with the `capture` clause, to allow new atomic update and capture forms, and to support  
16      sequentially consistent atomic operations with a new `seq_cst` clause.
- 17      • The `cancel` construct (see Section 2.13.1 on page 156), the `cancellation point`  
18      construct (see Section 2.13.2 on page 160), the `omp_get_cancellation` runtime routine  
19      (see Section 3.2.9 on page 217) and the `OMP_CANCELLATION` environment variable (see  
20      Section 4.11 on page 259) were added to support the concept of cancellation.
- 21      • The `OMP_DISPLAY_ENV` environment variable (see Section 4.12 on page 259) was added to  
22      display the value of ICVs associated with the OpenMP environment variables.
- 23      • Examples (previously Appendix A) were moved to a separate document.

## 24 E.3 Version 3.0 to 3.1 Differences

- 25      • The `final` and `mergeable` clauses (see Section 2.9.1 on page 80) were added to the `task`  
26      construct to support optimization of task data environments.
- 27      • The `taskyield` construct (see Section 2.9.4 on page 89) was added to allow user-defined task  
28      scheduling points.
- 29      • The `atomic` construct (see Section 2.12.6 on page 141) was extended to include `read`, `write`,  
30      and `capture` forms, and an `update` clause was added to apply the already existing form of the  
31      `atomic` construct.

- Data environment restrictions were changed to allow **intent (in)** and **const**-qualified types for the **firstprivate** clause (see Section 2.14.3.4 on page 179).
- Data environment restrictions were changed to allow Fortran pointers in **firstprivate** (see Section 2.14.3.4 on page 179) and **lastprivate** (see Section 2.14.3.5 on page 182).
- New reduction operators **min** and **max** were added for C and C++
- The nesting restrictions in Section 2.16 on page 205 were clarified to disallow closely-nested OpenMP regions within an **atomic** region. This allows an **atomic** region to be consistently defined with other OpenMP regions so that they include all the code in the atomic construct.
- The **omp\_in\_final** runtime library routine (see Section 3.2.21 on page 230) was added to support specialization of final task regions.
- The *nthreads*-var ICV has been modified to be a list of the number of threads to use at each nested parallel region level. The value of this ICV is still set with the **OMP\_NUM\_THREADS** environment variable (see Section 4.2 on page 252), but the algorithm for determining the number of threads used in a parallel region has been modified to handle a list (see Section 2.5.1 on page 48).
- The *bind*-var ICV has been added, which controls whether or not threads are bound to processors (see Section 2.3.1 on page 35). The value of this ICV can be set with the **OMP\_PROC\_BIND** environment variable (see Section 4.4 on page 253).
- Descriptions of examples (see Appendix Section A on page 261) were expanded and clarified.
- Replaced incorrect use of **omp\_integer\_kind** in Fortran interfaces (see Section C.3 on page 302 and Section C.4 on page 307) with **selected\_int\_kind(8)**.

## E.4 Version 2.5 to 3.0 Differences

The concept of tasks has been added to the OpenMP execution model (see Section 1.2.5 on page 9 and Section 1.3 on page 13).

- The **task** construct (see Section 2.9 on page 80) has been added, which provides a mechanism for creating tasks explicitly.
- The **taskwait** construct (see Section 2.12.4 on page 139) has been added, which causes a task to wait for all its child tasks to complete.
- The OpenMP memory model now covers atomicity of memory accesses (see Section 1.4.1 on page 17). The description of the behavior of **volatile** in terms of **flush** was removed.

- 1           • In Version 2.5, there was a single copy of the *nest-var*, *dyn-var*, *nthreads-var* and *run-sched-var*  
2           internal control variables (ICVs) for the whole program. In Version 3.0, there is one copy of  
3           these ICVs per task (see Section 2.3 on page 35). As a result, the **omp\_set\_num\_threads**,  
4           **omp\_set\_nested** and **omp\_set\_dynamic** runtime library routines now have specified  
5           effects when called from inside a **parallel** region (see Section 3.2.1 on page 208,  
6           Section 3.2.7 on page 214 and Section 3.2.10 on page 217).
- 7           • The definition of active **parallel** region has been changed: in Version 3.0 a **parallel**  
8           region is active if it is executed by a team consisting of more than one thread (see Section 1.2.2  
9           on page 2).
- 10          • The rules for determining the number of threads used in a **parallel** region have been modified  
11          (see Section 2.5.1 on page 48).
- 12          • In Version 3.0, the assignment of iterations to threads in a loop construct with a **static**  
13          schedule kind is deterministic (see Section 2.7.1 on page 55).
- 14          • In Version 3.0, a loop construct may be associated with more than one perfectly nested loop. The  
15          number of associated loops may be controlled by the **collapse** clause (see Section 2.7.1 on  
16          page 55).
- 17          • Random access iterators, and variables of unsigned integer type, may now be used as loop  
18          iterators in loops associated with a loop construct (see Section 2.7.1 on page 55).
- 19          • The schedule kind **auto** has been added, which gives the implementation the freedom to choose  
20          any possible mapping of iterations in a loop construct to threads in the team (see Section 2.7.1 on  
21          page 55).
- 22          • Fortran assumed-size arrays now have predetermined data-sharing attributes (see  
23          Section 2.14.1.1 on page 162).
- 24          • In Fortran, **firstprivate** is now permitted as an argument to the **default** clause (see  
25          Section 2.14.3.1 on page 173).
- 26          • For list items in the **private** clause, implementations are no longer permitted to use the storage  
27          of the original list item to hold the new list item on the master thread. If no attempt is made to  
28          reference the original list item inside the **parallel** region, its value is well defined on exit  
29          from the **parallel** region (see Section 2.14.3.3 on page 176).
- 30          • In Version 3.0, Fortran allocatable arrays may appear in **private**, **firstprivate**,  
31          **lastprivate**, **reduction**, **copyin** and **copyprivate** clauses. (see Section 2.14.2 on  
32          page 166, Section 2.14.3.3 on page 176, Section 2.14.3.4 on page 179, Section 2.14.3.5 on  
33          page 182,  
34          Section 2.14.3.6 on page 184, Section 2.14.4.1 on page 192 and Section 2.14.4.2 on page 193).
- 35          • In Version 3.0, static class members variables may appear in a **threadprivate** directive (see  
36          Section 2.14.2 on page 166).

- 1       • Version 3.0 makes clear where, and with which arguments, constructors and destructors of  
2       private and threadprivate class type variables are called (see Section 2.14.2 on page 166,  
3       Section 2.14.3.3 on page 176, Section 2.14.3.4 on page 179,  
4       Section 2.14.4.1 on page 192 and Section 2.14.4.2 on page 193).
- 5       • The runtime library routines `omp_set_schedule` and `omp_get_schedule` have been  
6       added; these routines respectively set and retrieve the value of the *run-sched-var* ICV (see  
7       Section 3.2.12 on page 220 and Section 3.2.13 on page 222).
- 8       • The *thread-limit-var* ICV has been added, which controls the maximum number of threads  
9       participating in the OpenMP program. The value of this ICV can be set with the  
10       **OMP\_THREAD\_LIMIT** environment variable and retrieved with the  
11       **omp\_get\_thread\_limit** runtime library routine (see Section 2.3.1 on page 35,  
12       Section 3.2.14 on page 223 and Section 4.10 on page 258).
- 13       • The *max-active-levels-var* ICV has been added, which controls the number of nested active  
14       **parallel** regions. The value of this ICV can be set with the **OMP\_MAX\_ACTIVE\_LEVELS**  
15       environment variable and the **omp\_set\_max\_active\_levels** runtime library routine, and  
16       it can be retrieved with the `omp_get_max_active_levels` runtime library routine (see Section 2.3.1  
17       on page 35, Section 3.2.15 on page 223, Section 3.2.16 on page 225 and Section 4.9 on page 258).
- 18       • The *stacksize-var* ICV has been added, which controls the stack size for threads that the OpenMP  
19       implementation creates. The value of this ICV can be set with the **OMP\_STACKSIZE**  
20       environment variable (see Section 2.3.1 on page 35 and Section 4.7 on page 256).
- 21       • The *wait-policy-var* ICV has been added, which controls the desired behavior of waiting threads.  
22       The value of this ICV can be set with the **OMP\_WAIT\_POLICY** environment variable (see  
23       Section 2.3.1 on page 35 and Section 4.8 on page 257).
- 24       • The `omp_get_level` runtime library routine has been added, which returns the number of  
25       nested **parallel** regions enclosing the task that contains the call (see Section 3.2.17 on  
26       page 226).
- 27       • The `omp_get_ancestor_thread_num` runtime library routine has been added, which  
28       returns, for a given nested level of the current thread, the thread number of the ancestor (see  
29       Section 3.2.18 on page 227).
- 30       • The `omp_get_team_size` runtime library routine has been added, which returns, for a given  
31       nested level of the current thread, the size of the thread team to which the ancestor belongs (see  
32       Section 3.2.19 on page 228).
- 33       • The `omp_get_active_level` runtime library routine has been added, which returns the  
34       number of nested, active **parallel** regions enclosing the task that contains the call (see  
35       Section 3.2.20 on page 229).
- 36       • In Version 3.0, locks are owned by tasks, not by threads (see Section 3.3 on page 239).

# Index

---

## Symbols

`_OPENMP` macro, [259](#)

`_OPENMP` macro, [32](#)

## A

affinity, [50](#)

array sections, [43](#)

**atomic**, [141](#)

**atomic** construct, [309](#)

attribute clauses, [172](#)

attributes, data-sharing, [162](#)

**auto**, [60](#)

## B

**barrier**, [138](#)

## C

C/C++ stub routines, [262](#)

**cancel**, [156](#)

cancellation constructs, [156](#)

**cancel**, [156](#)

**cancellation point**, [160](#)

**cancellation point**, [160](#)

canonical loop form, [52](#)

clauses

attribute data-sharing, [172](#)

**collapse**, [55](#), [57](#)

**copyin**, [192](#)

**copyprivate**, [193](#)

data copying, [191](#)

data-sharing, [172](#)

**default**, [173](#)

**depend**, [154](#)

**firstprivate**, [179](#)

**lastprivate**, [182](#)

**linear**, [190](#)

**map**, [195](#)

**private**, [176](#)

**reduction**, [184](#)

**schedule**, [57](#)

**shared**, [174](#)

combined constructs, [116](#)

parallel loop construct, [117](#)

parallel loop SIMD construct, [121](#)

**parallel sections**, [118](#)

**parallel workshare**, [120](#)

**target teams**, [123](#)

**target teams distribute**, [127](#)

target teams distribute parallel loop  
construct, [131](#)

target teams distribute parallel loop  
SIMD construct, [133](#)

**target teams distribute simd**,  
[128](#)

**teams distribute**, [125](#)

teams distribute parallel loop construct,  
[130](#)

teams distribute parallel loop SIMD  
construct, [132](#)

**teams distribute simd**, [126](#)

compilation sentinels, [33](#)

compliance, [21](#)

conditional compilation, [32](#)

constructs

**atomic**, [141](#)

**barrier**, [138](#)

**cancel**, [156](#)

cancellation constructs, [156](#)

**cancellation point**, [160](#)

combined constructs, [116](#)

**critical**, [136](#)

**declare simd**, [74](#)

**declare target**, [99](#)

device constructs, [92](#)

**distribute**, [104](#)

**distribute parallel do**, 109  
**distribute parallel do simd**,  
 110  
**distribute parallel for**, 109  
**distribute parallel for simd**,  
 110  
 distribute parallel loop, 109  
 distribute parallel loop SIMD, 110  
**distribute simd**, 107  
**do Fortran**, 55  
**flush**, 148  
**for**, C/C++, 55  
*loop*, 55  
 Loop SIMD, 78  
**master**, 135  
**ordered**, 152  
**parallel**, 44  
**parallel do Fortran**, 117  
**parallel for C/C++**, 117  
 parallel loop construct, 117  
 parallel loop SIMD construct, 121  
**parallel sections**, 118  
**parallel workshare**, 120  
**sections**, 63  
**simd**, 70  
**single**, 65  
**target**, 93  
**target data**, 92  
**target enter data**, 112  
**target exit data**, 114  
**target teams**, 123  
**target teams distribute**, 127  
 target teams distribute parallel loop  
 construct, 131  
 target teams distribute parallel loop  
 SIMD construct, 133  
**target teams distribute simd**,  
 128  
**target update**, 96  
**task**, 80  
**taskgroup**, 140  
 tasking constructs, 80  
**taskloop**, 83

**taskloop simd**, 88  
**taskwait**, 139  
**taskyield**, 89  
**teams**, 102  
**teams distribute**, 125  
 teams distribute parallel loop construct,  
 130  
 teams distribute parallel loop SIMD  
 construct, 132  
**teams distribute simd**, 126  
**workshare**, 67  
 worksharing, 54  
 controlling OpenMP thread affinity, 50  
**copyin**, 192  
**copyprivate**, 193  
**critical**, 136

## D

data copying clauses, 191  
 data environment, 162  
 data terminology, 10  
 data-sharing attribute clauses, 172  
 data-sharing attribute rules, 162  
**declare reduction**, 199  
**declare simd**, 74  
**declare target**, 99  
**default**, 173  
**depend**, 154  
 device constructs, 92  
**declare target**, 99  
 device constructs, 92  
**distribute**, 104  
 distribute parallel loop, 109  
 distribute parallel loop SIMD, 110  
**distribute simd**, 107  
**target**, 93  
**target update**, 96  
**teams**, 102  
 device data environments, 18, 112, 114  
 directive format, 26  
 directives, 25  
**declare reduction**, 199  
**declare target**, 99  
**threadprivate**, 166



**distribute**, 104  
distribute parallel loop construct, 109  
distribute parallel loop SIMD construct, 110  
**distribute simd**, 107  
**do**, *Fortran*, 55  
**do simd**, 78  
**dynamic**, 59  
dynamic thread adjustment, 308

## E

environment variables, 249  
    **OMP\_CANCELLATION**, 259  
    **OMP\_DEFAULT\_DEVICE**, 260  
    **OMP\_DISPLAY\_ENV**, 259  
    **OMP\_DYNAMIC**, 253  
    **OMP\_MAX\_ACTIVE\_LEVELS**, 258  
    **OMP\_NESTED**, 256  
    **OMP\_NUM\_THREADS**, 252  
    **OMP\_PLACES**, 254  
    **OMP\_PROC\_BIND**, 253  
    **OMP\_SCHEDULE**, 251  
    **OMP\_STACKSIZE**, 256  
    **OMP\_THREAD\_LIMIT**, 258  
    **OMP\_WAIT\_POLICY**, 257  
execution environment routines, 208  
execution model, 13

## F

features history, 312  
**firstprivate**, 179  
fixed source form conditional compilation  
    sentinels, 33  
fixed source form directives, 28  
**flush**, 148  
flush operation, 18  
**for**, *C/C++*, 55  
**for simd**, 78  
free source form conditional compilation  
    sentinel, 33  
free source form directives, 29

## G

glossary, 2  
grammar, 276

**guided**, 59

## H

header files, 207, 296  
history of features, 312

## I

ICVs (internal control variables), 35  
implementation, 308  
implementation terminology, 12  
include files, 207, 296  
interface declarations, 296  
internal control variables, 308  
internal control variables (ICVs), 35  
introduction, 1

## L

**lastprivate**, 182  
**linear**, 190  
lock routines, 239  
**loop**, 55  
loop SIMD Construct, 78  
loop terminology, 8

## M

**map**, 195  
**master**, 135  
master and synchronization constructs and  
    clauses, 135  
memory model, 17  
modifying and retrieving ICV values, 38  
modifying ICV's, 36

## N

nesting of regions, 205  
normative references, 21

## O

**omp\_get\_num\_teams**, 235  
**OMP\_CANCELLATION**, 259  
**OMP\_DEFAULT\_DEVICE**, 260  
**omp\_destroy\_lock**, 241  
**omp\_destroy\_nest\_lock**, 241  
**OMP\_DISPLAY\_ENV**, 259  
**OMP\_DYNAMIC**, 253

`omp_get_active_level`, 229  
`omp_get_ancestor_thread_num`, 227  
`omp_get_cancellation`, 217  
`omp_get_default_device`, 234  
`omp_get_dynamic`, 216  
`omp_get_level`, 226  
`omp_get_max_active_levels`, 225  
`omp_get_max_threads`, 210  
`omp_get_nested`, 219  
`omp_get_num_devices`, 235  
`omp_get_num_procs`, 213  
`omp_get_num_threads`, 209  
`omp_get_proc_bind`, 231  
`omp_get_schedule`, 222  
`omp_get_team_num`, 237  
`omp_get_team_size`, 228  
`omp_get_thread_limit`, 223  
`omp_get_thread_num`, 212  
`omp_get_wtick`, 248  
`omp_get_wtime`, 246  
`omp_in_final`, 230  
`omp_in_parallel`, 213  
`omp_init_lock`, 241  
`omp_init_nest_lock`, 241  
`omp_is_initial_device`, 238  
`OMP_MAX_ACTIVE_LEVELS`, 258  
`OMP_NESTED`, 256  
`OMP_NUM_THREADS`, 252  
`OMP_PLACES`, 254  
`OMP_PROC_BIND`, 253  
`OMP_SCHEDULE`, 251  
`omp_set_default_device`, 233  
`omp_set_dynamic`, 214  
`omp_set_lock`, 242  
`omp_set_max_active_levels`, 223  
`omp_set_nest_lock`, 242  
`omp_set_nested`, 217  
`omp_set_num_threads`, 208  
`omp_set_schedule`, 220  
`OMP_STACKSIZE`, 256  
`omp_test_lock`, 244  
`omp_test_nest_lock`, 244

`OMP_THREAD_LIMIT`, 258  
`omp_unset_lock`, 243  
`omp_unset_nest_lock`, 243  
`OMP_WAIT_POLICY`, 257  
OpenMP compliance, 21  
`ordered`, 152

## P

`parallel`, 44  
parallel loop construct, 117  
parallel loop SIMD construct, 121  
`parallel sections`, 118  
`parallel workshare`, 120  
`private`, 176

## R

`read, atomic`, 141  
`reduction`, 184  
runtime library definitions, 207  
runtime library routines, 206

## S

scheduling, 90  
`sections`, 63  
`shared`, 174  
`simd`, 70  
SIMD Constructs, 70  
Simple Lock Routines, 239  
`single`, 65  
stand-alone directives, 32  
stub routines, 262  
synchronization constructs, 135  
synchronization terminology, 8

## T

`target`, 93  
`target data`, 92  
`target teams`, 123  
`target teams distribute`, 127  
target teams distribute parallel loop construct, 131  
target teams distribute parallel loop SIMD construct, 133  
`target teams distribute simd`, 128

- target update**, [96](#)
- task**, [80](#)
- task scheduling, [90](#)
- taskgroup**, [140](#)
- tasking constructs, [80](#)
- tasking terminology, [9](#)
- taskloop**, [83](#)
- taskloop simd**, [88](#)
- taskwait**, [139](#)
- taskyield**, [89](#)
- teams**, [102](#)
- teams distribute**, [125](#)
- teams distribute parallel loop construct, [130](#)
- teams distribute parallel loop SIMD
  - construct, [132](#)
- teams distribute simd**, [126](#)
- thread affinity, [50](#)
- threadprivate**, [166](#)
- timer, [245](#)
- timing routines, [245](#)

## U

- update, atomic**, [141](#)

## V

- variables, environment, [249](#)

## W

- wall clock timer, [245](#)
- workshare**, [67](#)
- worksharing
  - constructs, [54](#)
  - parallel, [117](#)
  - scheduling, [62](#)
- worksharing constructs, [54](#)
- write, atomic**, [141](#)