



OpenMP Technical Report 7: Version 5.0 Public Comment Draft

This Technical Report is the public comment draft for the OpenMP Application Programming Specification version 5.0 that augments the OpenMP API Specification version 4.5 with support for C11, C++14/17, and Fortran 2008, for concurrent loops, improved worksharing constructs, task reductions, runtime interfaces for first-party (OMPT) and for third-party tools (OMPD), major extensions to the device constructs, memory allocation features, improved task dependencies, and several clarifications and corrections.

EDITORS

Bronis R. de Supinski

Michael Klemm

July 12, 2018

Expires November 8, 2018

We actively solicit comments. Please provide feedback on this document either to the Editors directly or in the OpenMP Forum at openmp.org

End of Public Comment Period: September 12, 2018

OpenMP Architecture Review Board – www.openmp.org – info@openmp.org

Ravi S. Rao, OpenMP, c/o Intel Corporation, 1300 MoPac Express Way, Austin, TX 78746, USA

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 provisions 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 5.0 Public Comment Draft, July 2018

Copyright ©1997-2018 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 3-TR7 (Third Official Draft) (12 July 2018) and includes the following internal tickets applied to the 4.5 LaTeX sources: 50, 134, 354, 389, 399, 408, 425, 426, 430, 445, 452, 458, 459, 461-463, 465-467, 484, 486, 489-504, 508, 510, 511, 514, 518, 520, 521, 523, 524, 530-533, 536, 539, 542, 545, 546, 548, 551, 555-562, 564, 568-578, 581-585, 586, 589-592, 594, 597, 598, 601, 603, 604, 606, 608-618, 620-646, 648-651, 654-656, 659, 661-666, 668-673, 675-677, 679, 680, 681, 684-686, 691-695, 696, 699-712, 715, 717-723, 725-736, 738, 740, 745, 748-763, 765-777, 780-787.

This is a draft; contents will change in official release.

Contents

1. Introduction	1
1.1. Scope	2
1.2. Glossary	2
1.2.1. Threading Concepts	2
1.2.2. OpenMP Language Terminology	3
1.2.3. Loop Terminology	9
1.2.4. Synchronization Terminology	9
1.2.5. Tasking Terminology	11
1.2.6. Data Terminology	13
1.2.7. Implementation Terminology	15
1.2.8. Tool Terminology	16
1.3. Execution Model	19
1.4. Memory Model	22
1.4.1. Structure of the OpenMP Memory Model	22
1.4.2. Device Data Environments	23
1.4.3. Memory Management	24
1.4.4. The Flush Operation	24
1.4.5. Flush Synchronization and Happens Before	26
1.4.6. OpenMP Memory Consistency	27
1.5. Tool Interface	28
1.5.1. OMPT	28
1.5.2. OMPD	29
1.6. OpenMP Compliance	30
1.7. Normative References	30
1.8. Organization of this Document	33

2. Directives	35
2.1. Directive Format	36
2.1.1. Fixed Source Form Directives	38
2.1.2. Free Source Form Directives	39
2.1.3. Stand-Alone Directives	40
2.2. Conditional Compilation	41
2.2.1. Fixed Source Form Conditional Compilation Sentinels	41
2.2.2. Free Source Form Conditional Compilation Sentinel	42
2.3. requires Directive	43
2.4. Internal Control Variables	47
2.4.1. ICV Descriptions	47
2.4.2. ICV Initialization	49
2.4.3. Modifying and Retrieving ICV Values	51
2.4.4. How ICVs are Scoped	54
2.4.4.1. How the Per-Data Environment ICVs Work	55
2.4.5. ICV Override Relationships	56
2.5. Array Shaping	58
2.6. Array Sections	59
2.7. Iterators	61
2.8. Variant Directives	63
2.8.1. OpenMP Context	63
2.8.2. Context Selectors	65
2.8.3. Matching and Scoring Context Selectors	66
2.8.4. declare variant Directive	67
2.8.5. Metadirective Meta-Directive	69
2.9. parallel Construct	72
2.9.1. Determining the Number of Threads for a parallel Region	77
2.9.2. Controlling OpenMP Thread Affinity	79
2.10. teams Construct	81
2.11. Worksharing Constructs	85
2.11.1. sections Construct	86
2.11.2. single Construct	89
2.11.3. workshare Construct	91

2.12. Loop Constructs	95
2.12.1. Canonical Loop Form	95
2.12.2. Worksharing-Loop Construct	100
2.12.2.1. Determining the Schedule of a Worksharing Loop	110
2.12.3. SIMD Constructs	111
2.12.3.1. simd Construct	111
2.12.3.2. Worksharing-Loop SIMD Construct	116
2.12.4. distribute Loop Constructs	117
2.12.4.1. distribute Construct	117
2.12.4.2. distribute simd Construct	121
2.12.4.3. Distribute Parallel Worksharing-Loop Construct	122
2.12.4.4. Distribute Parallel Worksharing-Loop SIMD Construct	124
2.12.5. loop Construct	126
2.12.6. scan Directive	129
2.13. Tasking Constructs	133
2.13.1. task Construct	133
2.13.2. taskloop Construct	138
2.13.3. taskloop simd Construct	144
2.13.4. taskyield Construct	145
2.13.5. Initial Task	146
2.13.6. Task Scheduling	147
2.14. Memory Management Directives	150
2.14.1. Memory Spaces	150
2.14.2. Memory Allocators	151
2.14.3. allocate Directive	154
2.14.4. allocate Clause	157
2.15. Device Constructs	158
2.15.1. Device Initialization	158
2.15.2. target data Construct	159
2.15.3. target enter data Construct	162
2.15.4. target exit data Construct	165
2.15.5. target Construct	168
2.15.6. target update Construct	174

2.15.7. declare target Directive	178
2.15.8. declare mapper Directive	183
2.16. Combined Constructs	185
2.16.1. Parallel Worksharing-Loop Construct	185
2.16.2. Parallel Loop Construct	187
2.16.3. parallel sections Construct	188
2.16.4. parallel workshare Construct	190
2.16.5. Parallel Worksharing-Loop SIMD Construct	191
2.16.6. target parallel Construct	192
2.16.7. Target Parallel Worksharing-Loop Construct	193
2.16.8. Target Parallel Worksharing-Loop SIMD Construct	195
2.16.9. target simd Construct	196
2.16.10. target teams Construct	198
2.16.11. teams distribute Construct	199
2.16.12. teams loop Construct	200
2.16.13. teams distribute simd Construct	201
2.16.14. target teams distribute Construct	203
2.16.15. target teams distribute simd Construct	204
2.16.16. Teams Distribute Parallel Worksharing-Loop Construct	205
2.16.17. Target Teams Distribute Parallel Worksharing-Loop Construct	206
2.16.18. Teams Distribute Parallel Worksharing-Loop SIMD Construct	208
2.16.19. Target Teams Distribute Parallel Worksharing-Loop SIMD Construct	209
2.17. Clauses on Combined and Composite Constructs	211
2.18. if Clause	213
2.19. master Construct	214
2.20. Synchronization Constructs and Clauses	216
2.20.1. critical Construct	216
2.20.2. barrier Construct	219
2.20.3. Implicit Barriers	221
2.20.4. Implementation-Specific Barriers	223
2.20.5. taskwait Construct	223
2.20.6. taskgroup Construct	225
2.20.7. atomic Construct	227

2.20.8.	flush Construct	235
	2.20.8.1. Implicit Flushes	240
2.20.9.	ordered Construct	243
2.20.10.	Depend Objects	246
	2.20.10.1. depobj Construct	247
2.20.11.	depend Clause	248
2.20.12.	Synchronization Hints	253
2.21.	Cancellation Constructs	256
	2.21.1. cancel Construct	256
	2.21.2. cancellation point Construct	260
2.22.	Data Environment	263
	2.22.1. Data-sharing Attribute Rules	263
	2.22.1.1. Data-sharing Attribute Rules for Variables Referenced in a Construct	263
	2.22.1.2. Data-sharing Attribute Rules for Variables Referenced in a Region but not in a Construct	266
	2.22.2. threadprivate Directive	268
	2.22.3. List Item Privatization	273
	2.22.4. Data-Sharing Attribute Clauses	276
	2.22.4.1. default Clause	277
	2.22.4.2. shared Clause	278
	2.22.4.3. private Clause	280
	2.22.4.4. firstprivate Clause	281
	2.22.4.5. lastprivate Clause	283
	2.22.4.6. linear Clause	286
	2.22.5. Reduction Clauses	290
	2.22.5.1. Properties Common To All Reduction Clauses	290
	2.22.5.2. Reduction Scoping Clauses	296
	2.22.5.3. Reduction Participating Clauses	296
	2.22.5.4. reduction Clause	297
	2.22.5.5. task_reduction Clause	299
	2.22.5.6. in_reduction Clause	300
	2.22.6. Data Copying Clauses	301
	2.22.6.1. copyin Clause	301

2.22.6.2. copyprivate Clause	303
2.22.7. Data-mapping Attribute Rules and Clauses	305
2.22.7.1. map Clause	307
2.22.7.2. defaultmap Clause	315
2.23. Declare Directives	317
2.23.1. declare simd Directive	317
2.23.2. declare reduction Directive	321
2.24. Nesting of Regions	327
3. Runtime Library Routines	329
3.1. Runtime Library Definitions	330
3.2. Execution Environment Routines	332
3.2.1. omp_set_num_threads	332
3.2.2. omp_get_num_threads	333
3.2.3. omp_get_max_threads	334
3.2.4. omp_get_thread_num	336
3.2.5. omp_get_num_procs	337
3.2.6. omp_in_parallel	337
3.2.7. omp_set_dynamic	338
3.2.8. omp_get_dynamic	339
3.2.9. omp_get_cancellation	340
3.2.10. omp_set_nested	341
3.2.11. omp_get_nested	342
3.2.12. omp_set_schedule	343
3.2.13. omp_get_schedule	345
3.2.14. omp_get_thread_limit	346
3.2.15. omp_set_max_active_levels	347
3.2.16. omp_get_max_active_levels	348
3.2.17. omp_get_level	349
3.2.18. omp_get_ancestor_thread_num	350
3.2.19. omp_get_team_size	351
3.2.20. omp_get_active_level	353
3.2.21. omp_in_final	353
3.2.22. omp_get_proc_bind	354

3.2.23.	<code>omp_get_num_places</code>	356
3.2.24.	<code>omp_get_place_num_procs</code>	357
3.2.25.	<code>omp_get_place_proc_ids</code>	358
3.2.26.	<code>omp_get_place_num</code>	359
3.2.27.	<code>omp_get_partition_num_places</code>	360
3.2.28.	<code>omp_get_partition_place_nums</code>	360
3.2.29.	<code>omp_set_affinity_format</code>	361
3.2.30.	<code>omp_get_affinity_format</code>	363
3.2.31.	<code>omp_display_affinity</code>	364
3.2.32.	<code>omp_capture_affinity</code>	365
3.2.33.	<code>omp_set_default_device</code>	367
3.2.34.	<code>omp_get_default_device</code>	368
3.2.35.	<code>omp_get_num_devices</code>	369
3.2.36.	<code>omp_get_device_num</code>	369
3.2.37.	<code>omp_get_num_teams</code>	370
3.2.38.	<code>omp_get_team_num</code>	371
3.2.39.	<code>omp_is_initial_device</code>	372
3.2.40.	<code>omp_get_initial_device</code>	373
3.2.41.	<code>omp_get_max_task_priority</code>	374
3.2.42.	<code>omp_pause_resource</code>	375
3.2.43.	<code>omp_pause_resource_all</code>	377
3.3.	Lock Routines	378
3.3.1.	<code>omp_init_lock</code> and <code>omp_init_nest_lock</code>	380
3.3.2.	<code>omp_init_lock_with_hint</code> and <code>omp_init_nest_lock_with_hint</code>	382
3.3.3.	<code>omp_destroy_lock</code> and <code>omp_destroy_nest_lock</code>	383
3.3.4.	<code>omp_set_lock</code> and <code>omp_set_nest_lock</code>	384
3.3.5.	<code>omp_unset_lock</code> and <code>omp_unset_nest_lock</code>	386
3.3.6.	<code>omp_test_lock</code> and <code>omp_test_nest_lock</code>	388
3.4.	Timing Routines	390
3.4.1.	<code>omp_get_wtime</code>	390
3.4.2.	<code>omp_get_wtick</code>	391

3.5.	Event Routines	392
3.5.1.	<code>omp_fulfill_event</code>	392
3.6.	Device Memory Routines	393
3.6.1.	<code>omp_target_alloc</code>	393
3.6.2.	<code>omp_target_free</code>	395
3.6.3.	<code>omp_target_is_present</code>	396
3.6.4.	<code>omp_target_memcpy</code>	396
3.6.5.	<code>omp_target_memcpy_rect</code>	398
3.6.6.	<code>omp_target_associate_ptr</code>	399
3.6.7.	<code>omp_target_disassociate_ptr</code>	401
3.7.	Memory Management Routines	403
3.7.1.	Memory Management Types	403
3.7.2.	<code>omp_init_allocator</code>	406
3.7.3.	<code>omp_destroy_allocator</code>	407
3.7.4.	<code>omp_set_default_allocator</code>	408
3.7.5.	<code>omp_get_default_allocator</code>	409
3.7.6.	<code>omp_alloc</code>	410
3.7.7.	<code>omp_free</code>	411
3.8.	Tool Control Routines	413
4.	Tool Support	417
4.1.	Tool Interfaces Definitions	418
4.2.	OMPT	418
4.2.1.	Activating an OMPT Tool	418
4.2.1.1.	Determining Whether an OMPT Tool Should be Initialized	419
4.2.1.2.	Initializing an OMPT Tool	420
4.2.1.3.	Monitoring Activity on the Host with OMPT	422
4.2.1.4.	Tracing Activity on Target Devices with OMPT	425
4.2.2.	Finalizing an OMPT Tool	429
4.2.3.	OMPT Data Types	429
4.2.3.1.	Tool Initialization and Finalization	429
4.2.3.2.	Callbacks	430
4.2.3.3.	Tracing	431
4.2.3.4.	Miscellaneous Type Definitions	433

4.2.4.	OMPT Tool Callback Signatures and Trace Records	443
4.2.4.1.	Initialization and Finalization Callback Signature	444
4.2.4.2.	Event Callback Signatures and Trace Records	446
4.2.5.	OMPT Runtime Entry Points for Tools	480
4.2.5.1.	Entry Points in the OMPT Callback Interface	481
4.2.5.2.	Entry Points in the OMPT Device Tracing Interface	502
4.2.5.3.	Lookup Entry Point	516
4.3.	OMPD	518
4.3.1.	Activating an OMPD Tool	519
4.3.1.1.	Enabling the Runtime for OMPD	519
4.3.1.2.	Finding the OMPD library	519
4.3.2.	OMPD Data Types	520
4.3.2.1.	Basic Types	520
4.3.2.2.	System Device Identifiers	522
4.3.2.3.	Thread Identifiers	522
4.3.2.4.	OMPD Handle Types	523
4.3.2.5.	OMPD Scope Types	523
4.3.2.6.	ICV ID Type	524
4.3.2.7.	Tool Context Types	524
4.3.2.8.	Return Code Types	525
4.3.2.9.	Primitive Types	526
4.3.3.	OMPD Tool Callback Interface	527
4.3.3.1.	Memory Management of OMPD Library	527
4.3.3.2.	Context Management and Navigation	529
4.3.3.3.	Accessing Memory in the OpenMP Program or Runtime	532
4.3.3.4.	Data Format Conversion	536
4.3.3.5.	Output	537
4.3.3.6.	The Callback Interface	538
4.3.4.	OMPD Tool Interface Routines	540
4.3.4.1.	Per OMPD Library Initialization and Finalization	540
4.3.4.2.	Per OpenMP Process Initialization and Finalization	543
4.3.4.3.	Thread and Signal Safety	546
4.3.4.4.	Address Space Information	546

4.3.4.5.	Thread Handles	548
4.3.4.6.	Parallel Region Handles	553
4.3.4.7.	Task Handles	558
4.3.4.8.	Display Control variables	568
4.3.4.9.	Accessing Scope Specific Information	570
4.3.5.	Runtime Entry Points for OMPD	577
4.3.5.1.	Beginning Parallel Regions	577
4.3.5.2.	Ending Parallel Regions	578
4.3.5.3.	Beginning Task Regions	580
4.3.5.4.	Ending Task Regions	581
4.3.5.5.	Beginning OpenMP Thread	582
4.3.5.6.	Ending OpenMP Thread	582
4.3.5.7.	Beginning OpenMP Device	583
4.3.5.8.	Ending OpenMP Device	584
4.4.	Tool Foundation	584
4.4.1.	Data Types	584
4.4.1.1.	Thread States	584
4.4.1.2.	Frames	589
4.4.1.3.	Wait Identifiers	591
4.4.2.	Global Symbols	591
4.4.2.1.	<code>ompt_start_tool</code>	592
4.4.2.2.	<code>ompd_dll_locations</code>	593
4.4.2.3.	<code>ompd_dll_locations_valid</code>	594
5.	Environment Variables	595
5.1.	<code>OMP_SCHEDULE</code>	596
5.2.	<code>OMP_NUM_THREADS</code>	597
5.3.	<code>OMP_DYNAMIC</code>	598
5.4.	<code>OMP_PROC_BIND</code>	598
5.5.	<code>OMP_PLACES</code>	599
5.6.	<code>OMP_NESTED</code>	601
5.7.	<code>OMP_STACKSIZE</code>	602
5.8.	<code>OMP_WAIT_POLICY</code>	603
5.9.	<code>OMP_MAX_ACTIVE_LEVELS</code>	603

5.10. OMP_THREAD_LIMIT	604
5.11. OMP_CANCELLATION	604
5.12. OMP_DISPLAY_ENV	605
5.13. OMP_DISPLAY_AFFINITY	606
5.14. OMP_AFFINITY_FORMAT	607
5.15. OMP_DEFAULT_DEVICE	609
5.16. OMP_MAX_TASK_PRIORITY	609
5.17. OMP_TARGET_OFFLOAD	610
5.18. OMP_TOOL	611
5.19. OMP_TOOL_LIBRARIES	611
5.20. OMP_DEBUG	612
5.21. OMP_ALLOCATOR	612
A. OpenMP Implementation-Defined Behaviors	613
B. Task Frame Management for the Tool Interface	621
C. Interaction Diagram of OMPD Components	623
D. Features History	625
D.1. Deprecated Features	625
D.2. Version 4.5 to 5.0 Differences	626
D.3. Version 4.0 to 4.5 Differences	629
D.4. Version 3.1 to 4.0 Differences	631
D.5. Version 3.0 to 3.1 Differences	632
D.6. Version 2.5 to 3.0 Differences	633
Index	637

This page intentionally left blank

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 parallelism in C, C++ and Fortran programs.

6 This specification provides a model for parallel programming that is portable across architectures
7 from different vendors. Compilers from numerous vendors support the OpenMP API. More
8 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 to manage parallel programs while permitting portability. The directives extend the C,
12 C++ 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, mapping 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.

1 1.1 Scope

2 The OpenMP API covers only user-directed parallelization, wherein the programmer explicitly
3 specifies the actions to be taken by the compiler and runtime system in order to execute the program
4 in parallel. OpenMP-compliant implementations are not required to check for data dependencies,
5 data conflicts, race conditions, or deadlocks, any of which may occur in conforming programs. In
6 addition, compliant implementations are not required to check for code sequences that cause a
7 program to be classified as non-conforming. Application developers are responsible for correctly
8 using the OpenMP API to produce a conforming program. The OpenMP API does not cover
9 compiler-generated automatic parallelization and directives to the compiler to assist such
10 parallelization.

11 1.2 Glossary

12 1.2.1 Threading Concepts

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

15 **OpenMP thread** A *thread* that is managed by the OpenMP implementation.

16 **idle thread** An *OpenMP thread* that is not currently part of any **parallel** region.

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

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

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

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

23 **host device** The *device* on which the *OpenMP program* begins execution.

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

25 **parent device** For a given **target** region, the device on which the corresponding **target**
26 construct was encountered.

1 1.2.2 OpenMP Language Terminology

2	base language	A programming language that serves as the foundation of the OpenMP specification.
3		COMMENT: See Section 1.7 on page 30 for a listing of current <i>base</i>
4		<i>languages</i> for the OpenMP API.
5	base program	A program written in a <i>base language</i> .
6	program order	An ordering of operations performed by the same thread as determined by the
7		execution sequence of operations specified by the <i>base language</i> .
8		COMMENT: For C11 and C++11, <i>program order</i> corresponds to the
9		sequenced before relation between operations performed by the same
10		thread.
11	structured block	For C/C++, an executable statement, possibly compound, with a single entry at the
12		top and a single exit at the bottom, or an OpenMP <i>construct</i> .
13		For Fortran, a block of executable statements with a single entry at the top and a
14		single exit at the bottom, or an OpenMP <i>construct</i> .
15		COMMENTS:
16		For all <i>base languages</i> :
17		• Access to the <i>structured block</i> must not be the result of a branch;
18		• The point of exit cannot be a branch out of the <i>structured block</i> ;
19		• Infinite loops where the point of exit is never reached are allowed in a
20		<i>structured block</i> ; and
21		• Halting caused by an IEEE exception is allowed in a <i>structured block</i> .
22		For C/C++:
23		• The point of entry must not be a call to <code>setjmp()</code> ;
24		• <code>longjmp()</code> and <code>throw()</code> must not violate the entry/exit criteria;
25		• A <i>structured block</i> may contain calls to <code>exit()</code> , <code>_Exit()</code> ,
26		<code>quick_exit()</code> , <code>abort()</code> or functions with <code>_Noreturn</code> specifier
27		(in C) or <code>noreturn</code> attribute (in C/C++); and
28		• An expression statement, iteration statement, selection statement, or try
29		block is considered to be a <i>structured block</i> if the corresponding
30		compound statement obtained by enclosing it in <code>{</code> and <code>}</code> would be a
31		<i>structured block</i> .
32		For Fortran:

- **STOP** statements are allowed in a *structured block*.

1		
2	compilation unit	For C/C++, a translation unit.
3		For Fortran, a program unit.
4	enclosing context	For C/C++, the innermost scope enclosing an OpenMP <i>directive</i> .
5		For Fortran, the innermost scoping unit enclosing an OpenMP <i>directive</i> .
6	directive	For C/C++, a #pragma , and for Fortran, a comment, that specifies <i>OpenMP program</i> behavior.
7		
8		COMMENT: See Section 2.1 on page 36 for a description of OpenMP
9		<i>directive</i> syntax.
10	meta-directive	A <i>directive</i> that conditionally resolves to another <i>directive</i> at compile time.
11	white space	A non-empty sequence of space and/or horizontal tab characters.
12	OpenMP program	A program that consists of a <i>base program</i> that is annotated with OpenMP <i>directives</i> or that calls OpenMP API runtime library routines
13		
14	conforming program	An <i>OpenMP program</i> that follows all rules and restrictions of the OpenMP specification.
15		
16	declarative directive	An OpenMP <i>directive</i> that may only be placed in a declarative context. A <i>declarative directive</i> results in one or more declarations only; it is not associated with the immediate execution of any user code.
17		
18		
19	executable directive	An OpenMP <i>directive</i> that is not declarative. That is, it may be placed in an executable context.
20		
21	stand-alone directive	An OpenMP <i>executable directive</i> that has no associated executable user code.
22	construct	An OpenMP <i>executable directive</i> (and for Fortran, the paired end directive , if any) and the associated statement, loop or <i>structured block</i> , if any, not including the code in any called routines. That is, the lexical extent of an <i>executable directive</i> .
23		
24		
25	combined construct	A construct that is a shortcut for specifying one construct immediately nested inside another construct. A combined construct is semantically identical to that of explicitly specifying the first construct containing one instance of the second construct and no other statements.
26		
27		
28		
29	composite construct	A construct that is composed of two constructs but does not have identical semantics to specifying one of the constructs immediately nested inside the other. A composite construct either adds semantics not included in the constructs from which it is composed or the nesting of the one construct inside the other is not conforming.
30		
31		
32		
33	combined target construct	A <i>combined construct</i> that is composed of a target construct and another construct.
34		

1	region	All code encountered during a specific instance of the execution of a given <i>construct</i> or of an OpenMP library routine. A <i>region</i> includes any code in called routines as well as any implicit code introduced by the OpenMP implementation. The generation of a <i>task</i> at the point where a <i>task generating construct</i> is encountered is a part of the <i>region</i> of the <i>encountering thread</i> , but an <i>explicit task region</i> corresponding to a <i>task generating construct</i> is not unless it is an <i>included task region</i> . The point where a target or teams directive is encountered is a part of the <i>region</i> of the <i>encountering thread</i> , but the <i>region</i> corresponding to the target or teams directive is not.
2		
3		
4		
5		
6		
7		
8		
9		
10		COMMENTS:
11		A <i>region</i> may also be thought of as the dynamic or runtime extent of a <i>construct</i> or of an OpenMP library routine.
12		
13		During the execution of an <i>OpenMP program</i> , a <i>construct</i> may give rise to many <i>regions</i> .
14		
15	active parallel region	A parallel <i>region</i> that is executed by a <i>team</i> consisting of more than one <i>thread</i> .
16	inactive parallel region	A parallel <i>region</i> that is executed by a <i>team</i> of only one <i>thread</i> .
17	active target region	A target <i>region</i> that is executed on a <i>device</i> other than the <i>device</i> that encountered the target <i>construct</i> .
18		
19	inactive target region	A target <i>region</i> that is executed on the same <i>device</i> that encountered the target <i>construct</i> .
20		
21	sequential part	All code encountered during the execution of an <i>initial task region</i> that is not part of a parallel <i>region</i> corresponding to a parallel <i>construct</i> or a task <i>region</i> corresponding to a task <i>construct</i> .
22		
23		
24		COMMENTS:
25		A <i>sequential part</i> is enclosed by an <i>implicit parallel region</i> .
26		Executable statements in called routines may be in both a <i>sequential part</i> and any number of explicit parallel <i>regions</i> at different points in the program execution.
27		
28		
29	master thread	An <i>OpenMP thread</i> that has <i>thread</i> number 0. A <i>master thread</i> may be an <i>initial thread</i> or the <i>thread</i> that encounters a parallel <i>construct</i> , creates a <i>team</i> , generates a set of <i>implicit tasks</i> , and then executes one of those <i>tasks</i> as <i>thread</i> number 0.
30		
31		
32		
33	parent thread	The <i>thread</i> that encountered the parallel <i>construct</i> and generated a parallel <i>region</i> is the <i>parent thread</i> of each of the <i>threads</i> in the <i>team</i> of that parallel
34		

1		<i>region</i> . The <i>master thread</i> of a parallel <i>region</i> is the same <i>thread</i> as its <i>parent thread</i> with respect to any resources associated with an <i>OpenMP thread</i> .
2		
3	child thread	When a thread encounters a parallel construct, each of the threads in the
4		generated parallel <i>region</i> 's team are <i>child threads</i> of the encountering <i>thread</i> .
5		The target or teams <i>region</i> 's <i>initial thread</i> is not a <i>child thread</i> of the thread that
6		encountered the target or teams construct.
7	ancestor thread	For a given <i>thread</i> , its <i>parent thread</i> or one of its <i>parent thread</i> 's <i>ancestor threads</i> .
8	descendent thread	For a given <i>thread</i> , one of its <i>child threads</i> or one of its <i>child threads</i> ' <i>descendent threads</i> .
9		
10	team	A set of one or more <i>threads</i> participating in the execution of a parallel <i>region</i> .
11		COMMENTS:
12		For an <i>active parallel region</i> , the team comprises the <i>master thread</i> and at
13		least one additional <i>thread</i> .
14		For an <i>inactive parallel region</i> , the <i>team</i> comprises only the <i>master thread</i> .
15	league	The set of <i>teams</i> created by a teams construct.
16	contention group	An <i>initial thread</i> and its <i>descendent threads</i> .
17	implicit parallel region	An <i>inactive parallel region</i> that is not generated from a parallel <i>construct</i> .
18		<i>Implicit parallel regions</i> surround the whole <i>OpenMP program</i> , all target <i>regions</i> ,
19		and all teams <i>regions</i> .
20	initial thread	A <i>thread</i> that executes an <i>implicit parallel region</i> .
21	initial team	A <i>team</i> that comprises an <i>initial thread</i> executing an <i>implicit parallel region</i> .
22	nested construct	A <i>construct</i> (lexically) enclosed by another <i>construct</i> .
23	closely nested construct	A <i>construct</i> nested inside another <i>construct</i> with no other <i>construct</i> nested between
24		them.
25	nested region	A <i>region</i> (dynamically) enclosed by another <i>region</i> . That is, a <i>region</i> generated from
26		the execution of another <i>region</i> or one of its <i>nested regions</i> .
27		COMMENT: Some nestings are <i>conforming</i> and some are not. See
28		Section 2.24 on page 327 for the restrictions on nesting.
29	closely nested region	A <i>region nested</i> inside another <i>region</i> with no parallel <i>region nested</i> between
30		them.
31	strictly nested region	A <i>region nested</i> inside another <i>region</i> with no other <i>region nested</i> between them.
32	all threads	All <i>OpenMP threads</i> participating in the <i>OpenMP program</i> .
33	current team	All <i>threads</i> in the <i>team</i> executing the innermost enclosing parallel <i>region</i> .

1	encountering thread	For a given <i>region</i> , the <i>thread</i> that encounters the corresponding <i>construct</i> .
2	all tasks	All <i>tasks</i> participating in the <i>OpenMP program</i> .
3	current team tasks	All <i>tasks</i> encountered by the corresponding <i>team</i> . The <i>implicit tasks</i> constituting the
4		parallel <i>region</i> and any <i>descendent tasks</i> encountered during the execution of
5		these <i>implicit tasks</i> are included in this set of tasks.
6	generating task	For a given <i>region</i> , the task for which execution by a <i>thread</i> generated the <i>region</i> .
7	binding thread set	The set of <i>threads</i> that are affected by, or provide the context for, the execution of a
8		<i>region</i> .
9		The <i>binding thread set</i> for a given <i>region</i> can be <i>all threads</i> on a <i>device</i> , <i>all threads</i>
10		in a <i>contention group</i> , <i>all master threads</i> executing an enclosing teams <i>region</i> , the
11		<i>current team</i> , or the <i>encountering thread</i> .
12		COMMENT: The <i>binding thread set</i> for a particular <i>region</i> is described in
13		its corresponding subsection of this specification.
14	binding task set	The set of <i>tasks</i> that are affected by, or provide the context for, the execution of a
15		<i>region</i> .
16		The <i>binding task set</i> for a given <i>region</i> can be <i>all tasks</i> , the <i>current team tasks</i> , the
17		<i>binding implicit task</i> or the <i>generating task</i> .
18		COMMENT: The <i>binding task set</i> for a particular <i>region</i> (if applicable) is
19		described in its corresponding subsection of this specification.
20	binding region	The enclosing <i>region</i> that determines the execution context and limits the scope of
21		the effects of the bound <i>region</i> is called the <i>binding region</i> .
22		<i>Binding region</i> is not defined for <i>regions</i> for which the <i>binding thread set</i> is <i>all</i>
23		<i>threads</i> or the <i>encountering thread</i> , nor is it defined for <i>regions</i> for which the <i>binding</i>
24		<i>task set</i> is <i>all tasks</i> .
25		COMMENTS:
26		The <i>binding region</i> for an ordered <i>region</i> is the innermost enclosing
27		<i>loop region</i> .
28		The <i>binding region</i> for a taskwait <i>region</i> is the innermost enclosing
29		<i>task region</i> .
30		The <i>binding region</i> for a cancel <i>region</i> is the innermost enclosing
31		<i>region</i> corresponding to the <i>construct-type-clause</i> of the cancel
32		<i>construct</i> .

1 The *binding region* for a **cancellation point** *region* is the
2 innermost enclosing *region* corresponding to the *construct-type-clause* of
3 the **cancellation point** construct.

4 For all other *regions* for which the *binding thread set* is the *current team*
5 or the *binding task set* is the *current team tasks*, the *binding region* is the
6 innermost enclosing **parallel** *region*.

7 For *regions* for which the *binding task set* is the *generating task*, the
8 *binding region* is the *region* of the *generating task*.

9 A **parallel** *region* need not be *active* nor explicit to be a *binding*
10 *region*.

11 A *task region* need not be explicit to be a *binding region*.

12 A *region* never binds to any *region* outside of the innermost enclosing
13 **parallel** *region*.

14 **orphaned construct** A *construct* that gives rise to a *region* for which the *binding thread set* is the *current*
15 *team*, but is not nested within another *construct* giving rise to the *binding region*.

16 **worksharing construct** A *construct* that defines units of work, each of which is executed exactly once by one
17 of the *threads* in the *team* executing the *construct*.

18 For C/C++, *worksharing constructs* are **for**, **sections**, and **single**.

19 For Fortran, *worksharing constructs* are **do**, **sections**, **single** and
20 **workshare**.

21 **place** Unordered set of *processors* on a device that is treated by the execution environment
22 as a location unit when dealing with OpenMP thread affinity.

23 **place list** The ordered list that describes all OpenMP *places* available to the execution
24 environment.

25 **place partition** An ordered list that corresponds to a contiguous interval in the OpenMP *place list*. It
26 describes the *places* currently available to the execution environment for a given
27 *parallel region*.

28 **place number** A number that uniquely identifies a *place* in the *place list*, with zero identifying the
29 first *place* in the *place list*, and each consecutive whole number identifying the next
30 *place* in the *place list*.

31 **SIMD instruction** A single machine instruction that can operate on multiple data elements.

32 **SIMD lane** A software or hardware mechanism capable of processing one data element from a
33 *SIMD instruction*.

34 **SIMD chunk** A set of iterations executed concurrently, each by a *SIMD lane*, by a single *thread* by
35 means of *SIMD instructions*.

1	memory	A storage resource to store and to retrieve variables accessible by OpenMP threads.
2	memory space	A representation of storage resources from which <i>memory</i> can be allocated or
3		deallocated.
4	memory allocator	An OpenMP object that fulfills requests to allocate and to deallocate <i>memory</i> for
5		program variables from the storage resources of its associated <i>memory space</i> .

6 1.2.3 Loop Terminology

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

20 1.2.4 Synchronization Terminology

21	barrier	A point in the execution of a program encountered by a <i>team of threads</i> , beyond
22		which no <i>thread</i> in the team may execute until all <i>threads</i> in the <i>team</i> have reached
23		the barrier and all <i>explicit tasks</i> generated by the <i>team</i> have executed to completion.
24		If <i>cancellation</i> has been requested, threads may proceed to the end of the canceled
25		<i>region</i> even if some threads in the team have not reached the <i>barrier</i> .
26	cancellation	An action that cancels (that is, aborts) an OpenMP <i>region</i> and causes executing
27		<i>implicit</i> or <i>explicit</i> tasks to proceed to the end of the canceled <i>region</i> .

1	cancellation point	A point at which implicit and explicit tasks check if cancellation has been requested.
2		If cancellation has been observed, they perform the <i>cancellation</i> .
3		COMMENT: For a list of cancellation points, see Section 2.21.1 on
4		page 256.
5	flush	An operation that a <i>thread</i> performs to enforce consistency between its view and
6		other <i>threads</i> ' view of memory.
7	flush property	Properties that determine the manner in which a <i>flush</i> operation enforces memory
8		consistency. These properties are:
9		• <i>strong</i> : flushes a set of variables from the current thread's temporary view of the
10		memory to the memory;
11		• <i>release</i> : orders memory operations that precede the flush before memory
12		operations performed by a different thread with which it synchronizes;
13		• <i>acquire</i> : orders memory operations that follow the flush after memory operations
14		performed by a different thread that synchronizes with it.
15		COMMENT: Any <i>flush</i> operation has one or more <i>flush properties</i> .
16	strong flush	A <i>flush</i> operation that has the <i>strong flush property</i> .
17	release flush	A <i>flush</i> operation that has the <i>release flush property</i> .
18	acquire flush	A <i>flush</i> operation that has the <i>acquire flush property</i> .
19	atomic operation	An operation that is specified by an atomic construct and atomically accesses
20		and/or modifies a specific storage location.
21	atomic read	An <i>atomic operation</i> that is specified by an atomic construct on which the read
22		clause is present.
23	atomic write	An <i>atomic operation</i> that is specified by an atomic construct on which the write
24		clause is present.
25	atomic update	An <i>atomic operation</i> that is specified by an atomic construct on which the
26		update clause is present.
27	atomic captured	An <i>atomic operation</i> that is specified by an atomic construct on which the
28	update	capture clause is present.
29	read-modify-write	An <i>atomic operation</i> that reads and writes to a given storage location.
30		COMMENT: All <i>atomic update</i> and <i>atomic captured update</i> operations
31		are <i>read-modify-write</i> operations.
32	sequentially consistent	An atomic construct for which the seq_cst clause is specified.
	atomic construct	

- 1 **non-sequentially** An **atomic** construct for which the **seq_cst** clause is not specified
 consistent atomic
 construct
- 2 **sequentially consistent** An *atomic operation* that is specified by a *sequentially consistent atomic construct*.
 atomic operation

3 1.2.5 Tasking Terminology

- 4 **task** A specific instance of executable code and its data environment that the OpenMP
 5 implementation can schedule for execution by threads.
- 6 **task region** A *region* consisting of all code encountered during the execution of a *task*.
 7 COMMENT: A **parallel** *region* consists of one or more implicit *task*
 8 *regions*.
- 9 **implicit task** A *task* generated by an *implicit parallel region* or generated when a **parallel**
 10 *construct* is encountered during execution.
- 11 **binding implicit task** The *implicit task* of the current thread team assigned to the encountering thread.
- 12 **explicit task** A *task* that is not an *implicit task*.
- 13 **initial task** An *implicit task* associated with an *implicit parallel region*.
- 14 **current task** For a given *thread*, the *task* corresponding to the *task region* in which it is executing.
- 15 **child task** A *task* is a *child task* of its generating *task region*. A *child task region* is not part of
 16 its generating *task region*.
- 17 **sibling tasks** *Tasks* that are *child tasks* of the same *task region*.
- 18 **descendent task** A *task* that is the *child task* of a *task region* or of one of its *descendent task regions*.
- 19 **task completion** *Task completion* occurs when the end of the *structured block* associated with the
 20 *construct* that generated the *task* is reached.
 21 COMMENT: Completion of the *initial task* that is generated when the
 22 program begins occurs at program exit.
- 23 **task scheduling point** A point during the execution of the current *task region* at which it can be suspended
 24 to be resumed later; or the point of *task completion*, after which the executing thread
 25 may switch to a different *task region*.
 26 COMMENT: For a list of *task scheduling points*, see Section 2.13.6 on
 27 page 147.

1	task switching	The act of a <i>thread</i> switching from the execution of one <i>task</i> to another <i>task</i> .
2	tied task	A <i>task</i> that, when its <i>task region</i> is suspended, can be resumed only by the same
3		<i>thread</i> that suspended it. That is, the <i>task</i> is tied to that <i>thread</i> .
4	untied task	A <i>task</i> that, when its <i>task region</i> is suspended, can be resumed by any <i>thread</i> in the
5		team. That is, the <i>task</i> is not tied to any <i>thread</i> .
6	undeferrred task	A <i>task</i> for which execution is not deferred with respect to its generating <i>task region</i> .
7		That is, its generating <i>task region</i> is suspended until execution of the <i>undeferrred task</i>
8		is completed.
9	included task	A <i>task</i> for which execution is sequentially included in the generating <i>task region</i> .
10		That is, an <i>included task</i> is <i>undeferrred</i> and executed immediately by the <i>encountering</i>
11		<i>thread</i> .
12	merged task	A <i>task</i> for which the <i>data environment</i> , inclusive of ICVs, is the same as that of its
13		generating <i>task region</i> .
14	mergeable task	A <i>task</i> that may be a <i>merged task</i> if it is an <i>undeferrred task</i> or an <i>included task</i> .
15	final task	A <i>task</i> that forces all of its <i>child tasks</i> to become <i>final</i> and <i>included tasks</i> .
16	task dependence	An ordering relation between two <i>sibling tasks</i> : the <i>dependent task</i> and a previously
17		generated <i>predecessor task</i> . The <i>task dependence</i> is fulfilled when the <i>predecessor</i>
18		<i>task</i> has completed.
19	dependent task	A <i>task</i> that because of a <i>task dependence</i> cannot be executed until its <i>predecessor</i>
20		<i>tasks</i> have completed.
21	mutually exclusive tasks	<i>Tasks</i> that may be executed in any order, but not at the same time.
22	predecessor task	A <i>task</i> that must complete before its <i>dependent tasks</i> can be executed.
23	task synchronization construct	A taskwait , taskgroup , or a barrier <i>construct</i> .
24	task generating construct	A <i>construct</i> that generates one or more <i>explicit tasks</i> .
25	target task	A <i>mergeable</i> and <i>untied task</i> that is generated by a target , target enter
26		data , target exit data , or target update <i>construct</i> .
27	taskgroup set	A set of tasks that are logically grouped by a taskgroup <i>region</i> .

1 1.2.6 Data Terminology

2 **variable** A named data storage block, for which the value can be defined and redefined during
3 the execution of a program.

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

5 **scalar variable** For C/C++: A scalar variable, as defined by the base language.

6 For Fortran: A scalar variable with intrinsic type, as defined by the base language,
7 excluding character type.

8 **aggregate variable** A variable, such as an array or structure, composed of other variables.

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

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

11 **base expression** For C/C++: The expression in an array section or array element that specifies the
12 address of the original array.

13 COMMENT: The *base expression* is *x* for *array element* $x[i]$ and for *array*
14 *section* $x[i:j]$.

15 **named array** For C/C++: An expression that is an array but not an array element and appears as
16 the array referred to by a given array item.

17 For Fortran: A variable that is an array and appears as the array referred to by a given
18 array item.

19 **named pointer** For C/C++: An lvalue expression that is a pointer and appears as a pointer to the
20 array implicitly referred to by a given array item.

21 For Fortran: A variable that has the **POINTER** attribute and appears as a pointer to
22 the array to which a given array item implicitly refers.

23 COMMENT: A given array item cannot have a *named pointer* if it has a
24 *named array*.

25 **attached pointer** A pointer variable in a device data environment to which the effect of a **map** clause
26 assigns the address of an array section. The pointer is an attached pointer for the
27 remainder of its lifetime in the device data environment.

28 **simply contiguous** An array section that statically can be determined to have contiguous storage or that
29 **array section** has the **CONTIGUOUS** attribute.

1	structure	A structure is a variable that contains one or more variables.
2		For C/C++: Implemented using struct types.
3		For C++: Implemented using class types.
4		For Fortran: Implemented using derived types.
5	private variable	With respect to a given set of <i>task regions</i> or <i>SIMD lanes</i> that bind to the same
6		parallel <i>region</i> , a <i>variable</i> for which the name provides access to a different
7		block of storage for each <i>task region</i> or <i>SIMD lane</i> .
8		A <i>variable</i> that is part of another variable (as an array or structure element) cannot be
9		made private independently of other components.
10	shared variable	With respect to a given set of <i>task regions</i> that bind to the same parallel <i>region</i> , a
11		<i>variable</i> for which the name provides access to the same block of storage for each
12		<i>task region</i> .
13		A <i>variable</i> that is part of another variable (as an array or structure element) cannot be
14		<i>shared</i> independently of the other components, except for static data members of
15		C++ classes.
16	threadprivate variable	A <i>variable</i> that is replicated, one instance per <i>thread</i> , by the OpenMP
17		implementation. Its name then provides access to a different block of storage for each
18		<i>thread</i> .
19		A <i>variable</i> that is part of another variable (as an array or structure element) cannot be
20		made <i>threadprivate</i> independently of the other components, except for static data
21		members of C++ classes.
22	threadprivate memory	The set of <i>threadprivate variables</i> associated with each <i>thread</i> .
23	data environment	The <i>variables</i> associated with the execution of a given <i>region</i> .
24	device data environment	The initial <i>data environment</i> associated with a device.
25	device address	An <i>implementation defined</i> reference to an address in a <i>device data environment</i> .
26	device pointer	A <i>variable</i> that contains a <i>device address</i> .
27	mapped variable	An original <i>variable</i> in a <i>data environment</i> with a corresponding <i>variable</i> in a device
28		<i>data environment</i> .
29		COMMENT: The original and corresponding <i>variables</i> may share storage.
30	map-type decay	The process used to determine the final map type used when mapping a variable with
31		a user defined mapper. The combination of the two map types determines the final
32		map type based on the following table.

	alloc	to	from	tofrom	release	delete
alloc	alloc	alloc	alloc	alloc	release	delete
to	alloc	to	alloc	to	release	delete
from	alloc	alloc	from	from	release	delete
tofrom	alloc	to	from	tofrom	release	delete

mappable type A type that is valid for a *mapped variable*. If a type is composed from other types (such as the type of an array or structure element) and any of the other types are not mappable then the type is not mappable.

COMMENT: Pointer types are *mappable* but the memory block to which the pointer refers is not *mapped*.

For C: The type must be a complete type.

For C++: The type must be a complete type.

In addition, for class types:

- All member functions accessed in any **target** region must appear in a **declare target** directive.

For Fortran: No restrictions on the type except that for derived types:

- All type-bound procedures accessed in any target region must appear in a **declare target** directive.

defined For *variables*, the property of having a valid value.

For C: For the contents of *variables*, the property of having a valid value.

For C++: For the contents of *variables* of POD (plain old data) type, the property of having a valid value.

For *variables* of non-POD class type, the property of having been constructed but not subsequently destructed.

For Fortran: For the contents of *variables*, the property of having a valid value. For the allocation or association status of *variables*, the property of having a valid status.

COMMENT: Programs that rely upon *variables* that are not *defined* are *non-conforming programs*.

class type For C++: *Variables* declared with one of the **class**, **struct**, or **union** keywords

1.2.7 Implementation Terminology

supporting *n* levels of parallelism Implies allowing an *active parallel region* to be enclosed by *n-1 active parallel regions*.

1	supporting the OpenMP API	Supporting at least one level of parallelism.
2	supporting nested parallelism	Supporting more than one level of parallelism.
3	internal control variable	A conceptual variable that specifies runtime behavior of a set of <i>threads</i> or <i>tasks</i> in an <i>OpenMP program</i> .
4		
5		COMMENT: The acronym ICV is used interchangeably with the term
6		<i>internal control variable</i> in the remainder of this specification.
7	compliant implementation	An implementation of the OpenMP specification that compiles and executes any <i>conforming program</i> as defined by the specification.
8		
9		COMMENT: A <i>compliant implementation</i> may exhibit <i>unspecified behavior</i> when compiling or executing a <i>non-conforming program</i> .
10		
11	unspecified behavior	A behavior or result that is not specified by the OpenMP specification or not known prior to the compilation or execution of an <i>OpenMP program</i> .
12		
13		Such <i>unspecified behavior</i> may result from:
14		• Issues documented by the OpenMP specification as having <i>unspecified behavior</i> .
15		• A <i>non-conforming program</i> .
16		• A <i>conforming program</i> exhibiting an <i>implementation defined</i> behavior.
17	implementation defined	Behavior that must be documented by the implementation, and is allowed to vary among different <i>compliant implementations</i> . An implementation is allowed to define this behavior as <i>unspecified</i> .
18		
19		
20		COMMENT: All features that have <i>implementation defined</i> behavior are
21		documented in Appendix A.
22	deprecated	Implies a construct, clause, or other feature is normative in the current specification but is considered obsolescent and will be removed in the future.
23		

24 1.2.8 Tool Terminology

25	tool	Executable code, distinct from application or runtime code, that can observe and/or
26		modify the execution of an application.
27	first-party tool	A tool that executes in the address space of the program it is monitoring.
28	third-party tool	A tool that executes as a separate process from that which it is monitoring and
29		potentially controlling.

1	activated tool	A first-party tool that successfully completed its initialization.
2	event	A point of interest in the execution of a thread where the condition defining that event
3		is true.
4	tool callback	A function provided by a tool to an OpenMP implementation that can be invoked
5		when needed.
6	registering a callback	Providing a callback function to an OpenMP implementation for a particular purpose.
7	dispatching a callback	Processing a callback when an associated event occurs in a manner consistent with
8	at an event	the return code provided when a <i>first-party</i> tool registered the callback.
9	thread state	An enumeration type that describes what an OpenMP thread is currently doing. A
10		thread can be in only one state at any time.
11	wait identifier	A unique opaque handle associated with each data object (e.g., a lock) used by the
12		OpenMP runtime to enforce mutual exclusion that may cause a thread to wait actively
13		or passively.
14	frame	A storage area on a thread's stack associated with a procedure invocation. A frame
15		includes space for one or more saved registers and often also includes space for saved
16		arguments, local variables, and padding for alignment.
17	canonical frame	An address associated with a procedure <i>frame</i> on a call stack defined as the value of
18	address	the stack pointer immediately prior to calling the procedure whose invocation the
19		frame represents.
20	runtime entry point	A function interface provided by an OpenMP runtime for use by a tool. A runtime
21		entry point is typically not associated with a global function symbol.
22	trace record	A data structure to store information associated with an occurrence of an <i>event</i> .
23	native trace record	A <i>trace record</i> for an OpenMP device that is in a device-specific format.
24	signal	A software interrupt delivered to a thread.
25	signal handler	A function called asynchronously when a <i>signal</i> is delivered to a thread.
26	async signal safe	Guaranteed not to interfere with operations that are being interrupted by <i>signal</i>
27		delivery. An async signal safe <i>runtime entry point</i> is safe to call from a <i>signal</i>
28		<i>handler</i> .
29	code block	A contiguous region of memory that contains code of an OpenMP program to be
30		executed on a device.
31	OMPT	An interface that helps a first-party tool monitor the execution of an OpenMP
32		program.

1	OMPD	An interface that helps a third-party tool inspect the OpenMP state of a program that
2		has begun execution.
3	OMPD library	A dynamically loadable library that implements the OMPD interface.
4	image file	An executable or shared library.
5	address space	A collection of logical, virtual, or physical memory address ranges containing code,
6		stack, and/or data. Address ranges within an address space need not be contiguous.
7		An address space consists of one or more <i>segments</i> .
8	segment	A region of an address space associated with a set of address ranges.
9	OpenMP architecture	The architecture on which an OpenMP region executes.
10	tool architecture	The architecture on which an OMPD tool executes.
11	OpenMP process	A collection of one or more threads and address spaces. A process may contain
12		threads and address spaces for multiple OpenMP architectures. At least one thread in
13		an OpenMP process is an OpenMP thread. A process may be live or a core file.
14	handle	An opaque reference provided by an OMPD library to a using tool. A handle
15		uniquely identifies an abstraction.
16	address space handle	A handle that refers to an address space within an OpenMP process.
17	thread handle	A handle that refers to an OpenMP thread.
18	parallel handle	A handle that refers to an OpenMP parallel region.
19	task handle	A handle that refers to an OpenMP task region.
20	descendent handle	An output handle that is returned from the OMPD library in a function that accepts
21		an input handle: the output handle is a descendent of the input handle.
22	ancestor handle	An input handle that is passed to the OMPD library in a function that returns an
23		output handle: the input handle is an ancestor of the output handle. For a given
24		handle, the ancestors of the handle are also the ancestors of the handle's descendent.
25		COMMENT: A handle cannot be used by the tool in an OMPD call if at
26		least one ancestor of the handle has been released, except for OMPD calls
27		that release the handle.
28	tool context	An opaque reference provided by a tool to an OMPD library. A tool context uniquely
29		identifies an abstraction.
30	address space context	A tool context that refers to an address space within a process.
31	thread context	A tool context that refers to a thread.
32	thread identifier	An identifier for a native thread defined by a thread implementation.

1 1.3 Execution Model

2 The OpenMP API uses the fork-join model of parallel execution. Multiple threads of execution
3 perform tasks defined implicitly or explicitly by OpenMP directives. The OpenMP API is intended
4 to support programs that will execute correctly both as parallel programs (multiple threads of
5 execution and a full OpenMP support library) and as sequential programs (directives ignored and a
6 simple OpenMP stubs library). However, it is possible and permitted to develop a program that
7 executes correctly as a parallel program but not as a sequential program, or that produces different
8 results when executed as a parallel program compared to when it is executed as a sequential
9 program. Furthermore, using different numbers of threads may result in different numeric results
10 because of changes in the association of numeric operations. For example, a serial addition
11 reduction may have a different pattern of addition associations than a parallel reduction. These
12 different associations may change the results of floating-point addition.

13 An OpenMP program begins as a single thread of execution, called an initial thread. An initial
14 thread executes sequentially, as if the code encountered is part of an implicit task region, called an
15 initial task region, that is generated by the implicit parallel region surrounding the whole program.

16 The thread that executes the implicit parallel region that surrounds the whole program executes on
17 the *host device*. An implementation may support other *target devices*. If supported, one or more
18 devices are available to the host device for offloading code and data. Each device has its own
19 threads that are distinct from threads that execute on another device. Threads cannot migrate from
20 one device to another device. The execution model is host-centric such that the host device offloads
21 **target** regions to target devices.

22 When a **target** construct is encountered, a new *target task* is generated. The *target task* region
23 encloses the **target** region. The *target task* is complete after the execution of the **target** region
24 is complete.

25 When a *target task* executes, the enclosed **target** region is executed by an initial thread. The
26 initial thread may execute on a *target device*. The initial thread executes sequentially, as if the target
27 region is part of an initial task region that is generated by an implicit parallel region. If the target
28 device does not exist or the implementation does not support the target device, all **target** regions
29 associated with that device execute on the host device.

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

33 The **teams** construct creates a *league of teams*, where each team is an initial team that comprises
34 an initial thread that executes the **teams** region. Each initial thread executes sequentially, as if the
35 code encountered is part of an initial task region that is generated by an implicit parallel region
36 associated with each team.

37 If a construct creates a data environment, the data environment is created at the time the construct is
38 encountered. Whether a construct creates a data environment is defined in the description of the

1 construct.

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

12 **parallel** regions may be arbitrarily nested inside each other. If nested parallelism is disabled, or
13 is not supported by the OpenMP implementation, then the new team that is created by a thread
14 encountering a **parallel** construct inside a **parallel** region will consist only of the
15 encountering thread. However, if nested parallelism is supported and enabled, then the new team
16 can consist of more than one thread. A **parallel** construct may include a **proc_bind** clause to
17 specify the places to use for the threads in the team within the **parallel** region.

18 When any team encounters a worksharing construct, the work inside the construct is divided among
19 the members of the team, and executed cooperatively instead of being executed by every thread.
20 There is a default barrier at the end of each worksharing construct unless the **nowait** clause is
21 present. Redundant execution of code by every thread in the team resumes after the end of the
22 worksharing construct.

23 When any thread encounters a *task generating construct*, one or more explicit tasks are generated.
24 Execution of explicitly generated tasks is assigned to one of the threads in the current team, subject
25 to the thread's availability to execute work. Thus, execution of the new task could be immediate, or
26 deferred until later according to task scheduling constraints and thread availability. Threads are
27 allowed to suspend the current task region at a task scheduling point in order to execute a different
28 task. If the suspended task region is for a tied task, the initially assigned thread later resumes
29 execution of the suspended task region. If the suspended task region is for an untied task, then any
30 thread may resume its execution. Completion of all explicit tasks bound to a given parallel region is
31 guaranteed before the master thread leaves the implicit barrier at the end of the region. Completion
32 of a subset of all explicit tasks bound to a given parallel region may be specified through the use of
33 task synchronization constructs. Completion of all explicit tasks bound to the implicit parallel
34 region is guaranteed by the time the program exits.

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

37 When a **loop** construct is encountered, the iterations of the loop associated with the construct are
38 executed in the context of its encountering thread(s), as determined according to its binding region.
39 If the **loop** region binds to a **teams** region, the region is encountered by the set of master threads
40 that execute the **teams** region. If the **loop** region binds to a **parallel** region, the region is

1 encountered by the team of threads executing the **parallel** region. Otherwise, the region is
2 encountered by a single thread.

3 If the **loop** region binds to a **teams** region, the encountering threads may continue execution
4 after the **loop** region without waiting for all iterations to complete; the iterations are guaranteed to
5 complete before the end of **teams** region. Otherwise, all iterations must complete before the
6 encountering thread(s) continue execution after the **loop** region. All threads that encounter the
7 **loop** construct may participate in the execution of the iterations. Only one of these threads may
8 execute any given iteration.

9 The **cancel** construct can alter the previously described flow of execution in an OpenMP region.
10 The effect of the **cancel** construct depends on its *construct-type-clause*. If a task encounters a
11 **cancel** construct with a **taskgroup** *construct-type-clause*, then the task activates cancellation
12 and continues execution at the end of its **task** region, which implies completion of that task. Any
13 other task in that **taskgroup** that has begun executing completes execution unless it encounters a
14 **cancellation point** construct, in which case it continues execution at the end of its **task**
15 region, which implies its completion. Other tasks in that **taskgroup** region that have not begun
16 execution are aborted, which implies their completion.

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

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

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

28 The OpenMP specification makes no guarantee that input or output to the same file is synchronous
29 when executed in parallel. In this case, the programmer is responsible for synchronizing input and
30 output statements (or routines) using the provided synchronization constructs or library routines.
31 For the case where each thread accesses a different file, no synchronization by the programmer is
32 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 corresponding to the directive is unspecified; see
20 Section 2.22.4.3 on page 280 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 2.22 on page 263.

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 is
4 the programmer's responsibility to ensure through synchronization that the lifetime of the variable
5 does not end before completion of the explicit task region sharing it. Any other access by one task
6 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.
16 Conversely the original variable becomes the new variable's corresponding variable in the device
17 data environment of the device that performs the mapping operation.

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 1.4.6 on page 27. 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. If a corresponding variable does not exist in
23 the device data environment then accesses to the original variable result in unspecified behavior
24 unless the `unified_shared_memory` requirement is specified.

25 The relationship between the value of the original variable and the initial or final value of the
26 corresponding variable depends on the *map-type*. Details of this issue, as well as other issues with
27 mapping a variable, are provided in Section 2.22.7.1 on page 307.

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

1 1.4.3 Memory Management

2 The host device, and target devices that an implementation may support, have attached storage
3 resources where program variables are stored. These resources can be of different kinds and have
4 different traits. A memory space in an OpenMP program represents a set of these storage resources.
5 Memory spaces are defined according to a set of traits and a single resource may be exposed as
6 multiple memory spaces with different traits or may be part of multiple memory spaces. In any
7 device at least one memory space is guaranteed to exist.

8 An OpenMP program can use a memory allocator to allocate *memory* to which store and from
9 which retrieve its variables. This *memory* will be allocated from the storage resources of the
10 *memory space* associated with the allocator. Memory allocators are also used to deallocate
11 previously allocated *memory*. When an OpenMP memory allocator is not used variables may be
12 allocated in any storage resource.

13 1.4.4 The Flush Operation

14 The memory model has relaxed-consistency because a thread's temporary view of memory is not
15 required to be consistent with memory at all times. A value written to a variable can remain in the
16 thread's temporary view until it is forced to memory at a later time. Likewise, a read from a variable
17 may retrieve the value from the thread's temporary view, unless it is forced to read from memory.
18 The OpenMP flush operation enforces consistency between multiple threads' view of memory.

19 If a flush operation is a strong flush, it enforces consistency between a thread's temporary view and
20 memory. A strong flush operation is applied to a set of variables called the *flush-set*. A strong flush
21 restricts reordering of memory operations that an implementation might otherwise do.
22 Implementations must not reorder the code for a memory operation for a given variable, or the code
23 for a flush operation for the variable, with respect to a strong flush operation that refers to the same
24 variable.

25 If a thread has performed a write to its temporary view of a shared variable since its last strong
26 flush of that variable, then when it executes another strong flush of the variable, the strong flush
27 does not complete until the value of the variable has been written to the variable in memory. If a
28 thread performs multiple writes to the same variable between two strong flushes of that variable,
29 the strong flush ensures that the value of the last write is written to the variable in memory. A
30 strong flush of a variable executed by a thread also causes its temporary view of the variable to be
31 discarded, so that if its next memory operation for that variable is a read, then the thread will read
32 from memory when it may again capture the value in the temporary view. When a thread executes a
33 strong flush, no later memory operation by that thread for a variable involved in that strong flush is
34 allowed to start until the strong flush completes. The completion of a strong flush of a set of
35 variables executed by a thread is defined as the point at which all writes to those variables

1 performed by the thread before the strong flush are visible in memory to all other threads and that
2 thread's temporary view of all variables involved is discarded.

3 A strong flush operation provides a guarantee of consistency between a thread's temporary view
4 and memory. Therefore, a strong flush can be used to guarantee that a value written to a variable by
5 one thread may be read by a second thread. To accomplish this, the programmer must ensure that
6 the second thread has not written to the variable since its last strong flush of the variable, and that
7 the following sequence of events are completed in the specified order:

- 8 1. The value is written to the variable by the first thread.
- 9 2. The variable is flushed, with a strong flush, by the first thread.
- 10 3. The variable is flushed, with a strong flush, by the second thread.
- 11 4. The value is read from the variable by the second thread.

12 If a flush operation is a release flush or acquire flush, it can enforce consistency between two
13 synchronizing threads' view of memory. A release flush guarantees that any prior operation that
14 writes or reads a shared variable will appear to be completed before any operation that writes or
15 reads the same shared variable that follows an acquire flush with which it synchronizes (see Section
16 Section 1.4.5 on page 26 for more details on flush synchronization). A release flush will propagate
17 the values of all shared variables in its temporary view to memory prior to the thread performing
18 any subsequent atomic operation that may establish a synchronization. An acquire flush will
19 discard any value of a shared variable in its temporary view to which the thread has not written
20 since last performing a release flush, so that it may subsequently read a value propagated by a
21 release flush that synchronizes with it. Therefore, release and acquire flushes may also be used to
22 guarantee that a value written to a variable by one thread may be read by a second thread. To
23 accomplish this, the programmer must ensure that the second thread has not written to the variable
24 since its last acquire flush, and that the following sequence of events happen in the specified order:

- 25 1. The value is written to the variable by the first thread.
- 26 2. The first thread performs a release flush.
- 27 3. The second thread performs an acquire flush.
- 28 4. The value is read from the variable by the second thread.

29 **Note** – OpenMP synchronization operations, described in Section 2.20 on page 216 and in
30 Section 3.3 on page 378, are recommended for enforcing this order. Synchronization through
31 variables is possible but is not recommended because the proper timing of flushes is difficult.

1 1.4.5 Flush Synchronization and Happens Before

2 OpenMP supports thread synchronization with the use of release flushes and acquire flushes. For
3 any such synchronization, a release flush is the source of the synchronization and an acquire flush is
4 the sink of the synchronization, such that the release flush *synchronizes with* the acquire flush.

5 A release flush has one or more associated *release sequences* that define the set of a modifications
6 that may be used to establish a synchronization. Any such release sequence starts with an atomic
7 operation that follows the release flush and modifies a shared variable and additionally includes any
8 read-modify-write atomic operations that read a value taken from some modification in the release
9 sequence. The atomic operations that start an associated release sequence are determined as
10 follows:

- 11 • If a release flush is performed on entry to an atomic operation, that atomic operation starts its
12 release sequence.
- 13 • If a release flush is performed by an implicit **flush** region, some atomic operation performed
14 by the implementation on an internal synchronization variable starts its release sequence.
- 15 • If a release flush is performed by an explicit **flush** region, any atomic operation that modifies a
16 shared variable and follows the **flush** region in its thread's program order starts an associated
17 release sequence.

18 An acquire flush is associated with one or more prior atomic operations that read a shared variable
19 and that may be used to establish a synchronization. The associated atomic operations that may
20 establish a synchronization are determined as follows:

- 21 • If an acquire flush is performed on exit from an atomic operation, that atomic operation is its
22 associated atomic operation.
- 23 • If an acquire flush is performed by an implicit **flush** region, some atomic operation performed
24 by the implementation that reads an internal synchronization variable is its associated atomic
25 operation.
- 26 • If an acquire flush is performed by an explicit **flush** region, any atomic operation that reads a
27 shared variable and precedes the **flush** region in its thread's program order is an associated
28 atomic operation.

29 A release flush synchronizes with an acquire flush if an atomic operation associated with the
30 acquire flush reads a value written by a modification from a release sequence associated with the
31 release flush.

32 An operation *X* *simply happens before* an operation *Y* if any of the following conditions are
33 satisfied:

- 34 1. *X* and *Y* are performed by the same thread, and *X* precedes *Y* in the thread's program order.
- 35 2. *X* synchronizes with *Y* according to the flush synchronization conditions explained above or
36 according to the base language's definition of *synchronizes with*, if such a definition exists.

1 3. There exists another operation Z , such that X simply happens before Z and Z simply happens
2 before Y .

3 An operation X *happens before* an operation Y if any of the following conditions are satisfied:

4 1. X happens before Y according to the base language's definition of *happens before*, if such a
5 definition exists.

6 2. X simply happens before Y .

7 A variable with an initial value is treated as if the value is stored to the variable by an operation that
8 happens before all operations that access or modify the variable in the program.

9 1.4.6 OpenMP Memory Consistency

10 The observable completion order of memory operations, as seen by all threads, is guaranteed
11 according to the following rules:

- 12 • If two operations performed by different threads are sequentially consistent atomic operations or
13 they are strong flushes that flush the same variable, then they must be completed as if in some
14 sequential order, seen by all threads.
- 15 • If two operations performed by the same thread are sequentially consistent atomic operations or
16 they access, modify, or, with a strong flush, flush the same variable, then they must be completed
17 as if in that thread's program order, as seen by all threads.
- 18 • If two operations are performed by different threads and one happens before the other, then they
19 must be completed as if in that happens before order, as seen by all threads, if:
 - 20 – both operations access or modify the same variable,
 - 21 – both operations are strong flushes that flush the same variable, or
 - 22 – both operations are sequentially consistent atomic operations.
- 23 • Any two atomic memory operations from different **atomic** regions must be completed as if in
24 the same order as the strong flushes implied in their respective regions, as seen by all threads.

25 The flush operation can be specified using the **flush** directive, and is also implied at various
26 locations in an OpenMP program: see Section 2.20.8 on page 235 for details.

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

- 1 OpenMP programs that:
- 2 • do not use non-sequentially consistent atomic directives,
 - 3 • do not rely on the accuracy of a *false* result from `omp_test_lock` and
 - 4 `omp_test_nest_lock`, and
 - 5 • correctly avoid data races as required in Section 1.4.1 on page 22
- 6 behave as though operations on shared variables were simply interleaved in an order consistent with
- 7 the order in which they are performed by each thread. The relaxed consistency model is invisible
- 8 for such programs, and any explicit flush operations in such programs are redundant.

9 1.5 Tool Interface

10 To enable development of high-quality, portable, tools that support monitoring, performance, or

11 correctness analysis and debugging of OpenMP programs developed using any implementation of

12 the OpenMP API, the OpenMP API includes two tool interfaces: OMPT and OMPD.

13 1.5.1 OMPT

14 The OMPT interface, which is intended for *first-party* tools, provides the following:

- 15 • a mechanism to initialize a first-party tool,
- 16 • routines that enable a tool to determine the capabilities of an OpenMP implementation,
- 17 • routines that enable a tool to examine OpenMP state information associated with a thread,
- 18 • mechanisms that enable a tool to map implementation-level calling contexts back to their
- 19 source-level representations,
- 20 • a callback interface that enables a tool to receive notification of OpenMP *events*,
- 21 • a tracing interface that enables a tool to trace activity on OpenMP target devices, and
- 22 • a runtime library routine that an application can use to control a tool.

23 OpenMP implementations may differ with respect to the *thread states* that they support, the mutual

24 exclusion implementations they employ, and the OpenMP events for which tool callbacks are

25 invoked. For some OpenMP events, OpenMP implementations must guarantee that a registered

26 callback will be invoked for each occurrence of the event. For other OpenMP events, OpenMP

1 implementations are permitted to invoke a registered callback for some or no occurrences of the
2 event; for such OpenMP events, however, OpenMP implementations are encouraged to invoke tool
3 callbacks on as many occurrences of the event as is practical to do so. Section 4.2.1.3 specifies the
4 subset of OMPT callbacks that an OpenMP implementation must support for a minimal
5 implementation of the OMPT interface.

6 An implementation of the OpenMP API may differ from the abstract execution model described by
7 its specification. The ability of tools using the OMPT interface to observe such differences does not
8 constrain implementations of the OpenMP API in any way.

9 With the exception of the `omp_control_tool` runtime library routine for tool control, all other
10 routines in the OMPT interface are intended for use only by tools and are not visible to
11 applications. For that reason, a Fortran binding is provided only for `omp_control_tool`; all
12 other OMPT functionality is described with C syntax only.

13 1.5.2 OMPD

14 The OMPD interface is intended for a *third-party* tool, which runs as a separate process. An
15 OpenMP implementation must provide an OMPD library that can be dynamically loaded and used
16 by a third-party tool. A third-party tool, such as a debugger, uses the OMPD library to access
17 OpenMP state of a program that has begun execution. OMPD defines the following:

- 18 • an interface that an OMPD library exports, which a tool can use to access OpenMP state of a
19 program that has begun execution;
- 20 • a callback interface that a tool provides to the OMPD library so that the library can use it to
21 access OpenMP state of a program that has begun execution; and
- 22 • a small number of symbols that must be defined by an OpenMP implementation to help the tool
23 find the correct OMPD library to use for that OpenMP implementation and to facilitate
24 notification of events.

25 OMPD is described in Chapter 4.

1 1.6 OpenMP Compliance

2 The OpenMP API defines constructs that operate in the context of the base language that is
3 supported by an implementation. If the implementation of the base language does not support a
4 language construct that appears in this document, a compliant OpenMP implementation is not
5 required to support it, with the exception that for Fortran, the implementation must allow case
6 insensitivity for directive and API routines names, and must allow identifiers of more than six
7 characters. An implementation of the OpenMP API is compliant if and only if it compiles and
8 executes all other conforming programs, and supports the tool interface, according to the syntax
9 and semantics laid out in Chapters 1, 2, 3, 4 and 5. Appendices A, B, C, D, and E, as well as
10 sections designated as Notes (see Section 1.8 on page 33) are for information purposes only and are
11 not part of the specification.

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

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

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

23 1.7 Normative References

- 24 • ISO/IEC 9899:1990, *Information Technology - Programming Languages - C*.

25 This OpenMP API specification refers to ISO/IEC 9899:1990 as C90.

- 26 • ISO/IEC 9899:1999, *Information Technology - Programming Languages - C*.

27 This OpenMP API specification refers to ISO/IEC 9899:1999 as C99.

- 28 • ISO/IEC 9899:2011, *Information Technology - Programming Languages - C*.

29 This OpenMP API specification refers to ISO/IEC 9899:2011 as C11. The following features are
30 not supported:

- 31 – Supporting the noreturn property

- 1 – Adding alignment support
- 2 – Creation of complex value
- 3 – Abandoning a process (adding `quick_exit` and `at_quick_exit`)
- 4 – Threads for the C standard library
- 5 – Thread-local storage
- 6 – Parallel memory sequencing model
- 7 – Atomic
- 8 ● ISO/IEC 14882:1998, *Information Technology - Programming Languages - C++*.
- 9 This OpenMP API specification refers to ISO/IEC 14882:1998 as C++.
- 10 ● ISO/IEC 14882:2011, *Information Technology - Programming Languages - C++*.
- 11 This OpenMP API specification refers to ISO/IEC 14882:2011 as C++11. The following features
- 12 are not supported:
 - 13 – Alignment support
 - 14 – Standard layout types
 - 15 – Allowing move constructs to throw
 - 16 – Defining move special member functions
 - 17 – Concurrency
 - 18 – Data-dependency ordering: atomics and memory model
 - 19 – Additions to the standard library
 - 20 – Thread-local storage
 - 21 – Dynamic initialization and destruction with concurrency
 - 22 – C++11 library
- 23 ● ISO/IEC 14882:2014, *Information Technology - Programming Languages - C++*.
- 24 This OpenMP API specification refers to ISO/IEC 14882:2014 as C++14. The following features
- 25 are not supported:
 - 26 – Sized deallocation
 - 27 – What signal handlers can do
- 28 ● ISO/IEC 14882:2017, *Information Technology - Programming Languages - C++*.
- 29 This OpenMP API specification refers to ISO/IEC 14882:2017 as C++17.

- 1 ● ISO/IEC 1539:1980, *Information Technology - Programming Languages - Fortran*.
2 This OpenMP API specification refers to ISO/IEC 1539:1980 as Fortran 77.
- 3 ● ISO/IEC 1539:1991, *Information Technology - Programming Languages - Fortran*.
4 This OpenMP API specification refers to ISO/IEC 1539:1991 as Fortran 90.
- 5 ● ISO/IEC 1539-1:1997, *Information Technology - Programming Languages - Fortran*.
6 This OpenMP API specification refers to ISO/IEC 1539-1:1997 as Fortran 95.
- 7 ● ISO/IEC 1539-1:2004, *Information Technology - Programming Languages - Fortran*.
8 This OpenMP API specification refers to ISO/IEC 1539-1:2004 as Fortran 2003.
- 9 ● ISO/IEC 1539-1:2010, *Information Technology - Programming Languages - Fortran*.
10 This OpenMP API specification refers to ISO/IEC 1539-1:2010 as Fortran 2008. The following
11 features are not supported:
- 12 – Submodules
 - 13 – Coarrays
 - 14 – DO CONCURRENT
 - 15 – Allocatable components of recursive type
 - 16 – Pointer initialization
 - 17 – Value attribute is permitted for any nonallocatable nonpointer nonarray
 - 18 – Simply contiguous arrays rank remapping to rank>1 target
 - 19 – Polymorphic assignment
 - 20 – Accessing real and imaginary parts
 - 21 – Pointer function reference is a variable
 - 22 – Recursive I/O
 - 23 – The BLOCK construct
 - 24 – EXIT statement
 - 25 – ERROR STOP
 - 26 – Internal procedure as an actual argument
 - 27 – Generic resolution by procedureness
 - 28 – Generic resolution by pointer vs. allocatable
 - 29 – Impure elemental procedures

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

3 1.8 Organization of this Document

4 The remainder of this document is structured as follows:

- 5 • Chapter 2 “Directives”
- 6 • Chapter 3 “Runtime Library Routines”
- 7 • Chapter 4 “Tool Support”
- 8 • Chapter 5 “Environment Variables”
- 9 • Appendix A “OpenMP Implementation-Defined Behaviors”
- 10 • Appendix B “Task Frame Management for the Tool Interface”
- 11 • Appendix C “Interaction Diagram of OMPD Components”
- 12 • Appendix D “Features History”

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

15  C / C++
C/C++ specific text...

16  C / C++

16 Text that applies only to programs for which the base language is C only is shown as follows:

17  C
C specific text...

18  C

18 Text that applies only to programs for which the base language is C90 only is shown as follows:

1 C90 specific text... C90

2 C90

3 Text that applies only to programs for which the base language is C99 only is shown as follows:

4 C99 specific text... C99

5 C99

6 Text that applies only to programs for which the base language is C++ only is shown as follows:

7 C++ specific text... C++

8 C++

9 Text that applies only to programs for which the base language is Fortran is shown as follows:

10 Fortran specific text..... Fortran

11 Fortran

12 Where an entire page consists of base language specific text, a marker is shown at the top of the page. For Fortran-specific text, the marker is:

13 Fortran (cont.)

14 For C/C++-specific text, the marker is:

15 C/C++ (cont.)

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



17 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 36)
- 6 • Mechanisms to control conditional compilation (Section 2.2 on page 41)
- 7 • Control of OpenMP API ICVs (Section 2.4 on page 47)
- 8 • How to specify and to use array sections for all base languages (Section 2.6 on page 59)
- 9 • Details of each OpenMP directive, including associated events and tool callbacks (Section 2.9 on
10 page 72 to Section 2.24 on page 327)

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

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

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

Fortran

Restrictions

The following restriction applies to all OpenMP directives:

- OpenMP directives, except SIMD and **declare target** directives, may not appear in pure procedures.

Fortran

2.1 Directive Format

C / C++

OpenMP directives for C/C++ are specified with the **pragma** preprocessing directive. The syntax of an OpenMP directive is as follows:

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

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

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

Directives are case-sensitive.

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

C / C++

C++

Directives may not appear in constexpr functions or in constant expressions. Variadic parameter packs cannot be expanded into a directive or its clauses except as part of an expression argument to be evaluated by the base language, such as into a function call inside an **if()** clause.

C++

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 38 and Section 2.1.2 on page 39.

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 2.16 on page 185). 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 2.22.4 on page 276), data copying clauses (Section 2.22.6 on page 301), the **threadprivate** directive (Section 2.22.2 on page 268), the **flush** directive (Section 2.20.8 on page 235), and the **link** clause of the **declare target** directive (Section 2.15.7 on page 178) accept a *list*. The **to** clause of the **declare target** directive (Section 2.15.7 on page 178) accepts an *extended-list*. The **depend** clause (Section 2.20.11 on page 248), when used to specify task dependences, accepts a *locator-list*. A *list* consists of a comma-separated collection of one or more *list items*. A *extended-list* consists of a comma-separated collection of one or more *extended list items*. A *locator-list* consists of a comma-separated collection of one or more *locator list items*.

C / C++

A *list item* is a variable or array section. An *extended list item* is a *list item* or a function name. A *locator list item* is any *lvalue* expression, including variables, or an array section.

C / C++

Fortran

1 A *list item* is a variable, array section or common block name (enclosed in slashes). An *extended*
2 *list item* is a *list item* or a procedure name. A *locator list item* is a *list item*.

3 When a named common block appears in a *list*, it has the same meaning as if every explicit member
4 of the common block appeared in the list. An explicit member of a common block is a variable that
5 is named in a **COMMON** statement that specifies the common block name and is declared in the same
6 scoping unit in which the clause appears.

7 Although variables in common blocks can be accessed by use association or host association,
8 common block names cannot. As a result, a common block name specified in a data-sharing
9 attribute, a data copying or a data-mapping attribute clause must be declared to be a common block
10 in the same scoping unit in which the clause appears.

11 If a list item that appears in a directive or clause is an optional dummy argument that is not present,
12 the directive or clause for that list item is ignored.

13 If the variable referenced inside a construct is an optional dummy argument that is not present, any
14 explicitly determined, implicitly determined, or predetermined data-sharing and data-mapping
15 attribute rules for that variable are ignored. Otherwise, if the variable is an optional dummy
16 argument that is present, it is present inside the construct.

Fortran

17 For all base languages, a *list item* or an *extended list item* is subject to the restrictions specified in
18 Section 2.6 on page 59 and in each of the sections describing clauses and directives for which the
19 *list* or *extended-list* appears.

Fortran

20 2.1.1 Fixed Source Form Directives

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

22 **!\$omp | c\$omp | *\$omp**

23 Sentinels must start in column 1 and appear as a single word with no intervening characters.
24 Fortran fixed form line length, white space, continuation, and column rules apply to the directive
25 line. Initial directive lines must have a space or zero in column 6, and continuation directive lines
26 must have a character other than a space or a zero in column 6.

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

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

```

3 c23456789
4 !$omp parallel do shared(a,b,c)
5
6 c$omp parallel do
7 c$omp+shared(a,b,c)
8
9 c$omp paralleldoshared(a,b,c)
```

10 2.1.2 Free Source Form Directives

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

```
12 !$omp
```

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

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

23 One or more blanks or horizontal tabs are optional to separate adjacent keywords in
 24 *directive-names* unless otherwise specified.

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

```
!23456789
    !$omp parallel do &
        !$omp shared(a,b,c)

    !$omp parallel &
    !$omp&do shared(a,b,c)

!$omp paralleldo shared(a,b,c)
```

Fortran

2.1.3 Stand-Alone Directives

Summary

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

Description

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

Restrictions

C / C++

For C/C++, a stand-alone directive may not be used in place of the statement following an **if**, **while**, **do**, **switch**, or **label**.

C / C++

Fortran

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

Fortran

3 2.2 Conditional Compilation

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

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

Fortran

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

11 2.2.1 Fixed Source Form Conditional Compilation 12 Sentinels

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

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

15 To enable conditional compilation, a line with a conditional compilation sentinel must satisfy the
16 following criteria:

- 17
- The sentinel must start in column 1 and appear as a single word with no intervening white space.
 - 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.
 - After the sentinel is replaced with two spaces, continuation lines must have a character other than a space or zero in column 6 and only white space in columns 1 through 5.

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

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

```

1  c23456789
2  !$ 10 iam = omp_get_thread_num() +
3  !$   &           index
4
5  #ifdef _OPENMP
6     10 iam = omp_get_thread_num() +
7     &           index
8  #endif
9
10

```

2.2.2 Free Source Form Conditional Compilation Sentinel

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

```
!$
```

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

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

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

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

```
3 c23456789  
4 !$ iam = omp_get_thread_num() + &  
5 !$& index  
6  
7 #ifdef _OPENMP  
8     iam = omp_get_thread_num() + &  
9     index  
10 #endif
```

Fortran

11 2.3 requires Directive

12 Summary

13 The **requires** directive specifies the features an implementation must provide in order for the
14 code to compile and to execute correctly. The **requires** directive is a declarative directive.

15 Syntax

C / C++

16 The syntax of the **requires** directive is as follows:

```
17 #pragma omp requires clause[ [ [, ] clause ] ... ] new-line
```

C / C++

Fortran

18 The syntax of the **requires** directive is as follows:

```
19 !$omp requires clause[ [ [, ] clause ] ... ]
```

Fortran

1 Where *clause* is either one of the requirement clauses listed below or a clause of the form
2 **ext_implementation-defined-requirement** for an implementation defined requirement clause.

```
3 reverse_offload  
4 unified_address  
5 unified_shared_memory  
6 atomic_default_mem_order(seq_cst | acq_rel | relaxed)  
7 dynamic_allocators
```

8 **Description**

9 The **requires** directive specifies features an implementation must support for correct execution.
10 The behavior specified by a requirement clause may override the normal behavior specified
11 elsewhere in this document.

▼ C / C++ ▼

12 The **requires** directive specifies requirements for the execution of all code in the current
13 translation unit.

▲ C / C++ ▲

▼ Fortran ▼

14 The **requires** directive specifies requirements for the execution of all code in the current
15 program unit.

▲ Fortran ▲

1 **Note** – Use of this directive makes your code less portable. Users should be aware that not all
2 devices or implementations support all requirements.

3 When the **reverse_offload** clause appears on a **requires** directive, the implementation
4 guarantees that a **target** region, for which the **target** construct specifies a **device** clause in
5 which the **ancestor** modifier appears, can execute on the parent device of an enclosing **target**
6 region.

7 When the **unified_address** clause appears on a **requires** directive, the implementation
8 guarantees that all devices accessible through OpenMP API routines and directives use a unified
9 address space. In this address space, a pointer will always refer to the same location in memory
10 from all devices accessible through OpenMP. The pointers returned by **omp_target_alloc** and
11 accessed through **use_device_ptr** are guaranteed to be pointer values that can support pointer
12 arithmetic while still being native device pointers. The **is_device_ptr** clause is not necessary
13 for device pointers to be translated in **target** regions, and pointers found not present are not set to
14 null but keep their original value. Memory local to a specific execution context may be exempt
15 from this, following the restrictions of locality to a given execution context, thread or contention
16 group. Target devices may still have discrete memories and dereferencing a device pointer on the
17 host device remains unspecified behavior.

18 The **unified_shared_memory** clause implies the **unified_address** requirement,
19 inheriting all of its behaviors. Additionally memory in the device data environment of any device
20 visible to OpenMP, including but not limited to the host, is considered part of the device data
21 environment of all devices accessible through OpenMP except as noted below. Every device address
22 allocated through OpenMP device memory routines is a valid host pointer. Memory local to an
23 execution context as defined in **unified_address** above may remain part of distinct device data
24 environments as long as the execution context is local to the device containing that environment.

25 The **unified_shared_memory** clause makes the **map** clause optional on **target** constructs
26 as well as the **declare target** directive on static lifetime variables accessed as part of
27 **declare target** functions. Scalar variables are still made **firstprivate** by default for
28 **target** regions. Values stored into memory by one device may not be visible to other devices
29 until those two devices synchronize with each other or both synchronize with the host.

30 The **atomic_default_mem_order** clause specifies the default memory ordering behavior for
31 **atomic** constructs that must be provided by an implementation. If the default memory ordering is
32 specified as **seq_cst**, all **atomic** constructs on which *memory-order-clause* is not specified
33 behave as if the **seq_cst** clause appears. If the default memory ordering is specified as
34 **relaxed**, all **atomic** constructs on which *memory-order-clause* is not specified behave as if the
35 **relaxed** clause appears.

36 If the default memory ordering is specified as **acq_rel**, **atomic** constructs on which
37 *memory-order-clause* is not specified behave in the following manner:

- as if the **release** clause is present if the construct specifies an atomic write or atomic update operation;
- as if the **acquire** clause is present if the construct specifies an atomic read operation;
- as if the **acq_rel** clause is present if the construct specifies an atomic captured update operation.

The **dynamic_allocators** clause has the following effects:

- makes the **uses_allocators** clause optional on **target** constructs for the purpose of using allocators in the corresponding **target** regions,
- allows the **omp_init_allocator** and **omp_destroy_allocator** API routines in **target** regions,
- allows default allocators to be used by **allocate** directives, **allocate** clauses and **omp_alloc** API routines in **target** regions.

Implementers are allowed to include additional implementation defined requirement clauses. Requirement names that do not start with **ext_** are reserved. All implementation-defined requirements should begin with **ext_**.

Restrictions

The restrictions for the **requires** directive are as follows:

- Each of the clauses can appear at most once on the directive.
- At most one **requires** directive with **atomic_default_mem_order** clause can appear in a single compilation unit.
- A **requires** directive with a **unified_address**, **unified_shared_memory** or **reverse_offload** clause shall appear lexically before any device constructs or device routines.
- A **requires** directive with the **unified_shared_memory** clause must appear in all *compilation units* of a program that contain device constructs or device routines or in none of them.
- A **requires** directive with the **reverse_offload** clause must appear in all *compilation units* of a program that contain device constructs or device routines or in none of them.
- The **requires** directive with **atomic_default_mem_order** clause may not appear lexically after any **atomic** construct on which *memory-order-clause* is not specified.

1 2.4 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 2.4.2 on page 49.

12 2.4.1 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. The *nest-var* ICV has been deprecated.
- 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.

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

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

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

10 The following ICVs store values that affect program execution.

- 11 • *stacksize-var* - controls the stack size for threads that the OpenMP implementation creates. There
12 is one copy of this ICV per device.
- 13 • *wait-policy-var* - controls the desired behavior of waiting threads. There is one copy of this ICV
14 per device.
- 15 • *display-affinity-var* - controls whether to display thread affinity. There is one copy of this ICV for
16 the whole program.
- 17 • *affinity-format-var* - controls the thread affinity format when displaying thread affinity. There is
18 one copy of this ICV per device.
- 19 • *cancel-var* - controls the desired behavior of the **cancel** construct and cancellation points.
20 There is one copy of this ICV for the whole program.
- 21 • *default-device-var* - controls the default target device. There is one copy of this ICV per data
22 environment.
- 23 • *target-offload-var* - controls the offloading behavior. There is one copy of this ICV for the whole
24 program.
- 25 • *max-task-priority-var* - controls the maximum priority value that can be specified in the
26 **priority** clause of the **task** construct. There is one copy of this ICV for the whole program.

27 The following ICVs store values that affect the operation of the first-party tool interface.

- 28 • *tool-var* - determines whether an OpenMP implementation will try to register a tool. There is
29 one copy of this ICV for the whole program.
- 30 • *tool-libraries-var* - specifies a list of absolute paths to tool libraries for OpenMP devices. There
31 is one copy of this ICV for the whole program.

32 The following ICVs store values that relate to the operation of the OMPD tool interface.

- 33 • *debug-var* - determines whether an OpenMP implementation will collect information that an
34 OMPD library can access to satisfy requests from a tool. There is one copy of this ICV for the
35 whole program.

- 1 The following ICVs store values that affect default memory allocation.
- 2 • *def-allocator-var* - determines the memory allocator to be used by memory allocation routines,
- 3 directives and clauses when a memory allocator is not specified by the user. There is one copy of
- 4 this ICV per implicit task.

5 2.4.2 ICV Initialization

6 Table 2.1 shows the ICVs, associated environment variables, and initial values.

TABLE 2.1: ICV Initial Values

ICV	Environment Variable	Initial value
<i>dyn-var</i>	OMP_DYNAMIC	See description below
<i>nest-var</i>	OMP_NESTED	Implementation defined
<i>nthreads-var</i>	OMP_NUM_THREADS	Implementation defined
<i>run-sched-var</i>	OMP_SCHEDULE	Implementation defined
<i>def-sched-var</i>	(none)	Implementation defined
<i>bind-var</i>	OMP_PROC_BIND	Implementation defined
<i>stacksize-var</i>	OMP_STACKSIZE	Implementation defined
<i>wait-policy-var</i>	OMP_WAIT_POLICY	Implementation defined
<i>thread-limit-var</i>	OMP_THREAD_LIMIT	Implementation defined
<i>max-active-levels-var</i>	OMP_MAX_ACTIVE_LEVELS	See description below
<i>active-levels-var</i>	(none)	<i>zero</i>
<i>levels-var</i>	(none)	<i>zero</i>
<i>place-partition-var</i>	OMP_PLACES	Implementation defined
<i>cancel-var</i>	OMP_CANCELLATION	<i>false</i>
<i>display-affinity-var</i>	OMP_DISPLAY_AFFINITY	<i>false</i>

table continued on next page

table continued from previous page

ICV	Environment Variable	Initial value
<i>affinity-format-var</i>	OMP_AFFINITY_FORMAT	Implementation defined
<i>default-device-var</i>	OMP_DEFAULT_DEVICE	Implementation defined
<i>target-offload-var</i>	OMP_TARGET_OFFLOAD	DEFAULT
<i>max-task-priority-var</i>	OMP_MAX_TASK_PRIORITY	<i>zero</i>
<i>tool-var</i>	OMP_TOOL	<i>enabled</i>
<i>tool-libraries-var</i>	OMP_TOOL_LIBRARIES	<i>empty string</i>
<i>debug-var</i>	OMP_DEBUG	<i>disabled</i>
<i>def-allocator-var</i>	OMP_ALLOCATOR	Implementation defined

Description

- 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 15 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 5.1 on page 596.
- **OMP_NUM_THREADS** environment variable, see Section 5.2 on page 597.
- **OMP_DYNAMIC** environment variable, see Section 5.3 on page 598.
- **OMP_PROC_BIND** environment variable, see Section 5.4 on page 598.
- **OMP_PLACES** environment variable, see Section 5.5 on page 599.

- 1 • **OMP_NESTED** environment variable, see Section 5.6 on page 601.
- 2 • **OMP_STACKSIZE** environment variable, see Section 5.7 on page 602.
- 3 • **OMP_WAIT_POLICY** environment variable, see Section 5.8 on page 603.
- 4 • **OMP_MAX_ACTIVE_LEVELS** environment variable, see Section 5.9 on page 603.
- 5 • **OMP_THREAD_LIMIT** environment variable, see Section 5.10 on page 604.
- 6 • **OMP_CANCELLATION** environment variable, see Section 5.11 on page 604.
- 7 • **OMP_DISPLAY_AFFINITY** environment variable, see Section 5.13 on page 606.
- 8 • **OMP_AFFINITY_FORMAT** environment variable, see Section 5.14 on page 607.
- 9 • **OMP_DEFAULT_DEVICE** environment variable, see Section 5.15 on page 609.
- 10 • **OMP_TARGET_OFFLOAD** environment variable, see Section 5.17 on page 610.
- 11 • **OMP_MAX_TASK_PRIORITY** environment variable, see Section 5.16 on page 609.
- 12 • **OMP_TOOL** environment variable, see Section 5.18 on page 611.
- 13 • **OMP_TOOL_LIBRARIES** environment variable, see Section 5.19 on page 611.
- 14 • **OMP_DEBUG** environment variable, see Section 5.20 on page 612.
- 15 • **OMP_ALLOCATOR** environment variable, see Section 5.21 on page 612.

16 2.4.3 Modifying and Retrieving ICV Values

17 Table 2.2 shows the method for modifying and retrieving the values of ICVs through OpenMP API
 18 routines.

TABLE 2.2: Ways to Modify and to Retrieve ICV Values

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>

table continued on next page

table continued from previous page

ICV	Ways to modify value	Ways to retrieve value
<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>	<code>thread_limit</code> 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_level()</code>
<i>levels-var</i>	(none)	<code>omp_get_level()</code>
<i>place-partition-var</i>	(none)	See description below
<i>cancel-var</i>	(none)	<code>omp_get_cancellation()</code>
<i>display-affinity-var</i>	(none)	(none)
<i>affinity-format-var</i>	<code>omp_set_affinity_format()</code>	<code>omp_get_affinity_format()</code>
<i>default-device-var</i>	<code>omp_set_default_device()</code>	<code>omp_get_default_device()</code>
<i>target-offload-var</i>	(none)	(none)
<i>max-task-priority-var</i>	(none)	<code>omp_get_max_task_priority()</code>
<i>tool-var</i>	(none)	(none)
<i>tool-libraries-var</i>	(none)	(none)
<i>debug-var</i>	(none)	(none)
<i>def-allocator-var</i>	<code>omp_set_default_allocator()</code>	<code>omp_get_default_allocator()</code>

Description

- 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.
- Detailed values in the *place-partition-var* ICV are retrieved using the runtime calls `omp_get_partition_num_places`, `omp_get_partition_place_nums`, `omp_get_place_num_procs`, and `omp_get_place_proc_ids`.

Cross References

- `thread_limit` clause of the `teams` construct, see Section 2.10 on page 81.
- `omp_set_num_threads` routine, see Section 3.2.1 on page 332.
- `omp_get_max_threads` routine, see Section 3.2.3 on page 334.
- `omp_set_dynamic` routine, see Section 3.2.7 on page 338.
- `omp_get_dynamic` routine, see Section 3.2.8 on page 339.
- `omp_get_cancellation` routine, see Section 3.2.9 on page 340.
- `omp_set_nested` routine, see Section 3.2.10 on page 341.
- `omp_get_nested` routine, see Section 3.2.11 on page 342.
- `omp_set_schedule` routine, see Section 3.2.12 on page 343.
- `omp_get_schedule` routine, see Section 3.2.13 on page 345.
- `omp_get_thread_limit` routine, see Section 3.2.14 on page 346.
- `omp_set_max_active_levels` routine, see Section 3.2.15 on page 347.
- `omp_get_max_active_levels` routine, see Section 3.2.16 on page 348.
- `omp_get_level` routine, see Section 3.2.17 on page 349.
- `omp_get_active_level` routine, see Section 3.2.20 on page 353.
- `omp_get_proc_bind` routine, see Section 3.2.22 on page 354.
- `omp_get_place_num_procs` routine, see Section 3.2.24 on page 357.
- `omp_get_place_proc_ids` routine, see Section 3.2.25 on page 358.
- `omp_get_partition_num_places` routine, see Section 3.2.27 on page 360.
- `omp_get_partition_place_nums` routine, see Section 3.2.28 on page 360.
- `omp_set_affinity_format` routine, see Section 3.2.29 on page 361.
- `omp_get_affinity_format` routine, see Section 3.2.30 on page 363.
- `omp_set_default_device` routine, see Section 3.2.33 on page 367.
- `omp_get_default_device` routine, see Section 3.2.34 on page 368.
- `omp_get_max_task_priority` routine, see Section 3.2.41 on page 374.
- `omp_set_default_allocator` routine, see Section 3.7.4 on page 408.
- `omp_get_default_allocator` routine, see Section 3.7.5 on page 409.

1 2.4.4 How ICVs are Scoped

2 Table 2.3 shows the ICVs and their scope.

TABLE 2.3: Scopes of ICVs

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>	global
<i>display-affinity-var</i>	global
<i>affinity-format-var</i>	device
<i>default-device-var</i>	data environment
<i>target-offload-var</i>	global
<i>max-task-priority-var</i>	global
<i>tool-var</i>	global
<i>tool-libraries-var</i>	global

table continued on next page

table continued from previous page

ICV	Scope
<i>debug-var</i>	global
<i>third-party-tool-var</i>	global
<i>def-allocator-var</i>	implicit task

1 Description

- 2 • There is one copy per device of each ICV with device scope
- 3 • Each data environment has its own copies of ICVs with data environment scope
- 4 • Each implicit task has its own copy of ICVs with implicit task scope

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

7 2.4.4.1 How the Per-Data Environment ICVs Work

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

10 When a **parallel** construct is encountered, the value of each ICV with implicit task scope is
11 inherited, unless otherwise specified, from the implicit binding task of the generating task unless
12 otherwise specified.

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

20 When a *target task* executes a **target** region, the generated initial task uses the values of the data
21 environment scoped ICVs from the device data environment ICV values of the device that will
22 execute the region.

23 If a **teams** construct with a **thread_limit** clause is encountered, the *thread-limit-var* ICV of
24 the construct's data environment is instead set to a value that is less than or equal to the value
25 specified in the clause.

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

4 2.4.5 ICV Override Relationships

5 Table 2.4 shows the override relationships among construct clauses and ICVs.

TABLE 2.4: ICV Override Relationships

ICV	construct clause, if used
<i>dyn-var</i>	(none)
<i>nest-var</i>	(none)
<i>nthreads-var</i>	num_threads
<i>run-sched-var</i>	schedule
<i>def-sched-var</i>	schedule
<i>bind-var</i>	proc_bind
<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)
<i>levels-var</i>	(none)
<i>place-partition-var</i>	(none)
<i>cancel-var</i>	(none)
<i>display-affinity-var</i>	(none)
<i>affinity-format-var</i>	(none)

table continued on next page

table continued from previous page

ICV	construct clause, if used
<i>default-device-var</i>	(none)
<i>target-offload-var</i>	(none)
<i>max-task-priority-var</i>	(none)
<i>tool-var</i>	(none)
<i>tool-libraries-var</i>	(none)
<i>debug-var</i>	(none)
<i>def-allocator-var</i>	allocator

1 **Description**

- 2 • The **num_threads** clause overrides the value of the first element of the *nthreads-var* ICV.
- 3 • If *bind-var* is not set to *false* then the **proc_bind** clause overrides the value of the first element
- 4 of the *bind-var* ICV; otherwise, the **proc_bind** clause has no effect.

5 **Cross References**

- 6 • **parallel** construct, see Section 2.9 on page 72.
- 7 • **proc_bind** clause, Section 2.9 on page 72.
- 8 • **num_threads** clause, see Section 2.9.1 on page 77.
- 9 • Worksharing-Loop construct, see Section 2.12.2 on page 100.
- 10 • **schedule** clause, see Section 2.12.2.1 on page 110.

2.5 Array Shaping

If an expression has a pointer to T type, then a shape-operator can be used to specify the extent of that pointer. In other words, the shape-operator is used to reinterpret, as an n-dimensional array, the region of memory to which that expression points.

Formally, the syntax of the shape-operator is as follows:

```
shaped-expression := ( [s1] [s2] . . . [sn] ) expression
```

The result of applying the shape-operator to an expression is an lvalue expression with an n-dimensional array type with dimensions $s_1 \times s_2 \dots \times s_n$ and element type T .

The precedence of the shape operator is the same as a type cast.

Each s_i is an integral type expression that must evaluate to a positive integer.

Restrictions

Restrictions on the shape-operator are as follows:

- The T type must be a complete type.
- The shape-operator can appear only in clauses where it is explicitly allowed.
- The type of the expression upon which a shape-operator is applied must be a pointer type.

C++

- If the T type is a reference to a type T' then the type will be considered to be T' for all purposes of the designated array.

C++

C / C++

1 2.6 Array Sections

2 An array section designates a subset of the elements in an array.

C / C++

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

```
5 [ lower-bound : length : stride ] or  
6 [ lower-bound : length : ] or  
7 [ lower-bound : length ] or  
8 [ lower-bound : : stride ] or  
9 [ lower-bound : : ] or  
10 [ lower-bound : ] or  
11 [ : length : stride ] or  
12 [ : length : ] or  
13 [ : length ] or  
14 [ : : stride ]  
15 [ : : ]  
16 [ : ]
```

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

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

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

```
22 { lower-bound, lower-bound + stride, lower-bound + 2 * stride, ... , lower-bound + ((length - 1) *  
23 stride) }
```

24 The *length* must evaluate to a non-negative integer.

25 The *stride* must evaluate to a positive integer.

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

27 When the *stride* is absent it defaults to 1.

28 When the *length* is absent, it defaults to $(size - lower-bound) / stride$ where *size* is the size of the
29 array dimension

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

1 The precedence of an array section is the same as the subscript operator.

2 Note – The following are examples of array sections:

```
3 a[0:6]
4 a[:6]
5 a[1:10]
6 a[1:]
7 b[10][:][:0]
8 c[1:10][42][0:6]
9 S.c[:100]
10 p->y[:10]
11 this->a[:N]
```

12 The first two examples are equivalent. If **a** is declared to be an eleven element array, the third and
13 fourth examples are equivalent. The fifth example is a zero-length array section. The sixth example
14 is not contiguous. The remaining examples show array sections that are formed from more general
15 base expressions.

▲ C / C++ ▼

▼ Fortran ▲

16 Fortran has built-in support for array sections although some restrictions apply to their use, as
17 enumerated in the following section.

▲ Fortran ▼

18 Restrictions

19 Restrictions to array sections are as follows:

- 20 • An array section can appear only in clauses where it is explicitly allowed.
- 21 • A *stride* expression may not be specified unless otherwise stated.

▼ C / C++ ▼

- 22 • An element of an array section with a non-zero size must have a complete type.
- 23 • The type of the base expression appearing in an array section must be an array or pointer type.

▲ C / C++ ▲

C++

- 1 • If the type of the base expression of an array section is a reference to a type *T* then the type will
- 2 be considered to be *T* for all purposes of the array section.
- 3 • An array section cannot be used in a C++ user-defined []-operator.

C++

Fortran

- 4 • If a stride expression is specified, it must be positive.
- 5 • The upper bound for the last dimension of an assumed-size dummy array must be specified.
- 6 • If a list item is an array section with vector subscripts, the first array element must be the lowest
- 7 in the array element order of the array section.

Fortran

8 2.7 Iterators

9 Iterators are identifiers that expand to multiple values in the clause on which they appear.

10 The syntax of an *iterators-definition* is the following:

11 `iterator-specifier [, iterators-definition]`

12 The syntax of an *iterator-specifier* is one of the following:

13 `[iterator-type] identifier = range-specification`

14 where:

- 15 • *identifier* is a base language identifier.

C / C++

- 16 • *iterator-type* is a type name.

C / C++

Fortran

- 17 • *iterator-type* is a type specifier.

Fortran

1 • *range-specification* is of the form *begin* : *end* [: *step*] where *begin*, *end* and *step* are expressions
2 for which their types can be converted to the *iterator-type* type.

▼ C / C++ ▼

3 • In an *iterator-specifier*, if the *iterator-type* is not specified then the type of that iterator is of **int**
4 type.

▲ C / C++ ▲

▼ Fortran ▼

5 • In an *iterator-specifier*, if the *iterator-type* is not specified then the type of that iterator is default
6 integer.

▲ Fortran ▲

7 In a *range-specification*, if the *step* is not specified its value is implicitly defined to be 1.

8 An iterator only exists in the context of the clause on which it appears. An iterator also hides all
9 accessible symbols with the same name in the context of the clause.

10 The use of a variable in an expression that appears in the *range-specification* causes an implicit
11 reference to the variable in all enclosing constructs.

▼ C / C++ ▼

12 The values of the iterator are the set of values $i_0 \dots i_{N-1}$ where $i_0 = \textit{begin}$, $i_j = i_{j-1} + \textit{step}$ and

13 • $i_0 < \textit{end}$ and $i_{N-1} < \textit{end}$ and $i_{N-1} + \textit{step} > \textit{end}$ if $\textit{step} > 0$.

14 • $i_0 > \textit{end}$ and $i_{N-1} > \textit{end}$ and $i_{N-1} + \textit{step} \leq \textit{end}$ if $\textit{step} < 0$.

▲ C / C++ ▲

▼ Fortran ▼

15 The values of the iterator are the set of values $i_1 \dots i_N$ where $i_1 = \textit{begin}$, $i_j = i_{j-1} + \textit{step}$ and

16 • $i_1 \leq \textit{end}$ and $i_N \leq \textit{end}$ and $i_N + \textit{step} > \textit{end}$ if $\textit{step} > 0$.

17 • $i_1 \geq \textit{end}$ and $i_N \geq \textit{end}$ and $i_N + \textit{step} < \textit{end}$ if $\textit{step} < 0$.

▲ Fortran ▲

18 The set of values will be empty if no possible value complies with the conditions above.

19 For those clauses that contain expressions containing iterator identifiers, the effect is as if the list
20 item is instantiated within the clause for each value of the iterator in the set defined above,
21 substituting each occurrence of the iterator identifier in the expression with the iterator value. If the
22 set of values of the iterator is empty then the effect is as if the clause was not specified.

Restrictions

- An expression containing an iterator identifier can only appear in clauses that explicitly allow expressions containing iterators.

▼ C / C++ ▼

- The *iterator-type* must be an integral or pointer type.

▲ C / C++ ▲

▼ Fortran ▼

- The *iterator-type* must be an integer type.

▲ Fortran ▲

- If the *step* expression of a *range-specification* equals zero the behavior is unspecified.

- Each iterator identifier can only be defined once in an *iterators-definition*.

- Iterators cannot appear in the *range-specification*.

2.8 Variant Directives

2.8.1 OpenMP Context

At any point in a program, an OpenMP context exists that defines traits describing the active OpenMP constructs, the execution devices, and functionality supported by the implementation. The traits are grouped in trait sets. The following trait sets exist: *construct*, *device* and *implementation*.

The *construct* set is composed of the directive names, each being a trait, of all enclosing executable directives at that point in the program up to a **target** directive. Combined and composite constructs will be added to the set as independent constructs in the same nesting order specified by the original construct. The set is ordered by their nesting level in increasing order. In addition, if the point in the program is not enclosed by a **target** directive, the following rules will be applied in order:

1. for functions with a **declare simd** directive, the *simd* trait will be added at the beginning of the set for the generated SIMD versions.
2. for functions with a **declare variant** directive, the selectors of the **construct** selector set will be added in the same order at the beginning of the set.

1 3. for functions within a **declare target** block, the *target* trait will be added at the beginning
2 of the set for the versions of the function being generated for **target** regions.

3 The *simd* trait can be further defined with properties that match the clauses accepted by the
4 **declare simd** directive with the same name and semantics. The *simd* trait will define at least the
5 *simdlen* property and one of the *inbranch* or *notinbranch* properties.

6 The *device* set includes traits that define the characteristics of the device being targeted by the
7 compiler at that point in the program. At least the following traits must be defined:

- 8 • The *kind(kind-name-list)* trait specifies the general kind of the device. The following *kind-name*
9 values are defined:
 - 10 – *host* specifies that the device is the host device.
 - 11 – *nohost* specifies that the devices is not the host device.
 - 12 – Values defined in the “OpenMP Context Definitions” document which is available on
13 <http://www.openmp.org/>.
- 14 • The *isa(isa-name-list)* trait specifies the Instruction Set Architectures supported by the device.
15 The accepted *isa-name* values are implementation defined.
- 16 • The *arch(arch-name-list)* trait specifies the architectures supported by the device. The accepted
17 *arch-name* values are implementation defined.

18 The *implementation* set includes traits that describe the functionality supported by the OpenMP
19 implementation at that point in the program. At least the following traits can be defined:

- 20 • The *vendor(vendor-name)* trait specifies the name of the vendor of the implementation. OpenMP
21 defined values for *vendor-name* are defined in the “OpenMP Context Definitions” document
22 which is available on <http://www.openmp.org/>.
- 23 • The *extension(extension-name-list)* trait specifies vendor specific extensions to the OpenMP
24 specification. The accepted *extension-name* values are implementation defined.
- 25 • A trait with the same name corresponding to each clause that can be supplied to the **requires**
26 directive.

1 Implementations can define further traits in the *device* and *implementation* sets. All implementation
2 defined traits must follow the following syntax:

```
3 identifier [ (context-element [, context-element [, ...]] ) ]  
4  
5 context-element :  
6   identifier [ (context-element [, context-element [, ...]] ) ]  
7   or  
8   context-value  
9  
10 context-value :  
11   string  
12   or  
13   integer expression
```

14 where *identifier* is a base language identifier.

15 2.8.2 Context Selectors

16 Context selectors allow to define the properties of an OpenMP context that a directive or clause
17 wants to match. OpenMP defines different sets of selectors, each containing different selectors.

18 The syntax to define a *context-selector-specification* is the following:

```
19 trait-set-selector [, trait-set-selector [, ...]]  
20  
21 trait-set-selector :  
22   trait-set-selector-name = { trait-selector [, trait-selector [, ...]] }  
23  
24 trait-selector :  
25   trait-selector-name [ (trait-property [, trait-property [, ...]] ) ]
```

26 The **construct** selector set defines which *construct* traits should be active in the OpenMP
27 context. The following selectors can be defined in the **construct** set: **target**, **teams**,
28 **parallel**, **for** (in C/C++), **do** (in Fortran), and **simd**. The properties of each selector are the
29 same defined for the corresponding trait. The **construct** selector is an ordered list.

30 The **device** and **implementation** selector sets define which traits should be active in the
31 corresponding trait set of the OpenMP context. The same traits defined in the corresponding traits
32 sets can be used as selectors with the same properties. The **kind** selector of the **device** selector
33 set can also be set to the value **any** which is as if no **kind** selector was specified.

34 The **user** selector set defines the **condition** selector that provides additional user-defined
35 conditions.

1 The **condition** (*boolean-expr*) selector defines a *constant expression* that must evaluate to true
2 for the selector to be true.

C

3 The **condition** (*boolean-expr*) selector defines a *constexpr* expression that must evaluate to true
4 for the selector to be true.

C++

5 The **condition** (*logical-expr*) selector defines a *constant expression* that must evaluate to true
6 for the selector to be true.

Fortran

7 Implementations can allow further selectors to be specified. Implementations can ignore specified
8 selectors that are not those described in this section.

9 Restrictions

- 10 • Each *trait-set-selector-name* can only be specified once.
- 11 • Each *trait-selector-name* can only be specified once.

12 2.8.3 Matching and Scoring Context Selectors

13 A given context selector is compatible with a given OpenMP context if:

- 14 • All selectors in the **user** set of the context selector are true,
- 15 • All selectors in the **construct**, **device** and **implementation** sets of the context selector
16 appear in the corresponding trait set of the OpenMP context,
- 17 • For each selector in the context selector, its properties are a subset of the properties of the
18 corresponding trait of the OpenMP context,
- 19 • Selectors in the **construct** set of the context selector appear in the same relative order as their
20 corresponding traits in the *construct* trait set of the OpenMP context.

21 Some properties of the **simd** selector have special rules to match the properties of the *simd* trait:

- 22 • The **simdlen** (*N*) property of the selector matches the *simdlen*(*M*) trait of the OpenMP context
23 $M\%N$ equals zero.

- 1 • The **aligned**(*list:N*) property of the selector matches the *aligned*(*list:M*) trait of the OpenMP
 2 context if $N\%M$ equals zero.
- 3 Among compatible context selectors a score will be computed using the following algorithm:
- 4 1. Each trait appearing in the *construct* trait set in the OpenMP context gets assigned the value
 5 2^{p-1} where p is the position of trait in the set.
- 6 2. The **kind**, **arch** and **isa** selectors will have the value 2^l , 2^{l+1} and 2^{l+2} respectively where l is
 7 the number of traits in the *construct* set.
- 8 3. Additional implementation allowed selector values are implementation defined.
- 9 4. Other selectors have a value of zero.
- 10 5. Context selectors which are a strict subset of another context selector have a score of zero. For
 11 other context selectors, the final score is the addition of the values of all the specified selectors
 12 plus 1. If the traits corresponding to the **construct** selectors appear multiple times in the
 13 OpenMP context, the highest valued subset of traits that contains all the selectors in the same
 14 order will be used.

15 2.8.4 declare variant Directive

16 Summary

17 The **declare variant** declares a function to be a specialized variant of another function and in
 18 which context it should be used.

19 Syntax

C / C++

20 The syntax of the **declare variant** directive is as follows:

```
21  | #pragma omp declare variant (base-func-name) [clause[ [, ] clause] ... ] new-line
22  |     function definition or declaration
```

23 where *clause* is one of the following:

```
24  |     match (context-selector-specification)
```

C / C++

Fortran

1 The syntax of the **declare variant** directive is as follows:

```
2 !$omp declare variant ([proc-name:]base-proc-name) [clause[ , ] clause] ... ]
```

3 where *clause* is one of the following:

```
4 match (context-selector-specification)
```

Fortran

5 Description

6 The use of a **declare variant** directive declares the function to be a function variant of the
7 *base-func-name* or *base-proc-name* function. If no **match** clause is specified then the context
8 selector for the variant is empty. If a **match** clause is specified then the context selector in the
9 clause will be associated to the variant.

10 At any point, after the declaration of variant for a given base function, where there is a direct call to
11 that base function the compiler will check if there is any variant that is compatible with OpenMP
12 context at that point. Among the compatible variants, the variant with the highest score according
13 to the algorithm described in Section 2.8.3 will be selected. If multiple variants have the highest
14 score, it is unspecified which one will be selected. If a compatible variant exists, the original call to
15 the base function will be replaced with a call to the selected variant function.

16 The prototype of the variant function shall, in general, match that of the base function. It is
17 implementation defined if for some specific OpenMP context the prototype of the variant should
18 differ, and how, from that of the base function.

19 Restrictions

20 Restrictions to the **declare variant** directive are as follows:

- 21 • At most one **match** clause can appear in a **declare variant** directive.
- 22 • If the function definition has a **declare variant** directive or if a declaration of the function
23 in the same compilation unit has a **declare variant**, then, calling the variant function
24 directly in an OpenMP context that is different than the one specified by the **construct** set of
25 the context selector is non-conforming.

C / C++

- 26 • If the function has any declarations, then the **declare variant** directive for any declaration
27 that has one must be equivalent. If the function definition has a **declare variant** it must
28 also be equivalent. Otherwise, the result is unspecified.

C / C++

C++

- 1 • *base-func-name* should not designate an overloaded function name. Otherwise, *base-func-name*
- 2 must be a function declaration without the return type.
- 3 • The *base-func-name* of a **declare variant** directive cannot be a template function.
- 4 • The *base-func-name* of a **declare variant** directive cannot be a virtual function.

C++

Fortran

- 5 • *proc-name* must not be a generic name, procedure pointer or entry name
- 6 • If *proc-name* is omitted, the **declare variant** directive must appear in the specification part
- 7 of a subroutine subprogram or a function subprogram.
- 8 • Any **declare variant** directive must appear in the specification part of a subroutine,
- 9 subroutine, function subprogram or interface body to which it applies.
- 10 • If a **declare variant** directive is specified in an interface block for a procedure, it must
- 11 match a **declare variant** directive in the definition of the procedure.
- 12 • If a procedure is declared via a procedure declaration statement, the procedure *proc-name* should
- 13 appear in the same specification.
- 14 • If a **declare variant** directive is specified for a procedure name with explicit interface and a
- 15 **declare variant** directive is also specified for the definition of the procedure the two
- 16 **declare variant** directives must match. Otherwise the result is unspecified.

Fortran

Cross References

- 17 • OpenMP Context Specification, see Section [2.8.1](#) on page [63](#).
- 18 • Context Selectors, see Section [2.8.2](#) on page [65](#).

20 2.8.5 Metadirective Meta-Directive

21 Summary

22 The metadirective meta-directive can specify multiple directive variants of which one may be
23 conditionally selected to replace the meta-directive based on the enclosing context.

Syntax

C / C++

The syntax of the metadirective meta-directive takes one of the following forms:

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

or

```
#pragma omp begin metadirective [clause[ [, ] clause] ... ] new-line  
    stmt(s)  
#pragma omp end metadirective
```

where *clause* is one of the following:

```
when (context-selector-specification: [directive-variant])  
default (directive-variant)
```

C / C++

Fortran

The syntax of the metadirective meta-directive takes one of the following forms:

```
!$omp metadirective [clause[ [, ] clause] ... ]
```

or

```
!$omp begin metadirective [clause[ [, ] clause] ... ]  
    stmt(s)  
!$omp end metadirective
```

where *clause* is one of the following:

```
when (context-selector-specification: [directive-variant])  
default (directive-variant)
```

Fortran

In the **when** clause, *context-selector-specification* specifies a context selector (see Section 2.8.2).

1 In the **when** and **default** clauses, *directive-variant* has the following form and specifies a
2 directive variant that is an OpenMP directive that has the same directive name and clauses.

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

4 Description

5 The metadirective directive is a meta-directive that behaves as if it is either ignored or replaced by
6 the directive variant specified in one of the **when** or **default** clauses that appears on the directive.

7 The OpenMP context for a given meta-directive is defined according to Section 2.8.1. For each
8 **when** clause that appears on the meta-directive, the specified directive variant, if present, is a
9 candidate to replace the meta-directive if the corresponding context selector is compatible with the
10 OpenMP context according to the matching rules defined in Section 2.8.3. If only one compatible
11 context selector specified by a **when** clause has the highest score and it specifies a directive variant,
12 the directive variant will replace the meta-directive. If more than one **when** clause specifies a
13 compatible context selector that has the highest computed score and at least one specifies a
14 directive variant, the first directive variant specified in the lexical order of those **when** clauses will
15 replace the meta-directive.

16 If no context selector from any **when** clause is compatible with the OpenMP context and a
17 **default** clause is present, the directive variant specified in the **default** clause will replace the
18 meta-directive.

19 If a directive variant is not selected to replace the meta-directive according to the above rules, the
20 meta-directive has no effect on the execution of program.

21 The **begin metadirective** directive behaves identically to the **metadirective** directive,
22 except that the directive syntax for the specified directive variants must accept a paired
23 **end directive**. For any directive variant that is selected to replace the **begin metadirective**
24 meta-directive, the **end metadirective** directive will be implicitly replaced by its paired
25 **end directive** to demarcate the statements that are affected by or are associated with the directive
26 variant. If no directive variant is selected to replace the meta-directive, its paired
27 **end metadirective** directive is ignored.

28 Restrictions

29 Restrictions for the metadirective directive are as follows:

- 30 • The directive variant appearing in a **when** or **default** clause must not specify a
31 **metadirective**, **begin metadirective**, or **end metadirective** directive.
- 32 • The context selector that appears in a **when** clause must not specify any properties for the **simd**
33 selector.
- 34 • Any replacement that occurs for a metadirective meta-directive must not result in a
35 non-conforming OpenMP program.

- 1 • Any directive variant that is specified by a **when** or **default** clause on a
2 **begin metadirective** meta-directive must be an OpenMP directive that has a paired
3 **end directive**, and the **begin metadirective** directive must have a paired
4 **end metadirective** directive.
- 5 • The **default** clause may appear at most once on the directive.

6 2.9 parallel Construct

7 Summary

8 This fundamental construct starts parallel execution. See Section 1.3 on page 19 for a general
9 description of the OpenMP execution model.

10 Syntax

▼ C / C++ ▼

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

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

14 where *clause* is one of the following:

```
15     if ([parallel :] scalar-expression)  
16     num_threads (integer-expression)  
17     default (shared | none)  
18     private (list)  
19     firstprivate (list)  
20     shared (list)  
21     copyin (list)  
22     reduction ([ reduction-modifier, ]reduction-identifier : list)  
23     proc_bind (master | close | spread)  
24     allocate ([allocator: ]list)
```

▲ C / C++ ▲

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

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

5 where *clause* is one of the following:

```
6   if ([parallel :] scalar-logical-expression)
7   num_threads (scalar-integer-expression)
8   default (private | firstprivate | shared | none)
9   private (list)
10  firstprivate (list)
11  shared (list)
12  copyin (list)
13  reduction ([ reduction-modifier , ] reduction-identifier : list)
14  proc_bind (master | close | spread)
15  allocate ([allocator: ]list)
```

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

17 Binding

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

20 Description

21 When a thread encounters a **parallel** construct, a team of threads is created to execute the
22 **parallel** region (see Section 2.9.1 on page 77 for more information about how the number of
23 threads in the team is determined, including the evaluation of the **if** and **num_threads** clauses).
24 The thread that encountered the **parallel** construct becomes the master thread of the new team,
25 with a thread number of zero for the duration of the new **parallel** region. All threads in the new
26 team, including the master thread, execute the region. Once the team is created, the number of
27 threads in the team remains constant for the duration of that **parallel** region.

28 The optional **proc_bind** clause, described in Section 2.9.2 on page 79, specifies the mapping of
29 OpenMP threads to places within the current place partition, that is, within the places listed in the
30 *place-partition-var* ICV for the implicit task of the encountering thread.

1 Within a **parallel** region, thread numbers uniquely identify each thread. Thread numbers are
2 consecutive whole numbers ranging from zero for the master thread up to one less than the number
3 of threads in the team. A thread may obtain its own thread number by a call to the
4 **omp_get_thread_num** library routine.

5 A set of implicit tasks, equal in number to the number of threads in the team, is generated by the
6 encountering thread. The structured block of the **parallel** construct determines the code that
7 will be executed in each implicit task. Each task is assigned to a different thread in the team and
8 becomes tied. The task region of the task being executed by the encountering thread is suspended
9 and each thread in the team executes its implicit task. Each thread can execute a path of statements
10 that is different from that of the other threads

11 The implementation may cause any thread to suspend execution of its implicit task at a task
12 scheduling point, and switch to execute any explicit task generated by any of the threads in the
13 team, before eventually resuming execution of the implicit task (for more details see Section 2.13
14 on page 133).

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

17 If a thread in a team executing a **parallel** region encounters another **parallel** directive, it
18 creates a new team, according to the rules in Section 2.9.1 on page 77, and it becomes the master of
19 that new team.

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

24 Execution Model Events

25 The *parallel-begin* event occurs in a thread encountering a **parallel** construct before any
26 implicit task is created for the corresponding **parallel** region.

27 Upon creation of each implicit task, an *implicit-task-begin* event occurs in the thread executing the
28 implicit task after the implicit task is fully initialized but before the thread begins to execute the
29 structured block of the **parallel** construct.

30 If the **parallel** region creates a thread, a *thread-begin* event occurs as the first event in the
31 context of the new thread prior to the *implicit-task-begin*.

32 Events associated with implicit barriers occur at the end of a **parallel** region. Section 2.20.3
33 describes events associated with implicit barriers.

34 When a thread finishes an implicit task, an *implicit-task-end* event occurs in the thread after events
35 associated with implicit barrier synchronization in the implicit task.

1 The *parallel-end* event occurs in the thread encountering the **parallel** construct after the thread
2 executes its *implicit-task-end* event but before resuming execution of the encountering task.
3 If a thread is destroyed at the end of a **parallel** region, a *thread-end* event occurs in the thread as
4 the last event prior to the thread's destruction.

5 **Tool Callbacks**

6 A thread dispatches a registered **ompt_callback_parallel_begin** callback for each
7 occurrence of a *parallel-begin* event in that thread. The callback occurs in the task encountering the
8 **parallel** construct. This callback has the type signature
9 **ompt_callback_parallel_begin_t**.

10 A thread dispatches a registered **ompt_callback_implicit_task** callback for each
11 occurrence of a *implicit-task-begin* and *implicit-task-end* event in that thread. The callback occurs
12 in the context of the implicit task. The callback has type signature
13 **ompt_callback_implicit_task_t**. The callback receives **ompt_scope_begin** or
14 **ompt_scope_end** as its *endpoint* argument, as appropriate.

15 A thread dispatches a registered **ompt_callback_parallel_end** callback for each
16 occurrence of a *parallel-end* event in that thread. The callback occurs in the task encountering the
17 **parallel** construct. This callback has the type signature
18 **ompt_callback_parallel_end_t**.

19 A thread dispatches a registered **ompt_callback_thread_begin** callback for the
20 *thread-begin* event in that thread. The callback occurs in the context of the thread. The callback has
21 type signature **ompt_callback_thread_begin_t**.

22 A thread dispatches a registered **ompt_callback_thread_end** callback for the *thread-end*
23 event in that thread. The callback occurs in the context of the thread. The callback has type
24 signature **ompt_callback_thread_end_t**.

25 **Restrictions**

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

- 27 • A program that branches into or out of a **parallel** region is non-conforming.
- 28 • A program must not depend on any ordering of the evaluations of the clauses of the **parallel**
29 directive, or on any side effects of the evaluations of the clauses.
- 30 • At most one **if** clause can appear on the directive.
- 31 • At most one **proc_bind** clause can appear on the directive.
- 32 • At most one **num_threads** clause can appear on the directive. The **num_threads**
33 expression must evaluate to a positive integer value.

C / C++

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

C / C++

Fortran

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

Fortran

Cross References

- **if** clause, see Section [2.18](#) on page [213](#).
- **default**, **shared**, **private**, **firstprivate**, and **reduction** clauses, see Section [2.22.4](#) on page [276](#).
- **copyin** clause, see Section [2.22.6](#) on page [301](#).
- **omp_get_thread_num** routine, see Section [3.2.4](#) on page [336](#).
- **ompt_scope_begin** and **ompt_scope_end**, see Section [4.2.3.4.10](#) on page [437](#).
- **ompt_callback_thread_begin_t**, see Section [4.2.4.2.1](#) on page [446](#).
- **ompt_callback_thread_end_t**, see Section [4.2.4.2.2](#) on page [447](#).
- **ompt_callback_parallel_begin_t**, see Section [4.2.4.2.3](#) on page [447](#).
- **ompt_callback_parallel_end_t**, see Section [4.2.4.2.4](#) on page [449](#).
- **ompt_callback_implicit_task_t**, see Section [4.2.4.2.10](#) on page [455](#).

1 2.9.1 Determining the Number of Threads for a `parallel` 2 Region

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

7 Using a variable in an `if` or `num_threads` clause expression of a `parallel` construct causes
8 an implicit reference to the variable in all enclosing constructs. The `if` clause expression and the
9 `num_threads` clause expression are evaluated in the context outside of the `parallel`
10 construct, and no ordering of those evaluations is specified. It is also unspecified whether, in what
11 order, or how many times any side effects of the evaluation of the `num_threads` or `if` clause
12 expressions occur.

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

15 Algorithm 2.1 16 17

```
18     let ThreadsBusy be the number of OpenMP threads currently executing in this  
19     contention group;  
20     let ActiveParRegions be the number of enclosing active parallel regions;  
21     if an if clause exists  
22     then let IfClauseValue be the value of the if clause expression;  
23     else let IfClauseValue = true;  
24     if a num_threads clause exists  
25     then let ThreadsRequested be the value of the num_threads clause expression;  
26     else let ThreadsRequested = value of the first element of nthreads-var;  
27     if a thread_limit clause exists on a teams construct corresponding to an  
28     enclosing teams region  
29     then let ThreadLimit be the value of the thread_limit clause expression;  
30     else let ThreadLimit = thread-limit-var;  
31     let ThreadsAvailable = (ThreadLimit - ThreadsBusy + 1);  
32     if (IfClauseValue = false)  
33     then number of threads = 1;
```

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

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

18 **Cross References**

- 19 • *nthreads-var*, *dyn-var*, *thread-limit-var*, *max-active-levels-var*, and *nest-var* ICVs, see
20 Section 2.4 on page 47.

1 2.9.2 Controlling OpenMP Thread Affinity

2 When a thread encounters a **parallel** directive without a **proc_bind** clause, the *bind-var* ICV
3 is 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 the
22 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 subpartition.
35 The master thread executes on the place of the parent thread and is assigned to the subpartition
36 that includes that place. The thread with the next smallest thread number is assigned to the first
37 place in the next subpartition, and so on, with wrap around with respect to the original place
38 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 ▼
11 Note – Wrap around is needed if the end of a place partition is reached before all thread
12 assignments are done. For example, wrap around may be needed in the case of **close** and $T \leq P$,
13 if the master thread is assigned to a place other than the first place in the place partition. In this
14 case, thread 1 is assigned to the place after the place of the master place, thread 2 is assigned to the
15 place after that, and so on. The end of the place partition may be reached before all threads are
 assigned. In this case, assignment of threads is resumed with the first place in the place partition.
 ▲

1 2.10 teams Construct

2 Summary

3 The **teams** construct creates a league of initial teams and the initial thread in each team executes
4 the region.

5 Syntax

C / C++

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

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

9 where *clause* is one of the following:

```
10     num_teams (integer-expression)  
11     thread_limit (integer-expression)  
12     default (shared | none)  
13     private (list)  
14     firstprivate (list)  
15     shared (list)  
16     reduction (reduction-identifier : list)  
17     allocate ([allocator: ]list)
```

C / C++

Fortran

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

```
19 !$omp teams [clause[ [, ] clause] ... ]  
20     structured-block  
21 !$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 allocate ([allocator: ]list)
```

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

Fortran

11 Binding

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

13 Description

14 When a thread encounters a **teams** construct, a league of teams is created. Each team is an initial
15 team, and the initial thread in each team executes the **teams** region.

16 The number of teams created is implementation defined, but is less than or equal to the value
17 specified in the **num_teams** clause. A thread may obtain the number of initial teams created by
18 the construct by a call to the **omp_get_num_teams** routine.

19 The maximum number of threads participating in the contention group that each team initiates is
20 implementation defined, but is less than or equal to the value specified in the **thread_limit**
21 clause.

22 On a combined or composite construct that includes **target** and **teams** constructs, the
23 expressions in **num_teams** and **thread_limit** clauses are evaluated on the host device on
24 entry to the **target** construct.

25 Once the teams are created, the number of initial teams remains constant for the duration of the
26 **teams** region.

27 Within a **teams** region, initial team numbers uniquely identify each initial team. Initial team
28 numbers are consecutive whole numbers ranging from zero to one less than the number of initial
29 teams. A thread may obtain its own initial team number by a call to the **omp_get_team_num**
30 library routine. The policy for assigning the initial threads to places is implementation defined. The

1 **team** construct sets the *place-partition-var* and *default-device-var* ICVs for each initial thread to
2 an implementation-defined value.

3 After the teams have completed execution of the **teams** region, the encountering task resumes
4 execution of the enclosing task region.

5 **Execution Model Events**

6 The *teams-begin* event occurs in a thread encountering a **teams** construct before any initial task is
7 created for the corresponding **teams** region.

8 Upon creation of each initial task, an *initial-task-begin* event occurs in the thread executing the
9 initial task after the initial task is fully initialized but before the thread begins to execute the
10 structured block of the **teams** construct.

11 If the **teams** region creates a thread, a *thread-begin* event occurs as the first event in the context of
12 the new thread prior to the *initial-task-begin*.

13 When a thread finishes an initial task, an *initial-task-end* event occurs in the thread.

14 The *teams-end* event occurs in the thread encountering the **teams** construct after the thread
15 executes its *initial-task-end* event but before resuming execution of the encountering task.

16 If a thread is destroyed at the end of a **teams** region, a *thread-end* event occurs in the thread as the
17 last event prior to the thread's destruction.

18 **Tool Callbacks**

19 A thread dispatches a registered **ompt_callback_parallel_begin** callback for each
20 occurrence of a *teams-begin* event in that thread. The callback occurs in the task encountering the
21 **parallel** construct. This callback has the type signature
22 **ompt_callback_parallel_begin_t**.

23 A thread dispatches a registered **ompt_callback_implicit_task** callback for each
24 occurrence of a *initial-task-begin* and *initial-task-end* event in that thread. The callback occurs in
25 the context of the initial task. The callback has type signature
26 **ompt_callback_implicit_task_t**. The callback receives **ompt_scope_begin** or
27 **ompt_scope_end** as its *endpoint* argument, as appropriate.

28 A thread dispatches a registered **ompt_callback_parallel_end** callback for each
29 occurrence of a *teams-end* event in that thread. The callback occurs in the task encountering the
30 **teams** construct. This callback has the type signature **ompt_callback_parallel_end_t**.

31 A thread dispatches a registered **ompt_callback_thread_begin** callback for the
32 *thread-begin* event in that thread. The callback occurs in the context of the thread. The callback has
33 type signature **ompt_callback_thread_begin_t**.

1 A thread dispatches a registered `ompt_callback_thread_end` callback for the *thread-end*
2 event in that thread. The callback occurs in the context of the thread. The callback has type
3 signature `ompt_callback_thread_end_t`.

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 • A `teams` region can only be strictly nested within the implicit parallel region or a `target`
14 region. If a `teams` construct is nested within a `target` construct, that `target` construct must
15 contain no statements, declarations or directives outside of the `teams` construct.
- 16 • `distribute`, `distribute simd`, distribute parallel worksharing-loop, distribute parallel
17 worksharing-loop SIMD, `parallel` regions, including any `parallel` regions arising from
18 combined constructs, `omp_get_num_teams()` regions, and `omp_get_team_num()` are
19 the only OpenMP regions that may be strictly nested inside the `teams` region.

20 Cross References

- 21 • `parallel` construct, see Section 2.9 on page 72.
- 22 • `distribute` construct, see Section 2.12.4.1 on page 117.
- 23 • `distribute simd` construct, see Section 2.12.4.2 on page 121.
- 24 • `target` construct, see Section 2.15.5 on page 168.
- 25 • Data-sharing attribute clauses, see Section 2.22.4 on page 276.
- 26 • `omp_get_num_teams` routine, see Section 3.2.37 on page 370.
- 27 • `omp_get_team_num` routine, see Section 3.2.38 on page 371.
- 28 • `ompt_callback_thread_begin_t`, see Section 4.2.4.2.1 on page 446.
- 29 • `ompt_callback_thread_end_t`, see Section 4.2.4.2.2 on page 447.
- 30 • `ompt_callback_parallel_begin_t`, see Section 4.2.4.2.3 on page 447.
- 31 • `ompt_callback_parallel_end_t`, see Section 4.2.4.2.4 on page 449.

1 • `ompt_callback_implicit_task_t`, see Section [4.2.4.2.10](#) on page 455.

2 2.11 Worksharing Constructs

3 A worksharing construct distributes the execution of the corresponding region among the members
4 of the team that encounters it. Threads execute portions of the region in the context of the implicit
5 tasks each one is executing. If the team consists of only one thread then the worksharing region is
6 not executed in parallel.

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

13 The OpenMP API defines the worksharing constructs that are described in this section. The
14 worksharing-loop construct is described in Section [2.12.2](#) on page 100.

15 **Restrictions**

16 The following restrictions apply to worksharing constructs:

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

1 2.11.1 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:

```
8 #pragma omp sections [clause [ , ] clause ... ] new-line  
9 {  
10     [#pragma omp section new-line]  
11     structured-block  
12     [#pragma omp section new-line]  
13     structured-block]  
14     ...  
15 }
```

16 where *clause* is one of the following:

```
17 private (list)  
18 firstprivate (list)  
19 lastprivate ([ lastprivate-modifier : ] list)  
20 reduction ([ reduction-modifier , ] reduction-identifier : list)  
21 nowait  
22 allocate ([ allocator : ] list)
```

C / C++

Fortran

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]
```

where *clause* is one of the following:

```
private (list)  
firstprivate (list)  
lastprivate ([ lastprivate-modifier: ] list)  
reduction ([reduction-modifier, ]reduction-identifier : list)  
allocate ([allocator: ]list)
```

Fortran

Binding

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

Description

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

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

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

1 **Execution Model Events**

2 The *sections-begin* event occurs after an implicit task encounters a **sections** construct but before
3 the task starts the execution of the structured block of the **sections** region.

4 The *sections-end* event occurs after a **sections** region finishes execution but before resuming
5 execution of the encountering task.

6 The *section-begin* event occurs before an implicit task starts executing a structured block in the
7 **sections** construct.

8 **Tool Callbacks**


9 A thread dispatches a registered **ompt_callback_work** callback for each occurrence of a
10 *sections-begin* and *sections-end* event in that thread. The callback occurs in the context of the
11 implicit task. The callback has type signature **ompt_callback_work_t**. The callback receives
12 **ompt_scope_begin** or **ompt_scope_end** as its *endpoint* argument, as appropriate, and
13 **ompt_work_sections** as its *wstype* argument.

14 A thread dispatches a registered **ompt_callback_dispatch** callback for each occurrence of a
15 *section-begin* event in that thread. The callback occurs in the context of the implicit task. The
16 callback has type signature **ompt_callback_dispatch_t**.

17 **Restrictions**

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

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



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



Cross References

- `private`, `firstprivate`, `lastprivate`, and `reduction` clauses, see Section 2.22.4 on page 276.
- `ompt_scope_begin` and `ompt_scope_end`, see Section 4.2.3.4.10 on page 437.
- `ompt_work_sections`, see Section 4.2.3.4.14 on page 438.
- `ompt_callback_work_t`, see Section 4.2.4.2.15 on page 461.
- `ompt_callback_dispatch_t`, see Section 4.2.4.2.17 on page 464.

2.11.2 `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
```

where *clause* is one of the following:

```
private (list)  
firstprivate (list)  
copyprivate (list)  
nowait  
allocate ([allocator: ]list)
```

C / C++

Fortran

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

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

5 where *clause* is one of the following:

```
6   private (list)  
7   firstprivate (list)  
8   allocate ([allocator: ]list)
```

9 and *end_clause* is one of the following:

```
10   copyprivate (list)  
11   nowait
```

Fortran

12 Binding

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

17 Description

18 Only one of the encountering threads will execute the structured block associated with the **single**
19 construct. The method of choosing a thread to execute the structured block each time the team
20 encounters the construct is implementation defined. There is an implicit barrier at the end of the
21 **single** construct unless a **nowait** clause is specified.

22 Execution Model Events

23 The *single-begin* event occurs after an **implicit task** encounters a **single** construct but
24 before the task starts the execution of the structured block of the **single** region.

25 The *single-end* event occurs after a **single** region finishes execution of the structured block but
26 before resuming execution of the encountering implicit task.

1 Tool Callbacks

2 A thread dispatches a registered `ompt_callback_work` callback for each occurrence of
3 *single-begin* and *single-end* events in that thread. The callback has type signature
4 `ompt_callback_work_t`. The callback receives `ompt_scope_begin` or
5 `ompt_scope_end` as its *endpoint* argument, as appropriate, and
6 `ompt_work_single_executor` or `ompt_work_single_other` as its *wstype* argument.

7 Restrictions

8 Restrictions to the `single` construct are as follows:

- 9
- The `copyprivate` clause must not be used with the `nowait` clause.
 - At most one `nowait` clause can appear on a `single` construct.
- 10

▼ C++ ▼

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

▲ C++ ▲

13 Cross References

- 14
- `private` and `firstprivate` clauses, see Section 2.22.4 on page 276.
 - 15
 - `copyprivate` clause, see Section 2.22.6.2 on page 303.
 - 16
 - `ompt_scope_begin` and `ompt_scope_end`, see Section 4.2.3.4.10 on page 437.
 - 17
 - `ompt_work_single_executor` and `ompt_work_single_other`, see
18 Section 4.2.3.4.14 on page 438.
 - 19
 - `ompt_callback_work_t`, Section 4.2.4.2.15 on page 461.

▼ Fortran ▼

20 2.11.3 workshare Construct

21 Summary

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

1 **Syntax**

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

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

6 The enclosed structured block must consist of only the following:

- 7 • array assignments
- 8 • scalar assignments
- 9 • **FORALL** statements
- 10 • **FORALL** constructs
- 11 • **WHERE** statements
- 12 • **WHERE** constructs
- 13 • **atomic** constructs
- 14 • **critical** constructs
- 15 • **parallel** constructs

16 Statements contained in any enclosed **critical** construct are also subject to these restrictions.
17 Statements in any enclosed **parallel** construct are not restricted.

18 **Binding**

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

23 **Description**

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

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

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

- 1 • For array expressions within each statement, including transformational array intrinsic functions
2 that compute scalar values from arrays:
- 3 – Evaluation of each element of the array expression, including any references to **ELEMENTAL**
4 functions, is a unit of work.
- 5 – Evaluation of transformational array intrinsic functions may be freely subdivided into any
6 number of units of work.
- 7 • For an array assignment statement, the assignment of each element is a unit of work.
- 8 • For a scalar assignment statement, the assignment operation is a unit of work.
- 9 • For a **WHERE** statement or construct, the evaluation of the mask expression and the masked
10 assignments are each a unit of work.
- 11 • For a **FORALL** statement or construct, the evaluation of the mask expression, expressions
12 occurring in the specification of the iteration space, and the masked assignments are each a unit
13 of work
- 14 • For an **atomic** construct, the atomic operation on the storage location designated as *x* is a unit
15 of work.
- 16 • For a **critical** construct, the construct is a single unit of work.
- 17 • For a **parallel** construct, the construct is a unit of work with respect to the **workshare**
18 construct. The statements contained in the **parallel** construct are executed by a new thread
19 team.
- 20 • If none of the rules above apply to a portion of a statement in the structured block, then that
21 portion is a unit of work.
- 22 The transformational array intrinsic functions are **MATMUL**, **DOT_PRODUCT**, **SUM**, **PRODUCT**,
23 **MAXVAL**, **MINVAL**, **COUNT**, **ANY**, **ALL**, **SPREAD**, **PACK**, **UNPACK**, **RESHAPE**, **TRANSPOSE**,
24 **EOSHIFT**, **CSHIFT**, **MINLOC**, and **MAXLOC**.
- 25 It is unspecified how the units of work are assigned to the threads executing a **workshare** region.
- 26 If an array expression in the block references the value, association status, or allocation status of
27 private variables, the value of the expression is undefined, unless the same value would be
28 computed by every thread.
- 29 If an array assignment, a scalar assignment, a masked array assignment, or a **FORALL** assignment
30 assigns to a private variable in the block, the result is unspecified.
- 31 The **workshare** directive causes the sharing of work to occur only in the **workshare** construct,
32 and not in the remainder of the **workshare** region.

1 Execution Model Events

2 The *workshare-begin* event occurs after an implicit task encounters a **workshare** construct but
3 before the task starts the execution of the structured block of the **workshare** region.

4 The *workshare-end* event occurs after a **workshare** region finishes execution but before resuming
5 execution of the encountering task.

6 Tool Callbacks

7 A thread dispatches a registered **ompt_callback_work** callback for each occurrence of a
8 *workshare-begin* and *workshare-end* event in that thread. The callback occurs in the context of the
9 implicit task. The callback has type signature **ompt_callback_work_t**. The callback receives
10 **ompt_scope_begin** or **ompt_scope_end** as its *endpoint* argument, as appropriate, and
11 **ompt_work_workshare** as its *wstype* argument.

12 Restrictions

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

- 14 • All array assignments, scalar assignments, and masked array assignments must be intrinsic
15 assignments.
- 16 • The construct must not contain any user defined function calls unless the function is
17 **ELEMENTAL**.

18 Cross References

- 19 • **ompt_scope_begin** and **ompt_scope_end**, see Section [4.2.3.4.10](#) on page [437](#).
- 20 • **ompt_work_workshare**, see Section [4.2.3.4.14](#) on page [438](#).
- 21 • **ompt_callback_work_t**, see Section [4.2.4.2.15](#) on page [461](#).


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

1 2.12 Loop Constructs

2 2.12.1 Canonical Loop Form

3  C++
A range-based for loop with random access iterator has a *canonical loop form*.

 C++

4  C / C++
5 The loops associated with a loop directive have *canonical loop form* if they conform to the following:

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.

continued on next page

	<p>For C, a variable of a pointer type.</p> <p>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 <i>for-loop</i> other than in <i>incr-expr</i>. Unless the variable is specified lastprivate or linear on the loop construct, its value after the loop is unspecified.</p>
<i>relational-op</i>	<p>One of the following:</p> <ul style="list-style-type: none"><<=>>=!=
<i>lb and b</i>	<p>Expressions of a type compatible with the type of <i>var</i> that are loop invariant with respect to the outermost associated loop or are one of the following (where <i>var-outer</i>, <i>a1</i>, and <i>a2</i> have a type compatible with the type of <i>var</i>, <i>var-outer</i> is <i>var</i> from an outer associated loop, and <i>a1</i> and <i>a2</i> are loop invariant integer expressions with respect to the outermost loop):</p> <ul style="list-style-type: none"><i>var-outer</i><i>var-outer</i> + <i>a2</i><i>a2</i> + <i>var-outer</i><i>var-outer</i> - <i>a2</i><i>a2</i> - <i>var-outer</i><i>a1</i> * <i>var-outer</i><i>a1</i> * <i>var-outer</i> + <i>a2</i><i>a2</i> + <i>a1</i> * <i>var-outer</i><i>a1</i> * <i>var-outer</i> - <i>a2</i><i>a2</i> - <i>a1</i> * <i>var-outer</i><i>var-outer</i> * <i>a1</i><i>var-outer</i> * <i>a1</i> + <i>a2</i><i>a2</i> + <i>var-outer</i> * <i>a1</i><i>var-outer</i> * <i>a1</i> - <i>a2</i><i>a2</i> - <i>var-outer</i> * <i>a1</i>
<i>incr</i>	<p>An integer expression that is loop invariant with respect to the outermost associated loop.</p>



Fortran

1 The loops associated with a loop directive have *canonical loop form* if each of them is a *do-loop*
2 that is a *do-construct* or an *inner-shared-do-construct* as defined by the Fortran standard. If an
3 **end do** directive follows a *do-construct* in which several loop statements share a **DO** termination
4 statement, then the directive can only be specified for the outermost of these **DO** statements.

5 The *do-stmt* for any *do-loop* must conform to the following:

DO [*label*] *var* = *lb* , *b* [, *incr*]


<i>var</i>	A variable of integer type. If this variable would otherwise be shared, it is implicitly made private in the loop construct. Unless the variable is specified lastprivate or linear on the loop construct, its value after the loop is unspecified.
<i>lb</i> and <i>b</i>	Expressions of a type compatible with the type of <i>var</i> that are loop invariant with respect to the outermost associated loop or are one of the following (where <i>var-outer</i> , <i>a1</i> , and <i>a2</i> have a type compatible with the type of <i>var</i> , <i>var-outer</i> is <i>var</i> from an outer associated loop, and <i>a1</i> and <i>a2</i> are loop invariant integer expressions with respect to the outermost loop): <i>var-outer</i> <i>var-outer</i> + <i>a2</i> <i>a2</i> + <i>var-outer</i> <i>var-outer</i> - <i>a2</i> <i>a2</i> - <i>var-outer</i> <i>a1</i> * <i>var-outer</i> <i>a1</i> * <i>var-outer</i> + <i>a2</i> <i>a2</i> + <i>a1</i> * <i>var-outer</i> <i>a1</i> * <i>var-outer</i> - <i>a2</i> <i>a2</i> - <i>a1</i> * <i>var-outer</i> <i>var-outer</i> * <i>a1</i> <i>var-outer</i> * <i>a1</i> + <i>a2</i> <i>a2</i> + <i>var-outer</i> * <i>a1</i> <i>var-outer</i> * <i>a1</i> - <i>a2</i> <i>a2</i> - <i>var-outer</i> * <i>a1</i>
<i>incr</i>	An integer expression that is loop invariant with respect to the outermost associated loop. If it is not explicitly specified, its value is assumed to be 1.

Fortran


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

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

- If *var* is of an integer type, then the type is the type of *var*.



- 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*.



- 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.



Restrictions

The following restrictions also apply:



- If *test-expr* is of the form *var relational-op b* and *relational-op* is `<` or `<=` then *incr-expr* must cause *var* to increase on each iteration of the loop. If *test-expr* is of the form *var relational-op b* and *relational-op* is `>` or `>=` then *incr-expr* must cause *var* to decrease on each iteration of the loop.
- If *test-expr* is of the form *b relational-op var* and *relational-op* is `<` or `<=` then *incr-expr* must cause *var* to decrease on each iteration of the loop. If *test-expr* is of the form *b relational-op var* and *relational-op* is `>` or `>=` then *incr-expr* must cause *var* to increase on each iteration of the loop.
- If *test-expr* is of the form *b != var* or *var != b* then *incr-expr* must cause *var* either to increase on each iteration of the loop or to decrease on each iteration of the loop.
- If *relational-op* is `!=` and *incr-expr* is of the form that has *incr* then *incr* must be a constant expression and evaluate to -1 or 1.



C++

- 1 • In the **simd** construct the only random access iterator types that are allowed for *var* are pointer
- 2 types.
- 3 • The iteration count of a range-based for loop must be loop invariant with respect to the outermost
- 4 associated loop.

C++

- 5 • The *b*, *lb*, and *incr* expressions may not reference *var* of any enclosed associated loop.
- 6 • For any associated loop where the *b* or *lb* expression is not loop invariant with respect to the
- 7 outermost loop, the *var-outer* that appears in the expression may not have a random access
- 8 iterator type.
- 9 • For any associated loop where *b* or *lb* is not loop invariant with respect to the outermost loop, the
- 10 expression $b - lb$ will have the form $c * var-outer + d$, where *c* and *d* are loop invariant integer
- 11 expressions. Let *incr-outer* be the *incr* expression of the outer loop referred to by *var-outer*. The
- 12 value of $c * incr-outer \bmod incr$ must be 0.

1 2.12.2 Worksharing-Loop Construct

2 Summary

3 The worksharing-loop construct specifies that the iterations of one or more associated loops will be
4 executed in parallel by threads in the team in the context of their implicit tasks. The iterations are
5 distributed across threads that already exist in the team executing the **parallel** region to which
6 the worksharing-loop region binds.

7 Syntax

C / C++

8 The syntax of the worksharing-loop construct is as follows:

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

11 where clause is one of the following:

```
12     private (list)  
13     firstprivate (list)  
14     lastprivate ([ lastprivate-modifier: ] list)  
15     linear (list[ : linear-step])  
16     reduction ([ reduction-modifier, ]reduction-identifier : list)  
17     schedule ([modifier [, modifier]:]kind[, chunk_size])  
18     collapse (n)  
19     ordered( (n) )  
20     nowait  
21     allocate ([allocator: ]list)  
22     order (concurrent)
```

23 The **for** directive places restrictions on the structure of all associated *for-loops*. Specifically, all
24 associated *for-loops* must have *canonical loop form* (see Section 2.12.1 on page 95).

C / C++

1 The syntax of the worksharing-loop construct is as follows:

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

5 where *clause* is one of the following:

```
6     private (list)
7     firstprivate (list)
8     lastprivate ([ lastprivate-modifier: ] list)
9     linear (list [ : linear-step ])
10    reduction ([ reduction-modifier, ] reduction-identifier : list)
11    schedule ([ modifier [ , modifier ]: ] kind [ , chunk_size ])
12    collapse (n)
13    ordered [ (n) ]
14    allocate ([ allocator: ] list)
15    order (concurrent)
```

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

17 The **do** directive places restrictions on the structure of all associated *do-loops*. Specifically, all
18 associated *do-loops* must have *canonical loop form* (see Section 2.12.1 on page 95).

19 Binding

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

Description

The worksharing-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 worksharing-loop construct unless a **nowait** clause is specified.

The **collapse** clause may be used to specify how many loops are associated with the worksharing-loop construct. The parameter of the **collapse** clause must be a constant positive integer expression. If a **collapse** clause is specified with a parameter value greater than 1, then the iterations of the associated loops to which the clause applies are collapsed into one larger iteration space that is then divided according to the **schedule** clause. The sequential execution of the iterations in these associated loops determines the order of the iterations in the collapsed iteration space. If no **collapse** clause is present or its parameter is 1, the only loop that is associated with the worksharing-loop construct for the purposes of determining how the iteration space is divided according to the **schedule** clause is the one that immediately follows the worksharing-loop directive.

If more than one loop is associated with the worksharing-loop construct then the number of times that any intervening code between any two associated loops will be executed is unspecified but will be at least once per iteration of the loop enclosing the intervening code and at most once per iteration of the innermost loop associated with the construct. If the iteration count of any loop that is associated with the worksharing-loop construct and does not enclose the intervening code is zero then the behavior is unspecified.

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 a set of associated loop(s) were executed sequentially. At the beginning of each logical iteration, the loop iteration variable of each associated loop has the value that it would have if the set of the associated loop(s) were executed sequentially. The **schedule** clause specifies how iterations of these 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 iterations of a given chunk are executed in sequential order by the assigned thread. The *chunk_size* expression is evaluated using the original list items of any variables that are made private in the worksharing-loop construct. It is unspecified whether, in what order, or how many times, any side effects of the evaluation of this expression occur. The use of a variable in a **schedule** clause expression of a worksharing-loop construct causes an implicit reference to the variable in all enclosing constructs.

1 Different worksharing-loop regions with the same schedule and iteration count, even if they occur
2 in the same parallel region, can distribute iterations among threads differently. The only exception
3 is for the **static** schedule as specified in Table 2.5. Programs that depend on which thread
4 executes a particular iteration under any other circumstances are non-conforming.

5 See Section 2.12.2.1 on page 110 for details of how the schedule for a worksharing
6 worksharing-loop is determined.

7 The schedule *kind* can be one of those specified in Table 2.5.

8 The schedule *modifier* can be one of those specified in Table 2.6. If the **static** schedule kind is
9 specified or if the **ordered** clause is specified, and if the **nonmonotonic** modifier is not
10 specified, the effect is as if the **monotonic** modifier is specified. Otherwise, unless the
11 **monotonic** modifier is specified, the effect is as if the **nonmonotonic** modifier is specified.

12 The **ordered** clause with the parameter may also be used to specify how many loops are
13 associated with the worksharing-loop construct. The parameter of the **ordered** clause must be a
14 constant positive integer expression if specified. The parameter of the **ordered** clause does not
15 affect how the logical iteration space is then divided. If an **ordered** clause with the parameter is
16 specified for the worksharing-loop construct, then those associated loops form a *doacross loop nest*.

17 If the value of the parameter in the **collapse** or **ordered** clause is larger than the number of
18 nested loops following the construct, the behavior is unspecified.

19 If an **order (concurrent)** clause is present, then after assigning the iterations of the associated
20 loops to their respective threads, as specified in Table 2.5, the iterations may be executed in any
21 order, including concurrently.

TABLE 2.5: `schedule` Clause *kind* Values

static	<p>When <code>schedule (static, chunk_size)</code> is specified, iterations are divided into chunks of size <code>chunk_size</code>, and the chunks are assigned to the threads in the team in a round-robin fashion in the order of the thread number. Each chunk contains <code>chunk_size</code> iterations, except for the chunk that contains the sequentially last iteration, which may have fewer iterations.</p> <p>When no <code>chunk_size</code> 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. The size of the chunks is unspecified in this case.</p> <p>A compliant implementation of the static schedule must ensure that the same assignment of logical iteration numbers to threads will be used in two worksharing-loop regions if the following conditions are satisfied: 1) both worksharing-loop regions have the same number of loop iterations, 2) both worksharing-loop regions have the same value of <code>chunk_size</code> specified, or both worksharing-loop regions have no <code>chunk_size</code> specified, 3) both worksharing-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 nowait clause.</p>
dynamic	<p>When <code>schedule (dynamic, chunk_size)</code> is specified, the iterations are distributed to threads in the team in chunks. Each thread executes a chunk of iterations, then requests another chunk, until no chunks remain to be distributed.</p> <p>Each chunk contains <code>chunk_size</code> iterations, except for the chunk that contains the sequentially last iteration, which may have fewer iterations.</p> <p>When no <code>chunk_size</code> is specified, it defaults to 1.</p>
guided	<p>When <code>schedule (guided, chunk_size)</code> is specified, the iterations are assigned to threads in the team in chunks. Each thread executes a chunk of iterations, then requests another chunk, until no chunks remain to be assigned.</p>

table continued on next page

table continued from previous page

For a *chunk_size* 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 *chunk_size* with value *k* (greater than 1), the size of each chunk is determined in the same way, with the restriction that the chunks do not contain fewer than *k* iterations (except for the chunk that contains the sequentially last iteration, which may have fewer than *k* iterations).

When no *chunk_size* is specified, it defaults to 1.

auto When **schedule (auto)** 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.

runtime When **schedule (runtime)** is specified, the decision regarding scheduling is deferred until run time, and the schedule and chunk size are taken from the *run-sched-var* ICV. If the ICV is set to **auto**, the schedule is implementation defined.

▼

1 **Note** – For a team of *p* threads and a loop of *n* iterations, let $\lceil n/p \rceil$ be the integer *q* that satisfies
2 $n = p * q - r$, with $0 \leq r < p$. One compliant implementation of the **static** schedule (with no
3 specified *chunk_size*) would behave as though *chunk_size* had been specified with value *q*. Another
4 compliant implementation would assign *q* iterations to the first *p* – *r* threads, and *q* – 1 iterations to
5 the remaining *r* threads. This illustrates why a conforming program must not rely on the details of a
6 particular implementation.

7 A compliant implementation of the **guided** schedule with a *chunk_size* value of *k* would assign
8 $q = \lceil n/p \rceil$ iterations to the first available thread and set *n* to the larger of *n* – *q* and $p * k$. It would
9 then repeat this process until *q* is greater than or equal to the number of remaining iterations, at
10 which time the remaining iterations form the final chunk. Another compliant implementation could
11 use the same method, except with $q = \lceil n/(2p) \rceil$, and set *n* to the larger of *n* – *q* and $2 * p * k$.

▲

TABLE 2.6: `schedule` Clause *modifier* Values

	monotonic	When the monotonic modifier is specified then each thread executes the chunks that it is assigned in increasing logical iteration order.
	nonmonotonic	When the nonmonotonic modifier is specified then chunks are assigned to threads in any order and the behavior of an application that depends on any execution order of the chunks is unspecified.
1	simd	When the simd modifier is specified and the loop is associated with a SIMD construct, the <i>chunk_size</i> for all chunks except the first and last chunks is $new_chunk_size = \lceil chunk_size / simd_width \rceil * simd_width$ where <i>simd_width</i> is an implementation-defined value. The first chunk will have at least <i>new_chunk_size</i> iterations except if it is also the last chunk. The last chunk may have fewer iterations than <i>new_chunk_size</i> . If the simd modifier is specified and the loop is not associated with a SIMD construct, the modifier is ignored.

2 **Execution Model Events**

3 The *ws-loop-begin* event occurs after an implicit task encounters a worksharing-loop construct but
4 before the task starts the execution of the structured block of the worksharing-loop region.

5 The *ws-loop-end* event occurs after a worksharing-loop region finishes execution but before
6 resuming execution of the encountering task.

7 The *ws-loop-iteration-begin* event occurs before an implicit task executes each iteration of a
8 parallel loop.

9 **Tool Callbacks**

10 A thread dispatches a registered **ompt_callback_work** callback for each occurrence of a
11 *ws-loop-begin* and *ws-loop-end* event in that thread. The callback occurs in the context of the
12 implicit task. The callback has type signature **ompt_callback_work_t**. The callback receives
13 **ompt_scope_begin** or **ompt_scope_end** as its *endpoint* argument, as appropriate, and
14 **work_loop** as its *wstype* argument.

15 A thread dispatches a registered **ompt_callback_dispatch** callback for each occurrence of a
16 *ws-loop-iteration-begin* event in that thread. The callback occurs in the context of the implicit task.
17 The callback has type signature **ompt_callback_dispatch_t**.

Restrictions

Restrictions to the worksharing-loop construct are as follows:

- There must be no OpenMP directive in the region between any associated loops.
- If a **collapse** clause is specified, exactly one loop must occur in the region at each nesting level up to the number of loops specified by the parameter of the **collapse** clause.
- If the **ordered** clause is present, all loops associated with the construct must be perfectly nested; that is there must be no intervening code between any two loops.
- If a **reduction** clause with the **inscan** modifier is specified, neither the **ordered** nor **schedule** clause may appear on the worksharing-loop directive.
- The values of the loop control expressions of the loops associated with the worksharing-loop construct must be the same for all threads in the team.
- Only one **schedule** clause can appear on a worksharing-loop directive.
- The **schedule** clause must not appear on the worksharing-loop directive if the associated loop(s) form a non-rectangular loop nest.
- The **ordered** clause must not appear on the worksharing-loop directive if the associated loop(s) form a non-rectangular loop nest.
- Only one **collapse** clause can appear on a worksharing-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.
- The value of the *run-sched-var* ICV must be the same for all threads in the team.
- When **schedule(runtime)** or **schedule(auto)** is specified, *chunk_size* must not be specified.
- A *modifier* may not be specified on a **linear** clause.
- Only one **ordered** clause can appear on a worksharing-loop directive.
- The **ordered** clause must be present on the worksharing-loop construct if any **ordered** region ever binds to a worksharing-loop region arising from the worksharing-loop construct.
- The **nonmonotonic** modifier cannot be specified if an **ordered** clause is specified.
- Either the **monotonic** modifier or the **nonmonotonic** modifier can be specified but not both.
- The loop iteration variable may not appear in a **threadprivate** directive.
- If both the **collapse** and **ordered** clause with a parameter are specified, the parameter of the **ordered** clause must be greater than or equal to the parameter of the **collapse** clause.

- 1 • A **linear** clause or an **ordered** clause with a parameter can be specified on a
2 worksharing-loop directive but not both.
- 3 • If an **order (concurrent)** clause is present, all restrictions from the **loop** construct with an
4 **order (concurrent)** clause also apply.
- 5 • If an **order (concurrent)** clause is present, an **ordered** clause may not appear on the
6 same directive.

▼ C / C++ ▼

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

- **order (concurrent)** clause, see Section 2.12.5 on page 126.
- **ordered** construct, see Section 2.20.9 on page 243.
- **depend** clause, see Section 2.20.11 on page 248.
- **private**, **firstprivate**, **lastprivate**, **linear**, and **reduction** clauses, see Section 2.22.4 on page 276.
- **ompt_scope_begin** and **ompt_scope_end**, see Section 4.2.3.4.10 on page 437.
- **ompt_work_loop**, see Section 4.2.3.4.14 on page 438.
- **ompt_callback_work_t**, see Section 4.2.4.2.15 on page 461.
- **OMP_SCHEDULE** environment variable, see Section 5.1 on page 596.

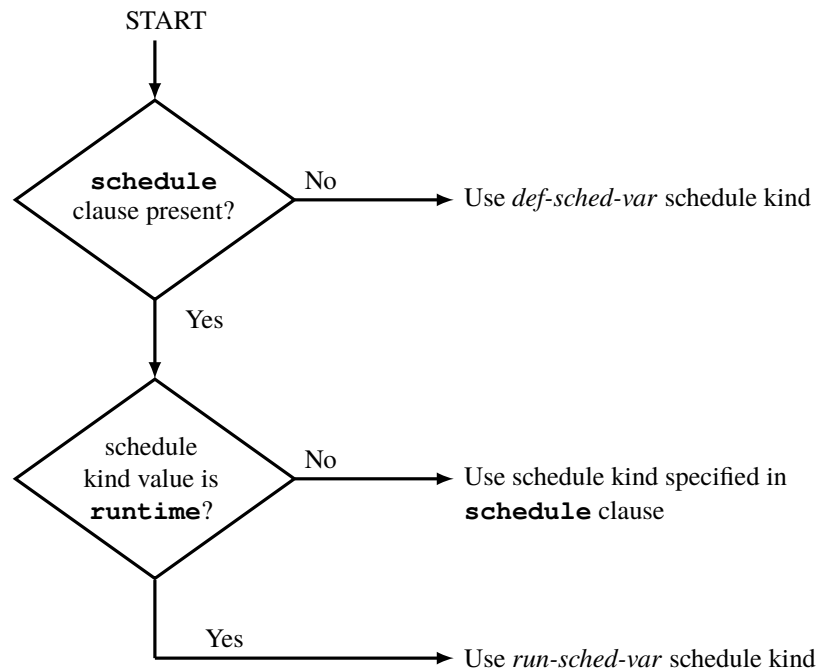


FIGURE 2.1: Determining the **schedule** for a Worksharing Loop

1 2.12.2.1 Determining the Schedule of a Worksharing Loop

2 When execution encounters a worksharing-loop directive, the **schedule** clause (if any) on the
3 directive, and the *run-sched-var* and *def-sched-var* ICVs are used to determine how loop iterations
4 are assigned to threads. See Section 2.4 on page 47 for details of how the values of the ICVs are
5 determined. If the worksharing-loop directive does not have a **schedule** clause then the current
6 value of the *def-sched-var* ICV determines the schedule. If the worksharing-loop directive has a
7 **schedule** clause that specifies the **runtime** schedule kind then the current value of the
8 *run-sched-var* ICV determines the schedule. Otherwise, the value of the **schedule** clause
9 determines the schedule. Figure 2.1 describes how the schedule for a worksharing loop is
10 determined.

11 Cross References

- 12 • ICVs, see Section 2.4 on page 47.

1 2.12.3 SIMD Constructs

2 2.12.3.1 `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++

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

11 where *clause* is one of the following:

```
12     if ([simd :] scalar-expression)  
13     safelen (length)  
14     simdlen (length)  
15     linear (list[ : linear-step])  
16     aligned (list[ : alignment])  
17     nontemporal (list)  
18     private (list)  
19     lastprivate ([ lastprivate-modifier : ] list)  
20     reduction ([reduction-modifier, ]reduction-identifier : list)  
21     collapse (n)  
22     order (concurrent)
```

23 The `simd` directive places restrictions on the structure of the associated *for-loops*. Specifically, all
24 associated *for-loops* must have *canonical loop form* (Section 2.12.1 on page 95).

C / C++

```

1  !$omp simd [clause[ [, ] clause ... ]
2      do-loops
3  [!$omp end simd]

```

where *clause* is one of the following:

```

5      if ([simd :] scalar-logical-expression)
6      safelen (length)
7      simdlen (length)
8      linear (list[ : linear-step])
9      aligned (list[ : alignment])
10     nontemporal (list)
11     private (list)
12     lastprivate ([ lastprivate-modifier : ] list)
13     reduction ([reduction-modifier, ]reduction-identifier : list)
14     collapse (n)
15     order (concurrent)

```

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

The **simd** directive places restrictions on the structure of all associated *do-loops*. Specifically, all associated *do-loops* must have *canonical loop form* (see Section 2.12.1 on page 95).

Binding

A **simd** region binds to the current task region. The binding thread set of the **simd** region is the 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 **simd** construct is the one that immediately follows the directive.

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

5 If more than one loop is associated with the **simd** construct then the number of times that any
6 intervening code between any two associated loops will be executed is unspecified but will be at
7 least once per iteration of the loop enclosing the intervening code and at most once per iteration of
8 the innermost loop associated with the construct. If the iteration count of any loop that is associated
9 with the **simd** construct and does not enclose the intervening code is zero then the behavior is
10 unspecified.

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

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

16 A SIMD loop has logical iterations numbered 0,1,...,N-1 where N is the number of loop iterations,
17 and the logical numbering denotes the sequence in which the iterations would be executed if the
18 associated loop(s) were executed with no SIMD instructions. At the beginning of each logical
19 iteration, the loop iteration variable of each associated loop has the value that it would have if the
20 set of the associated loop(s) were executed sequentially. The number of iterations that are executed
21 concurrently at any given time is implementation defined. Each concurrent iteration will be
22 executed by a different SIMD lane. Each set of concurrent iterations is a SIMD chunk. Lexical
23 forward dependencies in the iterations of the original loop must be preserved within each SIMD
24 chunk.

25 The **safelen** clause specifies that no two concurrent iterations within a SIMD chunk can have a
26 distance in the logical iteration space that is greater than or equal to the value given in the clause.
27 The parameter of the **safelen** clause must be a constant positive integer expression. The
28 **simdlen** clause specifies the preferred number of iterations to be executed concurrently unless an
29 **if** clause is present and evaluates to *false*, in which case the preferred number of iterations to be
30 executed concurrently is one. The parameter of the **simdlen** clause must be a constant positive
31 integer expression.

▼ C / C++ ▼

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

▲ C / C++ ▲

Fortran

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

Fortran

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

6 The **nontemporal** clause specifies that accesses to the storage locations to which the list items
7 refer have low temporal locality across the iterations in which those storage locations are accessed.

Restrictions

- 9 • There must be no OpenMP directive in the region between any associated loops.
- 10 • If a **collapse** clause is specified, exactly one loop must occur in the region at each nesting
11 level up to the number of loops specified by the parameter of the **collapse** clause.
- 12 • If the **ordered** clause is present, all loops associated with the construct must be perfectly
13 nested; that is there must be no intervening code between any two loops.
- 14 • The associated loops must be structured blocks.
- 15 • A program that branches into or out of a **simd** region is non-conforming.
- 16 • Only one **collapse** clause can appear on a **simd** directive.
- 17 • A *list-item* cannot appear in more than one **aligned** clause.
- 18 • A *list-item* cannot appear in more than one **nontemporal** clause.
- 19 • Only one **safelen** clause can appear on a **simd** directive.
- 20 • Only one **simdlen** clause can appear on a **simd** directive.
- 21 • If both **simdlen** and **safelen** clauses are specified, the value of the **simdlen** parameter
22 must be less than or equal to the value of the **safelen** parameter.
- 23 • A *modifier* may not be specified on a **linear** clause.
- 24 • The only OpenMP constructs that can be encountered during execution of a **simd** region are the
25 **atomic** construct, the **loop** construct, the **simd** construct and the **ordered** construct with
26 the **simd** clause.
- 27 • If an **order (concurrent)** clause is present, all restrictions from the **loop** construct with an
28 **order (concurrent)** clause also apply.

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.
- If a list item on the **aligned** clause has the **ALLOCATABLE** attribute, the allocation status must be allocated.
- If a list item on the **aligned** clause has the **POINTER** attribute, the association status must be associated.
- If the type of a list item on the **aligned** clause is either **C_PTR** or Cray pointer, the list item must be defined.

Fortran

Cross References

- **order (concurrent)** clause, see Section 2.12.5 on page 126.
- **if** Clause, see Section 2.18 on page 213.
- **private**, **lastprivate**, **linear** and **reduction** clauses, see Section 2.22.4 on page 276.

2.12.3.2 Worksharing-Loop SIMD Construct

Summary

The worksharing-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 worksharing-loop SIMD construct is a composite construct.

Syntax

C / C++

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

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

C / C++

Fortran

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

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

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

Fortran

1 **Description**

2 The worksharing-loop SIMD construct will first distribute the iterations of the associated loop(s)
3 across the implicit tasks of the parallel region in a manner consistent with any clauses that apply to
4 the worksharing-loop construct. The resulting chunks of iterations will then be converted to a
5 SIMD loop in a manner consistent with any clauses that apply to the **simd** construct.

6 **Execution Model Events**

7 This composite construct generates the same events as the worksharing-loop construct.

8 **Tool Callbacks**

9 This composite construct dispatches the same callbacks as the worksharing-loop construct.

10 **Restrictions**

11 All restrictions to the worksharing-loop construct and the **simd** construct apply to the
12 worksharing-loop SIMD construct. In addition, the following restrictions apply:

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

15 **Cross References**

- 16 • worksharing-loop construct, see Section [2.12.2](#) on page [100](#).
- 17 • **simd** construct, see Section [2.12.3.1](#) on page [111](#).
- 18 • Data attribute clauses, see Section [2.22.4](#) on page [276](#).
- 19 • Events and tool callbacks for the worksharing-loop construct, see Section [2.12.2](#) on page [100](#).

20 **2.12.4 distribute Loop Constructs**

21 **2.12.4.1 distribute Construct**

22 **Summary**

23 The **distribute** construct specifies that the iterations of one or more loops will be executed by
24 the initial teams in the context of their implicit tasks. The iterations are distributed across the initial
25 threads of all initial 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)  
lastprivate (list)  
collapse (n)  
dist_schedule (kind[ , chunk_size])  
allocate ([allocator: ]list)
```

All associated *for-loops* must have the canonical form described in Section 2.12.1 on page 95.

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)  
lastprivate (list)  
collapse (n)  
dist_schedule (kind[ , chunk_size])  
allocate ([allocator: ]list)
```

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

The **distribute** directive places restrictions on the structure of all associated *do-loops*. Specifically, all associated *do-loops* must have *canonical loop form* (see Section 2.12.1 on page 95).

Fortran

1 **Binding**

2 The binding thread set for a **distribute** region is the set of initial threads executing an
3 enclosing **teams** region. A **distribute** region binds to this **teams** region.

4 **Description**

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

7 There is no implicit barrier at the end of a **distribute** construct. To avoid data races the
8 original list items modified due to **lastprivate** or **linear** clauses should not be accessed
9 between the end of the **distribute** construct and the end of the **teams** region to which the
10 **distribute** binds.

11 The **collapse** clause may be used to specify how many loops are associated with the
12 **distribute** construct. The parameter of the **collapse** clause must be a constant positive
13 integer expression. If no **collapse** clause is present or its parameter is 1, the only loop that is
14 associated with the **distribute** construct is the one that immediately follows the **distribute**
15 construct. If a **collapse** clause is specified with a parameter value greater than 1 and more than
16 one loop is associated with the **distribute** construct, then the iteration of all associated loops
17 are collapsed into one larger iteration space. The sequential execution of the iterations in all
18 associated loops determines the order of the iterations in the collapsed iteration space.

19 A distribute loop has logical iterations numbered 0,1,...,N-1 where N is the number of loop
20 iterations, and the logical numbering denotes the sequence in which the iterations would be
21 executed if the set of associated loop(s) were executed sequentially. At the beginning of each
22 logical iteration, the loop iteration variable of each associated loop has the value that it would have
23 if the set of the associated loop(s) were executed sequentially.

24 If more than one loop is associated with the **distribute** construct then the number of times that
25 any intervening code between any two associated loops will be executed is unspecified but will be
26 at least once per iteration of the loop enclosing the intervening code and at most once per iteration
27 of the innermost loop associated with the construct. If the iteration count of any loop that is
28 associated with the **distribute** construct and does not enclose the intervening code is zero then
29 the behavior is unspecified.

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

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

35 If **dist_schedule** is specified, *kind* must be **static**. If specified, iterations are divided into
36 chunks of size *chunk_size*, chunks are assigned to the initial teams of the league in a round-robin
37 fashion in the order of the initial team number. When no *chunk_size* is specified, the iteration space

1 is divided into chunks that are approximately equal in size, and at most one chunk is distributed to
2 each initial team of the league. The size of the chunks is unspecified in this case.

3 When no **dist_schedule** clause is specified, the schedule is implementation defined.

4 Execution Model Events

5 The *distribute-begin* event occurs after an implicit task encounters a **distribute** construct but
6 before the task starts the execution of the structured block of the **distribute** region.

7 The *distribute-end* event occurs after a **distribute** region finishes execution but before
8 resuming execution of the encountering task.

9 Tool Callbacks

10 A thread dispatches a registered **ompt_callback_work** callback for each occurrence of a
11 *distribute-begin* and *distribute-end* event in that thread. The callback occurs in the context of the
12 implicit task. The callback has type signature **ompt_callback_work_t**. The callback receives
13 **ompt_scope_begin** or **ompt_scope_end** as its *endpoint* argument, as appropriate, and
14 **ompt_work_distribute** as its *wstype* argument.

15 Restrictions

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

- 17 • The **distribute** construct inherits the restrictions of the worksharing-loop construct.
- 18 • Each **distribute** region must be encountered by the initial threads of all initial teams in a
19 league or by none at all.
- 20 • The sequence of the **distribute** regions encountered must be the same for every initial thread
21 of every initial team in a league.
- 22 • The region corresponding to the **distribute** construct must be strictly nested inside a **teams**
23 region.
- 24 • A list item may appear in a **firstprivate** or **lastprivate** clause but not both.
- 25 • The **dist_schedule** clause must not appear on the **distribute** directive if the associated
26 loop(s) form a non-rectangular loop nest.

27 Cross References

- 28 • **teams** construct, see Section [2.10](#) on page [81](#)
- 29 • workshring-loop construct, see Section [2.12.2](#) on page [100](#).
- 30 • **ompt_work_distribute**, see Section [4.2.3.4.14](#) on page [438](#).
- 31 • **ompt_callback_work_t**, see Section [4.2.4.2.15](#) on page [461](#).



1 2.12.4.2 **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:

8  **C / C++** 
9 `#pragma omp distribute simd [clause[[,] clause] ...] newline`
10 `for-loops`

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

13  **C / C++** 

14  **Fortran** 
15 `!$omp distribute simd [clause[[,] clause] ...]`
16 `do-loops`
17 `[$omp end distribute simd]`

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

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

22  **Fortran** 

23 **Description**

24 The **distribute simd** construct will first distribute the iterations of the associated loop(s)
25 according to the semantics of the **distribute** construct and any clauses that apply to the
26 distribute construct. The resulting chunks of iterations will then be converted to a SIMD loop in a
27 manner consistent with any clauses that apply to the **simd** construct.

28 **Execution Model Events**

29 This composite construct generates the same events as the **distribute** construct.

1 Tool Callbacks

2 This composite construct dispatches the same callbacks as the **distribute** construct.

3 Restrictions

- 4 • The restrictions for the **distribute** and **simd** constructs apply.
- 5 • A list item may not appear in a **linear** clause, unless it is the loop iteration variable.
- 6 • The **conditional** modifier may not appear in a **lastprivate** clause.

7 Cross References

- 8 • **simd** construct, see Section [2.12.3.1](#) on page [111](#).
- 9 • **distribute** construct, see Section [2.12.4.1](#) on page [117](#).
- 10 • Data attribute clauses, see Section [2.22.4](#) on page [276](#).
- 11 • Events and tool callbacks for the **distribute** construct, see Section [2.12.3.1](#) on page [111](#).

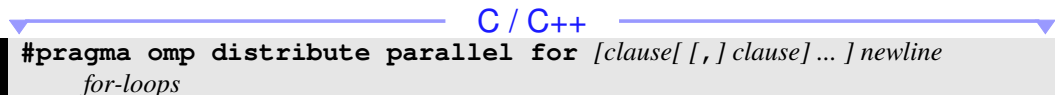
12 2.12.4.3 Distribute Parallel Worksharing-Loop Construct

13 Summary

14 The distribute parallel worksharing-loop construct specifies a loop that can be executed in parallel
15 by multiple threads that are members of multiple teams. The distribute parallel worksharing-loop
16 construct is a composite construct.

17 Syntax

18 The syntax of the distribute parallel worksharing-loop construct is as follows:

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

21 where *clause* can be any of the clauses accepted by the **distribute** or parallel worksharing-loop
22 directives with identical meanings and restrictions.

```
 C / C++
```

Fortran

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

4 where *clause* can be any of the clauses accepted by the **distribute** or parallel worksharing-loop
5 directives with identical meanings and restrictions.

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

Fortran

8 Description

9 The distribute parallel worksharing-loop construct will first distribute the iterations of the
10 associated loop(s) into chunks according to the semantics of the **distribute** construct and any
11 clauses that apply to the **distribute** construct. Each of these chunks will form a loop. Each
12 resulting loop will then be distributed across the threads within the **teams** region to which the
13 **distribute** construct binds in a manner consistent with any clauses that apply to the parallel
14 worksharing-loop construct.

15 Execution Model Events

16 This composite construct generates the same events as the **distribute** and parallel
17 worksharing-loop constructs.

18 Tool Callbacks

19 This composite construct dispatches the same callbacks as the **distribute** and parallel
20 worksharing-loop constructs.

21 Restrictions

- 22 • The restrictions for the **distribute** and parallel worksharing-loop constructs apply.
- 23 • No **ordered** clause can be specified.
- 24 • No **linear** clause can be specified.
- 25 • The **conditional** modifier may not appear in a **lastprivate** clause.

Cross References

- **distribute** construct, see Section 2.12.4.1 on page 117.
- Parallel worksharing-loop construct, see Section 2.16.1 on page 185.
- Data attribute clauses, see Section 2.22.4 on page 276.
- Events and tool callbacks for **distribute** construct, see Section 2.12.4.1 on page 117.
- Events and tool callbacks for parallel worksharing-loop construct, see Section 2.16.1 on page 185.

2.12.4.4 Distribute Parallel Worksharing-Loop SIMD Construct

Summary

The distribute parallel worksharing-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 worksharing-loop SIMD construct is a composite construct.

Syntax

C / C++

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

```
#pragma omp distribute parallel for simd \  
    [clause[ [, ] clause] ... ] newline  
    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++

Fortran

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

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

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

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

Fortran

1 **Description**

2 The distribute parallel worksharing-loop SIMD construct will first distribute the iterations of the
3 associated loop(s) according to the semantics of the **distribute** construct and any clauses that
4 apply to the **distribute** construct. The resulting loops will then be distributed across the
5 threads contained within the **teams** region to which the **distribute** construct binds in a
6 manner consistent with any clauses that apply to the parallel worksharing-loop construct. The
7 resulting chunks of iterations will then be converted to a SIMD loop in a manner consistent with
8 any clauses that apply to the **simd** construct.

9 **Execution Model Events**

10 This composite construct generates the same events as the **distribute** and parallel
11 worksharing-loop constructs.

12 **Tool Callbacks**

13 This composite construct dispatches the same callbacks as the **distribute** and parallel loop
14 constructs.

15 **Restrictions**

- 16 • The restrictions for the **distribute** and parallel worksharing-loop SIMD constructs apply.
- 17 • No **ordered** clause can be specified.
- 18 • A list item may not appear in a **linear** clause, unless it is the loop iteration variable.
- 19 • The **conditional** modifier may not appear in a **lastprivate** clause.

20 **Cross References**

- 21 • **distribute** construct, see Section [2.12.4.1](#) on page [117](#).
- 22 • Parallel worksharing-loop SIMD construct, see Section [2.16.5](#) on page [191](#).
- 23 • Data attribute clauses, see Section [2.22.4](#) on page [276](#).
- 24 • Events and tool callbacks for **distribute** construct, see Section [2.12.4.1](#) on page [117](#).
- 25 • Events and tool callbacks for parallel worksharing-loop construct, see Section [2.16.1](#) on
26 page [185](#).

1 2.12.5 loop Construct

2 Summary

3 A **loop** construct specifies that the iterations of the associated loops may execute concurrently and
4 permits the encountering thread(s) to execute the loop accordingly.

5 Syntax

C / C++

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

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

9 where *clause* is one of the following:

```
10 bind (binding)  
11 collapse (n)  
12 order (concurrent)  
13 private (list)  
14 reduction (reduction-identifier : list)
```

15 The **loop** directive places restrictions on the structure of all associated *for-loops*. Specifically, all
16 associated *for-loops* must have *canonical loop form* (see Section 2.12.1 on page 95).

C / C++

Fortran

17 The syntax of the **loop** construct is as follows:

```
18 !$omp loop [clause[ [, ] clause] ... ]  
19 do-loops  
20 [!$omp end loop]
```

21 where *clause* is one of the following:

```
22 bind (binding)  
23 collapse (n)  
24 order (concurrent)  
25 private (list)  
26 reduction (reduction-identifier : list)
```

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

3 The **loop** directive places restrictions on the structure of all associated *do-loops*. Specifically, all
4 associated *do-loops* must have *canonical loop form* (see Section 2.12.1 on page 95).

Fortran

5 Where *binding* is one of the following:

```
6     teams  
7     parallel  
8     thread
```

9 Binding

10 If the **bind** clause is present on the construct, the binding region is determined by *binding*.
11 Otherwise, if the **loop** construct is closely nested inside a **teams** or **parallel** construct, the
12 binding region is the corresponding **teams** or **parallel** region. If none of the above conditions
13 are true, the **loop** region does not have a binding region.

14 If the binding region is a **teams** region, then the binding thread set is the set of master threads
15 executing that region. If the binding region is a **parallel** region, then the binding thread set is
16 the team of threads executing that region. If the **loop** region does not have a binding region, then
17 the binding thread set is the encountering thread.

18 Description

19 The **loop** construct is associated with a loop nest consisting of one or more loops that follow the
20 directive. The directive asserts that the iterations may execute in any order, including concurrently.

21 If the **bind** clause is present then *binding* may be one of the following: **teams**, **parallel**, or
22 **thread**. If *binding* is **teams**, then the innermost enclosing **teams** region is the binding region.
23 If *binding* is **parallel**, then the innermost enclosing **parallel** region is the binding region. If
24 *binding* is **thread**, then the loop does not have a binding region and the binding thread set is the
25 encountering thread.

26 The **collapse** clause may be used to specify how many loops are associated with the **loop**
27 construct. The parameter of the **collapse** clause must be a constant positive integer expression.
28 If a **collapse** clause is specified with a parameter value greater than 1, then the iterations of the
29 associated loops to which the clause applies are collapsed into one larger iteration space with
30 unspecified ordering. If no **collapse** clause is present or its parameter is 1, the only loop that is
31 associated with the **loop** construct is the one that immediately follows the **loop** directive.

32 If more than one loop is associated with the **loop** construct then the number of times that any
33 intervening code between any two associated loops will be executed is unspecified but will be at

1 least once per iteration of the loop enclosing the intervening code and at most once per iteration of
2 the innermost loop associated with the construct. If the iteration count of any loop that is associated
3 with the **loop** construct and does not enclose the intervening code is zero then the behavior is
4 unspecified.

5 The iteration space of the associated loops correspond to logical iterations numbered 0,1,...,N-1
6 where N is the number of loop iterations, and the logical numbering denotes the sequence in which
7 the iterations would be executed if a set of associated loop(s) were executed sequentially. At the
8 beginning of each logical iteration, the loop iteration variable of each associated loop has the value
9 that it would have if the set of the associated loop(s) were executed sequentially.

10 Each logical iteration is executed once per instance of the **loop** region that is encountered by the
11 binding thread set.

12 If the **order (concurrent)** clause appears on the **loop** construct, the iterations of the
13 associated loops may execute in any order, including concurrently. If the **order** clause is not
14 present, the behavior is as if the **order (concurrent)** clause appeared on the construct.

15 The set of threads that may execute the iterations of the **loop** region is the binding thread set. Each
16 iteration is executed by one thread from this set.

17 If the **loop** region binds to a **teams** region, the threads in the binding thread set may continue
18 execution after the **loop** region without waiting for all iterations of the associated loop(s) to
19 complete. The iterations are guaranteed to complete before the end of the **teams** region.

20 If the **loop** region does not bind to a **teams** region, all iterations of the associated loop(s) must
21 complete before the encountering thread(s) continue execution after the **loop** region.

22 **Restrictions**

23 Restrictions to the **loop** construct are as follows:

- 24 • If the **collapse** clause exists there may be no intervening OpenMP directives between the
25 associated loops.
- 26 • At most one **collapse** clause can appear on a **loop** directive.
- 27 • If a **loop** construct is not nested inside another OpenMP construct and it appears in a procedure
28 called from the program, the **bind** clause must be present.
- 29 • If a **loop** region binds to a **teams** or **parallel** region, it must be encountered by all threads
30 in the binding thread set or by none of them.
- 31 • The only constructs that may be nested inside a **loop** region are the **loop** construct, the
32 **parallel** construct, the **simd** construct, and combined constructs for which the first construct
33 is a **parallel** construct.
- 34 • A **loop** region corresponding to a **loop** construct may not contain calls to procedures that
35 contain OpenMP directives.

1 • A **loop** region corresponding to a **loop** construct may not contain calls to the OpenMP
2 Runtime API.

3 • If a **threadprivate** variable is referenced inside a **loop** region, the behavior is unspecified.

▼ C / C++ ▼

4 • The associated for-loops must be structured blocks.

5 • No statement can branch to any associated **for** statement.

▲ C / C++ ▲

▼ Fortran ▼

6 • The associated do-loops must be structured blocks.

7 • No statement in the associated loops other than the DO statements can cause a branch out of the
8 loops.

▲ Fortran ▲

9 **Cross References**

10 • The Worksharing-Loop construct, see Section [2.12.2](#) on page [100](#).

11 • **distribute** construct, see Section [2.12.4.1](#) on page [117](#).

12 • SIMD constructs, see Section [2.12.3](#) on page [111](#).

13 • The **single** construct, see Section [2.11.2](#) on page [89](#).

14 **2.12.6 scan Directive**

15 **Summary**

16 The **scan** directive specifies that a scan computation is to be performed over the values used on
17 each iteration to update a list item.

Syntax

C / C++

The syntax of the **scan** directive is as follows:

```
loop-directive
for-loop-headers(s)
{
    structured-block
    #pragma omp scan clause new-line
    structured-block
}
```

where *clause* is one of the following:

```
inclusive (list)
exclusive (list)
```

and where *loop-directive* is a **for**, **for simd**, or **simd** directive.

C / C++

Fortran

The syntax of the **scan** directive is as follows:

```
loop-directive
do-loop-header(s)
    structured-block
    !$omp scan clause
    structured-block
do-termination-stmts(s)
[end-loop-directive]
```

where *clause* is one of the following:

```
inclusive (list)
exclusive (list)
```

and where *loop-directive* (*end-loop-directive*) is a **do** (**end do**), **do simd** (**end do simd**), or **simd** (**end simd**) directive.

Fortran

Description

The **scan** directive may appear in the body of a loop or loop nest associated with an enclosing worksharing-loop, worksharing-loop SIMD, or **simd** construct, to specify that one or more scan computations are to be performed by the loop. The directive specifies that either an inclusive scan computation is to be performed for each list item that appears in an **inclusive** clause on the directive, or an exclusive scan computation is to be performed for each list item that appears in an **exclusive** clause on the directive. For each list item for which a scan computation is specified, statements that lexically precede or follow the directive constitute one of two phases for a given logical iteration of the loop – an *input phase* or a *scan phase*.

If the list item appears in an **inclusive** clause, all statements in the structured block that lexically precede the directive constitute the input phase and all statements in the structured block that lexically follow the directive constitute the scan phase. If the list item appears in an **exclusive** clause and the iteration is not the last iteration, all statements in the structured block that lexically precede the directive constitute the scan phase and all statements in the structured block that lexically follow the directive constitute the input phase. If the list item appears in an **exclusive** clause and the iteration is the last iteration, there is no input phase for the iteration and all statements that lexically precede or follow the directive constitute the scan phase for the iteration. The input phase contains all computations that update the list item in the iteration, and the scan phase ensures that any statement that reads the list item will see the result of the scan computation for that iteration.

The result of a scan computation for a given iteration is calculated according to the last *generalized prefix sum* ($\text{PRESUM}_{\text{last}}$) applied over the sequence of values given by the original value of the list item prior to the loop and all preceding updates to the list item in the logical iteration space of the loop. The operation $\text{PRESUM}_{\text{last}}(op, a_1, \dots, a_N)$ is defined for a given binary operator op and a sequence of N values a_1, \dots, a_N as follows:

- if $N = 1$, a_1
- if $N > 1$, $op(\text{PRESUM}_{\text{last}}(op, a_1, \dots, a_K), \text{PRESUM}_{\text{last}}(op, a_L, \dots, a_N))$, where $1 \leq K + 1 = L \leq N$.

If the operator op is not a mathematically associative operation, the result of the $\text{PRESUM}_{\text{last}}$ operation is nondeterministic.

At the beginning of the input phase of each iteration, the list item is initialized with the initializer value of the *reduction-identifier* specified by the **reduction** clause on the innermost enclosing construct. The *update value* of a list item is, for a given iteration, the value of the list item on completion of its input phase.

Let *orig-val* be the value of the original list item on entry to enclosing worksharing-loop, worksharing-loop SIMD, or **simd** construct. Let *combiner* be the combiner for the *reduction-identifier* specified by the **reduction** clause on the construct. And let u_I be the update value of a list item for iteration I . For list items appearing in an **inclusive** clause on the **scan** directive, at the beginning of the scan phase for iteration I the list item is assigned the result of the

1 operation $\text{PRESUM}_{\text{last}}(\text{combiner}, \text{orig-val}, u_0, \dots, u_1)$. For list items appearing in an
2 **exclusive** clause on the **scan** directive, at the beginning of the scan phase for iteration $I = 0$
3 the list item is assigned the value *orig-val*, and at the beginning of the scan phase for iteration $I > 0$
4 the list item is assigned the result of the operation $\text{PRESUM}_{\text{last}}(\text{combiner}, \text{orig-val}, u_0, \dots, u_{I-1})$.

5 Restrictions

6 Restrictions to the **scan** directive are as follows:

- 7 • Exactly one **scan** directive must appear in the loop body of an enclosing worksharing-loop,
8 worksharing-loop SIMD, or **simd** construct on which a **reduction** clause with the **inscan**
9 modifier is present.
- 10 • A list item that appears in the **inclusive** or **exclusive** clause must appear in a
11 **reduction** clause with the **inscan** modifier on the enclosing worksharing-loop,
12 worksharing-loop SIMD, or **simd** construct.
- 13 • Cross-iteration dependences across different logical iterations must not exist, except for
14 dependences for the list items specified in an **inclusive** or **exclusive** clause.
- 15 • Intra-iteration dependences from a statement in the structured block preceding a **scan** directive
16 to a statement in the structured block following a **scan** directive must not exist, except for
17 dependences for the list items specified in an **inclusive** or **exclusive** clause.

18 Cross References

- 19 • worksharing-loop construct, see Section [2.12.2](#) on page [100](#).
- 20 • **simd** construct, see Section [2.12.3.1](#) on page [111](#).
- 21 • worksharing-loop SIMD construct, see Section [2.12.3.2](#) on page [116](#).
- 22 • **reduction** clause, see Section [2.22.5.4](#) on page [297](#).

1 2.13 Tasking Constructs

2 2.13.1 task Construct

3 Summary

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

5 Syntax

C / C++

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

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

9 where *clause* is one of the following:

```
10     if ([ task :] scalar-expression)  
11     final (scalar-expression)  
12     untied  
13     default (shared | none)  
14     mergeable  
15     private (list)  
16     firstprivate (list)  
17     shared (list)  
18     in_reduction (reduction-identifier : list)  
19     depend ([depend-modifier :][dependence-type :] locator-list)  
20     priority (priority-value)  
21     allocate ([allocator: ]list)  
22     affinity ([aff-modifier :] locator-list)  
23     detach (event-handler)
```

24 where *aff-modifier* is one of the following:

```
25     iterator (iterators-definition)
```

26 where *event-handler* is a variable of the **omp_event_t *** type.

C / C++

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

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

5 where *clause* is one of the following:

```
6     if([ task :] scalar-logical-expression)
7     final (scalar-logical-expression)
8     untied
9     default(private | firstprivate | shared | none)
10    mergeable
11    private (list)
12    firstprivate (list)
13    shared (list)
14    in_reduction (reduction-identifier : list)
15    depend ([depend-modifier :][dependence-type :] locator-list)
16    priority (priority-value)
17    allocate ([allocator :] list)
18    affinity ([aff-modifier :] locator-list)
19    detach (event-handler)
```

20 where *aff-modifier* is one of the following:

```
21    iterator (iterators-definition)
```

22 where *event-handler* is a variable of the **omp_event_kind** integer *kind*

23 **Binding**

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

Description

The **task** construct is a *task generating construct*. When a thread encounters a **task** construct, an explicit task is generated from the code for the associated *structured-block*. The data environment of the task is created according to the data-sharing attribute clauses on the **task** construct, per-data environment ICVs, and any defaults that apply. The data environment of the task is destroyed when the execution code of the associated *structured-block* is completed.

The encountering thread may immediately execute the task, or defer its execution. In the latter case, any thread in the team may be assigned the task. Completion of the task can be guaranteed using task synchronization constructs. If a **task** construct is encountered during execution of an outer task, the generated **task** region corresponding to this construct is not a part of the outer task region unless the generated task is an included task.

If a **detach** clause is present on a **task** construct a new event of type **omp_event_t**, *allow-completion-event*, is created. The *allow-completion-event* is connected to the completion of the associated **task** region. The original *event-handler* will be updated to point to the *allow-completion-event* event before the task data environment is created. The *event-handler* will be considered as if it was specified on a **firstprivate** clause. Note that the use of a variable in a **detach** clause expression of a **task** construct causes an implicit reference to the variable in all enclosing constructs.

If no **detach** clause is present on a **task** construct the generated **task** is completed when the execution of its associated *structured-block* is completed. If a **detach** clause is present on a **task** construct the task is completed when the execution of its associated *structured-block* is completed and the *allow-completion-event* is fulfilled.

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

When a **final** clause is present on a **task** construct and the **final** clause expression evaluates to *true*, the generated task will be a final task. All **task** constructs encountered during execution of a final task will generate final and included tasks. Note that the use of a variable in a **final** clause expression of a **task** construct causes an implicit reference to the variable in all enclosing constructs. Encountering a **task** construct with the **detach** clause during the execution of a final task results in unspecified behavior.

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

A thread that encounters a task scheduling point within the **task** region may temporarily suspend the **task** region. By default, a task is tied and its suspended **task** region can only be resumed by the thread that started its execution. If the **untied** clause is present on a **task** construct, any thread in the team can resume the **task** region after a suspension. The **untied** clause is ignored

1 if a **final** clause is present on the same **task** construct and the **final** clause expression
2 evaluates to *true*, or if a task is an included task.

3 The **task** construct includes a task scheduling point in the task region of its generating task,
4 immediately following the generation of the explicit task. Each explicit **task** region includes a
5 task scheduling point at the end of its associated *structured-block*.

6 When the **mergeable** clause is present on a **task** construct, the generated task is a *mergeable*
7 *task*.

8 The **priority** clause is a hint for the priority of the generated task. The *priority-value* is a
9 non-negative integer expression that provides a hint for task execution order. Among all tasks ready
10 to be executed, higher priority tasks (those with a higher numerical value in the **priority** clause
11 expression) are recommended to execute before lower priority ones. The default *priority-value*
12 when no **priority** clause is specified is zero (the lowest priority). If a value is specified in the
13 **priority** clause that is higher than the *max-task-priority-var* ICV then the implementation will
14 use the value of that ICV. A program that relies on task execution order being determined by this
15 *priority-value* may have unspecified behavior.

16 The **affinity** clause is a hint to indicate data affinity of the generated task. The task is
17 recommended to execute closely to the location of the list items. A program that relies on the task
18 execution location being determined by this list may have unspecified behavior.

19 The list items that appear in the **affinity** clause may reference iterators defined by an
20 *iterators-definition* appearing on the same clause. The list items that appear in the **affinity**
21 clause may include array sections.

▼ C / C++ ▼

22 The list items that appear in the **affinity** clause may use shape-operators.

▲ C / C++ ▲

23 If a list item appears in an **affinity** clause then data affinity refers to the original list item.

24 Note – When storage is shared by an explicit **task** region, the programmer must ensure, by
25 adding proper synchronization, that the storage does not reach the end of its lifetime before the
26 explicit **task** region completes its execution.

27 Execution Model Events

28 The *task-create* event occurs when a thread encounters a construct that causes a new task to be
29 created. The event occurs after the task is initialized but before it begins execution or is deferred.

1 Tool Callbacks

2 A thread dispatches a registered `ompt_callback_task_create` callback for each occurrence
3 of a `task-create` event in the context of the encountering task. This callback has the type signature
4 `ompt_callback_task_create_t`. In the dispatched callback, `(task_type &`
5 `ompt_task_explicit)` always evaluates to `true`. If the task is an undeferred task, then
6 `(task_type & ompt_task_undeferred)` evaluates to `true`. If the task is a final task,
7 `(task_type & ompt_task_final)` evaluates to `true`. If the task is an untied task,
8 `(task_type & ompt_task_untied)` evaluates to `true`. If the task is a mergeable task,
9 `(task_type & ompt_task_mergeable)` evaluates to `true`. If the task is a merged task,
10 `(task_type & ompt_task_merged)` evaluates to `true`.

11 Restrictions

12 Restrictions to the `task` construct are as follows:

- 13 • A program that branches into or out of a `task` region is non-conforming.
- 14 • A program must not depend on any ordering of the evaluations of the clauses of the `task`
15 directive, or on any side effects of the evaluations of the clauses.
- 16 • At most one `if` clause can appear on the directive.
- 17 • At most one `final` clause can appear on the directive.
- 18 • At most one `priority` clause can appear on the directive.
- 19 • At most one `detach` clause can appear on the directive.
- 20 • If a `detach` clause appears on the directive, then a `mergeable` clause cannot appear on the
21 same directive.

22  C / C++

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

 C / C++

 Fortran

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

 Fortran

Cross References

- Task scheduling constraints, see Section 2.13.6 on page 147.
- **if** Clause, see Section 2.18 on page 213.
- **depend** clause, see Section 2.20.11 on page 248.
- Data-sharing attribute clauses, Section 2.22.4 on page 276.
- **omp_fulfill_event**, see Section 3.5.1 on page 392.
- **ompt_callback_task_create_t**, see Section 4.2.4.2.6 on page 451.

2.13.2 taskloop Construct

Summary

The **taskloop** construct specifies that the iterations of one or more associated loops will be executed in parallel using explicit tasks. The iterations are distributed across tasks generated by the 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:


```
if ([ taskloop :] scalar-expr)  
shared (list)  
private (list)  
firstprivate (list)  
lastprivate (list)  
reduction (reduction-identifier : list)  
in_reduction (reduction-identifier : list)  
default (shared | none)  
grainsize (grain-size)  
num_tasks (num-tasks)
```

```

1  collapse(n)
2  final(scalar-expr)
3  priority(priority-value)
4  untied
5  mergeable
6  nogroup
7  allocate([allocator: list)

```

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 2.12.1 on page 95).



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

```

11  !$omp taskloop [clause[[, clause] ...]
12      do-loops
13  [!$omp end taskloop]

```

where *clause* is one of the following:

```

15  if([ taskloop :] scalar-logical-expr)
16  shared(list)
17  private(list)
18  firstprivate(list)
19  lastprivate(list)
20  reduction(reduction-identifier : list)
21  in_reduction(reduction-identifier : list)
22  default(private | firstprivate | shared | none)
23  grainsize(grain-size)
24  num_tasks(num-tasks)
25  collapse(n)
26  final(scalar-logical-expr)
27  priority(priority-value)
28  untied
29  mergeable
30  nogroup

```

1 **allocate** (*[allocator:]list*)

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

4 The **taskloop** directive places restrictions on the structure of all associated *do-loops*.
5 Specifically, all associated *do-loops* must have canonical loop form (see Section 2.12.1 on page 95).

Fortran

6 **Binding**

7 The binding thread set of the **taskloop** region is the current team. A **taskloop** region binds to
8 the innermost enclosing **parallel** region.

9 **Description**

10 The **taskloop** construct is a *task generating construct*. When a thread encounters a **taskloop**
11 construct, the construct partitions the associated loops into explicit tasks for parallel execution of
12 the loops' iterations. The data environment of each generated task is created according to the
13 data-sharing attribute clauses on the **taskloop** construct, per-data environment ICVs, and any
14 defaults that apply. The order of the creation of the loop tasks is unspecified. Programs that rely on
15 any execution order of the logical loop iterations are non-conforming.

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

20 If a **reduction** clause is present on the **taskloop** construct, the behavior is as if a
21 **task_reduction** clause with the same reduction operator and list items was applied to the
22 implicit **taskgroup** construct enclosing the **taskloop** construct. Furthermore, the **taskloop**
23 construct executes as if each generated task was defined by a **task** construct on which an
24 **in_reduction** clause with the same reduction operator and list items is present. Thus, the
25 generated tasks are participants of the reduction defined by the **task_reduction** clause that was
26 applied to the implicit **taskgroup** construct.

27 If an **in_reduction** clause is present on the **taskloop** construct, the behavior is as if each
28 generated task was defined by a **task** construct on which an **in_reduction** clause with the
29 same reduction operator and list items is present. Thus, the generated tasks are participants of a
30 reduction previously defined by a reduction scoping clause.

31 If a **grainsize** clause is present on the **taskloop** construct, the number of logical loop
32 iterations assigned to each generated task is greater than or equal to the minimum of the value of
33 the *grain-size* expression and the number of logical loop iterations, but less than two times the value
34 of the *grain-size* expression.

1 The parameter of the **grainsize** clause must be a positive integer expression. If **num_tasks** is
2 specified, the **taskloop** construct creates as many tasks as the minimum of the *num-tasks*
3 expression and the number of logical loop iterations. Each task must have at least one logical loop
4 iteration. The parameter of the **num_tasks** clause must evaluate to a positive integer. If neither a
5 **grainsize** nor **num_tasks** clause is present, the number of loop tasks generated and the
6 number of logical loop iterations assigned to these tasks is implementation defined.

7 The **collapse** clause may be used to specify how many loops are associated with the **taskloop**
8 construct. The parameter of the **collapse** clause must be a constant positive integer expression.
9 If no **collapse** clause is present or its parameter is 1, the only loop that is associated with the
10 **taskloop** construct is the one that immediately follows the **taskloop** directive. If a
11 **collapse** clause is specified with a parameter value greater than 1 and more than one loop is
12 associated with the **taskloop** construct, then the iterations of all associated loops are collapsed
13 into one larger iteration space that is then divided according to the **grainsize** and **num_tasks**
14 clauses. The sequential execution of the iterations in all associated loops determines the order of
15 the iterations in the collapsed iteration space.

16 If more than one loop is associated with the **taskloop** construct then the number of times that
17 any intervening code between any two associated loops will be executed is unspecified but will be
18 at least once per iteration of the loop enclosing the intervening code and at most once per iteration
19 of the innermost loop associated with the construct. If the iteration count of any loop that is
20 associated with the **taskloop** construct and does not enclose the intervening code is zero then the
21 behavior is unspecified.

22 A taskloop loop has logical iterations numbered 0,1,...,N-1 where N is the number of loop
23 iterations, and the logical numbering denotes the sequence in which the iterations would be
24 executed if the set of associated loop(s) were executed sequentially. At the beginning of each
25 logical iteration, the loop iteration variable of each associated loop has the value that it would have
26 if the set of the associated loop(s) were executed sequentially.

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



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



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

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

39 When a **priority** clause is present on a **taskloop** construct, the generated tasks have the

1 *priority-value* as if it was specified for each individual task. If the **priority** clause is not
2 specified, tasks generated by the **taskloop** construct have the default task priority (zero).
3 If the **untied** clause is specified, all tasks generated by the **taskloop** construct are untied tasks.
4 When the **mergeable** clause is present on a **taskloop** construct, each generated task is a
5 *mergeable task*.

6  C++ 
7 For **firstprivate** variables of class type, the number of invocations of copy constructors to
perform the initialization is implementation-defined.

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

11 Execution Model Events

12 The *taskloop-begin* event occurs after a task encounters a **taskloop** construct but before any
13 other events that may trigger as a consequence of executing the **taskloop**. Specifically, a
14 *taskloop-begin* event for a **taskloop** will precede the *taskgroup-begin* that occurs unless a
15 **nogroup** clause is present. Regardless of whether an implicit taskgroup is present, a
16 *taskloop-begin* will always precede any *task-create* events for generated tasks.

17 The *taskloop-end* event occurs after a **taskloop** region finishes execution but before resuming
18 execution of the encountering task.

19 The *taskloop-iteration-begin* event occurs before an implicit task executes each iteration of a
20 **taskloop**.

21 Tool Callbacks

22 A thread dispatches a registered **ompt_callback_work** callback for each occurrence of a
23 *taskloop-begin* and *taskloop-end* event in that thread. The callback occurs in the context of the
24 encountering task. The callback has type signature **ompt_callback_work_t**. The callback
25 receives **ompt_scope_begin** or **ompt_scope_end** as its *endpoint* argument, as appropriate,
26 and **ompt_work_taskloop** as its *wstype* argument.

27 A thread dispatches a registered **ompt_callback_dispatch** callback for each occurrence of a
28 *taskloop-iteration-begin* event in that thread. The callback occurs in the context of the implicit task.
29 The callback has type signature **ompt_callback_dispatch_t**.

Restrictions

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

- A program that branches into or out of a **taskloop** region is non-conforming.
- There must be no OpenMP directive in the region between any associated loops.
- If a **collapse** clause is specified, exactly one loop must occur in the region at each nesting level up to the number of loops specified by the parameter of the **collapse** clause.
- If the **ordered** clause is present, all loops associated with the construct must be perfectly nested; that is there must be no intervening code between any two loops.
- If a **reduction** clause is present on the **taskloop** directive, the **nogroup** clause must not be specified.
- The same list item cannot appear in both a **reduction** and an **in_reduction** clause.
- At most one **grainsize** clause can appear on a **taskloop** directive.
- At most one **num_tasks** clause can appear on a **taskloop** directive.
- The **grainsize** clause and **num_tasks** clause are mutually exclusive and may not appear on the same **taskloop** directive.
- At most one **collapse** clause can appear on a **taskloop** directive.
- At most one **if** clause can appear on the directive.
- At most one **final** clause can appear on the directive.
- At most one **priority** clause can appear on the directive.

Cross References

- **task** construct, Section [2.13.1](#) on page [133](#).
- **if** Clause, see Section [2.18](#) on page [213](#).
- **taskgroup** construct, Section [2.20.6](#) on page [225](#).
- Data-sharing attribute clauses, Section [2.22.4](#) on page [276](#).
- **ompt_scope_begin** and **ompt_scope_end**, see Section [4.2.3.4.10](#) on page [437](#).
- **ompt_work_taskloop**, see Section [4.2.3.4.14](#) on page [438](#).
- **ompt_callback_work_t**, see Section [4.2.4.2.15](#) on page [461](#).
- **ompt_callback_dispatch_t**, see Section [4.2.4.2.17](#) on page [464](#).

1 2.13.3 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 explicit tasks. The
5 **taskloop simd** construct is a composite construct.

6 Syntax

C / C++

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

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

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

C / C++

Fortran

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

```
13 !$omp taskloop simd [clause[[,] clause] ...]  
14     do-loops  
15 [!$omp end taskloop simd]
```

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

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

Fortran

20 Binding

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

1 **Description**

2 The **taskloop simd** construct will first distribute the iterations of the associated loop(s) across
3 tasks in a manner consistent with any clauses that apply to the **taskloop** construct. The resulting
4 tasks will then be converted to a SIMD loop in a manner consistent with any clauses that apply to
5 the **simd** construct, except for the **collapse** clause. For the purposes of each task’s conversion to
6 a SIMD loop, the **collapse** clause is ignored and the effect of any **in_reduction** clause is as
7 if a **reduction** clause with the same reduction operator and list items is present on the construct.

8 **Execution Model Events**

9 This composite construct generates the same events as the **taskloop** construct.

10 **Tool Callbacks**

11 This composite construct dispatches the same callbacks as the **taskloop** construct.

12 **Restrictions**

- 13 • The restrictions for the **taskloop** and **simd** constructs apply.

14 **Cross References**

- 15 • **simd** construct, see Section [2.12.3.1](#) on page [111](#).
- 16 • **taskloop** construct, see Section [2.13.2](#) on page [138](#).
- 17 • Data-sharing attribute clauses, see Section [2.22.4](#) on page [276](#).
- 18 • Events and tool callbacks for **taskloop** construct, see Section [2.13.2](#) on page [138](#).

19 **2.13.4 taskyield Construct**

20 **Summary**

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

Syntax

C / C++

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

```
#pragma omp taskyield new-line
```

C / C++

Fortran

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

```
!$omp taskyield
```

Fortran

Binding

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

Description

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

Cross References

- Task scheduling, see Section [2.13.6](#) on page [147](#).

2.13.5 Initial Task

Execution Model Events

No events are associated with the implicit parallel region in each initial thread.

The *initial-thread-begin* event occurs in an initial thread after the OpenMP runtime invokes the tool initializer but before the initial thread begins to execute the first OpenMP region in the initial task.

The *initial-task-create* event occurs after an *initial-thread-begin* event but before the first OpenMP region in the initial task begins to execute.

The *initial-thread-end* event occurs as the final event in an initial thread at the end of an initial task immediately prior to invocation of the tool finalizer.

1 Tool Callbacks

2 A thread dispatches a registered `ompt_callback_thread_begin` callback for the
3 *initial-thread-begin* event in an initial thread. The callback occurs in the context of the initial
4 thread. The callback has type signature `ompt_callback_thread_begin_t`. The callback
5 receives `ompt_thread_initial` as its *thread_type* argument.

6 A thread dispatches a registered `ompt_callback_task_create` callback for each occurrence
7 of a *initial-task-create* event in the context of the encountering task. This callback has the type
8 signature `ompt_callback_task_create_t`. The callback receives `ompt_task_initial`
9 as its *type* argument. The implicit parallel region does not dispatch a
10 `ompt_callback_parallel_begin` callback; however, the implicit parallel region can be
11 initialized within this `ompt_callback_task_create` callback.

12 A thread dispatches a registered `ompt_callback_thread_end` callback for the
13 *initial-thread-end* event in that thread. The callback occurs in the context of the thread. The
14 callback has type signature `ompt_callback_thread_end_t`. The implicit parallel region
15 does not dispatch a `ompt_callback_parallel_end` callback; however, the implicit parallel
16 region can be finalized within this `ompt_callback_thread_end` callback.

17 Cross References

- 18 • `ompt_task_initial`, see Section [4.2.3.4.17](#) on page 440.
- 19 • `ompt_callback_thread_begin_t`, see Section [4.2.4.2.1](#) on page 446.
- 20 • `ompt_callback_thread_end_t`, see Section [4.2.4.2.2](#) on page 447.
- 21 • `ompt_callback_task_create_t`, see Section [4.2.4.2.6](#) on page 451.

22 2.13.6 Task Scheduling

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

- 26 • the point immediately following the generation of an explicit task;
- 27 • after the point of completion of a `task` region;
- 28 • in a `taskyield` region;
- 29 • in a `taskwait` region;
- 30 • at the end of a `taskgroup` region;
- 31 • in an implicit and explicit `barrier` region;

- 1 • the point immediately following the generation of a **target** region;
- 2 • at the beginning and end of a **target data** region;
- 3 • in a **target update** region;
- 4 • in a **target enter data** region;
- 5 • in a **target exit data** region;
- 6 • in the **omp_target_memcpy** routine;
- 7 • in the **omp_target_memcpy_rect** routine;

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

- 10 • begin execution of a tied task bound to the current team
- 11 • resume any suspended task region, bound to the current team, to which it is tied
- 12 • begin execution of an untied task bound to the current team
- 13 • resume any suspended untied task region bound to the current team.

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

15 *Task Scheduling Constraints* are as follows:

- 16 1. An included task is executed immediately after generation of the task.
- 17 2. Scheduling of new tied tasks is constrained by the set of task regions that are currently tied to the
18 thread, and that are not suspended in a **barrier** region. If this set is empty, any new tied task
19 may be scheduled. Otherwise, a new tied task may be scheduled only if it is a descendent task of
20 every task in the set.
- 21 3. A dependent task shall not be scheduled until its task dependences are fulfilled.
- 22 4. A task shall not be scheduled while any task with which it is mutually exclusive has been
23 scheduled, but has not yet completed.
- 24 5. When an explicit task is generated by a construct containing an **if** clause for which the
25 expression evaluated to *false*, and the previous constraints are already met, the task is executed
26 immediately after generation of the task.

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

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

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

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

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

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

18 **Execution Model Events**

19 The *task-schedule* event occurs in a thread when the thread switches tasks at a task scheduling
20 point; no event occurs when switching to or from a merged task.

21 **Tool Callbacks**

22 A thread dispatches a registered **ompt_callback_task_schedule** callback for each
23 occurrence of a *task-schedule* event in the context of the task that begins or resumes. This callback
24 has the type signature **ompt_callback_task_schedule_t**. The argument *prior_task_status*
25 is used to indicate the cause for suspending the prior task. This cause may be the completion of the
26 prior task region, the encountering of a **taskyield** construct, or the encountering of an active
27 cancellation point.

28 **Cross References**

- 29 • **ompt_callback_task_schedule_t**, see Section [4.2.4.2.9](#) on page [454](#).

1 2.14 Memory Management Directives

2 2.14.1 Memory Spaces

3 OpenMP memory spaces represent storage resources where variables can be stored and retrieved.
4 Table 2.7 shows the list of predefined memory spaces. The selection of a given memory space
5 expresses an intent to use storage with certain traits for the allocations. The actual storage resources
6 that each memory space represents are implementation defined.

TABLE 2.7: Predefined Memory Spaces

Memory space name	Storage selection intent
<code>omp_default_mem_space</code>	Represents the system default storage.
<code>omp_large_cap_mem_space</code>	Represents storage with large capacity.
<code>omp_const_mem_space</code>	Represents storage optimized for variables with constant values. The result of writing to this storage is unspecified.
<code>omp_high_bw_mem_space</code>	Represents storage with high bandwidth.
<code>omp_low_lat_mem_space</code>	Represents storage with low latency.

7 Note – For variables allocated in the `omp_const_mem_space` memory space OpenMP
8 supports initializing constant memory either by means of the `firstprivate` clause or through
9 initialization with compile time constants for static and constant variables. Implementation-defined
10 mechanisms to provide the constant value of these variables may also be supported.

11 Cross References

- 12
- `omp_init_allocator` routine, see Section 3.7.2 on page 406.

1 2.14.2 Memory Allocators

2 OpenMP memory allocators can be used by a program to make allocation requests. When a
3 memory allocator receives a request to allocate storage of a certain size, it will try to return an
4 allocation of logically consecutive *memory* in the resources of its associated memory space of at
5 least the size being requested. This allocation will not overlap with any other existing allocation
6 from an OpenMP memory allocator.

7 The behavior of the allocation process can be affected by the allocator traits the user specifies.
8 Table 2.8 shows the allowed allocators traits, their possible values and the default value of each trait.

TABLE 2.8: Allocator traits

Allocator trait	Allowed values	Default value
<code>sync_hint</code>	<code>contended</code> , <code>uncontended</code> , <code>serialized</code> , <code>private</code>	<code>contended</code>
<code>alignment</code>	A positive integer value which is a power of 2	1 byte
<code>access</code>	<code>all</code> , <code>cgroup</code> , <code>pteam</code> , <code>thread</code>	<code>all</code>
9 <code>pool_size</code>	Positive integer value	Implementation defined
<code>fallback</code>	<code>default_mem_fb</code> , <code>null_fb</code> , <code>abort_fb</code> , <code>allocator_fb</code>	<code>default_mem_fb</code>
<code>fb_data</code>	an allocator handle	(none)
<code>pinned</code>	<code>true</code> , <code>false</code>	<code>false</code>
<code>partition</code>	<code>environment</code> , <code>nearest</code> , <code>blocked</code> , <code>interleaved</code>	<code>environment</code>

10 The `sync_hint` trait describes the expected manner in which multiple threads may use the
11 allocator. The values and their description are:

- 12 • **contended**: high contention is expected on the allocator; that is, many threads are expected to
13 request allocations simultaneously.
- 14 • **uncontended**: low contention is expected on the allocator; that is, few threads are expected to
15 request allocations simultaneously.
- 16 • **serialized**: only one thread at a time will request allocations with the allocator. Requesting
17 two allocations simultaneously when specifying **serialized** results in unspecified behavior.
- 18 • **private**: the same thread will request allocations with the allocator every time. Requesting an
19 allocation from different threads, simultaneously or not, when specifying **private** results in

1 unspecified behavior.

2 Memory allocated will be byte aligned to at least the value specified for the **alignment** trait of
3 the allocator.

4 Memory allocated by allocators with the **access** trait defined to be **all** must be accessible by all
5 threads in the device where the allocation was requested. Memory allocated by allocators with the
6 **access** trait defined to be **cgroup** will be memory accessible by all threads in the same
7 contention group of the thread requesting the allocation. Attempts to access the memory returned by
8 an allocator with the **access** trait defined to be **cgroup** from a thread that is not part of the same
9 contention group as the thread that allocated the memory result in unspecified behavior. Memory
10 allocated by allocators with the **access** trait defined to be **pteam** will be memory accessible by
11 all threads that bind to the same **parallel** region of the thread requesting the allocation.
12 Attempts to access the memory returned by an allocator with the **access** trait defined to be
13 **pteam** from a thread that does not bind to the same **parallel** region as the thread that allocated
14 the memory result in unspecified behavior. Memory allocated by allocator with the **access** trait
15 defined to be **thread** will be memory accessible by the *thread* requesting the allocation. Attempts
16 to access the memory returned by an allocator with the **access** trait defined to be **thread** from a
17 thread other than the one that allocated the memory result in unspecified behavior.

18 The total amount of storage in bytes that an allocator can use is limited by the **pool_size** trait.
19 For allocators with the **access** trait defined to be **all** this limit refers to allocations from all
20 threads accessing the allocator. For allocators with the **access** trait defined to be **cgroup** this
21 limit refers to allocations from threads accessing the allocator from the same contention group. For
22 allocators with the **access** trait defined to be **pteam** this limit refers to allocations from threads
23 accessing the allocator from the same parallel team. For allocators with the **access** trait defined
24 to be **thread** this limit refers to allocations from each thread accessing the allocator. Requests that
25 would result in using more storage than **pool_size** will not be fulfilled by the allocator.

26 The **fallback** trait specifies how the allocator behaves when it cannot fulfil an allocation request.
27 If the **fallback** trait is set to **null_fb** the allocator returns the value zero if it fails to allocate
28 the memory. If the **fallback** trait is set to **abort_fb** the program execution will be terminated
29 if the allocation fails. If the **fallback** trait is set to **allocator_fb** then when an allocation
30 fails the request will be delegated to the allocator specified in the **fb_data** trait. If the
31 **fallback** trait is set to **default_mem_fb** then when an allocation fails another allocation will
32 be tried in the **omp_default_mem_space** memory space assuming all allocator traits to be set
33 to their default values except for **fallback** trait which will be set to **null_fb**.

34 Allocators with the **pinned** trait defined to be **true** ensure that their allocations remain in the
35 same storage resource at the same location for their entire lifetime.

36 The **partition** trait describes the partitioning of allocated memory over the storage resources
37 represented by the memory space associated with the allocator. The partitioning will be done in
38 parts with a minimum size that is implementation defined. The values are:

- 39 • **environment**: the placement of allocated memory is determined by the execution
40 environment.

- 1 ● **nearest**: allocated memory is placed in the storage resource that is nearest to the thread that
- 2 requests the allocation.
- 3 ● **blocked**: allocated memory is partitioned into parts of approximately the same size with at
- 4 most one part per storage resource.
- 5 ● **interleaved**: allocated memory parts are distributed in a round-robin fashion across the
- 6 storage resources.
- 7 Table 2.9 shows the list of predefined memory allocators and their associated memory spaces. The
- 8 predefined memory allocators have default values for their allocator traits unless otherwise
- 9 specified.

TABLE 2.9: Predefined Allocators

Allocator name	Associated memory space	Non-default trait values
<code>omp_default_mem_alloc</code>	<code>omp_default_mem_space</code>	(none)
<code>omp_large_cap_mem_alloc</code>	<code>omp_large_cap_mem_space</code>	(none)
<code>omp_const_mem_alloc</code>	<code>omp_const_mem_space</code>	(none)
<code>omp_high_bw_mem_alloc</code>	<code>omp_high_bw_mem_space</code>	(none)
<code>omp_low_lat_mem_alloc</code>	<code>omp_low_lat_mem_space</code>	(none)
<code>omp_cgroup_mem_alloc</code>	Implementation defined	access:cgroup
<code>omp_pteam_mem_alloc</code>	Implementation defined	access:pteam
<code>omp_thread_mem_alloc</code>	Implementation defined	access:thread

11 ▼ Fortran ▼

12 If any operation of the base language causes a reallocation of an array that is allocated with a

13 memory allocator then that memory allocator will be used to release the current memory and to

 allocate the new memory.

 ▲ Fortran ▲

Cross References

- `omp_init_allocator` routine, see Section 3.7.2 on page 406.
- `omp_destroy_allocator` routine, see Section 3.7.3 on page 407.
- `omp_set_default_allocator` routine, see Section 3.7.4 on page 408.
- `omp_get_default_allocator` routine, see Section 3.7.5 on page 409.
- `OMP_ALLOCATOR` environment variable, see Section 5.21 on page 612.

2.14.3 `allocate` Directive

Summary

The `allocate` directive specifies how a set of variables are allocated. The `allocate` directive is a declarative directive if it is not associated with an allocation statement.

Syntax

C / C++

The syntax of the `allocate` directive is as follows:

```
#pragma omp allocate (list) [clause [ [,] clause] ... ]] new-line
```

where *clause* is one of the following:

```
allocator (allocator)
```

where *allocator* is an expression of `const omp_allocator_t *` type.

C / C++

Fortran

The syntax of the **allocate** directive is as follows:

```
!$omp allocate (list) [clause [ [, ] clause] ... ]
```

or

```
!$omp allocate[ (list) ] clause [ [, ] clause] ... ]  
[!$omp allocate (list) clause [ [, ] clause] ... ]  
[...]  
allocate statement
```

where *clause* is one of the following:

```
allocator (allocator)
```

where *allocator* is an integer expression of **omp_allocator_kind** *kind*.

Fortran

Description

If the directive is not associated with a statement, the storage for each *list item* that appears in the directive will be provided by an allocation through a memory allocator. If no clause is specified then the memory allocator specified by the *def-allocator-var* ICV will be used. If the **allocator** clause is specified, the memory allocator specified in the clause will be used. If a memory allocator is unable to fulfill the allocation request for any list item, the behavior is implementation defined.

The scope of this allocation is that of the list item in the base language. At the end of the scope for a given list item the memory allocator used to allocate that list item deallocates the storage.

Fortran

If the directive is associated with an **allocate** statement, the same list items appearing in the directive list and the **allocate** statement list are allocated with the memory allocator of the directive. If no list items are specified then all variables listed in the **allocate** statement are allocated with the memory allocator of the directive.

Fortran

For allocations that arise from this directive the **null_fb** value of the fallback allocator trait will behave as if the **abort_fb** had been specified.

Restrictions

- A variable that is part of another variable (as an array or structure element) cannot appear in an **allocate** directive.
- The directive must appear in the same scope of the *list item* declaration and before its first use.
- At most one **allocator** clause can appear on the **allocate** directive.
- **allocate** directives appearing in a **target** region must specify an **allocator** clause unless a **requires** directive with the **dynamic_allocators** clause is present in the same compilation unit.

C / C++

- If a list item has a static storage type, only predefined memory allocator variables can be used in the **allocator** clause.

C / C++

Fortran

- List items specified in the **allocate** directive must not have the **ALLOCATABLE** attribute unless the directive is associated with an **allocate** statement.
- List items specified in an **allocate** directive that is associated with an **allocate** statement must be variables that are allocated by the **allocate** statement.
- Multiple directives can only be associated with an **allocate** statement if list items are specified on each **allocate** directive.
- If a list item has the **SAVE** attribute, is a common block name, or is declared in the scope of a module, then only predefined memory allocator variables can be used in the **allocator** clause.
- A type parameter inquiry cannot appear in an **allocate** directive.

Fortran

Cross References

- *def-allocator-var* ICV, see Section 2.4.1 on page 47.
- Memory allocators, see Section 2.14.2 on page 151.
- **omp_allocator_t** and **omp_allocator_kind**, see Section 3.7.1 on page 403.

1 2.14.4 `allocate` Clause

2 Summary


3 The `allocate` clause specifies the memory allocator to be used to obtain storage for private
4 variables of a directive.

5 Syntax

6 The syntax of the `allocate` clause is as follows:

7 **`allocate`** (*[allocator:] list*)

8  `C / C++`
where *allocator* is an expression of the `const omp_allocator_t * type`.

9  `Fortran`
where *allocator* is an integer expression of the `omp_allocator_kind kind`.

10 Description

11 The storage for new list items that arise from list items that appear in the directive will be provided
12 through a memory allocator. If an *allocator* is specified in the clause this will be the memory
13 allocator used for allocations. For all directives except for the `target` directive, if no *allocator* is
14 specified in the clause then the memory allocator specified by the *def-allocator-var* ICV will be
15 used for the list items specified in the `allocate` clause. If a memory allocator is unable to fulfill
16 the allocation request for any list item, the behavior is implementation defined.

17 For allocations that arise from this clause the `null_fb` value of the fallback allocator trait will
18 behave as if the `abort_fb` had been specified.

19 Restrictions

- 20 • For any list item that is specified in the `allocate` clause on a directive, a data-sharing attribute
21 clause that may create a private copy of that list item must be specified on the same directive.
- 22 • For `task`, `taskloop` or `target` directives, allocation requests to memory allocators with the
23 trait `access` set to `thread` result in unspecified behavior.
- 24 • `allocate` clauses appearing in a `target` construct or in a `target` region must specify an
25 *allocator* expression unless a `requires` directive with the `dynamic_allocators` clause is
26 present in the same compilation unit.

1 Cross References

- 2 • *def-allocator-var* ICV, see Section 2.4.1 on page 47.
- 3 • Memory allocators, see Section 2.14.2 on page 151.
- 4 • `omp_allocator_t` and `omp_allocator_kind`, see Section 3.7.1 on page 403.

5 2.15 Device Constructs

6 2.15.1 Device Initialization

7 Execution Model Events

8 The *device-initialize* event occurs in a thread that encounters the first **target**, **target data**, or
9 **target enter data** construct associated with a particular target device after the thread
10 initiates initialization of OpenMP on the device and the device's OpenMP initialization, which may
11 include device-side tool initialization, completes.

12 The *device-load* event for a code block for a target device occurs in some thread before any thread
13 executes code from that code block on that target device.

14 The *device-unload* event for a target device occurs in some thread whenever a code block is
15 unloaded from the device.

16 The *device-finalize* event for a target device that has been initialized occurs in some thread before
17 an OpenMP implementation shuts down.

18 Tool Callbacks

19 A thread dispatches a registered `ompt_callback_device_initialize` callback for each
20 occurrence of a *device-initialize* event in that thread. This callback has type signature
21 `ompt_callback_device_initialize_t`.

22 A thread dispatches a registered `ompt_callback_device_load` callback for each occurrence
23 of a *device-load* event in that thread. This callback has type signature
24 `ompt_callback_device_load_t`.

25 A thread dispatches a registered `ompt_callback_device_unload` callback for each
26 occurrence of a *device-unload* event in that thread. This callback has type signature
27 `ompt_callback_device_unload_t`.

28 A thread dispatches a registered `ompt_callback_device_finalize` callback for each
29 occurrence of a *device-finalize* event in that thread. This callback has type signature
30 `ompt_callback_device_finalize_t`.

1 Restrictions

2 No thread may offload execution of an OpenMP construct to a device until a dispatched
3 `omp_callback_device_initialize` callback completes.

4 No thread may offload execution of an OpenMP construct to a device after a dispatched
5 `omp_callback_device_finalize` callback occurs.

6 Cross References

- 7 • `omp_callback_device_initialize_t`, see Section 4.2.4.2.28 on page 478.
- 8 • `omp_callback_device_load_t`, see Section 4.2.4.2.19 on page 466.
- 9 • `omp_callback_device_unload_t`, see Section 4.2.4.2.20 on page 467.
- 10 • `omp_callback_device_finalize_t`, see Section 4.2.4.2.29 on page 479.

11 2.15.2 target data Construct

12 Summary

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

14 Syntax

▼ C / C++ ▼

15 The syntax of the `target data` construct is as follows:

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

18 where *clause* is one of the following:

```
19     if([ target data :] scalar-expression)  
20     device(integer-expression)  
21     map([[map-type-modifier[, ] [map-type-modifier[, ] ...] map-type: ] list)  
22     use_device_ptr(ptr-list)  
23     use_device_addr(list)
```

▲ C / C++ ▲

Fortran

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

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

5 where *clause* is one of the following:

```
6 if ( [ target data : ] scalar-logical-expression )  
7 device ( scalar-integer-expression )  
8 map ( [ [ map-type-modifier [ , ] [ map-type-modifier [ , ] ... ] map-type : ] list )  
9 use_device_ptr ( ptr-list )  
10 use_device_addr ( list )
```

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

Fortran

12 Binding

13 The binding task set for a **target data** region is the generating task. The **target data** region
14 binds to the region of the generating task.

15 Description

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

21 Pointers that appear in a **use_device_ptr** clause are privatized and the device pointer to the
22 corresponding list items in the device data environment are assigned into the private versions.

23 List items that appear in a **use_device_addr** clause have the address of the corresponding
24 object in the device data environment inside the construct. For objects, any reference to the value of
25 the object will be to the corresponding object on the device, while references to the address will
26 result in a valid device address pointing to that object. Array sections privatize the base of the array
27 section and assign the private copy to the address of the corresponding array section in the device
28 data environment.

29 If one or more of the **use_device_ptr** or **use_device_addr** clauses and one or more **map**
30 clauses are present on the same construct, the address conversions of **use_device_addr** and
31 **use_device_ptr** clauses will occur as if performed after all variables are mapped according to
32 those **map** clauses.

1 **Execution Model Events**

2 The *target-data-begin* event occurs when a thread enters a **target data** region.

3 The *target-data-end* event occurs when a thread exits a **target data** region.

4 **Tool Callbacks**

5 A thread dispatches a registered **ompt_callback_target** callback for each occurrence of a
6 *target-data-begin* and *target-data-end* event in that thread in the context of the task encountering
7 the construct. The callback has type signature **ompt_callback_target_t**. The callback
8 receives **ompt_scope_begin** or **ompt_scope_end** as its *endpoint* argument, as appropriate,
9 and **ompt_target_enter_data** as its *kind* argument.

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** clause expression must
14 evaluate to a non-negative integer value less than the value of **omp_get_num_devices()**.
- 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**, **use_device_addr** or **use_device_ptr** clause must appear on the
18 directive.
- 19 • A list item in a **use_device_ptr** clause must hold the address of an object that has a
20 corresponding list item in the device data environment.
- 21 • A list item in a **use_device_addr** clause must have a corresponding list item in the device
22 data environment.
- 23 • A list item that specifies a given variable may not appear in more than one **use_device_ptr**
24 clause.
- 25 • A reference to a list item in a **use_device_addr** clause must be to the address of the list item.

26 **Cross References**

- 27 • *default-device-var*, see Section 2.4 on page 47.
- 28 • **if** Clause, see Section 2.18 on page 213.
- 29 • **map** clause, see Section 2.22.7.1 on page 307.
- 30 • **omp_get_num_devices** routine, see Section 3.2.35 on page 369.
- 31 • **ompt_callback_target_t**, see Section 4.2.4.2.18 on page 465.

1 2.15.3 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:

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

8 where *clause* is one of the following:

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

C / C++

Fortran

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

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

16 where *clause* is one of the following:

```
17 if ( [ target enter data : ] scalar-logical-expression )  
18 device ( scalar-integer-expression )  
19 map ( [ map-type-modifier [ , ] [ map-type-modifier [ , ] ... ] map-type : list )  
20 depend ( [ depend-modifier : ] [ dependence-type : ] locator-list )  
21 nowait
```

Fortran

1 **Binding**

2 The binding task set for a **target enter data** region is the generating task, which is the *target*
3 *task* generated by the **target enter data** construct. The **target enter data** region binds
4 to the corresponding *target task* region.

5 **Description**

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

8 The **target enter data** construct is a task generating construct. The generated task is a *target*
9 *task*. The generated task region encloses the **target enter data** region.

10 All clauses are evaluated when the **target enter data** construct is encountered. The data
11 environment of the *target task* is created according to the data-sharing attribute clauses on the
12 **target enter data** construct, per-data environment ICVs, and any default data-sharing
13 attribute rules that apply to the **target enter data** construct. A variable that is mapped in the
14 **target enter data** construct has a default data-sharing attribute of shared in the data
15 environment of the *target task*.

16 Assignment operations associated with mapping a variable (see Section [2.22.7.1](#) on page [307](#))
17 occur when the *target task* executes.

18 If the **nowait** clause is present, execution of the *target task* may be deferred. If the **nowait**
19 clause is not present, the *target task* is an included task.

20 If a **depend** clause is present, it is associated with the *target task*.

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

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

23 **Execution Model Events**

24 Events associated with a *target task* are the same as for the **task** construct defined in
25 Section [2.13.1](#) on page [133](#).

26 The *target-enter-data-begin* event occurs when a thread enters a **target enter data** region.

27 The *target-enter-data-end* event occurs when a thread exits a **target enter data** region.

1 Tool Callbacks

2 Callbacks associated with events for *target tasks* are the same as for the **task** construct defined in
3 Section 2.13.1 on page 133.

4 A thread dispatches a registered **ompt_callback_target** callback for each occurrence of a
5 *target-enter-data-begin* and *target-enter-data-end* event in that thread in the context of the target
6 task on the host. The callback has type signature **ompt_callback_target_t**. The callback
7 receives **ompt_scope_begin** or **ompt_scope_end** as its *endpoint* argument, as appropriate,
8 and **ompt_target_enter_data** as its *kind* argument.

9 Restrictions

- 10 • A program must not depend on any ordering of the evaluations of the clauses of the
- 11 **target enter data** directive, or on any side effects of the evaluations of the clauses.
- 12 • At least one **map** clause must appear on the directive.
- 13 • At most one **device** clause can appear on the directive. The **device** clause expression must
- 14 evaluate to a non-negative integer value less than the value of **omp_get_num_devices()**.
- 15 • At most one **if** clause can appear on the directive.
- 16 • A *map-type* must be specified in all **map** clauses and must be either **to** or **alloc**.
- 17 • At most one **nowait** clause can appear on the directive.

18 Cross References

- 19 • *default-device-var*, see Section 2.4.1 on page 47.
- 20 • **task**, see Section 2.13.1 on page 133.
- 21 • **task scheduling constraints**, see Section 2.13.6 on page 147.
- 22 • **target data**, see Section 2.15.2 on page 159.
- 23 • **target exit data**, see Section 2.15.4 on page 165.
- 24 • **if** Clause, see Section 2.18 on page 213.
- 25 • **map** clause, see Section 2.22.7.1 on page 307.
- 26 • **omp_get_num_devices** routine, see Section 3.2.35 on page 369.
- 27 • **ompt_callback_target_t**, see Section 4.2.4.2.18 on page 465.

1 2.15.4 target exit data Construct

2 Summary

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

5 Syntax

C / C++

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

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

8 where *clause* is one of the following:

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

C / C++

Fortran

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

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

16 where *clause* is one of the following:

```
17 if( [ target exit data : ] scalar-logical-expression )  
18 device( scalar-integer-expression )  
19 map( [ map-type-modifier [ , ] [ map-type-modifier [ , ] ... ] map-type : list )  
20 depend( [ depend-modifier : ] [ dependence-type : ] locator-list )  
21 nowait
```

Fortran

1 Binding

2 The binding task set for a **target exit data** region is the generating task, which is the *target*
3 *task* generated by the **target exit data** construct. The **target exit data** region binds to
4 the corresponding *target task* region.

5 Description

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

8 The **target exit data** construct is a task generating construct. The generated task is a *target*
9 *task*. The generated task region encloses the **target exit data** region.

10 All clauses are evaluated when the **target exit data** construct is encountered. The data
11 environment of the *target task* is created according to the data-sharing attribute clauses on the
12 **target exit data** construct, per-data environment ICVs, and any default data-sharing attribute
13 rules that apply to the **target exit data** construct. A variable that is mapped in the
14 **target exit data** construct has a default data-sharing attribute of shared in the data
15 environment of the *target task*.

16 Assignment operations associated with mapping a variable (see Section 2.22.7.1 on page 307)
17 occur when the *target task* executes.

18 If the **nowait** clause is present, execution of the *target task* may be deferred. If the **nowait**
19 clause is not present, the *target task* is an included task.

20 If a **depend** clause is present, it is associated with the *target task*.

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

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

23 Execution Model Events

24 Events associated with a *target task* are the same as for the **task** construct defined in
25 Section 2.13.1 on page 133.

26 The *target-exit-begin* event occurs when a thread enters a **target exit data** region.

27 The *target-exit-end* event occurs when a thread exits a **target exit data** region.

1 **Tool Callbacks**

2 Callbacks associated with events for *target tasks* are the same as for the **task** construct defined in
3 Section 2.13.1 on page 133.

4 A thread dispatches a registered **ompt_callback_target** callback for each occurrence of a
5 *target-exit-begin* and *target-exit-end* event in that thread in the context of the target task on the host.
6 The callback has type signature **ompt_callback_target_t**. The callback receives
7 **ompt_scope_begin** or **ompt_scope_end** as its *endpoint* argument, as appropriate, and
8 **ompt_target_exit_data** as its *kind* argument.

9 **Restrictions**

- 10 • A program must not depend on any ordering of the evaluations of the clauses of the
11 **target exit data** directive, or on any side effects of the evaluations of the clauses.
- 12 • At least one **map** clause must appear on the directive.
- 13 • At most one **device** clause can appear on the directive. The **device** clause expression must
14 evaluate to a non-negative integer value less than the value of **omp_get_num_devices()**.
- 15 • At most one **if** clause can appear on the directive.
- 16 • A *map-type* must be specified in all **map** clauses and must be either **from**, **release**, or
17 **delete**.
- 18 • At most one **nowait** clause can appear on the directive.

19 **Cross References**

- 20 • *default-device-var*, see Section 2.4.1 on page 47.
- 21 • **task**, see Section 2.13.1 on page 133.
- 22 • **task scheduling constraints**, see Section 2.13.6 on page 147.
- 23 • **target data**, see Section 2.15.2 on page 159.
- 24 • **target enter data**, see Section 2.15.3 on page 162.
- 25 • **if** Clause, see Section 2.18 on page 213.
- 26 • **map** clause, see Section 2.22.7.1 on page 307.
- 27 • **omp_get_num_devices** routine, see Section 3.2.35 on page 369.
- 28 • **ompt_callback_target_t**, see Section 4.2.4.2.18 on page 465.

1 2.15.5 target Construct

2 Summary

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

4 Syntax

C / C++

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

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

8 where *clause* is one of the following:

```
9 if ([ target : ] scalar-expression)  
10 device ([ device-modifier : ] integer-expression)  
11 private (list)  
12 firstprivate (list)  
13 in_reduction (reduction-identifier : list)  
14 map ([ [map-type-modifier [ , ] [map-type-modifier [ , ] ... ] map-type : ] list)  
15 is_device_ptr (list)  
16 defaultmap (implicit-behavior[:variable-category])  
17 nowait  
18 depend ([depend-modifier : ] [dependence-type : ] locator-list)  
19 allocate ( [ [allocator : ] list)  
20 uses_allocators (allocator [ (allocator-traits-array) ]  
21 [ , allocator [ (allocator-traits-array) ] ... ])
```

22 where *device-modifier* is one of the following:

```
23 ancestor  
24 device_num
```

25 where *allocator* is an identifier of **const omp_allocator_t *** type.

26 where *allocator-traits-array* is an identifier of **const omp_allocotrait_t *** type.

C / C++

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

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

5 where *clause* is one of the following:

```
6     if ([ target :] scalar-logical-expression)
7     device ([ device-modifier :] scalar-integer-expression)
8     private (list)
9     firstprivate (list)
10    in_reduction (reduction-identifier : list)
11    map ([[map-type-modifier [ , ] [map-type-modifier [ , ] ...] map-type : ] list)
12    is_device_ptr (list)
13    defaultmap (implicit-behavior[:variable-category])
14    nowait
15    depend ([depend-modifier :][dependence-type :] locator-list)
16    allocate ([allocator :]list)
17    uses_allocators (allocator[ (allocator-traits-array) ]
18                      [ , allocator[ (allocator-traits-array) ] ...])
```

19 where *device-modifier* is one of the following:

```
20    ancestor
21    device_num
```

22 where *allocator* is an integer expression of **omp_allocator_kind** *kind*.

23 where *allocator-traits-array* is an array of **type(omp_alloctrail)** type.

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

25 Binding

26 The binding task set for a **target** region is the generating task, which is the *target task* generated
27 by the **target** construct. The **target** region binds to the corresponding *target task* region.

Description

The **target** construct provides a superset of the functionality provided by the **target data** directive, except for the **use_device_ptr** clause.

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 is a task generating construct. The generated task is a *target task*. The generated task region encloses the **target** region.

All clauses are evaluated when the **target** construct is encountered. The data environment of the *target task* is created according to the data-sharing attribute clauses on the **target** construct, per-data environment ICVs, and any default data-sharing attribute rules that apply to the **target** construct. If a variable or part of a variable is mapped by the **target** construct and does not appear as a list item in an **in_reduction** clause on the construct, the variable has a default data-sharing attribute of shared in the data environment of the *target task*.

Assignment operations associated with mapping a variable (see Section 2.22.7.1 on page 307) occur when the *target task* executes.

If a **device** clause in which the **device_num** *device-modifier* appears is present on the construct, the **device** clause expression specifies the device number of the target device. If *device-modifier* does not appear in the clause, the behavior of the clause is as if *device-modifier* is **device_num**.

If a **device** clause in which the **ancestor** *device-modifier* appears is present on the **target** construct and the **device** clause expression evaluates to 1, execution of the **target** region occurs on the parent device of the enclosing **target** region. If the **target** construct is not encountered in a **target** region, the current device is treated as the parent device. The encountering thread waits for completion of the **target** region on the parent device before resuming. For any list item that appears in a **map** clause on the same construct, if the corresponding list item exists in the device data environment of the parent device, it is treated as if it has a reference count of positive infinity.

If the **nowait** clause is present, execution of the *target task* may be deferred. If the **nowait** clause is not present, the *target task* is an included task.

If a **depend** clause is present, it is associated with the *target task*.

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.

The **is_device_ptr** clause is used to indicate that a list item is a device pointer already in the device data environment and that it should be used directly. Support for device pointers created outside of OpenMP, specifically outside of the **omp_target_alloc** routine and the **use_device_ptr** clause, is implementation defined.

If a function (C, C++, Fortran) or subroutine (Fortran) is referenced in a **target** construct then that function or subroutine is treated as if its name had appeared in a **to** clause on a

1 **declare target** directive.

2 Each memory *allocator* specified in the **uses_allocators** clause will be made available in the
3 **target** region. For each non-predefined allocator that is specified, a new allocator handle will be
4 associated with an allocator that is created with the specified *traits* as if by a call to
5 **omp_init_allocator** at the beginning of the **target** region. Each non-predefined allocator
6 will be destroyed as if by a call to **omp_destroy_allocator** at the end of the **target** region.

C / C++

7 If an array section is a list item in a **map** clause and it has a named pointer that is a scalar variable
8 with a predetermined data-sharing attribute of *firstprivate* (see Section 2.22.1.1 on page 263) then
9 on entry to the **target** region:

- 10 • If the list item is not a zero-length array section, the corresponding private variable is initialized
11 relative to the address of the storage location of the corresponding array section in the device
12 data environment that is created by the **map** clause.
- 13 • If the list item is a zero-length array section, the corresponding private variable is initialized
14 relative to the address of the corresponding storage location in the device data environment. If
15 the corresponding storage location is not present in the device data environment, the
16 corresponding private variable is initialized to NULL.

C / C++

17 **Execution Model Events**

18 The *target-begin* event occurs when a thread enters a **target** region.

19 The *target-end* event occurs when a thread exits a **target** region.

20 The *target-submit* event occurs prior to creating an initial task on a target device for a target region.

21 **Tool Callbacks**

22 A thread dispatches a registered **ompt_callback_target** callback for each occurrence of a
23 *target-begin* and *target-end* event in that thread in the context of target task on the host. The
24 callback has type signature **ompt_callback_target_t**. The callback receives
25 **ompt_scope_begin** or **ompt_scope_end** as its *endpoint* argument, as appropriate, and
26 **ompt_target** as its *kind* argument.

27 A thread dispatches a registered **ompt_callback_target_submit** callback for each
28 occurrence of a *target-submit* event in that thread. The callback has type signature
29 **ompt_callback_target_submit_t**.

Restrictions

- If a **target**, **target update**, **target data**, **target enter data**, or **target exit data** construct is encountered during execution of 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.
- If a list item in a **map** clause is a structure element, any other element of that structure that is referenced in the **target** construct must also appear as a list item in a **map** clause.
- 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.
- At most one **defaultmap** clause for each category can appear on the directive.
- At most one **nowait** clause can appear on the directive.
- A *map-type* in a **map** clause must be **to**, **from**, **tofrom** or **alloc**.
- A list item that appears in an **is_device_ptr** clause must be a valid device pointer in the device data environment.
- At most one **device** clause can appear on the directive. The **device** clause expression must evaluate to a non-negative integer value less than the value of **omp_get_num_devices()**.
- If a **device** clause in which the **ancestor device-modifier** appears is present on the construct, then the following restrictions apply:
 - A **requires** directive with the **reverse_offload** clause must be specified.
 - The **device** clause expression must evaluate to 1.
 - Only the **device**, **firstprivate**, **private**, **defaultmap**, and **map** clauses may appear on the construct.
 - No OpenMP constructs or calls to OpenMP API runtime routines are allowed inside the corresponding **target** region.
- Memory allocators that do not appear in a **uses_allocators** clause cannot appear as an allocator in an **allocate** clause or be used in the **target** region unless a **requires** directive with the **dynamic_allocators** clause is present in the same compilation unit.
- Memory allocators that appear in a **uses_allocators** clause cannot appear in other data-sharing attribute clauses or data-mapping attribute clauses in the same construct.
- Predefined allocators appearing in a **uses_allocators** clause cannot have *traits* specified.
- Non-predefined allocators appearing in a **uses_allocators** clause must have *traits* specified.

- 1 • Arrays containing allocators traits that appear in a **uses_allocators** clause must be constant
2 arrays, have constant values and be defined in the same scope as the construct in which the clause
3 appears.
- 4 • Any IEEE floating-point exception status flag, halting mode, or rounding mode set prior to a
5 **target** region is unspecified in the region.
- 6 • Any IEEE floating-point exception status flag, halting mode, or rounding mode set in a **target**
7 region is unspecified upon exiting the region.

C / C++

- 8 • An attached pointer may not be modified in a **target** region.

C / C++

C

- 9 • A list item that appears in an **is_device_ptr** clause must have a type of pointer or array.

C

C++

- 10 • A list item that appears in an **is_device_ptr** clause must have a type of pointer, array,
11 reference to pointer or reference to array.
- 12 • The effect of invoking a virtual member function of an object on a device other than the device
13 on which the object was constructed is implementation defined.
- 14 • A throw executed inside a **target** region must cause execution to resume within the same
15 **target** region, and the same thread that threw the exception must catch it.

C++

Fortran

- 16 • A list item that appears in an **is_device_ptr** clause must be a dummy argument that does
17 not have the **ALLOCATABLE**, **POINTER** or **VALUE** attribute.
- 18 • If a list item in a **map** clause is an array section, and the array section is derived from a variable
19 with a **POINTER** or **ALLOCATABLE** attribute then the behavior is unspecified if the
20 corresponding list item's variable is modified in the region.

Fortran

Cross References

- *default-device-var*, see Section 2.4 on page 47.
- **task** construct, see Section 2.13.1 on page 133.
- **task** scheduling constraints, see Section 2.13.6 on page 147
- Memory allocators, see Section 2.14.2 on page 151.
- **target data** construct, see Section 2.15.2 on page 159.
- **if** Clause, see Section 2.18 on page 213.
- **private** and **firstprivate** clauses, see Section 2.22.4 on page 276.
- Data-mapping Attribute Rules and Clauses, see Section 2.22.7 on page 305.
- **omp_get_num_devices** routine, see Section 3.2.35 on page 369.
- **omp_set_default_allocator** routine, see Section 3.7.4 on page 408.
- **omp_get_default_allocator** routine, see Section 3.7.5 on page 409.
- **omp_alloctrail_t** and **omp_alloctrail** types, see Section 3.7.1 on page 403.
- **ompt_callback_target_t**, see Section 4.2.4.2.18 on page 465.
- **ompt_callback_target_submit_t**, Section 4.2.4.2.23 on page 472.

2.15.6 target update Construct

Summary

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

Syntax

C / C++

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

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

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

```
if( [ target update : ] scalar-expression )  
device ( integer-expression )  
nowait  
depend ( [ depend-modifier : ] [ dependence-type : ] locator-list )
```

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

```
to ( [ mapper ( mapper-identifier ) : ] list )  
from ( [ mapper ( mapper-identifier ) : ] list )
```

C / C++

Fortran

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

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

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

```
if( [ target update : ] scalar-logical-expression )  
device ( scalar-integer-expression )  
nowait  
depend ( [ depend-modifier : ] [ dependence-type : ] locator-list )
```

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

```
to ( [ mapper ( mapper-identifier ) : ] list )  
from ( [ mapper ( mapper-identifier ) : ] list )
```

Fortran

Binding

The binding task set for a **target update** region is the generating task, which is the *target task* generated by the **target update** construct. The **target update** region binds to the corresponding *target task* region.

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. If a **mapper ()** modifier appears in a **to** clause, each list item is replaced with the list items that the given mapper specifies are to be mapped with a **to** or **tofrom** map-type. If a **mapper ()** modifier appears in a **from** clause, each list item is replaced with the list items that the given mapper specifies are to be mapped with a **from** or **tofrom** map-type.

For each list item in a **from** or a **to** clause:

- For each part of the list item that is an attached pointer:
 - On exit from the region that part of the original list item will have the value it had on entry to the region;
 - On exit from the region that part of the corresponding list item will have the value it had on entry to the region;
- For each part of the list item that is not an attached pointer:
 - If the clause is **from**, the value of that part of the corresponding list item is assigned to that part of the original list item;
 - If the clause is **to**, the value of that part of the original list item is assigned to that part of the corresponding list item.
- To avoid race conditions:
 - Concurrent reads or updates of any part of the original list item must be synchronized with the update of the original list item that occurs as a result of the **from** clause;
 - Concurrent reads or updates of any part of the corresponding list item must be synchronized with the update of the corresponding list item that occurs as a result of the **to** clause.

▼ C / C++ ▼
The list items that appear in the **to** or **from** clauses may use shape-operators.
▲ C / C++ ▲

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

The **target update** construct is a task generating construct. The generated task is a *target task*. The generated task region encloses the **target update** region.

All clauses are evaluated when the **target update** construct is encountered. The data environment of the *target task* is created according to the data-sharing attribute clauses on the **target update** construct, per-data environment ICVs, and any default data-sharing attribute

1 rules that apply to the **target update** construct. A variable that is mapped in the
2 **target update** construct has a default data-sharing attribute of shared in the data environment
3 of the *target task*.

4 Assignment operations associated with mapping a variable (see Section 2.22.7.1 on page 307)
5 occur when the *target task* executes.

6 If the **nowait** clause is present, execution of the *target task* may be deferred. If the **nowait**
7 clause is not present, the *target task* is an included task.

8 If a **depend** clause is present, it is associated with the *target task*.

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

12 Execution Model Events

13 Events associated with a *target task* are the same as for the **task** construct defined in
14 Section 2.13.1 on page 133.

15 The *target-update-begin* event occurs when a thread enters a **target update** region.

16 The *target-update-end* event occurs when a thread exits a **target update** region.

17 Tool Callbacks

18 Callbacks associated with events for *target tasks* are the same as for the **task** construct defined in
19 Section 2.13.1 on page 133.

20 A thread dispatches a registered **ompt_callback_target** callback for each occurrence of a
21 *target-update-begin* and *target-update-end* event in that thread in the context of the target task on
22 the host. The callback has type signature **ompt_callback_target_t**. The callback receives
23 **ompt_scope_begin** or **ompt_scope_end** as its *endpoint* argument, as appropriate, and
24 **ompt_target_update** as its *kind* argument.

25 Restrictions

- 26 • A program must not depend on any ordering of the evaluations of the clauses of the
27 **target update** directive, or on any side effects of the evaluations of the clauses.

- 1 ● If a list item is an array section or it uses a shape-operator and the type of its base expression is a
2 pointer type, the base expression must be an lvalue expression.

- 3 ● At least one *motion-clause* must be specified.
4 ● A list item can only appear in a **to** or **from** clause, but not both.
5 ● A list item in a **to** or **from** clause must have a mappable type.
6 ● At most one **device** clause can appear on the directive. The **device** clause expression must
7 evaluate to a non-negative integer value less than the value of `omp_get_num_devices()`.
8 ● At most one **if** clause can appear on the directive.
9 ● At most one **nowait** clause can appear on the directive.

10 Cross References

- 11 ● *default-device-var*, see Section 2.4 on page 47.
12 ● Array shaping, Section 2.5 on page 58
13 ● Array sections, Section 2.6 on page 59
14 ● **task** construct, see Section 2.13.1 on page 133.
15 ● **task** scheduling constraints, see Section 2.13.6 on page 147
16 ● **target data**, see Section 2.15.2 on page 159.
17 ● **if** Clause, see Section 2.18 on page 213.
18 ● `omp_get_num_devices` routine, see Section 3.2.35 on page 369.
19 ● `ompt_callback_task_create_t`, see Section 4.2.4.2.6 on page 451.
20 ● `ompt_callback_target_t`, see Section 4.2.4.2.18 on page 465.

21 2.15.7 declare target Directive

22 Summary

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

Syntax

C / C++

The syntax of the **declare target** directive takes either of the following forms:

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

or

```
#pragma omp declare target (extended-list) new-line
```

or

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

where *clause* is one of the following:

```
to(extended-list)  
link(list)  
implements(function-name)  
device_type(host | nohost | any)
```

C / C++

Fortran

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

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

or

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

where *clause* is one of the following:

```
to(extended-list)  
link(list)  
implements(subroutine-name)  
device_type(host | nohost | any)
```

Fortran

Description

The **declare target** directive ensures that procedures and global variables can be executed or accessed on a device. Variables are mapped for all device executions, or for specific device executions through a **link** clause.

If an *extended-list* is present with no clause then the **to** clause is assumed.

The **implements** clause specifies that an alternate version of a procedure should be used.

The **device_type** clause specifies if a version of the procedure should be made available on host, device or both. If **host** is specified only host version of the procedure is made available. If **nohost** is specified then only device version of the procedure is made available. If **any** is specified then both device and host version of the procedure is made available.

C / C++

If a function is treated as if it appeared as a list item in a **to** clause on a **declare target** directive in the same translation unit in which the definition of the function occurs then a device-specific version of the function is created.

If a variable is treated as if it appeared as a list item in a **to** clause on a **declare target** directive in the same translation unit in which the definition of the variable occurs then the original list item is allocated a corresponding list item in the device data environment of all devices.

All calls in **target** constructs to the function in the **implements** clause are replaced by the function following the **declare target** constructs.

C / C++

Fortran

If an internal procedure is treated as if it appeared as a list item in a **to** clause on a **declare target** directive then a device-specific version of the procedure is created.

If a variable that is host associated is treated as if it appeared as a list item in a **to** clause on a **declare target** directive then the original list item is allocated a corresponding list item in the device data environment of all devices.

All calls in **target** constructs to the procedure in the **implements** clause are replaced by the procedure in which **declare target** construct appeared.

Fortran

If a variable is treated as if it appeared as a list item in a **to** clause on a **declare target** directive then the corresponding list item in the device data environment of each device 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 list item. The list item is never removed from those device data environments as if its reference count is initialized to positive infinity.

1 Including list items in a **link** clause supports compilation of functions called in a **target** region
2 that refer to the list items. They are not mapped by the **declare target** directive. Instead, they
3 are mapped according to the data mapping rules described in Section 2.22.7 on page 305.

▼ C / C++ ▼

4 If a function is referenced in a function that is treated as if it appeared as a list item in a **to** clause
5 on a **declare target** directive then the name of the referenced function is treated as if it had
6 appeared in a **to** clause on a **declare target** directive.

7 If a variable with static storage duration or a function (except *lambda* for C++) is referenced in the
8 initializer expression list of a variable with static storage duration that is treated as if it appeared as a
9 list item in a **to** clause on a **declare target** directive then the name of the referenced variable
10 or function is treated as if it had appeared in a **to** clause on a **declare target** directive.

11 The form of the **declare target** directive that has no clauses and requires a matching
12 **end declare target** directive defines an implicit *extended-list* to an implicit **to** clause. The
13 implicit *extended-list* consists of the variable names of any variable declarations at file or
14 namespace scope that appear between the two directives and of the function names of any function
15 declarations at file, namespace or class scope that appear between the two directives.

16 The *declaration-definition-seq* defined by a **declare target** directive and an
17 **end declare target** directive may contain **declare target** directives. If a
18 **device_type** clause is present on the contained **declare target** directive, then its argument
19 determines which versions are made available. If a list item appears both in an implicit and explicit
20 list, the explicit list determines which versions are made available.

▲ C / C++ ▲

▼ Fortran ▼

21 If a procedure is referenced in a procedure that is treated as if it appeared as a list item in a **to**
22 clause on a **declare target** directive then the name of the procedure is treated as if it had
23 appeared in a **to** clause on a **declare target** directive.

24 If a **declare target** does not have any clauses then an implicit *extended-list* to an implicit **to**
25 clause of one item is formed from the name of the enclosing subroutine subprogram, function
26 subprogram or interface body to which it applies.

27 If a **declare target** directive has an **implements** or **device_type** clause then any
28 enclosed internal procedures cannot contain any **declare target** directives. The enclosing
29 **device_type** clause implicitly applies to internal procedures.

▲ Fortran ▲

Restrictions

- A threadprivate variable cannot appear in a **declare target** directive.
- A variable declared in a **declare target** directive must have a mappable type.
- The same list item must not appear multiple times in clauses on the same directive.
- The same list item must not explicitly appear in both a **to** clause on one **declare target** directive and a **link** clause on another **declare target** directive.
- The **implements** clause can only appear with **device_type** clause.

C++

- The function names of overloaded functions or template functions may only be specified within an implicit *extended-list*.
- If a *lambda declaration and definition* appears between a **declare target** directive and the matching **end declare target** directive, all the variables that are captured by the *lambda* expression must also be variables that are treated as if they appear in a **to** clause.

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 clauses must appear in a specification part of a subroutine subprogram, function subprogram, program or module.
- Any **declare target** directive without clauses must 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, structure element or type parameter inquiry) 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.

- 1 • If a **declare target** directive specifying a common block name appears in one program unit,
- 2 then such a directive must also appear in every other program unit that contains a **COMMON**
- 3 statement specifying the same name. It must appear after the last such **COMMON** statement in the
- 4 program unit.
- 5 • If a list item is declared with the **BIND** attribute, the corresponding C entities must also be
- 6 specified in a **declare target** directive in the C program.
- 7 • A blank common block cannot appear in a **declare target** directive.
- 8 • A variable can only appear in a **declare target** directive in the scope in which it is declared.
- 9 It must not be an element of a common block or appear in an **EQUIVALENCE** statement.
- 10 • A variable that appears in a **declare target** directive must be declared in the Fortran scope
- 11 of a module or have the **SAVE** attribute, either explicitly or implicitly.

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

12 2.15.8 **declare mapper Directive**

13 **Summary**

14 The **declare mapper** directive declares a user-defined mapper for a given type, and may define

15 a *mapper-identifier* that can be used in a **map** clause. The **declare mapper** directive is a

16 declarative directive.

17 **Syntax**

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

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

```
19   | #pragma omp declare mapper ([mapper-identifier:]type var)
```

```
20   | [clause[ [, ] clause] ... ] new-line
```

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

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

21 The syntax of the **declare mapper** directive is as follows:

```
22   | !$omp declare mapper ([mapper-identifier:]type :: var)
```

```
23   | [clause[ [, ] clause] ... ]
```

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

1 where:

- 2 • *mapper-identifier* is a base-language identifier or **default**
- 3 • *type* is a valid type in scope (in Fortran, it must not be an abstract type)
- 4 • *var* is a valid base-language identifier
- 5 • *clause* is **map** (*[[map-type-modifier[,] [map-type-modifier[,] ...]] map-type:] list*)
- 6 • *map-type* is one of the following:
 - 7 – **alloc**
 - 8 – **to**
 - 9 – **from**
 - 10 – **tofrom**
- 11 • and *map-type-modifier* is one of the following:
 - 12 – **always**
 - 13 – **close**

14 Description

15 User-defined mappers can be defined using the **declare mapper** directive. The type and the
16 *mapper-identifier* uniquely identify the mapper for use in a **map** clause later in the program. If the
17 *mapper-identifier* is not specified, then **default** is used. The visibility and accessibility of this
18 declaration are the same as those of a variable declared at the same point in the program.

19 The variable declared by *var* is available for use in all **map** clauses on the directive, and no part of
20 the variable to be mapped is mapped by default.

21 The default mapper for all types *T*, designated by the pre-defined *mapper-identifier* **default**, is as
22 follows unless a user-defined mapper is specified for that type.

```
23 declare mapper(T v) map(tofrom: v)
```

25 Using the **default** *mapper-identifier* overrides the pre-defined default mapper for the given type,
26 making it the default for all variables of *type*. All **map** clauses with this construct in scope that map
27 a list item of *type* will use this mapper unless another is explicitly specified.

28 All **map** clauses on the directive are expanded into corresponding **map** clauses wherever this
29 mapper is invoked, either by matching type or by being explicitly named in a **map** clause. A **map**
30 clause with list item *var* maps *var* as though no mapper were specified.

1 The **declare mapper** directive can also appear at points in the program at which a static data
 2 member could be declared. In this case, the visibility and accessibility of the declaration are the
 3 same as those of a static data member declared at the same point in the program.

4 **Restrictions**

- 5 • No instance of the mapper type can be mapped as part of the mapper, either directly or indirectly
 6 through another type, except the instance passed as the list item. If a set of **declare mapper**
 7 directives results in a cyclic definition then the behavior is unspecified.
- 8 • The *type* must be of struct, union or class type in C and C++ or a non-intrinsic type in Fortran.
- 9 • At least one **map** clause that maps *var* or at least one element of *var* is required.
- 10 • List-items in **map** clauses on this construct may only refer to the declared variable *var* and
 11 entities that could be referenced by a procedure defined at the same location.
- 12 • Each *map-type-modifier* can appear at most once on the **map** clause.
- 13 • A *mapper-identifier* may not be re-declared in the current scope for the same type or for a type
 14 that is compatible according to the base language rules.

15 **2.16 Combined Constructs**

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

19 For combined constructs, tool callbacks shall be invoked as if the constructs were explicitly nested.

20 **2.16.1 Parallel Worksharing-Loop Construct**

21 **Summary**

22 The parallel worksharing-loop construct is a shortcut for specifying a **parallel** construct
 23 containing one worksharing-loop construct with one or more associated loops and no other
 24 statements.

Syntax

C / C++

The syntax of the parallel worksharing-loop construct is as follows:

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

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

C / C++

Fortran

The syntax of the parallel worksharing-loop construct is as follows:

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

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

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

Fortran

Description

The semantics are identical to explicitly specifying a **parallel** directive immediately followed by a worksharing-loop directive.

Restrictions

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

Cross References

- **parallel** construct, see Section 2.9 on page 72.
- worksharing-loop SIMD construct, see Section 2.12.3.2 on page 116.
- Data attribute clauses, see Section 2.22.4 on page 276.

1 2.16.2 Parallel Loop Construct

2 Summary

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

5 Syntax

C / C++

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

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

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

C / C++

Fortran

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

```
12 !$omp parallel loop [clause[ [, ] clause] ... ]  
13     do-loops  
14 [!$omp end parallel loop]
```

15 where *clause* can be any of the clauses accepted by the **parallel** or **loop** directives, with
16 identical meanings and restrictions.

17 If an **end parallel loop** directive is not specified, an **end parallel loop** directive is
18 assumed at the end of the *do-loops*. **nowait** may not be specified on an **end parallel loop**
19 directive.

Fortran

20 Description

21 The semantics are identical to explicitly specifying a **parallel** directive immediately followed
22 by a **loop** directive.

23 Restrictions

- 24 • The restrictions for the **parallel** construct and the **loop** construct apply.

Cross References

- **parallel** construct, see Section 2.9 on page 72.
- **loop** construct, see Section 2.12.5 on page 126.
- Data attribute clauses, see Section 2.22.4 on page 276.

2.16.3 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]
  ...
}
```

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

C / C++

Fortran

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
```

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

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

Fortran

Description

C / C++

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

C / C++

Fortran

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.

Fortran

Restrictions

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

Cross References

- **parallel** construct, see Section 2.9 on page 72.
- **sections** construct, see Section 2.11.1 on page 86.
- Data attribute clauses, see Section 2.22.4 on page 276.

Fortran

2.16.4 **parallel workshare Construct**

Summary

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

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 2.9 on page 72.
- **workshare** construct, see Section 2.11.3 on page 91.
- Data attribute clauses, see Section 2.22.4 on page 276.

Fortran

2.16.5 Parallel Worksharing-Loop SIMD Construct

Summary

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

Syntax

C / C++

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

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

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

C / C++

Fortran

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

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

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

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

Fortran

Description

The semantics of the parallel worksharing-loop SIMD construct are identical to explicitly specifying a **parallel** directive immediately followed by a worksharing-loop SIMD directive.

Restrictions

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

Cross References

- **parallel** construct, see Section 2.9 on page 72.
- worksharing-loop SIMD construct, see Section 2.12.3.2 on page 116.
- Data attribute clauses, see Section 2.22.4 on page 276.

2.16.6 target parallel Construct

Summary

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

Syntax

C / C++

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

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

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

C / C++

Fortran

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

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

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

Fortran

1 **Description**
2 The semantics are identical to explicitly specifying a **target** directive immediately followed by a
3 **parallel** directive.

4 **Restrictions**
5 The restrictions for the **target** and **parallel** constructs apply except for the following explicit
6 modifications:

- 7 • If any **if** clause on the directive includes a *directive-name-modifier* then all **if** clauses on the
8 directive must include a *directive-name-modifier*.
- 9 • At most one **if** clause without a *directive-name-modifier* can appear on the directive.
- 10 • At most one **if** clause with the **parallel** *directive-name-modifier* can appear on the directive.
- 11 • At most one **if** clause with the **target** *directive-name-modifier* can appear on the directive.
- 12 • If an **allocator** clause specifies an *allocator* it can only be a predefined allocator variable.

- 13 **Cross References**
- 14 • **parallel** construct, see Section 2.9 on page 72.
 - 15 • **target** construct, see Section 2.15.5 on page 168.
 - 16 • **if** Clause, see Section 2.18 on page 213.
 - 17 • Data attribute clauses, see Section 2.22.4 on page 276.

18 **2.16.7 Target Parallel Worksharing-Loop Construct**

19 **Summary**
20 The target parallel worksharing-loop construct is a shortcut for specifying a **target** construct
21 containing a parallel worksharing-loop construct and no other statements.

Syntax

C / C++

The syntax of the target parallel worksharing-loop construct is as follows:

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

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

C / C++

Fortran

The syntax of the target parallel worksharing-loop construct is as follows:

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

where *clause* can be any of the clauses accepted by the **target** or **parallel do** directives, except for **copyin**, with identical meanings and restrictions.

If an **end target parallel do** directive is not specified, an **end target parallel do** 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 parallel worksharing-loop directive.

Restrictions

The restrictions for the **target** and parallel worksharing-loop constructs apply except for the following explicit modifications:

- If any **if** clause on the directive includes a *directive-name-modifier* then all **if** clauses on the directive must include a *directive-name-modifier*.
- At most one **if** clause without a *directive-name-modifier* can appear on the directive.
- At most one **if** clause with the **parallel** *directive-name-modifier* can appear on the directive.
- At most one **if** clause with the **target** *directive-name-modifier* can appear on the directive.

Cross References

- **target** construct, see Section 2.15.5 on page 168.
- Parallel Worksharing-Loop construct, see Section 2.16.1 on page 185.
- **if** Clause, see Section 2.18 on page 213.
- Data attribute clauses, see Section 2.22.4 on page 276.

2.16.8 Target Parallel Worksharing-Loop SIMD Construct

Summary

The target parallel worksharing-loop SIMD construct is a shortcut for specifying a **target** construct containing a parallel worksharing-loop SIMD construct and no other statements.

Syntax

C / C++

The syntax of the target parallel worksharing-loop SIMD construct is as follows:

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

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

C / C++

Fortran

The syntax of the target parallel worksharing-loop SIMD construct is as follows:

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

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

If an **end target parallel do simd** directive is not specified, an **end target parallel do simd** 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 parallel worksharing-loop SIMD directive.

4 **Restrictions**

5 The restrictions for the **target** and parallel worksharing-loop SIMD constructs apply except for
6 the following explicit modifications:

- 7 • If any **if** clause on the directive includes a *directive-name-modifier* then all **if** clauses on the
8 directive must include a *directive-name-modifier*.
- 9 • At most one **if** clause without a *directive-name-modifier* can appear on the directive.
- 10 • At most one **if** clause with the **parallel** *directive-name-modifier* can appear on the directive.
- 11 • At most one **if** clause with the **target** *directive-name-modifier* can appear on the directive.

12 **Cross References**

- 13 • **target** construct, see Section [2.15.5](#) on page [168](#).
- 14 • Parallel worksharing-loop SIMD construct, see Section [2.16.5](#) on page [191](#).
- 15 • **if** Clause, see Section [2.18](#) on page [213](#).
- 16 • Data attribute clauses, see Section [2.22.4](#) on page [276](#).

17 **2.16.9 target simd Construct**

18 **Summary**

19 The **target simd** construct is a shortcut for specifying a **target** construct containing a **simd**
20 construct and no other statements.

Syntax

C / C++

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

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

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

C / C++

Fortran

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

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

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

If an **end target simd** directive is not specified, an **end target 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 **simd** directive.

Restrictions

The restrictions for the **target** and **simd** constructs apply.

Cross References

- **simd** construct, see Section [2.12.3.1](#) on page [111](#).
- **target** construct, see Section [2.15.5](#) on page [168](#).
- Data attribute clauses, see Section [2.22.4](#) on page [276](#).

1 2.16.10 target teams Construct

2 Summary

3 The **target teams** construct is a shortcut for specifying a **target** construct containing a
4 **teams** construct and no other statements.

5 Syntax

C / C++

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

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

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

C / C++

Fortran

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

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

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

Fortran

17 Description

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

20 Restrictions

21 The restrictions for the **target** and **teams** constructs apply except for the following explicit
22 modifications:

- 23 • If an **allocator** clause specifies an *allocator* it can only be a predefined allocator variable.

Cross References

- **teams** construct, see Section 2.10 on page 81.
- **target** construct, see Section 2.15.5 on page 168.
- Data attribute clauses, see Section 2.22.4 on page 276.

2.16.11 teams distribute Construct

Summary

The **teams distribute** construct is a shortcut for specifying a **teams** construct containing a **distribute** construct and no other statements.

Syntax

C / C++

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

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

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

C / C++

Fortran

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

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

where *clause* can be any of the clauses accepted by the **teams** or **distribute** 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

1 **Description**

2 The semantics are identical to explicitly specifying a **teams** directive immediately followed by a
3 **distribute** directive.

4 **Restrictions**

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

6 **Cross References**

- 7 • **teams** construct, see Section 2.10 on page 81.
8 • **distribute** construct, see Section 2.12.4.1 on page 117.
9 • Data attribute clauses, see Section 2.22.4 on page 276.

10 **2.16.12 teams loop Construct**

11 **Summary**

12 The **teams loop** construct is a shortcut for specifying a **teams** construct containing a **loop**
13 construct and no other statements.

14 **Syntax**

▼ C / C++ ▼

15 The syntax of the **teams loop** construct is as follows:

```
16   #pragma omp teams loop [clause [ , ] clause ... ] new-line  
17           for-loops
```

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

▲ C / C++ ▲

Fortran

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

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

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

If an **end teams loop** directive is not specified, an **end teams loop** 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 **loop** directive.

Restrictions

The restrictions for the **teams** and **loop** constructs apply.

Cross References

- **teams** construct, see Section [2.10](#) on page [81](#).
- **loop** construct, see Section [2.12.5](#) on page [126](#).
- Data attribute clauses, see Section [2.22.4](#) on page [276](#).

2.16.13 teams distribute simd Construct

Summary

The **teams distribute simd** construct is a shortcut for specifying a **teams** construct containing a **distribute simd** construct and no other statements.

Syntax

C / C++

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

```
#pragma omp teams distribute simd [clause[ [, ] clause] ... ] new-line  
    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

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

```
!$omp teams distribute simd [clause[ [, ] clause] ... ]  
    do-loops  
/$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 simd** directive is not specified, an **end teams distribute simd** 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.

Restrictions

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

Cross References

- **teams** construct, see Section [2.10](#) on page [81](#).
- **distribute simd** construct, see Section [2.12.4.2](#) on page [121](#).
- Data attribute clauses, see Section [2.22.4](#) on page [276](#).

1 2.16.14 target teams distribute Construct

2 Summary

3 The **target teams distribute** construct is a shortcut for specifying a **target** construct
4 containing a **teams distribute** construct and no other statements.

5 Syntax

C / C++

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

```
7 #pragma omp target teams distribute [clause[ [, ] clause] ... ] new-line  
8     for-loops
```

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

▲ C / C++

▼ Fortran

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

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

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

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

▲ Fortran

19 Description

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

22 Restrictions

23 The restrictions for the **target** and **teams distribute** constructs apply except for the
24 following explicit modifications:

- 25 • If an **allocator** clause specifies an *allocator* it can only be a predefined allocator variable.

Cross References

- **target** construct, see Section 2.15.2 on page 159.
- **teams distribute** construct, see Section 2.16.11 on page 199.
- Data attribute clauses, see Section 2.22.4 on page 276.

2.16.15 target teams distribute simd Construct

Summary

The **target teams distribute simd** construct is a shortcut for specifying a **target** construct containing a **teams distribute simd** construct and no other statements.

Syntax

C / C++

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

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

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

C / C++

Fortran

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

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

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

If an **end target teams distribute simd** directive is not specified, an **end target teams distribute 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 simd** directive.

Restrictions

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

Cross References

- **target** construct, see Section [2.15.2](#) on page [159](#).
- **teams distribute simd** construct, see Section [2.16.13](#) on page [201](#).
- Data attribute clauses, see Section [2.22.4](#) on page [276](#).

2.16.16 Teams Distribute Parallel Worksharing-Loop Construct

Summary

The **teams distribute parallel worksharing-loop** construct is a shortcut for specifying a **teams** construct containing a **distribute parallel worksharing-loop** construct and no other statements.

Syntax

C / C++

The syntax of the **teams distribute parallel worksharing-loop** construct is as follows:

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

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

C / C++

Fortran

1 The syntax of the teams distribute parallel worksharing-loop construct is as follows:

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

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

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

Fortran

9 Description

10 The semantics are identical to explicitly specifying a **teams** directive immediately followed by a
11 distribute parallel worksharing-loop directive.

12 Restrictions

13 The restrictions for the **teams** and distribute parallel worksharing-loop constructs apply.

14 Cross References

- 15 • **teams** construct, see Section [2.10](#) on page [81](#).
- 16 • Distribute parallel worksharing-loop construct, see Section [2.12.4.3](#) on page [122](#).
- 17 • Data attribute clauses, see Section [2.22.4](#) on page [276](#).

18 2.16.17 Target Teams Distribute Parallel Worksharing- 19 Loop Construct

20 Summary

21 The target teams distribute parallel worksharing-loop construct is a shortcut for specifying a
22 **target** construct containing a teams distribute parallel worksharing-loop construct and no other
23 statements.

Syntax

C / C++

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

```
#pragma omp target teams distribute parallel for \  
    [clause[ [, ] clause] ... ] new-line  
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

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

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

Description

The semantics are identical to explicitly specifying a **target** directive immediately followed by a teams distribute parallel worksharing-loop directive.

Restrictions

The restrictions for the **target** and teams distribute parallel worksharing-loop constructs apply except for the following explicit modifications:

- If any **if** clause on the directive includes a *directive-name-modifier* then all **if** clauses on the directive must include a *directive-name-modifier*.
- At most one **if** clause without a *directive-name-modifier* can appear on the directive.
- At most one **if** clause with the **parallel** *directive-name-modifier* can appear on the directive.
- At most one **if** clause with the **target** *directive-name-modifier* can appear on the directive.

Cross References

- **target** construct, see Section 2.15.5 on page 168.
- Teams distribute parallel worksharing-loop construct, see Section 2.16.16 on page 205.
- **if** Clause, see Section 2.18 on page 213.
- Data attribute clauses, see Section 2.22.4 on page 276.

2.16.18 Teams Distribute Parallel Worksharing-Loop SIMD Construct

Summary

The teams distribute parallel worksharing-loop SIMD construct is a shortcut for specifying a **teams** construct containing a distribute parallel worksharing-loop SIMD construct and no other statements.

Syntax

C / C++

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

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

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

C / C++

Fortran

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

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

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

If an **end teams distribute parallel do simd** directive is not specified, an **end 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 **teams** directive immediately followed by a distribute parallel worksharing-loop SIMD directive.

Restrictions

The restrictions for the **teams** and distribute parallel worksharing-loop SIMD constructs apply.

Cross References

- **teams** construct, see Section 2.10 on page 81.
- Distribute parallel worksharing-loop SIMD construct, see Section 2.12.4.4 on page 124.
- Data attribute clauses, see Section 2.22.4 on page 276.

2.16.19 Target Teams Distribute Parallel Worksharing-Loop SIMD Construct

Summary

The target teams distribute parallel worksharing-loop SIMD construct is a shortcut for specifying a **target** construct containing a teams distribute parallel worksharing-loop SIMD construct and no other statements.

Syntax

C / C++

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

```
#pragma omp target teams distribute parallel for simd \  
    [clause[ [, ] clause] ... ] new-line  
    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

1 The syntax of the target teams distribute parallel worksharing-loop SIMD construct is as follows:

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

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

7 If an **end target teams distribute parallel do simd** directive is not specified, an
8 **end target teams distribute parallel do simd** directive is assumed at the end of the
9 *do-loops*.

Fortran

10 Description

11 The semantics are identical to explicitly specifying a **target** directive immediately followed by a
12 teams distribute parallel worksharing-loop SIMD directive.

13 Restrictions

14 The restrictions for the **target** and teams distribute parallel worksharing-loop SIMD constructs
15 apply except for the following explicit modifications:

- 16 • If any **if** clause on the directive includes a *directive-name-modifier* then all **if** clauses on the
17 directive must include a *directive-name-modifier*.
- 18 • At most one **if** clause without a *directive-name-modifier* can appear on the directive.
- 19 • At most one **if** clause with the **parallel** *directive-name-modifier* can appear on the directive.
- 20 • At most one **if** clause with the **target** *directive-name-modifier* can appear on the directive.

21 Cross References

- 22 • **target** construct, see Section 2.15.5 on page 168.
- 23 • Teams distribute parallel worksharing-loop SIMD construct, see Section 2.16.18 on page 208.
- 24 • **if** Clause, see Section 2.18 on page 213.
- 25 • Data attribute clauses, see Section 2.22.4 on page 276.

1 2.17 Clauses on Combined and Composite 2 Constructs

3 This section specifies the handling of clauses on combined or composite constructs and handling of
4 implicit clauses from variables with predetermined data sharing if they are not predetermined only
5 on a particular construct. Some clauses are permitted only on a single construct from the constructs
6 that constitute the combined or composite construct, the effect is then as if the clause is applied to
7 that specific construct. Other clauses have the effect as if they are applied to one or more
8 constituent constructs as specified below:

- 9 • The **collapse** clause is applied once for the whole combined or composite construct.
- 10 • For the **private** clause the effect is as if it is applied to the innermost constituent construct only.
- 11 • For the **firstprivate** clause the effect is as if it is applied to one or more constructs as
12 follows:
 - 13 – to the **distribute** construct if it is among the constituent constructs,
 - 14 – to the **teams** construct if it is among the constituent constructs and **distribute** construct
15 is not,
 - 16 – to the worksharing-loop construct if it is among the constituent constructs,
 - 17 – to the **parallel** construct if it is among the constituent constructs and the worksharing-loop
18 construct is not,
 - 19 – to the outermost constituent construct if not already applied to it by the above rules and the
20 outermost constituent construct is neither **teams** nor **parallel** nor **target** construct,
 - 21 – to the **target** construct if it is among the constituent constructs and the same list item does
22 not appear in **lastprivate** or **map** clause.

23 If the **parallel** construct is among the constituent constructs and the effect is not as if the
24 **firstprivate** clause is applied to it by the above rules, then the effect is as if the **shared**
25 clause with the same list item is applied to the **parallel** construct.

26 If the **teams** construct is among the constituent constructs and the effect is not as if the
27 **firstprivate** clause is applied to it by the above rules, then the effect is as if the **shared**
28 clause with the same list item is applied to the **teams** construct.

- 29 • For the **lastprivate** clause the effect is as if it is applied to one or more constructs as follows:
 - 30 – to the worksharing-loop construct if it is among the constituent constructs,
 - 31 – to the **distribute** construct if it is among the constituent constructs,
 - 32 – to the innermost constituent construct that permits it unless it is a worksharing-loop or
33 **distribute** construct.

1 If the **parallel** construct is among the constituent constructs and the list item is not also
2 mentioned in the **firstprivate** clause, then the effect is as if the **shared** clause with the
3 same list item is applied to the **parallel** construct.

4 If the **teams** construct is among the constituent constructs and the list item is not also
5 mentioned in the **firstprivate** clause, then the effect is as if the **shared** clause with the
6 same list item is applied to the **teams** construct.

7 If the **target** construct is among the constituent constructs and the list item doesn't appear in a
8 **map** clause the effect is as if the same list item appears in a **map** clause with a *map-type* of
9 **tofrom**.

- 10 • For the **shared**, **default**, **order**, or **allocate** clauses the effect is as if it is applied to all
11 the constituent constructs that permit those clauses.
- 12 • For the **reduction** clause the effect is as if it is applied to all the constructs that permit the
13 clause, except for the following constructs:
 - 14 – the **parallel** construct, when combined with the worksharing-loop, **loop**, or **sections**
15 construct;
 - 16 – the **teams** construct, when combined with the **loop** construct.

17 For the **parallel** and **teams** constructs above, the behavior instead is as if each list item or,
18 for any list item that is an array item, its corresponding named array or named pointer appears in
19 a **shared** clause for the construct. If *reduction-modifier* is specified, the effect is as if it only
20 modifies the behavior of the **reduction** clause for the innermost construct that constitutes the
21 combined construct and accepts the modifier (see Section 2.22.5.4). If the construct is combined
22 with the **target** construct, the effect is also as if the same list item appears in a **map** clause
23 with a *map-type* of **tofrom**.

- 24 • The **in_reduction** clause is permitted on a single construct among the combined or
25 composite construct and the effect is as if it is applied to that construct, but if that construct is
26 **target**, the effect is also as if the same list item appears in a **map** clause with a *map-type* of
27 **tofrom** and a *map-type-modifier* of **always**.
- 28 • For the **if** clause the effect is described in the Section 2.18 on page 213 section.
- 29 • For the **linear** clause the effect is as if it is applied to the innermost constituent construct.
30 Additionally, if the list item is not the iteration variable of the **simd** or worksharing-loop SIMD
31 construct, the effect on the outer constituent constructs is as if the list item was present in the
32 **firstprivate** and **lastprivate** clauses on the combined or composite construct and the
33 rules specified above would apply. If the list item is the iteration variable of the **simd** or
34 worksharing-loop SIMD construct and it is not declared in the construct, the effect on the outer
35 constituent constructs is as if the list item was present in the **lastprivate** clause on the
36 combined or composite construct and the rules specified above would apply.
- 37 • For the **nowait** clause the effect is as if it is applied to the outermost constituent construct that
38 permits it.

1 If the clauses have expressions on them, such as for various clauses where the argument of the
2 clause is an expression, or *lower-bound*, *length*, or *stride* expressions inside array sections (or
3 *subscript* and *stride* expressions in *subscript-triple* for Fortran), or *linear-step* or *alignment*
4 expressions, the expressions are evaluated immediately before the construct to which the clause has
5 been split or duplicated per the above rules (therefore inside of the outer constituent constructs),
6 except that the expressions inside of the **num_teams** and **thread_limit** clauses are always
7 evaluated before the outermost constituent construct.

8 The restriction that a list item may not appear in more than one data sharing clause with the
9 exception of specifying a variable in both **firstprivate** and **lastprivate** clauses applies
10 after the clauses are split or duplicated per the above rules.

11 2.18 if Clause

12 Summary

13 The semantics of an **if** clause are described in the section on the construct to which it applies. The
14 **if** clause *directive-name-modifier* names the associated construct to which an expression applies,
15 and is particularly useful for composite and combined constructs.

16 Syntax

17  C / C++

The syntax of the **if** clause is as follows:

18 **if** ([*directive-name-modifier* :] *scalar-expression*)

19  C / C++

20  Fortran

The syntax of the **if** clause is as follows:

21 **if** ([*directive-name-modifier* :] *scalar-logical-expression*)

22  Fortran

Description

The effect of the **if** clause depends on the construct to which it is applied. For combined or composite constructs, the **if** clause only applies to the semantics of the construct named in the *directive-name-modifier* if one is specified. If no *directive-name-modifier* is specified for a combined or composite construct then the **if** clause applies to all constructs to which an **if** clause can apply.

2.19 master Construct

Summary

The **master** construct specifies a structured block that is executed by the master thread of the team.

Syntax

C / C++

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

```
#pragma omp master new-line
    structured-block
```

C / C++

Fortran

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

```
!$omp master
    structured-block
!$omp end master
```

Fortran

Binding

The binding thread set for a **master** region is the current team. A **master** region binds to the innermost enclosing **parallel** region. Only the master thread of the team executing the binding **parallel** region participates in the execution of the structured block of the **master** region.

1 **Description**

2 Other threads in the team do not execute the associated structured block. There is no implied
3 barrier either on entry to, or exit from, the **master** construct.

4 **Execution Model Events**

5 The *master-begin* event occurs in the thread encountering the **master** construct on entry to the
6 master region, if it is the master thread of the team.

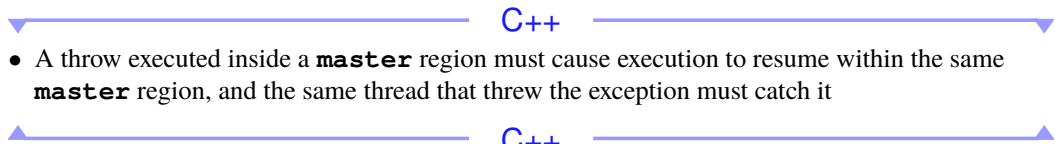
7 The *master-end* event occurs in the thread encountering the **master** construct on exit of the master
8 region, if it is the master thread of the team.

9 **Tool Callbacks**

10 A thread dispatches a registered **ompt_callback_master** callback for each occurrence of a
11 *master-begin* and a *master-end* event in that thread.

12 The callback occurs in the context of the task executed by the master thread. This callback has the
13 type signature **ompt_callback_master_t**. The callback receives **ompt_scope_begin** or
14 **ompt_scope_end** as its *endpoint* argument, as appropriate.

15 **Restrictions**

16  C++
17

- 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

18 **Cross References**

- 19
 - **ompt_scope_begin** and **ompt_scope_end**, see Section [4.2.3.4.10](#) on page [437](#).
 - **ompt_callback_master_t**, see Section [4.2.4.2.5](#) on page [450](#).

1 2.20 Synchronization Constructs and Clauses

2 A synchronization construct orders the completion of code executed by different threads. This
3 ordering is imposed by synchronizing flush operations that are executed as part of the region
4 corresponding to the construct.

5 Synchronization through the use of synchronizing flush operations and atomic operations is
6 described in Section 1.4.4 and Section 1.4.6. Section 2.20.8.1 defines the behavior of synchronizing
7 flush operations that are implied at various other locations in an OpenMP program.

8 The OpenMP API defines the following synchronization constructs, and these are described in the
9 sections that follow:

- 10 • the **critical** construct;
- 11 • the **barrier** construct;
- 12 • the **taskwait** construct;
- 13 • the **taskgroup** construct;
- 14 • the **atomic** construct;
- 15 • the **flush** construct;
- 16 • the **ordered** construct.

17 2.20.1 critical Construct

18 Summary

19 The **critical** construct restricts execution of the associated structured block to a single thread at
20 a time.

21 Syntax

22 C / C++

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

```
23 #pragma omp critical [ (name) [, hint (hint-expression)] ] new-line  
24     structured-block
```

25 where *hint-expression* is an integer constant expression that evaluates to a valid synchronization
26 hint (as described in Section 2.20.12 on page 253).

C / C++

Fortran

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

```
!$omp critical [ (name) [[,] hint (hint-expression) ] ]  
    structured-block  
!$omp end critical [ (name) ]
```

where *hint-expression* is a constant expression that evaluates to a scalar value with kind **omp_sync_hint_kind** and a value that is a valid synchronization hint (as described in Section 2.20.12 on page 253).

Fortran

Binding

The binding thread set for a **critical** region is all threads in the contention group. The region is executed as if only a single thread at a time among all threads in the contention group is entering the region for execution, without regard to the team(s) to which the threads belong.

Description

An optional *name* may be used to identify the **critical** construct. All **critical** constructs without a name are considered to have the same unspecified name.

C / C++

Identifiers used to identify a **critical** construct have external linkage and are in a name space that is separate from the name spaces used by labels, tags, members, and ordinary identifiers.

C / C++

Fortran

The names of **critical** constructs are global entities of the program. If a name conflicts with any other entity, the behavior of the program is unspecified.

Fortran

The threads of a contention group execute the **critical** region as if only one thread of the contention group is executing the **critical** region at a time. The **critical** construct enforces these execution semantics with respect to all **critical** constructs with the same name in all threads in the contention group, not just those threads in the current team.

If present, the **hint** clause gives the implementation additional information about the expected runtime properties of the **critical** region that can optionally be used to optimize the implementation. The presence of a **hint** clause does not affect the isolation guarantees provided by the **critical** construct. If no **hint** clause is specified, the effect is as if **hint(omp_sync_hint_none)** had been specified.

1 Execution Model Events

2 The *critical-acquire* event occurs in the thread encountering the **critical** construct on entry to
3 the critical region before initiating synchronization for the region.

4 The *critical-acquired* event occurs in the thread encountering the **critical** construct after
5 entering the region, but before executing the structured block of the **critical** region.

6 The *critical-release* event occurs in the thread encountering the **critical** construct after
7 completing any synchronization on exit from the **critical** region.

8 Tool Callbacks

9 A thread dispatches a registered **ompt_callback_mutex_acquire** callback for each
10 occurrence of a *critical-acquire* event in that thread. This callback has the type signature
11 **ompt_callback_mutex_acquire_t**.

12 A thread dispatches a registered **ompt_callback_mutex_acquired** callback for each
13 occurrence of a *critical-acquired* event in that thread. This callback has the type signature
14 **ompt_callback_mutex_t**.

15 A thread dispatches a registered **ompt_callback_mutex_released** callback for each
16 occurrence of a *critical-release* event in that thread. This callback has the type signature
17 **ompt_callback_mutex_t**. The callbacks occur in the task encountering the critical construct.
18 The callbacks should receive **ompt_mutex_critical** as their *kind* argument if practical, but a
19 less specific kind is acceptable.

20 Restrictions

- 21 • If the **hint** clause is specified, the **critical** construct must have a *name*.
- 22 • If the **hint** clause is specified, each of the **critical** constructs with the same *name* must
23 have a **hint** clause for which the *hint-expression* evaluates to the same value.

24 **C++**

- 25 • 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

26 The following restrictions apply to the critical construct:

- 27 • If a *name* is specified on a **critical** directive, the same *name* must also be specified on the
28 **end critical** directive.
- 29 • If no *name* appears on the **critical** directive, no *name* can appear on the **end critical**
30 directive.

Fortran

Cross References

- Synchronization Hints, see Section [2.20.12](#) on page [253](#).
- `omp_mutex_critical`, see Section [4.2.3.4.15](#) on page [439](#).
- `omp_callback_mutex_acquire_t`, see Section [4.2.4.2.12](#) on page [458](#).
- `omp_callback_mutex_t`, see Section [4.2.4.2.13](#) on page [459](#).

2.20.2 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

Binding

The binding thread set for a **barrier** region is the current team. A **barrier** region binds to the innermost enclosing **parallel** region.

Description

All threads of the team executing the binding **parallel** region must execute the **barrier** region and complete execution of all explicit tasks bound to this **parallel** region before any are allowed to continue execution beyond the barrier.

The **barrier** region includes an implicit task scheduling point in the current task region.

1 Execution Model Events

2 The *explicit-barrier-begin* event occurs in each thread encountering the **barrier** construct on
3 entry to the **barrier** region.

4 The *explicit-barrier-wait-begin* event occurs when a task begins an interval of active or passive
5 waiting in a **barrier** region.

6 The *explicit-barrier-wait-end* event occurs when a task ends an interval of active or passive waiting
7 and resumes execution in a **barrier** region.

8 The *explicit-barrier-end* event occurs in each thread encountering the **barrier** construct after the
9 barrier synchronization on exit from the **barrier** region.

10 A *cancellation* event occurs if cancellation is activated at an implicit cancellation point in an barrier
11 region.

12 Tool Callbacks

13 A thread dispatches a registered **ompt_callback_sync_region** callback for each occurrence
14 of a *explicit-barrier-begin* and *explicit-barrier-end* event in that thread. The callback occurs in the
15 task encountering the barrier construct. This callback has the type signature

16 **ompt_callback_sync_region_t**. The callback receives
17 **ompt_sync_region_barrier_explicit** — or **ompt_sync_region_barrier**, if the
18 implementation cannot make a distinction — as its *kind* argument and **ompt_scope_begin** or
19 **ompt_scope_end** as its *endpoint* argument, as appropriate.

20 A thread dispatches a registered **ompt_callback_sync_region_wait** callback for each
21 occurrence of a *barrier-wait-begin* and *barrier-wait-end* event. This callback has type signature

22 **ompt_callback_sync_region_t**. This callback executes in the context of the task that
23 encountered the **barrier** construct. The callback receives
24 **ompt_sync_region_barrier_explicit** — or **ompt_sync_region_barrier**, if the
25 implementation cannot make a distinction — as its *kind* argument and **ompt_scope_begin** or
26 **ompt_scope_end** as its *endpoint* argument, as appropriate.

27 A thread dispatches a registered **ompt_callback_cancel** callback for each occurrence of a
28 *cancellation* event in that thread. The callback occurs in the context of the encountering task. The
29 callback has type signature **ompt_callback_cancel_t**. The callback receives
30 **ompt_cancel_detected** as its *flags* argument.

31 Restrictions

32 The following restrictions apply to the **barrier** construct:

- 33 • Each **barrier** region must be encountered by all threads in a team or by none at all, unless
34 cancellation has been requested for the innermost enclosing parallel region.

- 1 • The sequence of worksharing regions and **barrier** regions encountered must be the same for
2 every thread in a team.

3 **Cross References**

- 4 • **ompt_scope_begin** and **ompt_scope_end**, see Section [4.2.3.4.10](#) on page [437](#).
5 • **ompt_sync_region_barrier**, see Section [4.2.3.4.12](#) on page [437](#).
6 • **ompt_callback_sync_region_t**, see Section [4.2.4.2.11](#) on page [457](#).
7 • **ompt_callback_cancel_t**, see Section [4.2.4.2.27](#) on page [477](#).

8 **2.20.3 Implicit Barriers**

9 Implicit tasks in a parallel region synchronize with one another using implicit barriers at the end of
10 worksharing constructs and at the end of the **parallel** region according to the premises defined
11 with the individual constructs. This section describes the OMPT events and tool callbacks
12 associated with implicit barriers.

13 Implicit barriers are task scheduling points. For a description of task scheduling points, associated
14 events, and tool callbacks, see Section [2.13.6](#) on page [147](#).

15 **Execution Model Events**

16 A *cancellation* event occurs if cancellation is activated at an implicit cancellation point in an
17 implicit barrier region.

18 The *implicit-barrier-begin* event occurs in each implicit task at the beginning of an implicit barrier.

19 The *implicit-barrier-wait-begin* event occurs when a task begins an interval of active or passive
20 waiting while executing in an implicit barrier region.

21 The *implicit-barrier-wait-end* event occurs when a task ends an interval of active or waiting and
22 resumes execution of an implicit barrier region.

23 The *implicit-barrier-end* event occurs in each implicit task at the end of an implicit barrier.

1 Tool Callbacks

2 A thread dispatches a registered `ompt_callback_sync_region` callback for each occurrence
3 of a *implicit-barrier-begin* and *implicit-barrier-end* event in that thread. The callback occurs in the
4 implicit task executing in a parallel region. This callback has the type signature

5 `ompt_callback_sync_region_t`. The callback receives
6 `ompt_sync_region_barrier_implicit` — or `ompt_sync_region_barrier`, if the
7 implementation cannot make a distinction — as its *kind* argument and `ompt_scope_begin` or
8 `ompt_scope_end` as its *endpoint* argument, as appropriate.

9 A thread dispatches a registered `ompt_callback_cancel` callback for each occurrence of a
10 *cancellation* event in that thread. The callback occurs in the context of the encountering task. The
11 callback has type signature `ompt_callback_cancel_t`. The callback receives
12 `ompt_cancel_detected` as its *flags* argument.

13 A thread dispatches a registered `ompt_callback_sync_region_wait` callback for each
14 occurrence of a *implicit-barrier-wait-begin* and *implicit-barrier-wait-end* event. This callback has
15 type signature `ompt_callback_sync_region_t`. The callback occurs in each implicit task
16 participating in an implicit barrier. The callback receives
17 `ompt_sync_region_barrier_implicit` — or `ompt_sync_region_barrier`, if the
18 implementation cannot make a distinction — as its *kind* argument and `ompt_scope_begin` or
19 `ompt_scope_end` as its *endpoint* argument, as appropriate.

20 Restrictions

21 If a thread is in the state `omp_state_wait_barrier_implicit_parallel`, a call to
22 `ompt_get_parallel_info` may return a pointer to a copy of the current parallel region's
23 *parallel_data* rather than a pointer to the data word for the region itself. This convention enables
24 the master thread for a parallel region to free storage for the region immediately after the region
25 ends, yet avoid having some other thread in the region's team potentially reference the region's
26 *parallel_data* object after it has been freed.

27 Cross References

- 28 • `ompt_scope_begin` and `ompt_scope_end`, see Section [4.2.3.4.10](#) on page [437](#).
- 29 • `ompt_sync_region_barrier`, see Section [4.2.3.4.12](#) on page [437](#)
- 30 • `ompt_cancel_detected`, see Section [4.2.3.4.23](#) on page [443](#).
- 31 • `ompt_callback_sync_region_t`, see Section [4.2.4.2.11](#) on page [457](#).
- 32 • `ompt_callback_cancel_t`, see Section [4.2.4.2.27](#) on page [477](#).

1 2.20.4 Implementation-Specific Barriers

2 An OpenMP implementation can execute implementation-specific barriers that are not implied by
3 the OpenMP specification; therefore, no *execution model events* are bound to these barriers.

4 The implementation can handle these barriers like implicit barriers and dispatch all events as for
5 implicit barriers.

6 In this case, the callbacks receive `ompt_sync_region_barrier_implementation` — or
7 `ompt_sync_region_barrier`, if the implementation cannot make a distinction — as its *kind*
8 argument.

9 2.20.5 taskwait Construct

10 Summary

11 The `taskwait` construct specifies a wait on the completion of child tasks of the current task. The
12 `taskwait` construct is a stand-alone directive.

13 Syntax

14  C / C++ 

The syntax of the `taskwait` construct is as follows:

```
15 | #pragma omp taskwait [clause[ [, ] clause] ... ] new-line
```

16 where *clause* is one of the following:

```
17 | depend ([depend-modifier:] [dependence-type :] locator-list)
```

18  C / C++ 

19  Fortran 

The syntax of the `taskwait` construct is as follows:

```
19 | !$omp taskwait [clause[ [, ] clause] ... ]
```

20 where *clause* is one of the following:

```
21 | depend ([depend-modifier:] [dependence-type :] locator-list)
```

22  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 If no **depend** clause is present on the **taskwait** construct, the current task region is suspended
6 at an implicit task scheduling point associated with the construct. The current task region remains
7 suspended until all child tasks that it generated before the **taskwait** region complete execution.

8 Otherwise, if one or more **depend** clauses are present on the **taskwait** construct, the behavior
9 is as if these clauses were applied to a **task** construct with an empty associated structured block
10 that generates a *mergeable* and *included task*. Thus, the current task region is suspended until the
11 *predecessor tasks* of this task complete execution.

12 Execution Model Events

13 The *taskwait-begin* event occurs in each thread encountering the **taskwait** construct on entry to
14 the **taskwait** region.

15 The *taskwait-wait-begin* event occurs when a task begins an interval of active or passive waiting in
16 a **taskwait** region.

17 The *taskwait-wait-end* event occurs when a task ends an interval of active or passive waiting and
18 resumes execution in a **taskwait** region.

19 The *taskwait-end* event occurs in each thread encountering the **taskwait** construct after the
20 taskwait synchronization on exit from the **taskwait** region.

21 Tool Callbacks

22 A thread dispatches a registered **ompt_callback_sync_region** callback for each occurrence
23 of a *taskwait-begin* and *taskwait-end* event in that thread. The callback occurs in the task
24 encountering the **taskwait** construct. This callback has the type signature
25 **ompt_callback_sync_region_t**. The callback receives
26 **ompt_sync_region_taskwait** as its *kind* argument and **ompt_scope_begin** or
27 **ompt_scope_end** as its *endpoint* argument, as appropriate.

28 A thread dispatches a registered **ompt_callback_sync_region_wait** callback for each
29 occurrence of a *taskwait-wait-begin* and *taskwait-wait-end* event. This callback has type signature
30 **ompt_callback_sync_region_t**. This callback executes in the context of the task that
31 encountered the **taskwait** construct. The callback receives **ompt_sync_region_taskwait**
32 as its *kind* argument and **ompt_scope_begin** or **ompt_scope_end** as its *endpoint* argument,
33 as appropriate.

Restrictions

The `mutexinoutset` *dependence-type* may not appear in a `depend` clause on a `taskwait` construct.

Cross References

- `task` construct, see Section 2.13.1 on page 133.
- Task scheduling, see Section 2.13.6 on page 147.
- `depend` clause, see Section 2.20.11 on page 248.
- `ompt_scope_begin` and `ompt_scope_end`, see Section 4.2.3.4.10 on page 437.
- `ompt_sync_region_taskwait`, see Section 4.2.3.4.12 on page 437.
- `ompt_callback_sync_region_t`, see Section 4.2.4.2.11 on page 457.

2.20.6 taskgroup Construct

Summary

The `taskgroup` construct specifies a wait on completion of child tasks of the current task and their descendent tasks.

Syntax

C / C++

The syntax of the `taskgroup` construct is as follows:

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

where *clause* is one of the following:

```
task_reduction (reduction-identifier : list)
allocate ([allocator: ]list)
```

C / C++

Fortran

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

```
2 !$omp taskgroup [clause [ [, ] clause] ...]  
3     structured-block  
4 !$omp end taskgroup
```

5 where *clause* is one of the following:

```
6     task_reduction (reduction-identifier : list)  
7     allocate ([allocator: ]list)
```

Fortran

8 Binding

9 A **taskgroup** region binds to the current task region. A **taskgroup** region binds to the
10 innermost enclosing **parallel** region.

11 Description

12 When a thread encounters a **taskgroup** construct, it starts executing the region. All child tasks
13 generated in the **taskgroup** region and all of their descendants that bind to the same **parallel**
14 region as the **taskgroup** region are part of the *taskgroup set* associated with the **taskgroup**
15 region.

16 There is an implicit task scheduling point at the end of the **taskgroup** region. The current task is
17 suspended at the task scheduling point until all tasks in the *taskgroup set* complete execution.

18 Execution Model Events

19 The *taskgroup-begin* event occurs in each thread encountering the **taskgroup** construct on entry
20 to the **taskgroup** region.

21 The *taskgroup-wait-begin* event occurs when a task begins an interval of active or passive waiting
22 in a **taskgroup** region.

23 The *taskgroup-wait-end* event occurs when a task ends an interval of active or passive waiting and
24 resumes execution in a **taskgroup** region.

25 The *taskgroup-end* event occurs in each thread encountering the **taskgroup** construct after the
26 taskgroup synchronization on exit from the **taskgroup** region.

1 Tool Callbacks

2 A thread dispatches a registered `ompt_callback_sync_region` callback for each occurrence
3 of a *taskgroup-begin* and *taskgroup-end* event in that thread. The callback occurs in the task
4 encountering the taskgroup construct. This callback has the type signature
5 `ompt_callback_sync_region_t`. The callback receives
6 `ompt_sync_region_taskgroup` as its *kind* argument and `ompt_scope_begin` or
7 `ompt_scope_end` as its *endpoint* argument, as appropriate.

8 A thread dispatches a registered `ompt_callback_sync_region_wait` callback for each
9 occurrence of a *taskgroup-wait-begin* and *taskgroup-wait-end* event. This callback has type
10 signature `ompt_callback_sync_region_t`. This callback executes in the context of the task
11 that encountered the `taskgroup` construct. The callback receives
12 `ompt_sync_region_taskgroup` as its *kind* argument and `ompt_scope_begin` or
13 `ompt_scope_end` as its *endpoint* argument, as appropriate.

14 Cross References

- 15 • Task scheduling, see Section [2.13.6](#) on page [147](#).
- 16 • `task_reduction` Clause, see Section [2.22.5.5](#) on page [299](#).
- 17 • `ompt_scope_begin` and `ompt_scope_end`, see Section [4.2.3.4.10](#) on page [437](#).
- 18 • `ompt_sync_region_taskgroup`, see Section [4.2.3.4.12](#) on page [437](#).
- 19 • `ompt_callback_sync_region_t`, see Section [4.2.4.2.11](#) on page [457](#).

20 2.20.7 atomic Construct

21 Summary

22 The `atomic` construct ensures that a specific storage location is accessed atomically, rather than
23 exposing it to the possibility of multiple, simultaneous reading and writing threads that may result
24 in indeterminate values.

25 Syntax

26 In the following syntax, *atomic-clause* is a clause that indicates the semantics for which atomicity is
27 enforced, *memory-order-clause* is a clause that indicates the memory ordering behavior of the
28 construct and *clause* is a clause other than *atomic-clause*. Specifically, *atomic-clause* is one of the
29 following:

```
1  read
2  write
3  update
4  capture
```

5 *memory-order-clause* is one of the following:

```
6  seq_cst
7  acq_rel
8  release
9  acquire
10 relaxed
```

11 and *clause* is either *memory-order-clause* or one of the following:

```
12  hint (hint-expression)
```

C / C++

13 The syntax of the **atomic** construct takes one of the following forms:

```
14 #pragma omp atomic [clause[[,] clause] ... ] [,] atomic-clause
15                    [[,] clause [[[,] clause] ... ]] new-line
16                    expression-stmt
```

17 or

```
18 #pragma omp atomic [clause[[,] clause] ... ] new-line
19                    expression-stmt
```

20 or

```
21 #pragma omp atomic [clause[[,] clause] ... ] [,] capture
22                    [[,] clause [[[,] clause] ... ]] new-line
23                    structured-block
```

24 where *expression-stmt* is an expression statement with one of the following forms:

- 25 • If *atomic-clause* is **read**:

```
26  v = x;
```

- 27 • If *atomic-clause* is **write**:

```
28  x = expr;
```

- 29 • If *atomic-clause* is **update** or not present:

```

1  x++;
2  x--;
3  ++x;
4  --x;
5  x binop= expr;
6  x = x binop expr;
7  x = expr binop x;

```

- If *atomic-clause* is **capture**:

```

9  v = x++;
10 v = x--;
11 v = ++x;
12 v = --x;
13 v = x binop= expr;
14 v = x = x binop expr;
15 v = x = expr binop x;

```

and where *structured-block* is a structured block with one of the following forms:

```

17 v = x; x binop= expr;
18 x binop= expr; v = x;
19 v = x; x = x binop expr;
20 v = x; x = expr binop x;
21 x = x binop expr; v = x;
22 x = expr binop x; v = x;
23 v = x; x = expr;
24 v = x; x++;
25 v = x; ++x;
26 ++x; v = x;
27 x++; v = x;
28 v = x; x--;
29 v = x; --x;
30 --x; v = x;
31 x--; v = x;

```

In the preceding expressions:

- x and v (as applicable) are both *l-value* expressions with scalar type.
- During the execution of an atomic region, multiple syntactic occurrences of x must designate the same storage location.
- Neither of v and $expr$ (as applicable) may access the storage location designated by x .
- Neither of x and $expr$ (as applicable) may access the storage location designated by v .
- $expr$ is an expression with scalar type.

- *binop* is one of **+**, *****, **-**, **/**, **&**, **^**, **|**, **<<**, or **>>**.
- *binop*, *binop=*, **++**, and **-** are not overloaded operators.
- The expression *x binop expr* must be numerically equivalent to *x binop (expr)*. This requirement is satisfied if the operators in *expr* have precedence greater than *binop*, or by using parentheses around *expr* or subexpressions of *expr*.
- The expression *expr binop x* must be numerically equivalent to *(expr) binop x*. This requirement is satisfied if the operators in *expr* have precedence equal to or greater than *binop*, or by using parentheses around *expr* or subexpressions of *expr*.
- For forms that allow multiple occurrences of *x*, the number of times that *x* is evaluated is unspecified.

▲ C / C++ ▲

▼ Fortran ▼

The syntax of the **atomic** construct takes any of the following forms:

```
!$omp atomic [clause[[, ] clause] ... ] [, ] read [[, ] clause [[[, ] clause] ... ]]
    capture-statement
!$omp end atomic
```

or

```
!$omp atomic [clause[[, ] clause] ... ] [, ] write [[, ] clause [[[, ] clause] ... ]]
    write-statement
!$omp end atomic
```

or

```
!$omp atomic [clause[[, ] clause] ... ] [, ] update [[, ] clause [[[, ] clause] ... ]]
    update-statement
!$omp end atomic
```

or

```
!$omp atomic [clause[[, ] clause] ... ]
    update-statement
!$omp end atomic
```

or

```
!$omp atomic [clause[[, ] clause] ... ] [, ] capture [[, ] clause [[[, ] clause] ... ]]
    update-statement
    capture-statement
!$omp end atomic
```

or

```

1  !$omp atomic [clause[[[,] clause] ... ] [, ]] capture [[[,] clause [[[,] clause] ... ]]
2      capture-statement
3      update-statement
4  !$omp end atomic

```

5 or

```

6  !$omp atomic [clause[[[,] clause] ... ] [, ]] capture [[[,] clause [[[,] clause] ... ]]
7      capture-statement
8      write-statement
9  !$omp end atomic

```

10 where *write-statement* has the following form (if *atomic-clause* is **capture** or **write**):

```

11      x = expr

```

12 where *capture-statement* has the following form (if *atomic-clause* is **capture** or **read**):

```

13      v = x

```

14 and where *update-statement* has one of the following forms (if *atomic-clause* is **update**,
15 **capture**, or not present):

```

16      x = x operator expr
17
18      x = expr operator x
19
20      x = intrinsic_procedure_name (x, expr_list)
21
22      x = intrinsic_procedure_name (expr_list, x)

```

23 In the preceding statements:

- 24 • *x* and *v* (as applicable) are both scalar variables of intrinsic type.
- 25 • *x* must not have the **ALLOCATABLE** attribute.
- 26 • During the execution of an atomic region, multiple syntactic occurrences of *x* must designate the
27 same storage location.
- 28 • None of *v*, *expr*, and *expr_list* (as applicable) may access the same storage location as *x*.
- 29 • None of *x*, *expr*, and *expr_list* (as applicable) may access the same storage location as *v*.
- 30 • *expr* is a scalar expression.
- 31 • *expr_list* is a comma-separated, non-empty list of scalar expressions. If
32 *intrinsic_procedure_name* refers to **IAND**, **IOR**, or **IEOR**, exactly one expression must appear in
33 *expr_list*.

- 1 • *intrinsic_procedure_name* is one of **MAX**, **MIN**, **IAND**, **IOR**, or **IEOR**.
- 2 • *operator* is one of **+**, *****, **-**, **/**, **.AND.**, **.OR.**, **.EQV.**, or **.NEQV.**
- 3 • The expression *x operator expr* must be numerically equivalent to *x operator (expr)*. This
4 requirement is satisfied if the operators in *expr* have precedence greater than *operator*, or by
5 using parentheses around *expr* or subexpressions of *expr*.
- 6 • The expression *expr operator x* must be numerically equivalent to *(expr) operator x*. This
7 requirement is satisfied if the operators in *expr* have precedence equal to or greater than
8 *operator*, or by using parentheses around *expr* or subexpressions of *expr*.
- 9 • *intrinsic_procedure_name* must refer to the intrinsic procedure name and not to other program
10 entities.
- 11 • *operator* must refer to the intrinsic operator and not to a user-defined operator.
- 12 • All assignments must be intrinsic assignments.
- 13 • For forms that allow multiple occurrences of *x*, the number of times that *x* is evaluated is
14 unspecified.

Fortran

Binding

If the size of *x* is 8, 16, 32, or 64 bits and *x* is aligned to a multiple of its size, the binding thread set for the **atomic** region is all threads on the device. Otherwise, the binding thread set for the **atomic** region is all threads in the contention group. **atomic** regions enforce exclusive access with respect to other **atomic** regions that access the same storage location *x* among all threads in the binding thread set without regard to the teams to which the threads belong.

Description

If *atomic-clause* is not present on the construct, the behavior is as if the **update** clause is specified.

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. 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 captured update — an atomic update of the location designated by *x* using the designated operator or intrinsic while also capturing

1 the original or final value of the location designated by x with respect to the atomic update. The
2 original or final value of the location designated by x is written in the location designated by v
3 depending on the form of the **atomic** construct structured block or statements following the usual
4 language semantics. Only the read and write of the location designated by x are performed
5 mutually atomically. Neither the evaluation of $expr$ or $expr_list$, nor the write to the location
6 designated by v , need be atomic with respect to the read or write of the location designated by x . No
7 task scheduling points are allowed between the read and the write of the location designated by x .

8 The **atomic** construct may be used to enforce memory consistency between threads, based on the
9 guarantees provided by Section 1.4.6 on page 27. A strong flush on the location designated by x is
10 performed on entry to and exit from the atomic operation, ensuring that the set of all atomic
11 operations in the program applied to the same location has a total completion order. If the **write**,
12 **update**, or **capture** clause is specified and the **release**, **acq_rel**, or **seq_cst** clause is
13 specified, the flush on entry to the atomic operation is a release flush. If the **read** or **capture**
14 clause is specified and the **acquire**, **acq_rel**, or **seq_cst** clause is specified, the flush on exit
15 from the atomic operation is an acquire flush. Therefore, if *memory-order-clause* is specified and is
16 not **relaxed**, release and/or acquire flush operations are implied and permit synchronization
17 between the threads without the use of explicit **flush** directives.

18 For all forms of the **atomic** construct, any combination of two or more of these **atomic**
19 constructs enforces mutually exclusive access to the locations designated by x among threads in the
20 binding thread set. To avoid race conditions, all accesses of the locations designated by x that could
21 potentially occur in parallel must be protected with an **atomic** construct.

22 **atomic** regions do not guarantee exclusive access with respect to any accesses outside of
23 **atomic** regions to the same storage location x even if those accesses occur during a **critical**
24 or **ordered** region, while an OpenMP lock is owned by the executing task, or during the
25 execution of a **reduction** clause.

26 However, other OpenMP synchronization can ensure the desired exclusive access. For example, a
27 barrier following a series of atomic updates to x guarantees that subsequent accesses do not form a
28 race with the atomic accesses.

29 A compliant implementation may enforce exclusive access between **atomic** regions that update
30 different storage locations. The circumstances under which this occurs are implementation defined.

31 If the storage location designated by x is not size-aligned (that is, if the byte alignment of x is not a
32 multiple of the size of x), then the behavior of the **atomic** region is implementation defined.

33 If present, the **hint** clause gives the implementation additional information about the expected
34 properties of the atomic operation that can optionally be used to optimize the implementation. The
35 presence of a **hint** clause does not affect the semantics of the **atomic** construct, and it is legal to
36 ignore all hints. If no **hint** clause is specified, the effect is as if
37 **hint(omp_sync_hint_none)** had been specified.

1 Execution Model Events

2 The *atomic-acquire* event occurs in the thread encountering the **atomic** construct on entry to the
3 atomic region before initiating synchronization for the region.

4 The *atomic-acquired* event occurs in the thread encountering the **atomic** construct after entering
5 the region, but before executing the structured block of the **atomic** region.

6 The *atomic-release* event occurs in the thread encountering the **atomic** construct after completing
7 any synchronization on exit from the **atomic** region.

8 Tool Callbacks

9 A thread dispatches a registered **ompt_callback_mutex_acquire** callback for each
10 occurrence of an *atomic-acquire* event in that thread. This callback has the type signature
11 **ompt_callback_mutex_acquire_t**.

12 A thread dispatches a registered **ompt_callback_mutex_acquired** callback for each
13 occurrence of an *atomic-acquired* event in that thread. This callback has the type signature
14 **ompt_callback_mutex_t**.

15 A thread dispatches a registered **ompt_callback_mutex_released** callback for each
16 occurrence of an *atomic-release* event in that thread. This callback has the type signature
17 **ompt_callback_mutex_t**. The callbacks occur in the task encountering the atomic construct.
18 The callbacks should receive **ompt_mutex_atomic** as their *kind* argument if practical, but a
19 less specific kind is acceptable.

20 Restrictions

21 The following restrictions apply to the **atomic** construct:

- 22 • At most one *memory-order-clause* may appear on the construct.
- 23 • At most one **hint** clause may appear on the construct.
- 24 • If *atomic-clause* is **read** then *memory-order-clause* must not be **acq_rel** or **release**.
- 25 • If *atomic-clause* is **write** then *memory-order-clause* must not be **acq_rel** or **acquire**.
- 26 • If *atomic-clause* is **update** or not present then *memory-order-clause* must not be **acq_rel** or
27 **acquire**.

▼ C / C++ ▼

- 28 • All atomic accesses to the storage locations designated by *x* throughout the program are required
29 to have a compatible type.

▲ C / C++ ▲

Fortran

- All atomic accesses to the storage locations designated by *x* throughout the program are required to have the same type and type parameters.

Fortran

- OpenMP constructs may not be encountered during execution of an **atomic** region.

Cross References

- **critical** construct, see Section 2.20.1 on page 216.
- **barrier** construct, see Section 2.20.2 on page 219.
- **flush** construct, see Section 2.20.8 on page 235.
- **ordered** construct, see Section 2.20.9 on page 243.
- **reduction** clause, see Section 2.22.5.4 on page 297.
- Synchronization Hints, see Section 2.20.12 on page 253.
- lock routines, see Section 3.3 on page 378.
- **ompt_mutex_atomic**, see Section 4.2.3.4.15 on page 439.
- **ompt_callback_mutex_acquire_t**, see Section 4.2.4.2.12 on page 458.
- **ompt_callback_mutex_t**, see Section 4.2.4.2.13 on page 459.

2.20.8 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 22 for more details. The **flush** construct is a stand-alone directive.

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19

Syntax

C / C++

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

```
#pragma omp flush [memory-order-clause] [(list)] new-line
```

where *memory-order-clause* is one of the following:

```
acq_rel  
release  
acquire
```

C / C++

Fortran

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

```
!$omp flush [memory-order-clause] [(list)]
```

where *memory-order-clause* is one of the following:

```
acq_rel  
release  
acquire
```

Fortran

Binding

The binding thread set for a **flush** region is the encountering thread. Execution of a **flush** region affects the memory and the temporary view of memory of only the thread that executes the region. It does not affect the temporary view of other threads. Other threads must themselves execute a flush operation in order to be guaranteed to observe the effects of the encountering thread's flush operation

Description

If *memory-order-clause* is not specified then the **flush** construct results in a strong flush operation with the following behavior. A **flush** construct without a list, executed on a given thread, operates as if the whole thread-visible data state of the program, as defined by the base language, is flushed. A **flush** construct with a list applies the flush operation to the items in the list, and does not return until the operation is complete for all specified list items. An implementation may implement a **flush** with a list by ignoring the list, and treating it the same as a **flush** without a list.

If no list items are specified, the flush operation has the release and/or acquire flush properties:

- If *memory-order-clause* is not specified or is **acq_rel**, the flush operation is both a release flush and an acquire flush.
- If *memory-order-clause* is **release**, the flush operation is a release flush.
- If *memory-order-clause* is **acquire**, the flush operation is an acquire flush.

▼ C / C++ ▼

If a pointer is present in the list, the pointer itself is flushed, not the memory block to which the pointer refers.

▲ C / C++ ▲

▼ Fortran ▼

If the list item or a subobject of the list item has the **POINTER** attribute, the allocation or association status of the **POINTER** item is flushed, but the pointer target is not. If the list item is a Cray pointer, the pointer is flushed, but the object to which it points is not. If the list item is of type **C_PTR**, the variable is flushed, but the storage that corresponds to that address is not flushed. If the list item or the subobject of the list item has the **ALLOCATABLE** attribute and has an allocation status of allocated, the allocated variable is flushed; otherwise the allocation status is flushed.

▲ Fortran ▲

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. Execution of the protected section by neither
19 thread is considered correct in this example. This occurs if both flushes complete prior to either
20 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.



Execution Model Events

The *flush* event occurs in a thread encountering the `flush` construct.

Tool Callbacks

A thread dispatches a registered `ompt_callback_flush` callback for each occurrence of a *flush* event in that thread. This callback has the type signature `ompt_callback_flush_t`.

Restrictions

The following restrictions apply to the `flush` construct:

- If *memory-order-clause* is `release`, `acquire`, or `acq_rel`, list items must not be specified on the `flush` directive.

Cross References

- `ompt_callback_flush_t`, see Section 4.2.4.2.16 on page 463.

1 2.20.8.1 Implicit Flushes

2 Flush operations implied when executing an **atomic** region are described in Section [2.20.7](#).

3 A **flush** region that corresponds to a **flush** directive with the **release** clause present is
4 implied at the following locations:

- 5 • During a barrier region.
- 6 • At entry to **parallel** regions.
- 7 • At entry to **teams** regions.
- 8 • At exit from **critical** regions.
- 9 • During **omp_unset_lock** and **omp_unset_nest_lock** regions.
- 10 • Immediately before every task scheduling point.
- 11 • At exit from the task region of each implicit task.
- 12 • At exit from an **ordered** region, if a **threads** clause or a **depend** clause with a **source**
13 dependence type is present, or if no clauses are present.
- 14 • During a **cancel** region, if the *cancel-var* ICV is *true*.

15 A **flush** region that corresponds to a **flush** directive with the **acquire** clause present is
16 implied at the following locations:

- 17 • During a barrier region.
- 18 • At exit from **teams** regions.
- 19 • At entry to **critical** regions.
- 20 • During **omp_set_lock**, **omp_test_lock**, **omp_set_nest_lock**, and
21 **omp_test_nest_lock** regions, if the region causes the lock to be set.
- 22 • Immediately after every task scheduling point.
- 23 • At entry to the task region of each implicit task.
- 24 • At entry to an **ordered** region, if a **threads** clause or a **depend** clause with a **sink**
25 dependence type is present, or if no clauses are present.
- 26 • Immediately before a cancellation point, if the *cancel-var* ICV is *true* and cancellation has been
27 activated.

1 Note – A **flush** region is not implied at the following locations:

- 2
- At entry to worksharing regions.
 - At entry to or exit from a **master** region.
- 3
-

4 The synchronization behavior of implicit flushes is as follows:

- 5
- When a thread executes a **critical** region that has a given name, the behavior is as if the release flush performed on exit from the region synchronizes with the acquire flush performed on entry to the next **critical** region with the same name that is performed by a different thread, if it exists.
 - When a thread team executes a **barrier** region, the behavior is as if the release flush performed by each thread within the region synchronizes with the acquire flush performed by all other threads within the region.
 - When a thread executes a **taskwait** region that does not result in the creation of a dependent task, the behavior is as if each thread that executes a remaining child task performs a release flush upon completion of the child task that synchronizes with an acquire flush performed in the **taskwait** region.
 - When a thread executes a **taskgroup** region, the behavior is as if each thread that executes a remaining descendant task performs a release flush upon completion of the descendant task that synchronizes with an acquire flush performed on exit from the **taskgroup** region.
 - When a thread executes an **ordered** region that does not arise from a stand-alone **ordered** directive, the behavior is as if the release flush performed on exit from the region synchronizes with the acquire flush performed on entry to an **ordered** region encountered in the next logical iteration to be executed by a different thread, if it exists.
 - When a thread executes an **ordered** region that arises from a stand-alone **ordered** directive, the behavior is as if the release flush performed in the **ordered** region from a given source iteration synchronizes with the acquire flush performed in all **ordered** regions executed by a different thread that are waiting for dependences on that iteration to be satisfied.
 - When a thread team begins execution of a **parallel** region, the behavior is as if the release flush performed by the master thread on entry to the **parallel** region synchronizes with the acquire flush performed on entry to each implicit task that is assigned to a different thread.
 - When an initial thread begins execution of a **target** region that is generated by a different thread from a target task, the behavior is as if the release flush performed by the generating thread in the target task synchronizes with the acquire flush performed by the initial thread on entry to its initial task region.
- 30
- 31
- 32
- 33

- 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
- 27
- 28
- 29
- 30
- 31
- 32
- 33
- 34
- 35
- 36
- 37
- 38
- When an initial thread completes execution of a **target** region that is generated by a different thread from a target task, the behavior is as if the release flush performed by the initial thread on exit from its initial task region synchronizes with the acquire flush performed by the generating thread in the target task.
 - When a thread encounters a **teams** construct, the behavior is as if the release flush performed by the thread on entry to the **teams** region synchronizes with the acquire flush performed on entry to each initial task that is executed by a different initial thread that participates in the execution of the **teams** region.
 - When a thread that encounters a **teams** construct reaches the end of the **teams** region, the behavior is as if the release flush performed by each different participating initial thread at exit from its initial task synchronizes with the acquire flush performed by the thread at exit from the **teams** region.
 - When a task generates an explicit task that begins execution on a different thread, the behavior is as if the thread that is executing the generating task performs a release flush that synchronizes with the acquire flush performed by the thread that begins to execute the explicit task.
 - When a dependent task with one or more predecessor tasks begins execution on a given thread, the behavior is as if each release flush performed by a different thread on completion of a predecessor task synchronizes with the acquire flush performed by the thread that begins to execute the dependent task.
 - When a task begins execution on a given thread and it is mutually exclusive with respect to another sibling task that is executed by a different thread, the behavior is as if each release flush performed on completion of the sibling task synchronizes with the acquire flush performed by the thread that begins to execute the task.
 - When a thread executes a **cancel** region, the *cancel-var* ICV is *true*, and cancellation is not already activated for the specified region, the behavior is as if the release flush performed during the **cancel** region synchronizes with the acquire flush performed by a different thread immediately before a cancellation point in which that thread observes cancellation was activated for the region.
 - When a thread executes an **omp_unset_lock** region that causes the specified lock to be unset, the behavior is as if a release flush is performed during the **omp_unset_lock** region that synchronizes with an acquire flush that is performed during the next **omp_set_lock** or **omp_test_lock** region to be executed by a different thread that causes the specified lock to be set.
 - When a thread executes an **omp_unset_nest_lock** region that causes the specified nested lock to be unset, the behavior is as if a release flush is performed during the **omp_unset_nest_lock** region that synchronizes with an acquire flush that is performed during the next **omp_set_nest_lock** or **omp_test_nest_lock** region to be executed by a different thread that causes the specified nested lock to be set.

1 2.20.9 ordered Construct

2 Summary

3 The **ordered** construct either specifies a structured block in a worksharing-loop, **simd**, or
4 worksharing-loop SIMD region that will be executed in the order of the loop iterations, or it is a
5 stand-alone directive that specifies cross-iteration dependences in a doacross loop nest. The
6 **ordered** construct sequentializes and orders the execution of **ordered** regions while allowing
7 code outside the region to run in parallel.

8 Syntax

C / C++

9 The syntax of the **ordered** construct is as follows:

```
10 #pragma omp ordered [clause[ [, ] clause] ] new-line  
11     structured-block
```

12 where *clause* is one of the following:

```
13     threads  
14     simd
```

15 or

```
16 #pragma omp ordered clause [[[ , ] clause] ... ] new-line
```

17 where *clause* is one of the following:

```
18     depend(source)  
19     depend(sink : vec)
```

C / C++

Fortran

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

```
21 !$omp ordered [clause[ [, ] clause] ]  
22     structured-block  
23 !$omp end ordered
```

24 where *clause* is one of the following:

```
25     threads  
26     simd
```

27 or

1 **!\$omp ordered** *clause* [[[,] *clause*] ...]

2 where *clause* is one of the following:

3 **depend**(*source*)
4 **depend**(*sink* : *vec*)

Fortran

5 If the **depend** clause is specified, the **ordered** construct is a stand-alone directive.

6 Binding

7 The binding thread set for an **ordered** region is the current team. An **ordered** region binds to
8 the innermost enclosing **simd** or worksharing-loop SIMD region if the **simd** clause is present, and
9 otherwise it binds to the innermost enclosing worksharing-loop region. **ordered** regions that bind
10 to different regions execute independently of each other.

11 Description

12 If no clause is specified, the **ordered** construct behaves as if the **threads** clause had been
13 specified. If the **threads** clause is specified, the threads in the team executing the
14 worksharing-loop region execute **ordered** regions sequentially in the order of the loop iterations.
15 If any **depend** clauses are specified then those clauses specify the order in which the threads in the
16 team execute **ordered** regions. If the **simd** clause is specified, the **ordered** regions
17 encountered by any thread will execute one at a time in the order of the loop iterations.

18 When the thread executing the first iteration of the loop encounters an **ordered** construct, it can
19 enter the **ordered** region without waiting. When a thread executing any subsequent iteration
20 encounters an **ordered** construct without a **depend** clause, it waits at the beginning of the
21 **ordered** region until execution of all **ordered** regions belonging to all previous iterations has
22 completed. When a thread executing any subsequent iteration encounters an **ordered** construct
23 with one or more **depend**(*sink:vec*) clauses, it waits until its dependences on all valid
24 iterations specified by the **depend** clauses are satisfied before it completes execution of the
25 **ordered** region. A specific dependence is satisfied when a thread executing the corresponding
26 iteration encounters an **ordered** construct with a **depend**(*source*) clause.

27 Execution Model Events

28 The *ordered-acquire* event occurs in the thread encountering the **ordered** construct on entry to
29 the ordered region before initiating synchronization for the region.

30 The *ordered-acquired* event occurs in the thread encountering the **ordered** construct after
31 entering the region, but before executing the structured block of the **ordered** region.

32 The *ordered-release* event occurs in the thread encountering the **ordered** construct after
33 completing any synchronization on exit from the **ordered** region.

1 Tool Callbacks

2 A thread dispatches a registered `ompt_callback_mutex_acquire` callback for each
3 occurrence of an *ordered-acquire* event in that thread. This callback has the type signature
4 `ompt_callback_mutex_acquire_t`.

5 A thread dispatches a registered `ompt_callback_mutex_acquired` callback for each
6 occurrence of an *ordered-acquired* event in that thread. This callback has the type signature
7 `ompt_callback_mutex_t`.

8 A thread dispatches a registered `ompt_callback_mutex_released` callback for each
9 occurrence of an *ordered-release* event in that thread. This callback has the type signature
10 `ompt_callback_mutex_t`. The callbacks occur in the task encountering the ordered
11 construct. The callbacks should receive `ompt_mutex_ordered` as their *kind* argument if
12 practical, but a less specific kind is acceptable.

13 Restrictions

14 Restrictions to the `ordered` construct are as follows:

- 15 • At most one `threads` clause can appear on an `ordered` construct.
- 16 • At most one `simd` clause can appear on an `ordered` construct.
- 17 • At most one `depend(source)` clause can appear on an `ordered` construct.
- 18 • The construct corresponding to the binding region of an `ordered` region must not specify a
19 `reduction` clause with the `inscan` modifier.
- 20 • Either `depend(sink:vec)` clauses or `depend(source)` clauses may appear on an
21 `ordered` construct, but not both.
- 22 • The worksharing-loop or worksharing-loop SIMD region to which an `ordered` region
23 corresponding to an `ordered` construct without a `depend` clause binds must have an
24 `ordered` clause without the parameter specified on the corresponding worksharing-loop or
25 worksharing-loop SIMD directive.
- 26 • The worksharing-loop region to which an `ordered` region corresponding to an `ordered`
27 construct with any `depend` clauses binds must have an `ordered` clause with the parameter
28 specified on the corresponding worksharing-loop directive.
- 29 • An `ordered` construct with the `depend` clause specified must be closely nested inside a
30 worksharing-loop (or parallel worksharing-loop) construct.
- 31 • An `ordered` region corresponding to an `ordered` construct without the `simd` clause
32 specified must be closely nested inside a loop region.
- 33 • An `ordered` region corresponding to an `ordered` construct with the `simd` clause specified
34 must be closely nested inside a `simd` or worksharing-loop SIMD region.

- 1 • An **ordered** region corresponding to an **ordered** construct with both the **simd** and
2 **threads** clauses must be closely nested inside a worksharing-loop SIMD region or must be
3 closely nested inside a worksharing-loop and **simd** region.
- 4 • During execution of an iteration of a worksharing-loop or a loop nest within a worksharing-loop,
5 **simd**, or worksharing-loop SIMD region, a thread must not execute more than one **ordered**
6 region corresponding to an **ordered** construct without a **depend** clause.

▼ C++ ▼

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

▲ C++ ▲

9 Cross References

- 10 • worksharing-loop construct, see Section [2.12.2](#) on page [100](#).
- 11 • **simd** construct, see Section [2.12.3.1](#) on page [111](#).
- 12 • parallel Worksharing-loop construct, see Section [2.16.1](#) on page [185](#).
- 13 • **depend** Clause, see Section [2.20.11](#) on page [248](#)
- 14 • **ompt_mutex_ordered**, see Section [4.2.3.4.15](#) on page [439](#).
- 15 • **ompt_callback_mutex_acquire_t**, see Section [4.2.4.2.12](#) on page [458](#).
- 16 • **ompt_callback_mutex_t**, see Section [4.2.4.2.13](#) on page [459](#).

17 2.20.10 Depend Objects

18 This section describes constructs that support OpenMP depend objects that can be used to supply
19 user computed dependences to **depend** clauses. OpenMP depend objects must be accessed only
20 through the **depobj** construct or through the **depend** clause; programs that otherwise access
21 OpenMP depend objects are non-conforming.

22 An OpenMP dependence object can be in one of the following states: *uninitialized* or *initialized*.
23 Initially OpenMP depend objects are in the *uninitialized* state.

1 2.20.10.1 depobj Construct

2 Summary

3 The **depobj** construct allows to initialize, update and destroy depend objects that represent data
4 dependences. The **depobj** construct is a stand-alone directive.

5 Syntax

C / C++

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

```
7 #pragma omp depobj (depobj) clause new-line
```

8 where *depobj* is an lvalue expression of type **omp_depend_t**.

9 where *clause* is one of the following:

```
10 depend (dependence-type : locator)  
11 destroy  
12 update (dependence-type)
```

C / C++

Fortran

13 The syntax of the **depobj** construct is as follows:

```
14 !$omp depobj (depobj) clause
```

15 where *depobj* is a scalar integer variable of the **omp_depend_kind** kind.

16 where *clause* is one of the following:

```
17 depend (dependence-type : locator)  
18 destroy  
19 update (dependence-type)
```

Fortran

20 Binding

21 The binding thread set for **depobj** regions is the encountering thread.

Description

A **depobj** construct with a **depend** clause present initializes the *depobj* to represent the dependence specified by the **depend** clause.

A **depobj** construct with a **destroy** clause present changes the state of the *depobj* to uninitialized.

A **depobj** construct with a **update** clause present changes the dependence type of the dependence represented by *depobj* to the one specified by *update* clause.

Restrictions

- A **depend** clause on a **depobj** construct must not have **source** or **sink** as *dependence-type*.
- A **depend** clause on a **depobj** construct must specify the *dependence-type*.
- A **depend** clause on a **depobj** construct can only specify one locator.
- The *depobj* of a **depobj** construct with the **depend** clause present must be in the uninitialized state.
- The *depobj* of a **depobj** construct with the **destroy** clause present must be in the initialized state.
- The *depobj* of a **depobj** construct with the **update** clause present must be in the initialized state.

Cross References

- **depend** clause, see Section [2.20.11](#) on page [248](#).

2.20.11 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.

Syntax

The syntax of the **depend** clause is as follows:

```
depend ([depend-modifier : ] [dependence-type : ] locator-list)
```

where *dependence-type* is one of the following:

```
in  
out  
inout  
mutexinoutset
```

where *depend-modifier* is one of the following:

```
iterator (iterators-definition)
```

or

```
depend (dependence-type)
```

where *dependence-type* is:

```
source
```

or

```
depend (dependence-type : vec)
```

where *dependence-type* is:

```
sink
```

and where *vec* is the iteration vector, which has the form:

$$x_1 [\pm d_1], x_2 [\pm d_2], \dots, x_n [\pm d_n]$$

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.

Description

Task dependences are derived from the *dependence-type* of a **depend** clause and its list items when *dependence-type* is **in**, **out**, **inout**, or **mutexinoutset**. When the **depend** clause has no *dependence-type* specified, the task dependences are derived from the **depend** clause specified in the **depobj** constructs that initialized the depend objects specified on the **depend** clause as if the **depend** clauses of the **depobj** constructs were specified in the current construct.

1 For the **in** *dependence-type*, if the storage location of at least one of the list items is the same as the
2 storage location of a list item appearing in a **depend** clause with an **out**, **inout**, or
3 **mutexinoutset** *dependence-type* on a construct from which a sibling task was previously
4 generated, then the generated task will be a dependent task of that sibling task.

5 For the **out** and **inout** *dependence-types*, if the storage location of at least one of the list items is
6 the same as the storage location of a list item appearing in a **depend** clause with an **in**, **out**,
7 **inout**, or **mutexinoutset** *dependence-type* on a construct from which a sibling task was
8 previously generated, then the generated task will be a dependent task of that sibling task.

9 For the **mutexinoutset** *dependence-type*, if the storage location of at least one of the list items
10 is the same as the storage location of a list item appearing in a **depend** clause with an **in**, **out**, or
11 **inout** *dependence-type* on a construct from which a sibling task was previously generated, then
12 the generated task will be a dependent task of that sibling task.

13 If a list item appearing in a **depend** clause with a **mutexinoutset** *dependence-type* on a
14 task-generating construct has the same storage location as a list item appearing in a **depend** clause
15 with a **mutexinoutset** *dependence-type* on a different task generating construct, and both
16 constructs generate sibling tasks, the sibling tasks will be mutually exclusive tasks.

17 The list items that appear in the **depend** clause may reference iterators defined by an
18 *iterators-definition* appearing on an **iterator** modifier.

▼ Fortran ▲

19 If a list item has the **ALLOCATABLE** attribute and its allocation status is unallocated, the behavior
20 is unspecified. If a list item has the **POINTER** attribute and its association status is disassociated or
21 undefined, the behavior is unspecified.

▲ Fortran ▼

22 The list items that appear in the **depend** clause may include array sections.

▼ C / C++ ▼

23 The list items that appear in the **depend** clause may use shape-operators.

▲ C / C++ ▲

24 **Note** – The enforced task dependence establishes a synchronization of memory accesses
25 performed by a dependent task with respect to accesses performed by the predecessor tasks.
26 However, it is the responsibility of the programmer to synchronize properly with respect to other
27 concurrent accesses that occur outside of those tasks.



1 The **source** *dependence-type* specifies the satisfaction of cross-iteration dependences that arise
2 from the current iteration.

3 The **sink** *dependence-type* specifies a cross-iteration dependence, where the iteration vector *vec*
4 indicates the iteration that satisfies the dependence.

5 If the iteration vector *vec* does not occur in the iteration space, the **depend** clause is ignored. If all
6 **depend** clauses on an **ordered** construct are ignored then the construct is ignored.

7 **Note** – If the iteration vector *vec* does not indicate a lexicographically earlier iteration, it can cause
8 a deadlock.

9 Execution Model Events

10 The *task-dependences* event occurs in a thread encountering a tasking construct or a taskwait
11 construct with a **depend** clause immediately after the *task-create* event for the new task or the
12 *taskwait-begin* event.

13 The *task-dependence* event indicates an unfulfilled dependence for the generated task. This event
14 occurs in a thread that observes the unfulfilled dependence before it is satisfied.

15 Tool Callbacks

16 A thread dispatches the **ompt_callback_task_dependences** callback for each occurrence
17 of the *task-dependences* event to announce its dependences with respect to the list items in the
18 **depend** clause. This callback has type signature **ompt_callback_task_dependences_t**.

19 A thread dispatches the **ompt_callback_task_dependence** callback for a *task-dependence*
20 event to report a dependence between a predecessor task (*src_task_data*) and a dependent task
21 (*sink_task_data*). This callback has type signature **ompt_callback_task_dependence_t**.

22 Restrictions

23 Restrictions to the **depend** clause are as follows:

- 24 • List items used in **depend** clauses of the same task or sibling tasks must indicate identical
25 storage locations or disjoint storage locations.
- 26 • List items used in **depend** clauses cannot be zero-length array sections.
- 27 • Array sections cannot be specified in **depend** clauses with no *dependence-type* specified.
- 28 • List items used in **depend** clauses with no *dependence-type* specified must be depend objects in
29 the initialized state.

C / C++

1
2
3
4

- List items used in **depend** clauses with no *dependence-type* specified must be expressions of the **omp_depend_t** type.
- List items used in **depend** clauses with the **in**, **out**, **inout** or **mutexinoutset** dependence types cannot be expressions of the **omp_depend_t** type.

C / C++

Fortran

5
6
7

- A common block name cannot appear in a **depend** clause.
- List items used in **depend** clauses with no *dependence-type* specified must be integer expressions of the **omp_depend_kind** *kind*.

Fortran

8
9
10
11

- For a *vec* element of **sink** *dependence-type* of the form $x_i + d_i$ or $x_i - d_i$ if the loop iteration variable x_i has an integral or pointer type, the expression $x_i + d_i$ or $x_i - d_i$ for any value of the loop iteration variable x_i that can encounter the **ordered** construct must be computable in the loop iteration variable's type without overflow.

C++

12
13
14
15
16

- For a *vec* element of **sink** *dependence-type* of the form $x_i + d_i$ or $x_i - d_i$ if the loop iteration variable x_i is of a random access iterator type other than pointer type, the expression $(x_i - lb_i) + d_i$ or $(x_i - lb_i) - d_i$ for any value of the loop iteration variable x_i that can encounter the **ordered** construct must be computable in the type that would be used by *std::distance* applied to variables of the type of x_i without overflow.

C++

C / C++

17

- A bit-field cannot appear in a **depend** clause.

C / C++

Cross References

- Array sections, see Section 2.6 on page 59.
- Iterators, see Section 2.7 on page 61.
- **task** construct, see Section 2.13.1 on page 133.
- Task scheduling constraints, see Section 2.13.6 on page 147.
- **target enter data** construct, see Section 2.15.3 on page 162.
- **target exit data** construct, see Section 2.15.4 on page 165.
- **target** construct, see Section 2.15.5 on page 168.
- **target update** construct, see Section 2.15.6 on page 174.
- **ordered** construct, see Section 2.20.9 on page 243.
- **omp_callback_task_dependences_t**, see Section 4.2.4.2.7 on page 453.
- **omp_callback_task_dependence_t**, see Section 4.2.4.2.8 on page 454.

2.20.12 Synchronization Hints

Hints about the expected dynamic behavior or suggested implementation can be provided by the programmer to locks (by using the **omp_init_lock_with_hint** or **omp_init_nest_lock_with_hint** functions to initialize the lock), and to **atomic** and **critical** directives by using the **hint** clause. The effect of a hint is implementation defined. The OpenMP implementation is free to ignore the hint since doing so cannot change program semantics.

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 hint constants. The valid constants must include the following, which can be extended with implementation-defined values:

C / C++

```

1  typedef enum omp_sync_hint_t {
2      omp_sync_hint_none = 0x0,
3      omp_lock_hint_none = omp_sync_hint_none,
4      omp_sync_hint_uncontended = 0x1,
5      omp_lock_hint_uncontended = omp_sync_hint_uncontended,
6      omp_sync_hint_contended = 0x2,
7      omp_lock_hint_contended = omp_sync_hint_contended,
8      omp_sync_hint_nonspeculative = 0x4,
9      omp_lock_hint_nonspeculative = omp_sync_hint_nonspeculative,
10     omp_sync_hint_speculative = 0x8
11     omp_lock_hint_speculative = omp_sync_hint_speculative
12 } omp_sync_hint_t;
13
14 typedef omp_sync_hint_t omp_lock_hint_t;

```

C / C++

Fortran

```

15 integer, parameter :: omp_lock_hint_kind = omp_sync_hint_kind
16
17 integer (kind=omp_sync_hint_kind), &
18     parameter :: omp_sync_hint_none = Z'0'
19 integer (kind=omp_lock_hint_kind), &
20     parameter :: omp_lock_hint_none = omp_sync_hint_none
21 integer (kind=omp_sync_hint_kind), &
22     parameter :: omp_sync_hint_uncontended = Z'1'
23 integer (kind=omp_lock_hint_kind), &
24     parameter :: omp_lock_hint_uncontended = &
25         omp_sync_hint_uncontended
26 integer (kind=omp_sync_hint_kind), &
27     parameter :: omp_sync_hint_contended = Z'2'
28 integer (kind=omp_lock_hint_kind), &
29     parameter :: omp_lock_hint_contended = &
30         omp_sync_hint_contended
31 integer (kind=omp_sync_hint_kind), &
32     parameter :: omp_sync_hint_nonspeculative = Z'4'
33 integer (kind=omp_lock_hint_kind), &
34     parameter :: omp_lock_hint_nonspeculative = &
35         omp_sync_hint_nonspeculative
36 integer (kind=omp_sync_hint_kind), &
37     parameter :: omp_sync_hint_speculative = Z'8'
38 integer (kind=omp_lock_hint_kind), &

```

```
1 parameter :: omp_lock_hint_speculative = &
2 omp_sync_hint_speculative
```

Fortran

3 The hints can be combined by using the `+` or `|` operators in C/C++ or the `+` operator in Fortran.
4 The effect of the combined hint is implementation defined and can be ignored by the
5 implementation. Combining `omp_sync_hint_none` with any other hint is equivalent to
6 specifying the other hint. The following restrictions apply to combined hints; violating these
7 restrictions results in unspecified behavior:

- 8 • the hints `omp_sync_hint_uncontended` and `omp_sync_hint_contended` cannot be
9 combined,
- 10 • the hints `omp_sync_hint_nonspeculative` and `omp_sync_hint_speculative`
11 cannot be combined.

12 The rules for combining multiple values of `omp_sync_hint` apply equally to the corresponding
13 values of `omp_lock_hint`, and expressions mixing the two types.

14 The intended meaning of hints is

- 15 • `omp_sync_hint_uncontended`: low contention is expected in this operation, that is, few
16 threads are expected to be performing the operation simultaneously in a manner that requires
17 synchronization.
- 18 • `omp_sync_hint_contended`: high contention is expected in this operation, that is, many
19 threads are expected to be performing the operation simultaneously in a manner that requires
20 synchronization.
- 21 • `omp_sync_hint_speculative`: the programmer suggests that the operation should be
22 implemented using speculative techniques such as transactional memory.
- 23 • `omp_sync_hint_nonspeculative`: the programmer suggests that the operation should
24 not be implemented using speculative techniques such as transactional memory.

25 Note – Future OpenMP specifications may add additional hints to the `omp_sync_hint_t` type
26 and the `omp_sync_hint_kind` kind. Implementers are advised to add implementation-defined
27 hints starting from the most significant bit of the `omp_sync_hint_t` type and
28 `omp_sync_hint_kind` kind and to include the name of the implementation in the name of the
29 added hint to avoid name conflicts with other OpenMP implementations.

30 The `omp_sync_hint_t` and `omp_lock_hint_t` enumeration types and the equivalent types
31 in Fortran are synonyms for each other. The type `omp_lock_hint_t` has been deprecated.

Cross References

- **critical** construct, see Section 2.20.1 on page 216.
- **atomic** construct, see Section 2.20.7 on page 227
- **omp_init_lock_with_hint** and **omp_init_nest_lock_with_hint**, see Section 3.3.2 on page 382.

2.21 Cancellation Constructs

2.21.1 `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
```

where *construct-type-clause* is one of the following:

```
parallel  
sections  
for  
taskgroup
```

and *if-clause* is

```
if ([ cancel : ] scalar-expression)
```

C / C++

Fortran

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

```
!$omp cancel construct-type-clause [ [, ] if-clause ]
```

where *construct-type-clause* is one of the following:

```
parallel  
sections  
do  
taskgroup
```

and *if-clause* is

```
if ( [ cancel : ] scalar-logical-expression )
```

Fortran

Binding

The binding thread set of the **cancel** region is the current team. The binding region of the **cancel** region is the innermost enclosing region of the type corresponding to the *construct-type-clause* specified in the directive (that is, the innermost **parallel**, **sections**, loop, or **taskgroup** region).

Description

The **cancel** construct activates cancellation of the binding region only if the *cancel-var* ICV is *true*, in which case the **cancel** construct causes the encountering task to continue execution at the end of the binding region if *construct-type-clause* is **parallel**, **for**, **do**, or **sections**. If the *cancel-var* ICV is *true* and *construct-type-clause* is **taskgroup**, the encountering task continues execution at the end of the current task region. If the *cancel-var* ICV is *false*, the **cancel** construct is ignored.

Threads check for active cancellation only at cancellation points that are implied at the following locations:

- **cancel** regions;
- **cancellation point** regions;
- **barrier** regions;
- implicit barriers regions.

When a thread reaches one of the above cancellation points and if the *cancel-var* ICV is *true*, then:

- If the thread is at a **cancel** or **cancellation point** region and *construct-type-clause* is **parallel**, **for**, **do**, or **sections**, the thread continues execution at the end of the canceled region if cancellation has been activated for the innermost enclosing region of the type specified.
- If the thread is at a **cancel** or **cancellation point** region and *construct-type-clause* is **taskgroup**, the encountering task checks for active cancellation of all of the *taskgroup sets* to which the encountering task belongs, and continues execution at the end of the current task region if cancellation has been activated for any of the *taskgroup sets*.
- If the encountering task is at a barrier region, the encountering task checks for active cancellation of the innermost enclosing **parallel** region. If cancellation has been activated, then the encountering task continues execution at the end of the canceled region.

Note – If one thread activates cancellation and another thread encounters a cancellation point, the 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 tasks that belong to the *taskgroup set* of the innermost enclosing **taskgroup** region 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**, or worksharing-loop 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**, or worksharing-loop region, no task cancellation occurs.

C++

The usual C++ rules for object destruction are followed when cancellation is performed.

C++

Fortran

1 All private objects or subobjects with **ALLOCATABLE** attribute that are allocated inside the
2 canceled construct are deallocated.

Fortran

3 If the canceled construct contains a **reduction** or **lastprivate** clause, the final value of the
4 **reduction** or **lastprivate** variable is undefined.

5 When an **if** clause is present on a **cancel** construct and the **if** expression evaluates to *false*, the
6 **cancel** construct does not activate cancellation. The cancellation point associated with the
7 **cancel** construct is always encountered regardless of the value of the **if** expression.

8 **Note** – The programmer is responsible for releasing locks and other synchronization data
9 structures that might cause a deadlock when a **cancel** construct is encountered and blocked
10 threads cannot be canceled. The programmer is also responsible for ensuring proper
11 synchronizations to avoid deadlocks that might arise from cancellation of OpenMP regions that
12 contain OpenMP synchronization constructs.

13 Execution Model Events

14 If a task encounters a **cancel** construct that will activate cancellation then a *cancel* event occurs.
15 A *discarded-task* event occurs for any discarded tasks.

16 Tool Callbacks

17 A thread dispatches a registered **ompt_callback_cancel** callback for each occurrence of a
18 *cancel* event in that thread. The callback occurs in the context of the encountering task. The
19 callback has type signature **ompt_callback_cancel_t**. The callback receives
20 **ompt_cancel_activated** as its *flags* argument.

21 A thread dispatches a registered **ompt_callback_cancel** callback for each occurrence of a
22 *discarded-task* event. The callback occurs in the context of the task that discards the task. The
23 callback has type signature **ompt_callback_cancel_t**. The callback receives the
24 *ompt_data_t* associated with the discarded task as its *task_data* argument. The callback receives
25 **ompt_cancel_discarded_task** as its *flags* argument.

1 **Restrictions**

2 The restrictions to the **cancel** construct are as follows:

- 3 • The behavior for concurrent cancellation of a region and a region nested within it is unspecified.
- 4 • If *construct-type-clause* is **taskgroup**, the **cancel** construct must be closely nested inside a
5 **task** construct and the **cancel** region must be closely nested inside a **taskgroup** region. If
6 *construct-type-clause* is **sections**, the **cancel** construct must be closely nested inside a
7 **sections** or **section** construct. Otherwise, the **cancel** construct must be closely nested
8 inside an OpenMP construct that matches the type specified in *construct-type-clause* of the
9 **cancel** construct.
- 10 • A worksharing construct that is canceled must not have a **nowait** clause.
- 11 • A worksharing-loop construct that is canceled must not have an **ordered** clause.
- 12 • During execution of a construct that may be subject to cancellation, a thread must not encounter
13 an orphaned cancellation point. That is, a cancellation point must only be encountered within
14 that construct and must not be encountered elsewhere in its region.

15 **Cross References**

- 16 • *cancel-var* ICV, see Section 2.4.1 on page 47.
- 17 • **if** Clause, see Section 2.18 on page 213.
- 18 • **cancellation point** construct, see Section 2.21.2 on page 260.
- 19 • **omp_get_cancellation** routine, see Section 3.2.9 on page 340.
- 20 • **ompt_callback_cancel_t**, see Section 4.2.4.2.27 on page 477.
- 21 • **omp_cancel_flag_t** enumeration type, see Section 4.2.3.4.23 on page 443.

22 **2.21.2 cancellation point Construct**

23 **Summary**

24 The **cancellation point** construct introduces a user-defined cancellation point at which
25 implicit or explicit tasks check if cancellation of the innermost enclosing region of the type
26 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++

Fortran

The syntax of the **cancellation point** construct is as follows:

```
!$omp cancellation point construct-type-clause
```

where *construct-type-clause* is one of the following:

```
parallel  
sections  
do  
taskgroup
```

Fortran

Binding

The binding thread set of the **cancellation point** construct is the current team. The binding region of the **cancellation point** region is the innermost enclosing region of the type corresponding to the *construct-type-clause* specified in the directive (that is, the innermost **parallel**, **sections**, loop, or **taskgroup** region).

Description

This directive introduces a user-defined cancellation point at which an implicit or explicit task must check if cancellation of the innermost enclosing region of the type specified in the clause has been requested. This construct does not implement any synchronization between threads or tasks.

When an implicit or explicit task reaches a user-defined cancellation point and if the *cancel-var* ICV is *true*, then:

- If the *construct-type-clause* of the encountered **cancellation point** construct is **parallel**, **for**, **do**, or **sections**, the thread continues execution at the end of the canceled region if cancellation has been activated for the innermost enclosing region of the type specified.
- If the *construct-type-clause* of the encountered **cancellation point** construct is **taskgroup**, the encountering task checks for active cancellation of all *taskgroup sets* to which the encountering task belongs and continues execution at the end of the current task region if cancellation has been activated for any of them.

Execution Model Events

The *cancellation* event occurs if a task encounters a cancellation point and detected the activation of cancellation.

Tool Callbacks

A thread dispatches a registered **ompt_callback_cancel** callback for each occurrence of a *cancellation* event in that thread. The callback occurs in the context of the encountering task. The callback has type signature **ompt_callback_cancel_t**. The callback receives **ompt_cancel_detected** as its *flags* argument.

Restrictions

- A **cancellation point** construct for which *construct-type-clause* is **taskgroup** must be closely nested inside a **task** construct, and the **cancellation point** region must be closely nested inside a **taskgroup** region. A **cancellation point** construct for which *construct-type-clause* is **sections** must be closely nested inside a **sections** or **section** construct. Otherwise, a **cancellation point** construct must be closely nested inside an OpenMP construct that matches the type specified in *construct-type-clause*.

Cross References

- *cancel-var* ICV, see Section 2.4.1 on page 47.
- **cancel** construct, see Section 2.21.1 on page 256.
- **omp_get_cancellation** routine, see Section 3.2.9 on page 340.
- **ompt_callback_cancel_t**, see Section 4.2.4.2.27 on page 477.

1 2.22 Data Environment

2 This section presents a directive and several clauses for controlling data environments.

3 2.22.1 Data-sharing Attribute Rules

4 This section describes how the data-sharing attributes of variables referenced in data environments
5 are determined. The following two cases are described separately:

- 6 • Section 2.22.1.1 on page 263 describes the data-sharing attribute rules for variables referenced in
7 a construct.
- 8 • Section 2.22.1.2 on page 266 describes the data-sharing attribute rules for variables referenced in
9 a region, but outside any construct.

10 2.22.1.1 Data-sharing Attribute Rules for Variables Referenced 11 in a Construct

12 The data-sharing attributes of variables that are referenced in a construct can be *predetermined*,
13 *explicitly determined*, or *implicitly determined*, according to the rules outlined in this section.

14 Specifying a variable on a **firstprivate**, **lastprivate**, **linear**, **reduction**, or
15 **copyprivate** clause of an enclosed construct causes an implicit reference to the variable in the
16 enclosing construct. Specifying a variable on a **map** clause of an enclosed construct may cause an
17 implicit reference to the variable in the enclosing construct. Such implicit references are also
18 subject to the data-sharing attribute rules outlined in this section.

19 Certain variables and objects have *predetermined* data-sharing attributes as follows:



- 20 • Variables appearing in **threadprivate** directives are threadprivate.
- 21 • Variables with automatic storage duration that are declared in a scope inside the construct are
22 private.
- 23 • Objects with dynamic storage duration are shared.
- 24 • Static data members are shared.
- 25 • The loop iteration variable(s) in the associated *for-loop(s)* of a **for**, **parallel for**,
26 **taskloop**, or **distribute** construct is (are) private.

- The loop iteration variable in the associated *for-loop* of a **simd** or **loop** construct with just one associated *for-loop* is linear with a *linear-step* that is the increment of the associated *for-loop*.
- The loop iteration variables in the associated *for-loops* of a **simd** or **loop** construct with multiple associated *for-loops* are lastprivate.
- Variables with static storage duration that are declared in a scope inside the construct are shared.
- If an array section with a named pointer is a list item in a **map** clause on the **target** construct and the named pointer is a scalar variable that does not appear in a **map** clause on the construct, the named pointer is firstprivate.

C / C++

Fortran

- Variables and common blocks appearing in **threadprivate** directives are threadprivate.
- The loop iteration variable(s) in the associated *do-loop(s)* of a **do**, **parallel do**, **taskloop**, or **distribute** construct is (are) private.
- The loop iteration variable in the associated *do-loop* of a **simd** or **loop** construct with just one associated *do-loop* is linear with a *linear-step* that is the increment of the associated *do-loop*.
- The loop iteration variables in the associated *do-loops* of a **simd** or **loop** construct with multiple associated *do-loops* are lastprivate.
- A loop iteration variable for a sequential loop in a **parallel** or task generating construct is private in the innermost such construct that encloses the loop.
- Implied-do indices and **forall** indices are private.
- Cray pointees have the same the data-sharing attribute as the storage with which their Cray pointers are associated.
- Assumed-size arrays are shared.
- An associate name preserves the association with the selector established at the **ASSOCIATE** or **SELECT TYPE** statement.

Fortran

Variables with predetermined data-sharing attributes may not be listed in data-sharing attribute clauses, except for the cases listed below. For these exceptions only, listing a predetermined variable in a data-sharing attribute clause is allowed and overrides the variable's predetermined data-sharing attributes.

C / C++

- 1 • The loop iteration variable(s) in the associated *for-loop(s)* of a **for**, **parallel for**,
2 **taskloop**, or **distribute** construct may be listed in a **private** or **lastprivate** clause.
- 3 • The loop iteration variable in the associated *for-loop* of a **simd** construct with just one
4 associated *for-loop* may be listed in a **private**, **lastprivate**, or **linear** clause with a
5 *linear-step* that is the increment of the associated *for-loop*.
- 6 • The loop iteration variables in the associated *for-loops* of a **simd** construct with multiple
7 associated *for-loops* may be listed in a **private** or **lastprivate** clause.
- 8 • The loop iteration variable(s) in the associated *for-loop(s)* of a **loop** construct may be listed in a
9 **private** clause.
- 10 • Variables with **const**-qualified type having no mutable member may be listed in a
11 **firstprivate** clause, even if they are static data members.

C / C++

Fortran

- 12 • The loop iteration variable(s) in the associated *do-loop(s)* of a **do**, **parallel do**, **taskloop**,
13 or **distribute** construct may be listed in a **private** or **lastprivate** clause.
- 14 • The loop iteration variable in the associated *do-loop* of a **simd** construct with just one
15 associated *do-loop* may be listed in a **private**, **lastprivate**, or **linear** clause with a
16 *linear-step* that is the increment of the associated loop.
- 17 • The loop iteration variables in the associated *do-loops* of a **simd** construct with multiple
18 associated *do-loops* may be listed in a **private** or **lastprivate** clause.
- 19 • The loop iteration variable(s) in the associated *do-loop(s)* of a **loop** construct may be listed in a
20 **private** clause.
- 21 • Variables used as loop iteration variables in sequential loops in a **parallel** or task generating
22 construct may be listed in data-sharing attribute clauses on the construct itself, and on enclosed
23 constructs, subject to other restrictions.
- 24 • Assumed-size arrays may be listed in a **shared** clause.

Fortran

25 Additional restrictions on the variables that may appear in individual clauses are described with
26 each clause in Section [2.22.4](#) on page [276](#).

27 Variables with *explicitly determined* data-sharing attributes are those that are referenced in a given
28 construct and are listed in a data-sharing attribute clause on the construct.

29 Variables with *implicitly determined* data-sharing attributes are those that are referenced in a given
30 construct, do not have predetermined data-sharing attributes, and are not listed in a data-sharing
31 attribute clause on the construct.

Rules for variables with *implicitly determined* data-sharing attributes are as follows:

- In a **parallel**, **teams**, or task generating construct, the data-sharing attributes of these variables are determined by the **default** clause, if present (see Section 2.22.4.1 on page 277).
- In a **parallel** construct, if no **default** clause is present, these variables are shared.
- For constructs other than task generating constructs, if no **default** clause is present, these variables reference the variables with the same names that exist in the enclosing context.
- In a **target** construct, variables that are not mapped after applying data-mapping attribute rules (see Section 2.22.7 on page 305) are firstprivate.

▼ C++ ▼

- In an orphaned task generating construct, if no **default** clause is present, formal arguments passed by reference are firstprivate.

▲ C++ ▲

▼ Fortran ▼

- In an orphaned task generating construct, if no **default** clause is present, dummy arguments are firstprivate.

▲ Fortran ▲

- In a task generating construct, if no **default** clause is present, a variable for which the data-sharing attribute is not determined by the rules above and that in the enclosing context is determined to be shared by all implicit tasks bound to the current team is shared.
- In a task generating construct, if no **default** clause is present, a variable for which the data-sharing attribute is not determined by the rules above is firstprivate.

Additional restrictions on the variables for which data-sharing attributes cannot be implicitly determined in a task generating construct are described in Section 2.22.4.4 on page 281.

2.22.1.2 Data-sharing Attribute Rules for Variables Referenced in a Region but not in a Construct

The data-sharing attributes of variables that are referenced in a region, but not in a construct, are determined as follows:

C / C++

- 1 • Variables with static storage duration that are declared in called routines in the region are shared.
- 2 • File-scope or namespace-scope variables referenced in called routines in the region are shared
- 3 unless they appear in a **threadprivate** directive.
- 4 • Objects with dynamic storage duration are shared.
- 5 • Static data members are shared unless they appear in a **threadprivate** directive.
- 6 • In C++, formal arguments of called routines in the region that are passed by reference have the
- 7 same data-sharing attributes as the associated actual arguments.
- 8 • Other variables declared in called routines in the region are private.

C / C++

Fortran

- 9 • Local variables declared in called routines in the region and that have the **save** attribute, or that
- 10 are data initialized, are shared unless they appear in a **threadprivate** directive.
- 11 • Variables belonging to common blocks, or accessed by host or use association, and referenced in
- 12 called routines in the region are shared unless they appear in a **threadprivate** directive.
- 13 • Dummy arguments of called routines in the region that have the **VALUE** attribute are private.
- 14 • Dummy arguments of called routines in the region that do not have the **VALUE** attribute are
- 15 private if the associated actual argument is not shared.
- 16 • Dummy arguments of called routines in the region that do not have the **VALUE** attribute are
- 17 shared if the actual argument is shared and it is a scalar variable, structure, an array that is not a
- 18 pointer or assumed-shape array, or a simply contiguous array section. Otherwise, the
- 19 data-sharing attribute of the dummy argument is implementation-defined if the associated actual
- 20 argument is shared.
- 21 • Cray pointees have the same data-sharing attribute as the storage with which their Cray pointers
- 22 are associated.
- 23 • Implied-do indices, **forall** indices, and other local variables declared in called routines in the
- 24 region are private.

Fortran

1 2.22.2 threadprivate Directive

2 Summary

3 The **threadprivate** directive specifies that variables are replicated, with each thread having its
4 own copy. The **threadprivate** directive is a declarative directive.

5 Syntax

▼ C / C++ ▼

6 The syntax of the **threadprivate** directive is as follows:

7 **#pragma omp threadprivate** (*list*) *new-line*

8 where *list* is a comma-separated list of file-scope, namespace-scope, or static block-scope variables
9 that do not have incomplete types.

▲ C / C++ ▲

▼ Fortran ▼

10 The syntax of the **threadprivate** directive is as follows:

11 **!\$omp threadprivate** (*list*)

12 where *list* is a comma-separated list of named variables and named common blocks. Common
13 block names must appear between slashes.

▲ Fortran ▲

14 Description

15 Each copy of a threadprivate variable is initialized once, in the manner specified by the program,
16 but at an unspecified point in the program prior to the first reference to that copy. The storage of all
17 copies of a threadprivate variable is freed according to how static variables are handled in the base
18 language, but at an unspecified point in the program.

19 A program in which a thread references another thread's copy of a threadprivate variable is
20 non-conforming.

21 The content of a threadprivate variable can change across a task scheduling point if the executing
22 thread switches to another task that modifies the variable. For more details on task scheduling, see
23 Section 1.3 on page 19 and Section 2.13 on page 133.

24 In **parallel** regions, references by the master thread will be to the copy of the variable in the
25 thread that encountered the **parallel** region.

1 During a sequential part references will be to the initial thread's copy of the variable. The values of
2 data in the initial thread's copy of a threadprivate variable are guaranteed to persist between any
3 two consecutive references to the variable in the program.

4 The values of data in the threadprivate variables of non-initial threads are guaranteed to persist
5 between two consecutive active **parallel** regions only if all of the following conditions hold:

- 6 • Neither **parallel** region is nested inside another explicit **parallel** region.
- 7 • The number of threads used to execute both **parallel** regions is the same.
- 8 • The thread affinity policies used to execute both **parallel** regions are the same.
- 9 • The value of the *dyn-var* internal control variable in the enclosing task region is *false* at entry to
10 both **parallel** regions.
- 11 • Neither the **omp_pause_resource** nor **omp_pause_resource_all** routine is called.

12 If these conditions all hold, and if a threadprivate variable is referenced in both regions, then
13 threads with the same thread number in their respective regions will reference the same copy of that
14 variable.

▼ C / C++ ▼

15 If the above conditions hold, the storage duration, lifetime, and value of a thread's copy of a
16 threadprivate variable that does not appear in any **copyin** clause on the second region will be
17 retained. Otherwise, the storage duration, lifetime, and value of a thread's copy of the variable in
18 the second region is unspecified.

19 If the value of a variable referenced in an explicit initializer of a threadprivate variable is modified
20 prior to the first reference to any instance of the threadprivate variable, then the behavior is
21 unspecified.

▲ C / C++ ▲

▼ C++ ▼

22 The order in which any constructors for different threadprivate variables of class type are called is
23 unspecified. The order in which any destructors for different threadprivate variables of class type
24 are called is unspecified.

▲ C++ ▲

Fortran

1 A variable is affected by a **copyin** clause if the variable appears in the **copyin** clause or it is in a
2 common block that appears in the **copyin** clause.

3 If the above conditions hold, the definition, association, or allocation status of a thread's copy of a
4 threadprivate variable or a variable in a threadprivate common block, that is not affected by any
5 **copyin** clause that appears on the second region, will be retained. Otherwise, the definition and
6 association status of a thread's copy of the variable in the second region are undefined, and the
7 allocation status of an allocatable variable will be implementation defined.

8 If a threadprivate variable or a variable in a threadprivate common block is not affected by any
9 **copyin** clause that appears on the first **parallel** region in which it is referenced, the thread's
10 copy of the variable inherits the declared type parameter and the default parameter values from the
11 original variable. The variable or any subobject of the variable is initially defined or undefined
12 according to the following rules:

- 13 • If it has the **ALLOCATABLE** attribute, each copy created will have an initial allocation status of
14 unallocated.
- 15 • If it has the **POINTER** attribute:
 - 16 – if it has an initial association status of disassociated, either through explicit initialization or
17 default initialization, each copy created will have an association status of disassociated;
 - 18 – otherwise, each copy created will have an association status of undefined.
- 19 • If it does not have either the **POINTER** or the **ALLOCATABLE** attribute:
 - 20 – if it is initially defined, either through explicit initialization or default initialization, each copy
21 created is so defined;
 - 22 – otherwise, each copy created is undefined.

Fortran

Restrictions

23 The restrictions to the **threadprivate** directive are as follows:

- 24 • A threadprivate variable must not appear in any clause except the **copyin**, **copyprivate**,
25 **schedule**, **num_threads**, **thread_limit**, and **if** clauses.
- 26 • A program in which an untied task accesses threadprivate storage is non-conforming.
- 27

C / C++

- 1 • A variable that is part of another variable (as an array or structure element) cannot appear in a
2 **threadprivate** clause unless it is a static data member of a C++ class.
- 3 • A **threadprivate** directive for file-scope variables must appear outside any definition or
4 declaration, and must lexically precede all references to any of the variables in its list.
- 5 • A **threadprivate** directive for namespace-scope variables must appear outside any
6 definition or declaration other than the namespace definition itself, and must lexically precede all
7 references to any of the variables in its list.
- 8 • Each variable in the list of a **threadprivate** directive at file, namespace, or class scope must
9 refer to a variable declaration at file, namespace, or class scope that lexically precedes the
10 directive.
- 11 • A **threadprivate** directive for static block-scope variables must appear in the scope of the
12 variable and not in a nested scope. The directive must lexically precede all references to any of
13 the variables in its list.
- 14 • Each variable in the list of a **threadprivate** directive in block scope must refer to a variable
15 declaration in the same scope that lexically precedes the directive. The variable declaration must
16 use the static storage-class specifier.
- 17 • If a variable is specified in a **threadprivate** directive in one translation unit, it must be
18 specified in a **threadprivate** directive in every translation unit in which it is declared.
- 19 • The address of a threadprivate variable is not an address constant.

C / C++

C++

- 20 • A **threadprivate** directive for static class member variables must appear in the class
21 definition, in the same scope in which the member variables are declared, and must lexically
22 precede all references to any of the variables in its list.
- 23 • A threadprivate variable must not have an incomplete type or a reference type.
- 24 • A threadprivate variable with class type must have:
 - 25 – an accessible, unambiguous default constructor in case of default initialization without a given
26 initializer;
 - 27 – an accessible, unambiguous constructor accepting the given argument in case of direct
28 initialization;
 - 29 – an accessible, unambiguous copy constructor in case of copy initialization with an explicit
30 initializer

C++

Fortran

- 1 • A variable that is part of another variable (as an array, structure element or type parameter
2 inquiry) cannot appear in a **threadprivate** clause.
- 3 • The **threadprivate** directive must appear in the declaration section of a scoping unit in
4 which the common block or variable is declared. Although variables in common blocks can be
5 accessed by use association or host association, common block names cannot. This means that a
6 common block name specified in a **threadprivate** directive must be declared to be a
7 common block in the same scoping unit in which the **threadprivate** directive appears.
- 8 • If a **threadprivate** directive specifying a common block name appears in one program unit,
9 then such a directive must also appear in every other program unit that contains a **COMMON**
10 statement specifying the same name. It must appear after the last such **COMMON** statement in the
11 program unit.
- 12 • If a **threadprivate** variable or a **threadprivate** common block is declared with the **BIND** attribute,
13 the corresponding C entities must also be specified in a **threadprivate** directive in the C
14 program.
- 15 • A blank common block cannot appear in a **threadprivate** directive.
- 16 • A variable can only appear in a **threadprivate** directive in the scope in which it is declared.
17 It must not be an element of a common block or appear in an **EQUIVALENCE** statement.
- 18 • A variable that appears in a **threadprivate** directive must be declared in the scope of a
19 module or have the **SAVE** attribute, either explicitly or implicitly.

Fortran

Cross References

- 20 • *dyn-var* ICV, see Section 2.4 on page 47.
- 21 • Number of threads used to execute a **parallel** region, see Section 2.9.1 on page 77.
- 22 • **copyin** clause, see Section 2.22.6.1 on page 301.
- 23

1 2.22.3 List Item Privatization

2 For any construct, a list item that appears in a data-sharing attribute clause, including a reduction
3 clause, may be privatized. Each task that references a privatized list item in any statement in the
4 construct receives at least one new list item if the construct has one or more associated loops, and
5 otherwise each such task receives one new list item. Each SIMD lane used in a **simd** construct that
6 references a privatized list item in any statement in the construct receives at least one new list item.
7 Language-specific attributes for new list items are derived from the corresponding original list item.
8 Inside the construct, all references to the original list item are replaced by references to a new list
9 item received by the task or SIMD lane.

10 If the construct has one or more associated loops, within the same logical iteration of the loop(s)
11 the same new list item replaces all references to the original list item. For any two logical iterations,
12 if the references to the original list item are replaced by the same list item then the logical iterations
13 must execute in some sequential order.

14 In the rest of the region, it is unspecified whether references are to a new list item or the original list
15 item. Therefore, if an attempt is made to reference the original item, its value after the region is also
16 unspecified. If a task or a SIMD lane does not reference a privatized list item, it is unspecified
17 whether the task or SIMD lane receives a new list item.

18 The value and/or allocation status of the original list item will change only:

- 19
- 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
- 21

22

▼ Fortran ▲

- if accessed and modified via construct association.

▲ Fortran ▼

23

▼ C++ ▼

24 If the construct is contained in a member function, it is unspecified anywhere in the region if
accesses through the implicit **this** pointer refer to the new list item or the original list item.

▲ C++ ▲

C / C++

1 A new list item of the same type, with automatic storage duration, is allocated for the construct.
2 The storage and thus lifetime of these list items lasts until the block in which they are created exits.
3 The size and alignment of the new list item are determined by the type of the variable. This
4 allocation occurs once for each task generated by the construct and once for each SIMD lane used
5 by the construct.

6 The new list item is initialized, or has an undefined initial value, as if it had been locally declared
7 without an initializer.

C / C++

C++

8 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
9 purposes of this clause.

10 The order in which any default constructors for different private variables of class type are called is
11 unspecified. The order in which any destructors for different private variables of class type are
12 called is unspecified.

C++

Fortran

13 If any statement of the construct references a list item, a new list item of the same type and type
14 parameters is allocated. This allocation occurs once for each task generated by the construct and
15 once for each SIMD lane used by the construct. The initial value of the new list item is undefined.
16 The initial status of a private pointer is undefined.

17 For a list item or the subobject of a list item with the **ALLOCATABLE** attribute:

- 18 • if the allocation status is unallocated, the new list item or the subobject of the new list item will
19 have an initial allocation status of unallocated.
- 20 • if the allocation status is allocated, the new list item or the subobject of the new list item will
21 have an initial allocation status of allocated.
- 22 • If the new list item or the subobject of the new list item is an array, its bounds will be the same as
23 those of the original list item or the subobject of the original list item.

24 A privatized list item may be storage-associated with other variables when the data-sharing
25 attribute clause is encountered. Storage association may exist because of constructs such as
26 **EQUIVALENCE** or **COMMON**. If A is a variable that is privatized by a construct and B is a variable
27 that is storage-associated with A , then:

- 28 • The contents, allocation, and association status of B are undefined on entry to the region.
- 29 • Any definition of A , or of its allocation or association status, causes the contents, allocation, and
30 association status of B to become undefined.

- 1 • Any definition of *B*, or of its allocation or association status, causes the contents, allocation, and
2 association status of *A* to become undefined.

3 A privatized list item clause may be a selector of an **ASSOCIATE** or **SELECT TYPE** construct. If
4 the construct association is established prior to a **parallel** region, the association between the
5 associate name and the original list item will be retained in the region.

6 Finalization of a list item of a finalizable type or subobjects of a list item of a finalizable type occurs
7 at the end of the region. The order in which any final subroutines for different variables of a
8 finalizable type are called is unspecified.



Fortran

9 **Restrictions**

10 The following restrictions apply to any list item that is privatized unless otherwise stated for a given
11 data-sharing attribute clause:



C

- 12 • A variable that is part of another variable (as an array or structure element) cannot be privatized.



C



C++

- 13 • A variable that is part of another variable (as an array or structure element) cannot be privatized
14 except if the data-sharing attribute clause is associated with a construct within a class non-static
15 member function and the variable is an accessible data member of the object for which the
16 non-static member function is invoked.

- 17 • A variable of class type (or array thereof) that is privatized requires an accessible, unambiguous
18 default constructor for the class type.



C++



C / C++

- 19 • A variable that is privatized must not have a **const**-qualified type unless it is of class type with
20 a **mutable** member. This restriction does not apply to the **firstprivate** clause.
- 21 • A variable that is privatized must not have an incomplete type or be a reference to an incomplete
22 type.



C / C++

Fortran

- 1 • A variable that is part of another variable (as an array or structure element) cannot be privatized.
- 2 • A variable that is privatized must either be definable, or an allocatable variable. This restriction
- 3 does not apply to the **firstprivate** clause.
- 4 • Variables that appear in namelist statements, in variable format expressions, and in expressions
- 5 for statement function definitions, may not be privatized.
- 6 • Pointers with the **INTENT (IN)** attribute may not appear be privatized. This restriction does not
- 7 apply to the **firstprivate** clause.
- 8 • Assumed-size arrays may not be privatized in a **target**, **teams**, or **distribute** construct.

Fortran

9 2.22.4 Data-Sharing Attribute Clauses

10 Several constructs accept clauses that allow a user to control the data-sharing attributes of variables
11 referenced in the construct. Data-sharing attribute clauses apply only to variables for which the
12 names are visible in the construct on which the clause appears.

13 Not all of the clauses listed in this section are valid on all directives. The set of clauses that is valid
14 on a particular directive is described with the directive.

15 Most of the clauses accept a comma-separated list of list items (see Section 2.1 on page 36). All list
16 items appearing in a clause must be visible, according to the scoping rules of the base language.
17 With the exception of the **default** clause, clauses may be repeated as needed. A list item that
18 specifies a given variable may not appear in more than one clause on the same directive, except that
19 a variable may be specified in both **firstprivate** and **lastprivate** clauses.

20 The reduction data-sharing attribute clauses are explained in Section 2.22.5.

C++

21 If a variable referenced in a data-sharing attribute clause has a type derived from a template, and
22 there are no other references to that variable in the program, then any behavior related to that
23 variable is unspecified.

C++

Fortran

1 When a named common block appears in a **private**, **firstprivate**, **lastprivate**, or
2 **shared** clause of a directive, none of its members may be declared in another data-sharing
3 attribute clause in that directive. When individual members of a common block appear in a
4 **private**, **firstprivate**, **lastprivate**, **reduction**, or **linear** clause of a directive,
5 the storage of the specified variables is no longer Fortran associated with the storage of the common
6 block itself.

Fortran

7 2.22.4.1 default Clause

8 Summary

9 The **default** clause explicitly determines the data-sharing attributes of variables that are
10 referenced in a **parallel**, **teams**, or task generating construct and would otherwise be implicitly
11 determined (see Section 2.22.1.1 on page 263).

12 Syntax

C / C++

13 The syntax of the **default** clause is as follows:

```
14 | default(shared | none)
```

C / C++

Fortran

15 The syntax of the **default** clause is as follows:

```
16 | default(private | firstprivate | shared | none)
```

Fortran

1 **Description**

2 The **default (shared)** clause causes all variables referenced in the construct that have
3 implicitly determined data-sharing attributes to be shared.



4 The **default (firstprivate)** clause causes all variables in the construct that have implicitly
5 determined data-sharing attributes to be firstprivate.

6 The **default (private)** clause causes all variables referenced in the construct that have
7 implicitly determined data-sharing attributes to be private.



8 The **default (none)** clause requires that each variable that is referenced in the construct, and
9 that does not have a predetermined data-sharing attribute, must have its data-sharing attribute
10 explicitly determined by being listed in a data-sharing attribute clause.

11 **Restrictions**

12 The restrictions to the **default** clause are as follows:

- 13 • Only a single **default** clause may be specified on a **parallel**, **task**, **taskloop** or
14 **teams** directive.

15 **2.22.4.2 shared Clause**

16 **Summary**

17 The **shared** clause declares one or more list items to be shared by tasks generated by a
18 **parallel**, **teams**, or task generating construct.

19 **Syntax**

20 The syntax of the **shared** clause is as follows:

21 **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.

The programmer must 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**, **teams**, or task generating construct if it is associated with a target or a subobject of a target that is in a **private**, **firstprivate**, **lastprivate**, or **reduction** clause in the construct.

Note – Passing a shared variable to a procedure may result in the use of temporary storage in place of the actual argument when the corresponding dummy argument does not have the **VALUE** or **CONTIGUOUS** attribute and its data-sharing attribute is implementation-defined as per the rules in Section 2.22.1.2 on page 266. These conditions effectively result in references to, and definitions of, the temporary storage during the procedure reference. Furthermore, the value of the shared variable is copied into the intervening temporary storage before the procedure reference when the dummy argument does not have the **INTENT (OUT)** attribute, and back out of the temporary storage into the shared variable when the dummy argument does not have the **INTENT (IN)** attribute. Any references to (or definitions of) the shared storage that is associated with the dummy argument by any other task must be synchronized with the procedure reference to avoid possible race conditions.

Fortran

Restrictions

The restrictions for the **shared** clause are as follows:

C

- A variable that is part of another variable (as an array or structure element) cannot appear in a **shared** clause.

C

C++

- A variable that is part of another variable (as an array or structure element) cannot appear in a **shared** clause except if the **shared** clause is associated with a construct within a class non-static member function and the variable is an accessible data member of the object for which the non-static member function is invoked.

C++

- 1 • A variable that is part of another variable (as an array, structure element or type parameter
2 inquiry) cannot appear in a shared clause.

3 2.22.4.3 private Clause

4 Summary

5 The **private** clause declares one or more list items to be private to a task or to a SIMD lane.

6 Syntax

7 The syntax of the private clause is as follows:

8 **private** (*list*)

9 Description

10 The **private** clause specifies that its list items are to be privatized according to Section 2.22.3 on
11 page 273. Each task or SIMD lane that references a list item in the construct receives only one new
12 list item, unless the construct has one or more associated loops and the **order (concurrent)**
13 clause is also present.

14 List items that appear in a **private**, **firstprivate**, or **reduction** clause in a **parallel**
15 construct may also appear in a **private** clause in an enclosed **parallel**, worksharing, **loop**,
16 **task**, **taskloop**, **simd**, or **target** construct.

17 List items that appear in a **private** or **firstprivate** clause in a **task** or **taskloop**
18 construct may also appear in a **private** clause in an enclosed **parallel**, **loop**, **task**,
19 **taskloop**, **simd**, or **target** construct.

20 List items that appear in a **private**, **firstprivate**, **lastprivate**, or **reduction** clause
21 in a worksharing construct may also appear in a **private** clause in an enclosed **parallel**,
22 **loop**, **task**, **simd**, or **target** construct.

23 List items that appear in a **private** clause on a **loop** construct may also appear in a **private**
24 clause in an enclosed **loop**, **parallel**, or **simd** construct.

25 Restrictions

26 The restrictions to the **private** clause are as specified in Section 2.22.3.

Cross References

- List Item Privatization, see Section [2.22.3](#) on page [273](#).

2.22.4.4 `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 [2.22.4.3](#) on page [280](#), 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`, `taskloop`, `target`, 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 unless otherwise specified. 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 unless otherwise specified.

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++

1 For variables of non-array type, the initialization occurs by copy assignment. For an array of
2 elements of non-array type, each element is initialized as if by assignment from an element of the
3 original array to the corresponding element of the new array.

C / C++

C++

4 For each variable of class type:

- 5 • If the **firstprivate** clause is not on a **target** construct then a copy constructor is invoked
6 to perform the initialization;
- 7 • If the **firstprivate** clause is on a **target** construct then it is unspecified how many copy
8 constructors, if any, are invoked.

9 If copy constructors are called, the order in which copy constructors for different variables of class
10 type are called is unspecified.

C++

Fortran

11 If the original list item does not have the **POINTER** attribute, initialization of the new list items
12 occurs as if by intrinsic assignment unless the list item has a type bound procedure as a defined
13 assignment. If the original list item that does not have the **POINTER** attribute has the allocation
14 status of unallocated, the new list items will have the same status.

15 If the original list item has the **POINTER** attribute, the new list items receive the same association
16 status of the original list item as if by pointer assignment.

Fortran

Restrictions

17 The restrictions to the **firstprivate** clause are as follows:

- 18 • A list item that is private within a **parallel** region must not appear in a **firstprivate**
19 clause on a worksharing construct if any of the worksharing regions arising from the worksharing
20 construct ever bind to any of the **parallel** regions arising from the **parallel** construct.
- 21 • A list item that is private within a **teams** region must not appear in a **firstprivate** clause
22 on a **distribute** construct if any of the **distribute** regions arising from the
23 **distribute** construct ever bind to any of the **teams** regions arising from the **teams**
24 construct.
25

- 1 • A list item that appears in a **reduction** clause of a **parallel** construct must not appear in a
2 **firstprivate** clause on a worksharing, **task**, or **taskloop** construct if any of the
3 worksharing or task regions arising from the worksharing, **task**, or **taskloop** construct ever
4 bind to any of the **parallel** regions arising from the **parallel** construct.
- 5 • A list item that appears in a **reduction** clause of a **teams** construct must not appear in a
6 **firstprivate** clause on a **distribute** construct if any of the **distribute** regions
7 arising from the **distribute** construct ever bind to any of the **teams** regions arising from the
8 **teams** construct.
- 9 • A list item that appears in a **reduction** clause of a worksharing construct must not appear in a
10 **firstprivate** clause in a **task** construct encountered during execution of any of the
11 worksharing regions arising from the worksharing construct.

C++

- 12 • A variable of class type (or array thereof) that appears in a **firstprivate** clause requires an
13 accessible, unambiguous copy constructor for the class type.

C++

C / C++

- 14 • A variable that appears in a **firstprivate** clause must not have an incomplete C/C++ type or
15 be a reference to an incomplete type.
- 16 • If a list item in a **firstprivate** clause on a worksharing construct has a reference type then it
17 must bind to the same object for all threads of the team.

C / C++

Fortran

- 18 • Variables that appear in namelist statements, in variable format expressions, or in expressions for
19 statement function definitions, may not appear in a **firstprivate** clause.
- 20 • Assumed-size arrays may not appear in the **firstprivate** clause in a **target**, **teams**, or
21 **distribute** construct.
- 22 • If the list item is a polymorphic variable with the **ALLOCATABLE** attribute, the behavior is
23 unspecified.

Fortran

24 2.22.4.5 lastprivate Clause

25 Summary

26 The **lastprivate** clause declares one or more list items to be private to an implicit task or to a
27 SIMD lane, and causes the corresponding original list item to be updated after the end of the region.

Syntax

The syntax of the **lastprivate** clause is as follows:

```
lastprivate ([ lastprivate-modifier : ] list)
```

where *lastprivate-modifier* is:

```
conditional
```

Description

The **lastprivate** clause provides a superset of the functionality provided by the **private** clause.

A list item that appears in a **lastprivate** clause is subject to the **private** clause semantics described in Section 2.22.4.3 on page 280. In addition, when a **lastprivate** clause without the **conditional** modifier appears on a directive, the value of each new list item from the sequentially last iteration of the associated loops, or the lexically last **section** construct, is assigned to the original list item. When the **conditional** modifier appears on the clause, if an assignment to a list item is encountered in the construct then the original list item is assigned the value that is assigned to the new list item in the sequentially last iteration or lexically last section in which such an assignment is encountered.

C / C++

For an array of elements of non-array type, each element is assigned to the corresponding element of the original array.

C / C++

Fortran

If the original list item does not have the **POINTER** attribute, its update occurs as if by intrinsic assignment unless it has a type bound procedure as a defined assignment.

If the original list item has the **POINTER** attribute, its update occurs as if by pointer assignment.

Fortran

When the **conditional** modifier does not appear on the **lastprivate** clause, list items that are not assigned a value by the sequentially last iteration of the loops, or by the lexically last **section** construct, have unspecified values after the construct. Unassigned subcomponents also have unspecified values after the construct.

If the **lastprivate** clause is used on a construct to which neither the **nowait** nor the **nogroup** clauses are applied, the original list item becomes defined at the end of the construct. To avoid race conditions, concurrent reads or updates of the original list item must be synchronized with the update of the original list item that occurs as a result of the **lastprivate** clause.

1 Otherwise, If the **lastprivate** clause is used on a construct to which the **nowait** or the
2 **nogroup** clauses are applied, accesses to the original list item may create a data race. To avoid
3 this, if an assignment to the original list item occurs then synchronization must be inserted to
4 ensure that the assignment completes and the original list item is flushed to memory.

5 If a list item appears in both **firstprivate** and **lastprivate** clauses, the update required
6 for **lastprivate** occurs after all initializations for **firstprivate**.

7 Restrictions

8 The restrictions to the **lastprivate** clause are as follows:

- 9 • A list item that is private within a **parallel** region, or that appears in the **reduction** clause
10 of a **parallel** construct, must not appear in a **lastprivate** clause on a worksharing
11 construct if any of the corresponding worksharing regions ever binds to any of the corresponding
12 **parallel** regions.
- 13 • If a list item that appears in a **lastprivate** clause with the **conditional** modifier is
14 modified in the region by an assignment outside the construct or not to the list item then the value
15 assigned to the original list item is unspecified.
- 16 • A list item that appears in a **lastprivate** clause with the **conditional** modifier must be a
17 scalar variable.

▼ C++ ▼

- 18 • A variable of class type (or array thereof) that appears in a **lastprivate** clause requires an
19 accessible, unambiguous default constructor for the class type, unless the list item is also
20 specified in a **firstprivate** clause.
- 21 • A variable of class type (or array thereof) that appears in a **lastprivate** clause requires an
22 accessible, unambiguous copy assignment operator for the class type. The order in which copy
23 assignment operators for different variables of class type are called is unspecified.

▲ C++ ▲

▼ C / C++ ▼

- 24 • A variable that appears in a **lastprivate** clause must not have a **const**-qualified type unless
25 it is of class type with a **mutable** member.
- 26 • A variable that appears in a **lastprivate** clause must not have an incomplete C/C++ type or
27 be a reference to an incomplete type.
- 28 • If a list item in a **lastprivate** clause on a worksharing construct has a reference type then it
29 must bind to the same object for all threads of the team.

▲ C / C++ ▲

Fortran

- 1 • A variable that appears in a **lastprivate** clause must be definable.
- 2 • If the original list item has the **ALLOCATABLE** attribute, the corresponding list item whose value
- 3 is assigned to the original list item must have an allocation status of allocated upon exit from the
- 4 sequentially last iteration or lexically last **section** construct.
- 5 • Variables that appear in namelist statements, in variable format expressions, or in expressions for
- 6 statement function definitions, may not appear in a **lastprivate** clause.
- 7 • If the list item is a polymorphic variable with the **ALLOCATABLE** attribute, the behavior is
- 8 unspecified.

Fortran

9 2.22.4.6 linear Clause

10 Summary

11 The **linear** clause declares one or more list items to be private to a SIMD lane and to have a
12 linear relationship with respect to the iteration space of a loop.

13 Syntax

14 The syntax of the **linear** clause is as follows:

```
15 linear (linear-list [ : linear-step ])
```

16 where *linear-list* is one of the following

```
17 list  
18 modifier (list)
```

19 where *modifier* is one of the following:

```
20 val
```

C++

1 The syntax of the **linear** clause is as follows:

2 **linear** (*linear-list* [: *linear-step*])

3 where *linear-list* is one of the following

4 *list*
5 *modifier* (*list*)

6 where *modifier* is one of the following:

7 **ref**
8 **val**
9 **uval**

C++

Fortran

10 The syntax of the **linear** clause is as follows:

11 **linear** (*linear-list* [: *linear-step*])

12 where *linear-list* is one of the following

13 *list*
14 *modifier* (*list*)

15 where *modifier* is one of the following:

16 **ref**
17 **val**
18 **uval**

Fortran

Description





The **linear** clause provides a superset of the functionality provided by the **private** clause. A list item that appears in a **linear** clause is subject to the **private** clause semantics described in Section 2.22.4.3 on page 280 except as noted. If *linear-step* is not specified, it is assumed to be 1.

When a **linear** clause is specified on a construct, the value of the new list item on each iteration of the associated loop(s) corresponds to the value of the original list item before entering the construct plus the logical number of the iteration times *linear-step*. The value corresponding to the sequentially last iteration of the associated loop(s) is assigned to the original list item.

When a **linear** clause is specified on a declarative directive, all list items must be formal parameters (or, in Fortran, dummy arguments) of a function that will be invoked concurrently on each SIMD lane. If no *modifier* is specified or the **val** or **uval** modifier is specified, the value of each list item on each lane corresponds to the value of the list item upon entry to the function plus the logical number of the lane times *linear-step*. If the **uval** modifier is specified, each invocation uses the same storage location for each SIMD lane; this storage location is updated with the final value of the logically last lane. If the **ref** modifier is specified, the storage location of each list item on each lane corresponds to an array at the storage location upon entry to the function indexed by the logical number of the lane times *linear-step*.

Restrictions

- The *linear-step* expression must be invariant during the execution of the region corresponding to 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.
- Only a loop iteration variable of a loop that is associated with the construct may appear as a *list-item* in a **linear** clause if a **reduction** clause with the **inscan** modifier also appears on the construct.

- ▼  C  ▼
- A *list-item* that appears in a **linear** clause must be of integral or pointer type.
- ▲  C  ▲

C++

- 1 • A *list-item* that appears in a **linear** clause without the **ref** modifier must be of integral or
2 pointer type, or must be a reference to an integral or pointer type.
- 3 • The **ref** or **uval** modifier can only be used if the *list-item* is of a reference type.
- 4 • If a list item in a **linear** clause on a worksharing construct has a reference type then it must
5 bind to the same object for all threads of the team.
- 6 • If the list item is of a reference type and the **ref** modifier is not specified and if any write to the
7 list item occurs before any read of the list item then the result is unspecified.

C++

Fortran

- 8 • A *list-item* that appears in a **linear** clause without the **ref** modifier must be of type
9 **integer**.
- 10 • The **ref** or **uval** modifier can only be used if the *list-item* is a dummy argument without the
11 **VALUE** attribute.
- 12 • Variables that have the **POINTER** attribute and Cray pointers may not appear in a linear clause.
- 13 • If the list item has the **ALLOCATABLE** attribute and the **ref** modifier is not specified, the
14 allocation status of the list item in the sequentially last iteration must be allocated upon exit from
15 that iteration.
- 16 • If the **ref** modifier is specified, variables with the **ALLOCATABLE** attribute, assumed-shape
17 arrays and polymorphic variables may not appear in the **linear** clause.
- 18 • If the list item is a dummy argument without the **VALUE** attribute and the **ref** modifier is not
19 specified and if any write to the list item occurs before any read of the list item then the result is
20 unspecified.
- 21 • A common block name cannot appear in a **linear** clause.

Fortran

1 2.22.5 Reduction Clauses

2 The reduction clauses are data-sharing attribute clauses that can be used to perform some forms of
3 recurrence calculations (involving mathematically associative and commutative operators) in
4 parallel.


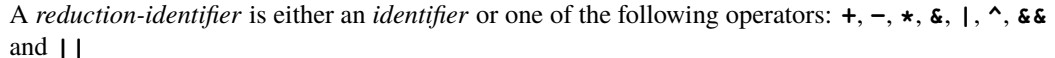
5 Reduction clauses include reduction scoping clauses and reduction participating clauses. Reduction
6 scoping clauses define the region in which a reduction is computed. Reduction participating clauses
7 define the participants in the reduction.


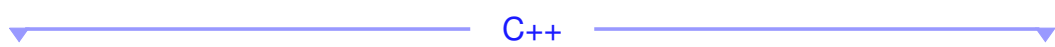
8 Reduction clauses specify a *reduction-identifier* and one or more list items.

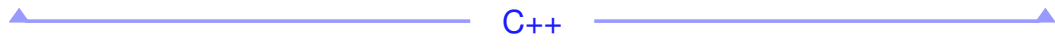

9 2.22.5.1 Properties Common To All Reduction Clauses



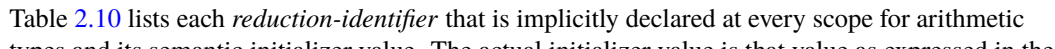
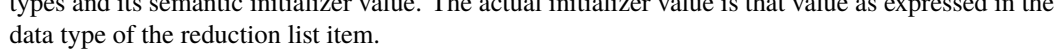
10 Syntax

11 The syntax of a *reduction-identifier* is defined as follows:

12  **C** 
A *reduction-identifier* is either an *identifier* or one of the following operators: +, -, *, &, |, ^, &&
13 and ||

14  **C++** 
A *reduction-identifier* is either an *id-expression* or one of the following operators: +, -, *, &, |, ^,
15 && and ||

16  **Fortran** 
A *reduction-identifier* is either a base language identifier, or a user-defined operator, or one of the
17 following operators: +, -, *, **.and.**, **.or.**, **.eqv.**, **.neqv.**, or one of the following intrinsic
18 procedure names: **max**, **min**, **iand**, **ior**, **ieor**.

19  **Fortran** 
20  **C / C++** 

21 Table 2.10 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.

TABLE 2.10: Implicitly Declared C/C++ *reduction-identifiers*

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 &= 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 && 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 > omp_out ? omp_in : omp_out</code>
min	<code>omp_priv = Largest representable number in the reduction list item type</code>	<code>omp_out = omp_in < omp_out ? omp_in : omp_out</code>

C / C++

Fortran

1
2
3

Table 2.11 lists each *reduction-identifier* that is implicitly declared for numeric and logical types and its semantic initializer value. The actual initializer value is that value as expressed in the data type of the reduction list item.

TABLE 2.11: Implicitly Declared Fortran *reduction-identifiers*

Identifier	Initializer	Combiner
+	<code>omp_priv = 0</code>	<code>omp_out = omp_in + omp_out</code>
*	<code>omp_priv = 1</code>	<code>omp_out = omp_in * omp_out</code>
-	<code>omp_priv = 0</code>	<code>omp_out = omp_in + omp_out</code>

table continued on next page

table continued from previous page

Identifier	Initializer	Combiner
<code>.and.</code>	<code>omp_priv = .true.</code>	<code>omp_out = omp_in .and. omp_out</code>
<code>.or.</code>	<code>omp_priv = .false.</code>	<code>omp_out = omp_in .or. omp_out</code>
<code>.eqv.</code>	<code>omp_priv = .true.</code>	<code>omp_out = omp_in .eqv. omp_out</code>
<code>.neqv.</code>	<code>omp_priv = .false.</code>	<code>omp_out = omp_in .neqv. omp_out</code>
<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

1 In the above tables, `omp_in` and `omp_out` correspond to two identifiers that refer to storage of the
2 type of the list item. `omp_out` holds the final value of the combiner operation.

3 Any *reduction-identifier* that is defined with the **declare reduction** directive is also valid. In
4 that case, the initializer and combiner of the *reduction-identifier* are specified by the
5 *initializer-clause* and the *combiner* in the **declare reduction** directive.

Description

6 A reduction clause specifies a *reduction-identifier* and one or more list items.

7 The *reduction-identifier* specified in a reduction clause must match a previously declared
8 *reduction-identifier* of the same name and type for each of the list items. This match is done by
9 means of a name lookup in the base language.

10 The list items that appear in a **reduction** clause may include array sections.

1 If the type is a derived class, then any *reduction-identifier* that matches its base classes is also a
 2 match, if there is no specific match for the type.

3 If the *reduction-identifier* is not an *id-expression*, then it is implicitly converted to one by
 4 prepending the keyword operator (for example, `+` becomes *operator+*).

5 If the *reduction-identifier* is qualified then a qualified name lookup is used to find the declaration.

6 If the *reduction-identifier* is unqualified then an *argument-dependent name lookup* must be
 7 performed using the type of each list item.

8 If the list item is an array or array section, it will be treated as if a reduction clause would be applied
 9 to each separate element of the array section.

10 Any copies associated with the reduction are initialized with the initializer value of the
 11 *reduction-identifier*.

12 Any copies are combined using the combiner associated with the *reduction-identifier*.

13 Restrictions

14 The restrictions common to reduction clauses are as follows:

- 15 • Any number of reduction clauses can be specified on the directive, but a list item (or any array
 16 element in an array section) can appear only once in reduction clauses for that directive.
- 17 • For a *reduction-identifier* declared with the **declare reduction** construct, the directive must
 18 appear before its use in a reduction clause.

- 19 • If a list item is an array section, its base expression must be a base language identifier.

- 20 • If a list item is an array section, it must specify contiguous storage and it cannot be a zero-length
 21 array section.
- 22 • If a list item is an array section, accesses to the elements of the array outside the specified array
 23 section result in unspecified behavior.

C

- 1 • A variable that is part of another variable, with the exception of array elements, cannot appear in
2 a reduction clause.

C

- 3 • A variable that is part of another variable, with the exception of array elements, cannot appear in
4 a reduction clause except if the reduction clause is associated with a construct within a class
5 non-static member function and the variable is an accessible data member of the object for which
6 the non-static member function is invoked.

C++

- 7 • The type of a list item that appears in a reduction clause must be valid for the
8 *reduction-identifier*. For a **max** or **min** reduction in C, the type of the list item must be an
9 allowed arithmetic data type: **char**, **int**, **float**, **double**, or **_Bool**, possibly modified with
10 **long**, **short**, **signed**, or **unsigned**. For a **max** or **min** reduction in C++, the type of the
11 list item must be an allowed arithmetic data type: **char**, **wchar_t**, **int**, **float**, **double**, or
12 **bool**, possibly modified with **long**, **short**, **signed**, or **unsigned**.
- 13 • A list item that appears in a reduction clause must not be **const**-qualified.
- 14 • The *reduction-identifier* for any list item must be unambiguous and accessible.

C / C++

Fortran

- 15 • A variable that is part of another variable, with the exception of array elements, cannot appear in
16 a reduction clause.
- 17 • A type parameter inquiry cannot appear in a reduction clause.
- 18 • The type, type parameters and rank of a list item that appears in a reduction clause must be valid
19 for the *combiner* and *initializer*.
- 20 • A list item that appears in a reduction clause must be definable.
- 21 • A procedure pointer may not appear in a reduction clause.
- 22 • A pointer with the **INTENT (IN)** attribute may not appear in the reduction clause.
- 23 • An original list item with the **POINTER** attribute or any pointer component of an original list
24 item that is referenced in the *combiner* must be associated at entry to the construct that contains
25 the reduction clause. Additionally, the list item or the pointer component of the list item must not
26 be deallocated, allocated, or pointer assigned within the region.

- 1 • An original list item with the **ALLOCATABLE** attribute or any allocatable component of an
2 original list item that corresponds to the special variable identifier in the *combiner* or the
3 *initializer* must be in the allocated state at entry to the construct that contains the reduction
4 clause. Additionally, the list item or the allocatable component of the list item must be neither
5 deallocated nor allocated, explicitly or implicitly, within the region.
- 6 • If the *reduction-identifier* is defined in a **declare reduction** directive, the
7 **declare reduction** directive must be in the same subprogram, or accessible by host or use
8 association.
- 9 • If the *reduction-identifier* is a user-defined operator, the same explicit interface for that operator
10 must be accessible as at the **declare reduction** directive.
- 11 • If the *reduction-identifier* is defined in a **declare reduction** directive, any subroutine or
12 function referenced in the initializer clause or combiner expression must be an intrinsic function,
13 or must have an explicit interface where the same explicit interface is accessible as at the
14 **declare reduction** directive.

Fortran

15 **Execution Model Events**

16 The *reduction-begin* event occurs before a task begins to perform loads and stores that belong to the
17 implementation of a reduction and the *reduction-end* event occurs after the task has completed
18 loads and stores associated with the reduction. If a task participates in multiple reductions, each
19 reduction may be bracketed by its own pair of *reduction-begin/reduction-end* events or multiple
20 reductions may be bracketed by a single pair of events. The interval defined by a pair of
21 *reduction-begin/reduction-end* events may not contain a task scheduling point.

22 **Tool Callbacks**

23 A thread dispatches a registered **ompt_callback_reduction** for each occurrence of a
24 *reduction-begin* or *reduction-end* event in that thread. The callback occurs in the context of the task
25 performing the reduction. This callback has the type signature
26 **ompt_callback_sync_region_t**. The callback receives
27 **ompt_sync_region_reduction** in its *kind* argument and **ompt_scope_begin** or
28 **ompt_scope_end** as its *endpoint* argument, as appropriate.

29 **Cross References**

- 30 • **ompt_scope_begin** and **ompt_scope_end**, see Section [4.2.3.4.10](#) on page [437](#).
- 31 • **ompt_sync_region_reduction**, see Section [4.2.3.4.12](#) on page [437](#).
- 32 • **ompt_callback_sync_region_t**, see Section [4.2.4.2.11](#) on page [457](#).

1 2.22.5.2 Reduction Scoping Clauses

2 Reduction scoping clauses define the region in which a reduction is computed by tasks or SIMD
3 lanes. All properties common to all reduction clauses, which are defined in Section 2.22.5.1, apply
4 to reduction scoping clauses.

5 The number of copies created for each list item and the time at which those copies are initialized
6 are determined by the particular reduction scoping clause that appears on the construct.

7 The time at which the original list item contains the result of the reduction is determined by the
8 particular reduction scoping clause.

▼ Fortran ▼

9 If the original list item has the **POINTER** attribute, copies of the list item are associated with
10 private targets.

▲ Fortran ▲

11 If the list item is an array section, the elements of any copy of the array section will be allocated
12 contiguously.

13 The location in the OpenMP program at which values are combined and the order in which values
14 are combined are unspecified. Therefore, when comparing sequential and parallel runs, or when
15 comparing one parallel run to another (even if the number of threads used is the same), there is no
16 guarantee that bit-identical results will be obtained or that side effects (such as floating-point
17 exceptions) will be identical or take place at the same location in the OpenMP program.

18 To avoid race conditions, concurrent reads or updates of the original list item must be synchronized
19 with the update of the original list item that occurs as a result of the reduction computation.

20 2.22.5.3 Reduction Participating Clauses

21 A reduction participating clause specifies a task or a SIMD lane as a participant in a reduction
22 defined by a reduction scoping clause. All properties common to all reduction clauses, which are
23 defined in Section 2.22.5.1, apply to reduction participating clauses.

24 Accesses to the original list item may be replaced by accesses to copies of the original list item
25 created by a region corresponding to a construct with a reduction scoping clause.

26 In any case, the final value of the reduction must be determined as if all tasks or SIMD lanes that
27 participate in the reduction are executed sequentially in some arbitrary order.

1 2.22.5.4 reduction Clause

2 Summary

3 The **reduction** clause specifies a *reduction-identifier* and one or more list items. For each list
4 item, a private copy is created in each implicit task or SIMD lane and is initialized with the
5 initializer value of the *reduction-identifier*. After the end of the region, the original list item is
6 updated with the values of the private copies using the combiner associated with the
7 *reduction-identifier*.

8 Syntax

```
9 | reduction ([reduction-modifier, ] reduction-identifier : list)
```

10 Where *reduction-identifier* is defined in Section 2.22.5.1, and *reduction-modifier* is one of the
11 following:

```
12 | inscan  
13 | task  
14 | default
```

15 Description

16 The **reduction** clause is a reduction scoping clause and a reduction participating clause, as
17 described in Sections 2.22.5.2 and 2.22.5.3.

18 If *reduction-modifier* is not present or the **default** *reduction-modifier* is present, the behavior is
19 as follows. For **parallel** and worksharing constructs, one or more private copies of each list
20 item are created for each implicit task, as if the **private** clause had been used. For the **simd**
21 construct, one or more private copies of each list item are created for each SIMD lane, as if the
22 **private** clause had been used. For the **taskloop** construct, private copies are created
23 according to the rules of the reduction scoping clauses. For the **teams** construct, one or more
24 private copies of each list item are created for the initial task of each team in the league, as if the
25 **private** clause had been used. For the **loop** construct, private copies are created and used in the
26 construct according to Section 2.22.3. At the end of a region corresponding to an above construct
27 for which the **reduction** clause was specified, the original list item is updated by combining its
28 original value with the final value of each of the private copies, using the combiner of the specified
29 *reduction-identifier*.

30 If the **inscan** *reduction-modifier* is present, a scan computation is performed over updates to the
31 list item performed in each logical iteration of the loop associated with the worksharing-loop,
32 worksharing-loop SIMD, or **simd** construct (see Section 2.12.6). The list items are privatized in
33 the construct according to Section 2.22.3. At the end of the region, each original list item is
34 assigned the value of the private copy from the last logical iteration of the loops associated with the
35 construct.

1 If the **task reduction-modifier** is present for a **parallel** or worksharing construct, then each list
2 item is privatized according to Section 2.22.3, and an unspecified number of additional private
3 copies are created to support task reductions. Any copies associated with the reduction are
4 initialized before they are accessed by the tasks participating in the reduction, which include all
5 implicit tasks in the corresponding region and all participating explicit tasks that specify an
6 **in_reduction** clause (see Section 2.22.5.6). After the end of the region, the original list item
7 contains the result of the reduction.

8 If **nowait** is not specified for the construct, the reduction computation will be complete at the end
9 of the construct; however, if the **reduction** clause is used on a construct to which **nowait** is
10 also applied, accesses to the original list item will create a race and, thus, have unspecified effect
11 unless synchronization ensures that they occur after all threads have executed all of their iterations
12 or **section** constructs, and the reduction computation has completed and stored the computed
13 value of that list item. This can most simply be ensured through a barrier synchronization.

14 Restrictions

15 The restrictions to the **reduction** clause are as follows:

- 16 • All the common restrictions to all reduction clauses, which are listed in Section 2.22.5.1, apply to
17 this clause.
- 18 • A list item that appears in a **reduction** clause of a worksharing construct must be shared in
19 the **parallel** region to which a corresponding worksharing region binds.
- 20 • A list item that appears in a **reduction** clause with the **inscan reduction-modifier** must
21 appear as a list item in an **inclusive** or **exclusive** clause on a **scan** directive enclosed by
22 the construct.
- 23 • A **reduction** clause without the **inscan reduction-modifier** may not appear on a construct
24 on which a **reduction** clause with the **inscan reduction-modifier** appears.
- 25 • A **reduction** clause with the **task reduction-modifier** may only appear on a **parallel**
26 construct, a worksharing construct or a combined construct for which any of the aforementioned
27 constructs is a constituent construct.
- 28 • A **reduction** clause with the **inscan reduction-modifier** may only appear on a
29 worksharing-loop construct, a worksharing-loop SIMD construct, a **simd** construct or a
30 combined construct for which any of the aforementioned constructs is a constituent construct.
- 31 • A list item that appears in a **reduction** clause of the innermost enclosing worksharing or
32 **parallel** construct may not be accessed in an explicit task generated by a construct for which
33 an **in_reduction** clause over the same list item does not appear.
- 34 • The **task reduction-modifier** may not appear in a **reduction** clause if the **nowait** clause is
35 specified on the same construct.

- 1 • If a list item in a **reduction** clause on a worksharing construct has a reference type then it
2 must bind to the same object for all threads of the team.
- 3 • A variable of class type (or array thereof) that appears in a **reduction** clause with the
4 **inscan** *reduction-modifier* requires an accessible, unambiguous default constructor for the
5 class type. The number of calls to the default constructor while performing the scan computation
6 is unspecified.
- 7 • A variable of class type (or array thereof) that appears in a **reduction** clause with the
8 **inscan** *reduction-modifier* requires an accessible, unambiguous copy assignment operator for
9 the class type. The number of calls to the copy assignment operator while performing the scan
10 computation is unspecified.

11 Cross References

- 12 • List Item Privatization, see Section [2.22.3](#) on page [273](#).
- 13 • **private** clause, see Section [2.22.4.3](#) on page [280](#).
- 14 • **scan** directive, see Section [2.12.6](#) on page [129](#).

15 2.22.5.5 **task_reduction** Clause

16 Summary

17 The **task_reduction** clause specifies a reduction among tasks.

18 Syntax

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

20 Where *reduction-identifier* is defined in Section [2.22.5.1](#).

21 Description

22 The **task_reduction** clause is a reduction scoping clause, as described in [2.22.5.2](#).

23 For each list item, the number of copies is unspecified. Any copies associated with the reduction
24 are initialized before they are accessed by the tasks participating in the reduction. After the end of
25 the region, the original list item contains the result of the reduction.

1 **Restrictions**

2 The restrictions to the **task_reduction** clause are as follows:

- 3 • All the common restrictions to all reduction clauses, which are listed in Section [2.22.5.1](#), apply to
4 this clause.

5 **2.22.5.6 in_reduction Clause**

6 **Summary**

7 The **in_reduction** clause specifies that a task participates in a reduction.

8 **Syntax**

9 **in_reduction** (*reduction-identifier* : list)

10 Where *reduction-identifier* is defined in Section [2.22.5.1](#)

11 **Description**

12 The **in_reduction** clause is a reduction participating clause, as described in Section [2.22.5.3](#).

13 For a given a list item, the **in_reduction** clause defines a task to be a participant in a task
14 reduction that is defined by an enclosing region for a matching list item that appears in a
15 **task_reduction** clause or a **reduction** clause with the **task** modifier, where either:

- 16 1. the matching list item has the same storage location as the list item in the **in_reduction**
17 clause; or
- 18 2. a private copy, derived from the matching list item, that is used to perform the task reduction has
19 the same storage location as the list item in the **in_reduction** clause.

20 For the **task** construct, the generated task becomes the participating task. For each list item, a
21 private copy may be created as if the **private** clause had been used.

22 For the **target** construct, the target task becomes the participating task. For each list item, a
23 private copy will be created in the data environment of the target task as if the **private** clause had
24 been used, and this private copy will be implicitly mapped into the device data environment of the
25 target device.

26 At the end of the task region, if a private copy was created its value is combined with a copy created
27 by a reduction scoping clause or with the original list item.

1 Restrictions

2 The restrictions to the **in_reduction** clause are as follows:

- 3 • All the common restrictions to all reduction clauses, which are listed in Section 2.22.5.1, apply to
4 this clause.
- 5 • For each list item, there must exist a matching list item that appears in a **task_reduction**
6 clause or a **reduction** clause with the **task** modifier that is specified on a construct
7 corresponding to a region in which the region of the participating task is closely nested. The
8 construct corresponding to the innermost enclosing region that meets this condition must specify
9 the same *reduction-identifier* for the matching list item as the **in_reduction** clause.

10 2.22.6 Data Copying Clauses

11 This section describes the **copyin** clause (allowed on the **parallel** directive and combined
12 parallel worksharing directives) and the **copyprivate** clause (allowed on the **single** directive).

13 These clauses support the copying of data values from private or threadprivate variables on one
14 implicit task or thread to the corresponding variables on other implicit tasks or threads in the team.

15 The clauses accept a comma-separated list of list items (see Section 2.1 on page 36). All list items
16 appearing in a clause must be visible, according to the scoping rules of the base language. Clauses
17 may be repeated as needed, but a list item that specifies a given variable may not appear in more
18 than one clause on the same directive.

Fortran

19 An associate name preserves the association with the selector established at the **ASSOCIATE**
20 statement. A list item that appears in a data copying clause may be a selector of an **ASSOCIATE**
21 construct. If the construct association is established prior to a parallel region, the association
22 between the associate name and the original list item will be retained in the region.

Fortran

23 2.22.6.1 copyin Clause

24 Summary

25 The **copyin** clause provides a mechanism to copy the value of the master thread's threadprivate
26 variable to the threadprivate variable of each other member of the team executing the **parallel**
27 region.

Syntax

The syntax of the **copyin** clause is as follows:

```
copyin (list)
```

Description

C / C++

The copy is done after the team is formed and prior to the start of execution of the associated structured block. For variables of non-array type, the copy occurs by copy assignment. For an array of elements of non-array type, each element is copied as if by assignment from an element of the master thread's array to the corresponding element of the other thread's array.

C / C++

C++

For class types, the copy assignment operator is invoked. The order in which copy assignment operators for different variables of class type are called is unspecified.

C++

Fortran

The copy is done, as if by assignment, after the team is formed and prior to the start of execution of the associated structured block.

On entry to any **parallel** region, each thread's copy of a variable that is affected by a **copyin** clause for the **parallel** region will acquire the type parameters, allocation, association, and definition status of the master thread's copy, according to the following rules:

- If the original list item has the **POINTER** attribute, each copy receives the same association status of the master thread's copy as if by pointer assignment.
- If the original list item does not have the **POINTER** attribute, each copy becomes defined with the value of the master thread's copy as if by intrinsic assignment unless the list item has a type bound procedure as a defined assignment. If the original list item that does not have the **POINTER** attribute has the allocation status of unallocated, each copy will have the same status.
- If the original list item is unallocated or unassociated, the thread's copy of the variable inherits the declared type parameters and the default type parameter values from the original list item.

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.
- If the list item is a polymorphic variable with the **ALLOCATABLE** attribute, the behavior is unspecified.

Fortran

2.22.6.2 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)
```

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 2.11.2 on page 89), and before any of the threads in the team have left the barrier at the end of the construct.

C / C++

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 associated with the thread that 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

C / C++

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.

C++

Fortran

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 a type bound procedure as a defined assignment, the assignment is performed by the defined assignment.

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.

Fortran

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).

1 Restrictions

2 The restrictions to the **copyprivate** clause are as follows:

- 3 • All list items that appear in the **copyprivate** clause must be either **threadprivate** or **private** in
4 the enclosing context.
- 5 • A list item that appears in a **copyprivate** clause may not appear in a **private** or
6 **firstprivate** clause on the **single** construct.

▼ C++ ▼

- 7 • A variable of class type (or array thereof) that appears in a **copyprivate** clause requires an
8 accessible unambiguous copy assignment operator for the class type.

▲ C++ ▲

▼ Fortran ▼

- 9 • A common block that appears in a **copyprivate** clause must be **threadprivate**.
- 10 • Pointers with the **INTENT (IN)** attribute may not appear in the **copyprivate** clause.
- 11 • The list item with the **ALLOCATABLE** attribute must have the allocation status of **allocated** when
12 the intrinsic assignment is performed.
- 13 • If the list item is a polymorphic variable with the **ALLOCATABLE** attribute, the behavior is
14 unspecified.

▲ Fortran ▲

15 2.22.7 Data-mapping Attribute Rules and Clauses

16 This section describes how the data-mapping and data-sharing attributes of any variable referenced
17 in a **target** region are determined. When specified, explicit data-sharing attributes, **map** or
18 **is_device_ptr** clauses on **target** directives determine these attributes. Otherwise, the first
19 matching rule from the following implicit data-mapping rules applies for variables referenced in a
20 **target** construct that are not declared in the construct and do not appear in data-sharing attribute,
21 **map** or **is_device_ptr** clauses:

- 22 • If a variable appears in a **to** or **link** clause on a **declare target** directive then it is treated
23 as if it had appeared in a **map** clause with a *map-type* of **tofrom**.
- 24 • If a list item appears in a **reduction**, **lastprivate** or **linear** clause on a combined
25 **target** construct then it is treated as if it also appears in a **map** clause with a *map-type* of
26 **tofrom**.

- 1 • If a list item appears in an **in_reduction** clause on a **target** construct then it is treated as if
2 it also appears in a **map** clause with a *map-type* of **tofrom** and a *map-type-modifier* of
3 **always**.
- 4 • If a **defaultmap** clause is present for the category of the variable and specifies an implicit
5 behavior other than **default**, the data-mapping attribute is determined by that clause.

▼ C++ ▼

- 6 • If the **target** construct is within a class non-static member function, and a variable is an
7 accessible data member of the object for which the non-static data member function is invoked,
8 the variable is treated as if the **this[:1]** expression had appeared in a **map** clause with a
9 *map-type* of **tofrom**. Additionally, if the variable is of a type pointer or reference to pointer, it
10 is also treated as if it has appeared in a **map** clause as a zero-length array section.
- 11 • If the **this** keyword is referenced inside a **target** construct within a class non-static member
12 function, it is treated as if the **this[:1]** expression had appeared in a **map** clause with a
13 *map-type* of **tofrom**.

▲ C++ ▲

▼ C / C++ ▼

- 14 • A variable that is of type pointer is treated as if it is the named pointer of a zero-length array
15 section that appeared as a list item in a **map** clause.

▲ C / C++ ▲

▼ C++ ▼

- 16 • A variable that is of type reference to pointer is treated as if it had appeared in a **map** clause as a
17 zero-length array section.

▲ C++ ▲

- 18 • If a variable is not a scalar then it is treated as if it had appeared in a **map** clause with a *map-type*
19 of **tofrom**.

▼ Fortran ▼

- 20 • If a scalar variable has the **TARGET**, **ALLOCATABLE** or **POINTER** attribute then it is treated as
21 if it has appeared in a **map** clause with a *map-type* of **tofrom**.

▲ Fortran ▲

- 22 • If none of the above rules applies then a scalar variable is not mapped, but instead has an implicit
23 data-sharing attribute of **firstprivate** (see Section 2.22.1.1 on page 263).

1 2.22.7.1 map Clause

2 Summary

3 The **map** clause specifies how an original list item is mapped from the current task's data
4 environment to a corresponding list item in the device data environment of the device identified by
5 the construct.

6 Syntax

7 The syntax of the map clause is as follows:

```
8 map ([ [map-type-modifier[, ] [map-type-modifier[, ] ...] map-type : ] list)
```

9 where *map-type* is one of the following:

```
10 to  
11 from  
12 tofrom  
13 alloc  
14 release  
15 delete
```

16 and *map-type-modifier* is one of the following:

```
17 always  
18 close  
19 mapper (mapper-identifier)
```

20 Description

21 The list items that appear in a **map** clause may include array sections and structure elements.

22 The *map-type* and *map-type-modifier* specify the effect of the **map** clause, as described below.

23 For a given construct, the effect of a **map** clause with the **to**, **from**, or **tofrom** *map-type* is
24 ordered before the effect of a **map** clause with the **alloc**, **release**, or **delete** *map-type*. If a
25 mapper is specified for the type being mapped, or explicitly specified with the **mapper**
26 *map-type-modifier*, then the effective **map-type** of a list item will be determined according to the
27 rules of map-type decay.

28 If a mapper is specified for the type being mapped, or explicitly specified with the **mapper**
29 *map-type-modifier*, then all map clauses that appear on the **declare mapper** directive are
30 treated as though they appeared on the construct with the **map** clause. Array sections of a mapper

1 type are mapped as normal, then each element in the array section is mapped according to the rules
2 of the mapper.

▼ C / C++ ▼

3 If a list item in a **map** clause is a variable of structure type then it is treated as if each structure
4 element contained in the variable is a list item in the clause.

▲ C / C++ ▲

▼ Fortran ▼

5 If a list item in a **map** clause is a derived type variable then it is treated as if each nonpointer
6 component is a list item in the clause.

▲ Fortran ▲

7 If a list item in a **map** clause is a structure element then all other structure elements (except pointer
8 component, for Fortran) of the containing structure variable form a *structure sibling list*. The **map**
9 clause and the structure sibling list are associated with the same construct. If a corresponding list
10 item of the structure sibling list item is present in the device data environment when the construct is
11 encountered then:

12 • If the structure sibling list item does not appear in a **map** clause on the construct then:

13 – If the construct is a **target**, **target data**, or **target enter data** construct then the
14 structure sibling list item is treated as if it is a list item in a **map** clause on the construct with a
15 *map-type* of **alloc**.

16 – If the construct is **target exit data** construct, then the structure sibling list item is treated
17 as if it is a list item in a **map** clause on the construct with a *map-type* of **release**.

18 • If the **map** clause in which the structure element appears as a list item has a *map-type* of
19 **delete** and the structure sibling list item does not appear as a list item in a **map** clause on the
20 construct with a *map-type* of **delete** then the structure sibling list item is treated as if it is a list
21 item in a **map** clause on the construct with a *map-type* of **delete**.

▼ Fortran ▼

22 If a list item in a **map** clause has the **POINTER** attribute and if the association status of the list item
23 is associated, then it is treated as if the pointer target is a list item in the clause.

▲ Fortran ▲

C / C++

1 If $item_1$ is a list item in a **map** clause, and $item_2$ is another list item in a **map** clause on the same
2 construct that has a named pointer that is, or is part of, $item_1$, then:

- 3 • If the **map** clause(s) appear on a **target**, **target data**, or **target enter data** construct,
4 then on entry to the corresponding region the effect of the **map** clause on $item_1$ is ordered to
5 occur before the effect of the **map** clause on $item_2$.
- 6 • If the **map** clause(s) appear on a **target**, **target data**, or **target exit data** construct,
7 then on exit from the corresponding region the effect of the **map** clause on $item_2$ is ordered to
8 occur before the effect of the **map** clause on $item_1$.

9 If an array section with a named pointer is a list item in a **map** clause and a pointer variable is
10 present in the device data environment that corresponds to the named pointer when the effect of the
11 **map** clause occurs, then if the corresponding array section is created in the device data
12 environment:

- 13 1. The corresponding pointer variable is assigned the address of the corresponding array section.
- 14 2. The corresponding pointer variable becomes an attached pointer for the corresponding array
15 section.
- 16 3. If the original named pointer and the corresponding attached pointer share storage, then the
17 original array section and the corresponding array section must share storage.

C / C++

C++

18 If a *lambda* is mapped explicitly or implicitly, variables that are captured by the *lambda* behave as
19 follows:

- 20 • the variables that are of pointer type are treated as if they had appeared in a **map** clause as
21 zero-length array sections
- 22 • the variables that are of reference type are treated as if they had appeared in a **map** clause.

23 If a member variable is captured by a *lambda* in class scope, and the *lambda* is later mapped
24 explicitly or implicitly with its full static type, the *this* pointer is treated as if it had appeared on a
25 **map** clause.

C++

26 The original and corresponding list items may share storage such that writes to either item by one
27 task followed by a read or write of the other item by another task without intervening
28 synchronization can result in data races.

29 If the **map** clause appears on a **target**, **target data**, or **target enter data** construct then
30 on entry to the region the following sequence of steps occurs as if performed as a single atomic
31 operation:

- 1 1. If a corresponding list item of the original list item is not present in the device data environment,
2 then:
- 3 a) A new list item with language-specific attributes is derived from the original list item and
4 created in the device data environment.
- 5 b) The new list item becomes the corresponding list item to the original list item in the device
6 data environment.
- 7 c) The corresponding list item has a reference count that is initialized to zero.
- 8 d) The value of the corresponding list item is undefined.
- 9 2. If the corresponding list item's reference count was not already incremented because of the
10 effect of a **map** clause on the construct then:
- 11 a) The corresponding list item's reference count is incremented by one
- 12 3. If the corresponding list item's reference count is one or the **always map-type-modifier** is
13 present, then:
- 14 a) If the *map-type* is **to** or **tofrom**, then:
- 15 • For each part of the list item that is an attached pointer:
 - 16 – That part of the corresponding list item will have the value it had immediately prior to
17 the effect of the **map** clause;
 - 18 • For each part of the list item that is not an attached pointer:
 - 19 – The value of that part of the original list item is assigned to that part of the
20 corresponding list item.
 - 21 • Concurrent reads or updates of any part of the corresponding list item must be
22 synchronized with the update of the corresponding list item that occurs as a result of the
23 **map** clause.

24 **Note** – If the effect of the **map** clauses on a construct would assign the value of an original list
25 item to a corresponding list item more than once, then an implementation is allowed to ignore
26 additional assignments of the same value to the corresponding list item.

27 If the **map** clause appears on a **target**, **target data**, or **target exit data** construct then
28 on exit from the region the following sequence of steps occurs as if performed as a single atomic
29 operation:

- 30 1. If a corresponding list item of the original list item is not present in the device data environment,
31 then the list item is ignored.

- 1 2. If a corresponding list item of the original list item is present in the device data environment,
2 then:
- 3 a) If the corresponding list item's reference count is finite, then:
- 4 i. If the corresponding list item's reference count was not already decremented because
5 of the effect of a **map** clause on the construct then:
- 6 A. If the *map-type* is not **delete**, then the corresponding list item's reference count
7 is decremented by one.
- 8 ii. If the *map-type* is **delete**, then the corresponding list item's reference count is set to
9 zero.
- 10 b) If the corresponding list item's reference count is zero or the **always map-type-modifier** is
11 present, then:
- 12 i. If the *map-type* is **from** or **tofrom** then:
- 13 • For each part of the list item that is an attached pointer:
- 14 – That part of the original list item will have the value it had immediately prior to
15 the effect of the **map** clause;
- 16 • For each part of the list item that is not an attached pointer:
- 17 – The value of that part of the corresponding list item is assigned to that part of the
18 original list item;
- 19 • To avoid race conditions:
- 20 – Concurrent reads or updates of any part of the original list item must be
21 synchronized with the update of the original list item that occurs as a result of the
22 **map** clause;
- 23 c) If the corresponding list item's reference count is zero, then the corresponding list item is
24 removed from the device data environment

25 **Note** – If the effect of the **map** clauses on a construct would assign the value of a corresponding
26 list item to an original list item more than once, then an implementation is allowed to ignore
27 additional assignments of the same value to the original list item.

28 If a single contiguous part of the original storage of a list item with an implicit data-mapping
29 attribute has corresponding storage in the device data environment prior to a task encountering the
30 construct associated with the **map** clause, only that part of the original storage will have
31 corresponding storage in the device data environment as a result of the **map** clause.

C / C++

1 If a new list item is created then a new list item of the same type, with automatic storage duration, is
2 allocated for the construct. The size and alignment of the new list item are determined by the static
3 type of the variable. This allocation occurs if the region references the list item in any statement.

C / C++

Fortran

4 If a new list item is created then a new list item of the same type, type parameter, and rank is
5 allocated. The new list item inherits all default values for the type parameters from the original list
6 item.

7 If the allocation status of the original list item with the **ALLOCATABLE** attribute is changed in the
8 host device data environment and the corresponding list item is already present in the device data
9 environment, the allocation status of the corresponding list item is unspecified until a mapping
10 operation is performed with a **map** clause on entry to a **target**, **target data**, or
11 **target enter data** region.

Fortran

12 The *map-type* determines how the new list item is initialized.

13 If a *map-type* is not specified, the *map-type* defaults to **tofrom**.

14 The **close map-type-modifier** is a hint to the runtime to allocate memory close to the target device.

Execution Model Events

15 The *target-map* event occurs when a thread maps data to or from a target device.

16 The *target-data-op* event occurs when a thread initiates a data operation on a target device.

Tool Callbacks

17 A thread dispatches a registered **ompt_callback_target_map** callback for each occurrence
18 of a *target-map* event in that thread. The callback occurs in the context of the target task. The
19 callback has type signature **ompt_callback_target_map_t**.

20 A thread dispatches a registered **ompt_callback_target_data_op** callback for each
21 occurrence of a *target-data-op* event in that thread. The callback occurs in the context of the target
22 task. The callback has type signature **ompt_callback_target_data_op_t**.

Restrictions

- A list item cannot appear in both a **map** clause and a data-sharing attribute clause on the same construct, unless the the construct is a combined construct.

- Each of the *map-type-modifier* modifiers can appear at most once on the **map** clause.

C / C++

- If a list item is an array section and the type of its base expression is a pointer type, the base expression must be an lvalue expression.

C / C++

- If a list item is an array section, it must specify contiguous storage.
- If more than one list item of the **map** clauses on the same construct are, or are part of, array items that have the same named array, they must indicate identical original storage.
- List items of the **map** clauses on the same construct must not share original storage unless they are the same variable or array section.
- If any part of the original storage of a list item with an explicit data-mapping attribute has corresponding storage in the device data environment prior to a task encountering the construct associated with the map clause, all of the original storage must have corresponding storage in the device data environment prior to the task encountering the construct.
- If a list item is an element of a structure, and a different element of the structure has a corresponding list item in the device data environment prior to a task encountering the construct associated with the **map** clause, then the list item must also have a corresponding list item in the device data environment prior to the task encountering the construct.
- If a list item is an element of a structure, only the rightmost symbol of the variable reference can be an array section.
- A list item must have a mappable type.
- **threadprivate** variables cannot appear in a **map** clause.
- If a **mapper** map-type-modifier is specified, its type must match the type of the list-items passed to that map clause.
- Memory spaces and memory allocators cannot appear as a list item in a **map** clause.

C++

- If the type of a list item is a reference to a type T then the reference in the device data environment is initialized to refer to the object in the device data environment that corresponds to the object referenced by the list item. If mapping occurs, it occurs as though the object were mapped through a pointer with an array section of type T and length one.
- No type mapped through a reference can contain a reference to its own type, or any cycle of references to types that could produce a cycle of references.
- If the list item is a *lambda*, any pointers and references captured by the *lambda* must have the corresponding list item in the device data environment prior to the task encountering the construct.

C++

C / C++

- Initialization and assignment are through bitwise copy.
- A list item cannot be a variable that is a member of a structure with a union type.
- A bit-field cannot appear in a **map** clause.
- A pointer that has a corresponding attached pointer may not be modified for the duration of the lifetime of the array section to which the corresponding pointer is attached in the device data environment.

C / C++

Fortran

- The value of the new list item becomes that of the original list item in the map initialization and assignment.
- If the allocation status of a list item or any subobject of the list item with the **ALLOCATABLE** attribute is unallocated upon entry to a **target** region, the list item or any subobject of the corresponding list item must be unallocated upon exit from the region.
- If the allocation status of a list item or any subobject of the list item with the **ALLOCATABLE** attribute is allocated upon entry to a **target** region, the allocation status of the corresponding list item or any subobject of the corresponding list item must not be changed and must not be reshaped in the region.
- If an array section is mapped and the size of the section is smaller than that of the whole array, the behavior of referencing the whole array in the **target** region is unspecified.
- A list item must not be a whole array of an assumed-size array.
- If the association status of a list item with the **POINTER** attribute is associated upon entry to a **target** region, the list item remains associated with the same pointer target upon exit from the region.

- 1 • If the association status of a list item with the **POINTER** attribute is disassociated upon entry to a
2 **target** region, the list item must be disassociated upon exit from the region.
- 3 • If the association status of a list item with the **POINTER** attribute is undefined upon entry to a
4 **target** region, the list item must be undefined upon exit from the region.
- 5 • If the association status of a list item with the **POINTER** attribute is disassociated or undefined
6 on entry and if the list item is associated with a pointer target inside a **target** region, then the
7 pointer association status must become disassociated before the end of the region; otherwise the
8 behavior is unspecified.

Fortran

Cross References

- 9 • `ompt_callback_target_map_t`, see Section [4.2.4.2.22](#) on page [471](#).
- 10 • `ompt_callback_target_data_op_t`, see Section [4.2.4.2.21](#) on page [468](#).

2.22.7.2 defaultmap Clause

Summary

14 The **defaultmap** clause redefines the implicit data-mapping attributes of variables that are
15 referenced in a **target** construct and are implicitly determined.

Syntax

17 The syntax of the **defaultmap** clause is as follows:

```
18 defaultmap (implicit-behavior[:variable-category])
```

19 Where *implicit-behavior* is one of:

```
20     alloc  
21     to  
22     from  
23     tofrom  
24     firstprivate  
25     none  
26     default
```

C / C++

1 and *variable-category* is one of:

```
2     scalar  
3     aggregate  
4     pointer
```

C / C++

Fortran

5 and *variable-category* is one of:

```
6     scalar  
7     aggregate  
8     allocatable  
9     pointer
```

Fortran

10 Description

11 The **defaultmap** clause sets the implicit data-mapping attribute for all variables referenced in the
12 construct. If *variable-category* is specified, the effect of the **defaultmap** clause is as follows:

- 13 • If *variable-category* is **scalar**, all scalar variables of non-pointer type or all non-pointer
14 non-allocatable scalar variables that have an implicitly determined data-mapping or data-sharing
15 attribute will have a data-mapping or data-sharing attribute specified by *implicit-behavior*.
- 16 • If *variable-category* is **aggregate** or **allocatable**, all aggregate or allocatable variables
17 that have an implicitly determined data-mapping or data-sharing attribute will have a
18 data-mapping or data-sharing attribute specified by *implicit-behavior*.
- 19 • If *variable-category* is **pointer**, all variables of pointer type or with the POINTER attribute
20 that have implicitly determined data-mapping or data-sharing attributes will have a data-mapping
21 or data-sharing attribute specified by *implicit-behavior*. The zero-length array section and
22 attachment an implicitly mapped pointer normally gets is only provided for the **default**
23 behavior.

1 If no *variable-category* is specified in the clause then *implicit-behavior* specifies the implicitly
2 determined data-mapping or data-sharing attribute for all variables referenced in the construct. If
3 *implicit-behavior* is **none**, each variable referenced in the construct that does not have a
4 predetermined data-sharing attribute and does not appear in a **to** or **link** clause on a
5 **declare target** directive must be listed in a data-mapping attribute clause, a data-sharing
6 attribute clause (including data-sharing attribute clause on a combined construct where **target** is
7 one of the constituent constructs), or an **is_device_ptr** clause. If *implicit-behavior* is
8 **default**, then the clause has no effect for the variables in the category specified by
9 *variable-category*.

10 2.23 Declare Directives

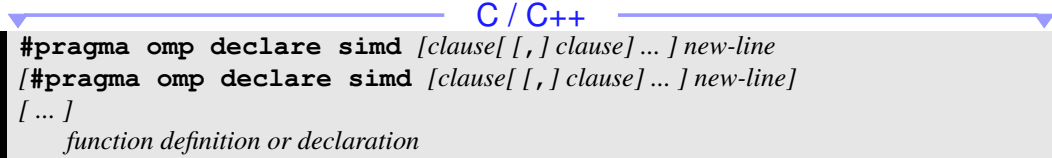
11 2.23.1 declare simd Directive

12 Summary

13 The **declare simd** directive can be applied to a function (C, C++ and Fortran) or a subroutine
14 (Fortran) to enable the creation of one or more versions that can process multiple arguments using
15 SIMD instructions from a single invocation in a SIMD loop. The **declare simd** directive is a
16 declarative directive. There may be multiple **declare simd** directives for a function (C, C++,
17 Fortran) or subroutine (Fortran).

18 Syntax

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

20 
21 **#pragma omp declare simd** [*clause*[[,] *clause*] ...] *new-line*
22 **[/#pragma omp declare simd** [*clause*[[,] *clause*] ...] *new-line*]
23 [...]
 function definition or declaration

24 where *clause* is one of the following:

25 **simdlen** (*length*)
26 **linear** (*linear-list*[: *linear-step*])
27 **aligned** (*argument-list*[: *alignment*])
28 **uniform** (*argument-list*)
29 **inbranch**

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23

```
notinbranch
```

C / C++

Fortran

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

where *clause* is one of the following:

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

Fortran

Description

C / C++

The use of one or more **declare simd** directives immediately prior to a function declaration or definition enables the creation of corresponding SIMD versions of the associated function that can be used to process multiple arguments from a single invocation in a SIMD loop concurrently.

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

C / C++

Fortran

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

Fortran

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

C++

1 The special *this* pointer can be used as if was one of the arguments to the function in any of the
2 **linear**, **aligned**, or **uniform** clauses.

C++

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

C / C++

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

C / C++

Fortran

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

Fortran

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

12 The **inbranch** clause specifies that the SIMD version of the function will always be called from
13 inside a conditional statement of a SIMD loop. The **notinbranch** clause specifies that the SIMD
14 version of the function will never be called from inside a conditional statement of a SIMD loop. If
15 neither clause is specified, then the SIMD version of the function may or may not be called from
16 inside a conditional statement of a SIMD loop.

Restrictions

- 17 • Each argument can appear in at most one **uniform** or **linear** clause.
- 18 • At most one **simdlen** clause can appear in a **declare simd** directive.
- 19 • Either **inbranch** or **notinbranch** may be specified, but not both.
- 20 • When a *linear-step* expression is specified in a **linear** clause it must be either a constant integer
21 expression or an integer-typed parameter that is specified in a **uniform** clause on the directive.
- 22 • The function or subroutine body must be a structured block.
- 23 • The execution of the function or subroutine, when called from a SIMD loop, cannot result in the
24 execution of an OpenMP construct except for an **ordered** construct with the **simd** clause or an
25 **atomic** construct.
- 26

- 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 **set jmp** 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++

Fortran

- *proc-name* must not be a generic name, procedure pointer or entry name.
- If *proc-name* is omitted, the **declare simd** directive must appear in the specification part of a subroutine subprogram or a function subprogram for which creation of the SIMD versions is enabled.
- Any **declare simd** directive must appear in the specification part of a subroutine subprogram, function subprogram or interface body to which it applies.
- If a **declare simd** directive is specified in an interface block for a procedure, it must match a **declare simd** directive in the definition of the procedure.
- If a procedure is declared via a procedure declaration statement, the procedure *proc-name* should appear in the same specification.
- If a **declare simd** directive is specified for a procedure name with explicit interface and a **declare simd** directive is also specified for the definition of the procedure then the two **declare simd** directives must match. Otherwise the result is unspecified.
- Procedure pointers may not be used to access versions created by the **declare simd** directive.

- 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

- **reduction** clause, see Section 2.22.5.4 on page 297.
- **linear** clause, see Section 2.22.4.6 on page 286.

2.23.2 declare reduction Directive

Summary

The following section describes the directive for declaring user-defined reductions. The **declare reduction** directive declares a *reduction-identifier* that can be used in a **reduction** clause. The **declare reduction** directive is a declarative directive.

Syntax

```
#pragma omp declare reduction(reduction-identifier : typename-list :  
combiner ) [initializer-clause] new-line
```

where:

- *reduction-identifier* is either a base language identifier or one of the following operators: **+**, **-**, *****, **&**, **|**, **^**, **&&** and **||**
- *typename-list* is a list of type names
- *combiner* is an expression
- *initializer-clause* is **initializer** (*initializer-expr*) where *initializer-expr* is **omp_priv = initializer** or **function-name** (*argument-list*)

C++

```
#pragma omp declare reduction(reduction-identifier : typename-list :  
combiner) [initializer-clause] new-line
```

where:

- *reduction-identifier* is either an *id-expression* or one of the following operators: +, -, *, &, |, ^, && and ||
- *typename-list* is a list of type names
- *combiner* is an expression
- *initializer-clause* is **initializer**(*initializer-expr*) where *initializer-expr* is **omp_priv** *initializer* or *function-name* (*argument-list*)

C++

Fortran

```
!$omp declare reduction(reduction-identifier : type-list : combiner)  
[initializer-clause]
```

where:

- *reduction-identifier* is either a base language identifier, or a user-defined operator, or one of the following operators: +, -, *, .and., .or., .eqv., .neqv., or one of the following intrinsic procedure names: **max**, **min**, **iand**, **ior**, **ieor**.
- *type-list* is a list of type specifiers that must not be **CLASS (*)** and abstract type
- *combiner* is either an assignment statement or a subroutine name followed by an argument list
- *initializer-clause* is **initializer**(*initializer-expr*), where *initializer-expr* is **omp_priv** = *expression* or *subroutine-name* (*argument-list*)

Fortran

Description

Custom reductions can be defined using the **declare reduction** directive; the *reduction-identifier* and the type identify the **declare reduction** directive. The *reduction-identifier* can later be used in a **reduction** clause using variables of the type or types specified in the **declare reduction** directive. If the directive applies to several types then it is considered as if there were multiple **declare reduction** directives, one for each type.

Fortran

1 If a type with deferred or assumed length type parameter is specified in a **declare reduction**
2 directive, the *reduction-identifier* of that directive can be used in a **reduction** clause with any
3 variable of the same type and the same kind parameter, regardless of the length type Fortran
4 parameters with which the variable is declared.

Fortran

5 The visibility and accessibility of this declaration are the same as those of a variable declared at the
6 same point in the program. The enclosing context of the *combiner* and of the *initializer-expr* will be
7 that of the **declare reduction** directive. The *combiner* and the *initializer-expr* must be correct
8 in the base language as if they were the body of a function defined at the same point in the program.

Fortran

9 If the *reduction-identifier* is the same as the name of a user-defined operator or an extended
10 operator, or the same as a generic name that is one of the allowed intrinsic procedures, and if the
11 operator or procedure name appears in an accessibility statement in the same module, the
12 accessibility of the corresponding **declare reduction** directive is determined by the
13 accessibility attribute of the statement.

14 If the *reduction-identifier* is the same as a generic name that is one of the allowed intrinsic
15 procedures and is accessible, and if it has the same name as a derived type in the same module, the
16 accessibility of the corresponding **declare reduction** directive is determined by the
17 accessibility of the generic name according to the base language.

Fortran

C++

18 The **declare reduction** directive can also appear at points in the program at which a static
19 data member could be declared. In this case, the visibility and accessibility of the declaration are
20 the same as those of a static data member declared at the same point in the program.

C++

21 The *combiner* specifies how partial results can be combined into a single value. The *combiner* can
22 use the special variable identifiers **omp_in** and **omp_out** that are of the type of the variables
23 being reduced with this *reduction-identifier*. Each of them will denote one of the values to be
24 combined before executing the *combiner*. It is assumed that the special **omp_out** identifier will
25 refer to the storage that holds the resulting combined value after executing the *combiner*.

26 The number of times the *combiner* is executed, and the order of these executions, for any
27 **reduction** clause is unspecified.

Fortran

1 If the *combiner* is a subroutine name with an argument list, the *combiner* is evaluated by calling the
2 subroutine with the specified argument list.

3 If the *combiner* is an assignment statement, the *combiner* is evaluated by executing the assignment
4 statement.

Fortran

5 As the *initializer-expr* value of a user-defined reduction is not known *a priori* the *initializer-clause*
6 can be used to specify one. Then the contents of the *initializer-clause* will be used as the initializer
7 for private copies of reduction list items where the **omp_priv** identifier will refer to the storage to
8 be initialized. The special identifier **omp_orig** can also appear in the *initializer-clause* and it will
9 refer to the storage of the original variable to be reduced.

10 The number of times that the *initializer-expr* is evaluated, and the order of these evaluations, is
11 unspecified.

C / C++

12 If the *initializer-expr* is a function name with an argument list, the *initializer-expr* is evaluated by
13 calling the function with the specified argument list. Otherwise, the *initializer-expr* specifies how
14 **omp_priv** is declared and initialized.

C / C++

C

15 If no *initializer-clause* is specified, the private variables will be initialized following the rules for
16 initialization of objects with static storage duration.

C

C++

17 If no *initializer-expr* is specified, the private variables will be initialized following the rules for
18 *default-initialization*.

C++

Fortran

1 If the *initializer-expr* is a subroutine name with an argument list, the *initializer-expr* is evaluated by
2 calling the subroutine with the specified argument list.

3 If the *initializer-expr* is an assignment statement, the *initializer-expr* is evaluated by executing the
4 assignment statement.

5 If no *initializer-clause* is specified, the private variables will be initialized as follows:

- 6 • For **complex**, **real**, or **integer** types, the value 0 will be used.
- 7 • For **logical** types, the value **.false.** will be used.
- 8 • For derived types for which default initialization is specified, default initialization will be used.
- 9 • Otherwise, not specifying an *initializer-clause* results in unspecified behavior.

Fortran

C / C++

10 If *reduction-identifier* is used in a **target** region then a **declare target** construct must be
11 specified for any function that can be accessed through the *combiner* and *initializer-expr*.

C / C++

Fortran

12 If *reduction-identifier* is used in a **target** region then a **declare target** construct must be
13 specified for any function or subroutine that can be accessed through the *combiner* and
14 *initializer-expr*.

Fortran

Restrictions

- 15 • The only variables allowed in the *combiner* are **omp_in** and **omp_out**.
- 16 • The only variables allowed in the *initializer-clause* are **omp_priv** and **omp_orig**.
- 17 • If the variable **omp_orig** is modified in the *initializer-clause*, the behavior is unspecified.
- 18 • If execution of the *combiner* or the *initializer-expr* results in the execution of an OpenMP
19 construct or an OpenMP API call, then the behavior is unspecified.
- 20 • A *reduction-identifier* may not be re-declared in the current scope for the same type or for a type
21 that is compatible according to the base language rules.
- 22 • At most one *initializer-clause* can be specified.
- 23

C / C++

- 1
- 2
- A type name in a **declare reduction** directive cannot be a function type, an array type, a reference type, or a type qualified with **const**, **volatile** or **restrict**.

C / C++

- 3
- 4
- If the *initializer-expr* is a function name with an argument list, then one of the arguments must be the address of **omp_priv**.

C

- 5
- 6
- If the *initializer-expr* is a function name with an argument list, then one of the arguments must be **omp_priv** or the address of **omp_priv**.

C++

- 7
- 8
- If the *initializer-expr* is a subroutine name with an argument list, then one of the arguments must be **omp_priv**.

Fortran

- 9
- 10
- 11
- 12
- If the **declare reduction** directive appears in the specification part of a module and the corresponding reduction clause does not appear in the same module, the *reduction-identifier* must be the same as the name of a user-defined operator, one of the allowed operators that is extended or a generic name that is the same as the name of one of the allowed intrinsic procedures.
- 13
- 14
- 15
- 16
- 17
- 18
- If the **declare reduction** directive appears in the specification of a module, if the corresponding **reduction** clause does not appear in the same module, and if the *reduction-identifier* is the same as the name of a user-defined operator or an extended operator, or the same as a generic name that is the same as one of the allowed intrinsic procedures then the interface for that operator or the generic name must be defined in the specification of the same module, or must be accessible by use association.
- 19
- 20
- Any subroutine or function used in the **initializer** clause or *combiner* expression must be an intrinsic function, or must have an accessible interface.
- 21
- 22
- Any user-defined operator, defined assignment or extended operator used in the **initializer** clause or *combiner* expression must have an accessible interface.
- 23
- 24
- 25
- If any subroutine, function, user-defined operator, defined assignment or extended operator is used in the **initializer** clause or *combiner* expression, it must be accessible to the subprogram in which the corresponding **reduction** clause is specified.
- 26
- If the length type parameter is specified for a type, it must be a constant, a colon or an *****.

- If a type with deferred or assumed length parameter is specified in a **declare reduction** directive, no other **declare reduction** directive with the same type, the same kind parameters and the same *reduction-identifier* is allowed in the same scope.
- Any subroutine used in the **initializer** clause or *combiner* expression must not have any alternate returns appear in the argument list.

Fortran

Cross References

- **reduction** clause, Section [2.22.5.4](#) on page [297](#).

2.24 Nesting of Regions

This section describes a set of restrictions on the nesting of regions. The restrictions on nesting are as follows:

- A worksharing region may not be closely nested inside a worksharing, **task**, **taskloop**, **critical**, **ordered**, **atomic**, or **master** region.
- A **barrier** region may not be closely nested inside a worksharing, **task**, **taskloop**, **critical**, **ordered**, **atomic**, or **master** region.
- A **master** region may not be closely nested inside a worksharing, **atomic**, **task**, or **taskloop** region.
- An **ordered** region corresponding to an **ordered** construct without any clause or with the **threads** or **depend** clause may not be closely nested inside a **critical**, **ordered**, **atomic**, **task**, or **taskloop** region.
- An **ordered** region corresponding to an **ordered** construct without the **simd** clause specified must be closely nested inside a worksharing-loop region.
- An **ordered** region corresponding to an **ordered** construct with the **simd** clause specified must be closely nested inside a **simd** or worksharing-loop SIMD region.
- An **ordered** region corresponding to an **ordered** construct with both the **simd** and **threads** clauses must be closely nested inside a worksharing-loop SIMD region or closely nested inside a worksharing-loop and **simd** region.
- A **critical** region may not be nested (closely or otherwise) inside a **critical** region with the same name. This restriction is not sufficient to prevent deadlock.

- 1 • OpenMP constructs may not be encountered during execution of an **atomic** region.
- 2 • The only OpenMP constructs that can be encountered during execution of a **simd** (or
- 3 worksharing-loop SIMD) region are the **atomic** construct, the **loop** construct, the **simd**
- 4 construct and the **ordered** construct with the **simd** clause.
- 5 • If a **target update**, **target data**, **target enter data**, or **target exit data**
- 6 construct is encountered during execution of a **target** region, the behavior is unspecified.
- 7 • If a **target** construct is encountered during execution of a **target** region and a **device**
- 8 clause in which the **ancestor device-modifier** appears is not present on the construct, the
- 9 behavior is unspecified.
- 10 • A **teams** region can only be strictly nested within the implicit parallel region or a **target**
- 11 region. If a **teams** construct is nested within a **target** construct, that **target** construct must
- 12 contain no statements, declarations or directives outside of the **teams** construct.
- 13 • **distribute**, **distribute simd**, distribute parallel worksharing-loop, distribute parallel
- 14 worksharing-loop SIMD, **loop**, **parallel** regions, including any **parallel** regions arising
- 15 from combined constructs, **omp_get_num_teams()** regions, and **omp_get_team_num()**
- 16 regions are the only OpenMP regions that may be strictly nested inside the **teams** region.
- 17 • The region corresponding to the **distribute** construct must be strictly nested inside a **teams**
- 18 region.
- 19 • If *construct-type-clause* is **taskgroup**, the **cancel** construct must be closely nested inside a
- 20 **task** construct and the **cancel** region must be closely nested inside a **taskgroup** region. If
- 21 *construct-type-clause* is **sections**, the **cancel** construct must be closely nested inside a
- 22 **sections** or **section** construct. Otherwise, the **cancel** construct must be closely nested
- 23 inside an OpenMP construct that matches the type specified in *construct-type-clause* of the
- 24 **cancel** construct.
- 25 • A **cancellation point** construct for which *construct-type-clause* is **taskgroup** must be
- 26 closely nested inside a **task** construct, and the **cancellation point** region must be closely
- 27 nested inside a **taskgroup** region. A **cancellation point** construct for which
- 28 *construct-type-clause* is **sections** must be closely nested inside a **sections** or **section**
- 29 construct. Otherwise, a **cancellation point** construct must be closely nested inside an
- 30 OpenMP construct that matches the type specified in *construct-type-clause*.
- 31 • The only constructs that may be nested inside a **loop** region are the **loop** construct, the
- 32 **parallel** construct, the **simd** construct, and combined constructs for which the first construct
- 33 is a **parallel** construct.
- 34 • A **loop** region corresponding to a **loop** construct may not contain calls to procedures that
- 35 contain OpenMP directives.
- 36 • A **loop** region may not contain calls to the OpenMP Runtime API.

Runtime Library Routines

This chapter describes the OpenMP API runtime library routines and queryable runtime states, and is divided into the following sections:

- Runtime library definitions (Section 3.1 on page 330).

- Execution environment routines that can be used to control and to query the parallel execution environment (Section 3.2 on page 332).

- Lock routines that can be used to synchronize access to data (Section 3.3 on page 378).

- Portable timer routines (Section 3.4 on page 390).

- Device memory routines that can be used to allocate memory and to manage pointers on target devices (Section 3.6 on page 393).

- Execution routines to control the application monitoring (Section 3.8 on page 413)

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 ▲

Restrictions

The following restriction applies to all OpenMP runtime library routines:

- OpenMP runtime library routines may not be called from **PURE** or **ELEMENTAL** procedures.

3.1 Runtime Library Definitions

For each base language, a compliant implementation must supply a set of definitions for the OpenMP API runtime library routines and the special data types of their parameters. The set of definitions must contain a declaration for each OpenMP API runtime library routine and a declaration for the *simple lock*, *nestable lock*, *schedule*, and *thread affinity policy* data types. In addition, each set of definitions may specify other implementation specific values.

The library routines are external functions with “C” linkage.

Prototypes for the C/C++ runtime library routines described in this chapter shall be provided in a header file named **omp.h**. This file defines the following:

- The prototypes of all the routines in the chapter.
- The type **omp_lock_t**.
- The type **omp_nest_lock_t**.
- The type **omp_sync_hint_t**.
- The type **omp_lock_hint_t** (deprecated).
- The type **omp_sched_t**.
- The type **omp_proc_bind_t**.
- The type **omp_control_tool_t**.
- The type **omp_control_tool_result_t**.
- The type **omp_depend_t**.
- The type **omp_memspace_t**.
- The type **omp_allocator_t**.

- 1 • The type `omp_uintptr_t` which is an unsigned integer type capable of holding a pointer.
- 2 • A global variable of type `const omp_memspace_t *` for each predefined memory space in
- 3 Table 2.7 on page 150.
- 4 • A global variable of type `const omp_allocator_t *` for each predefined memory
- 5 allocator in Table 2.9 on page 153.
- 6 • The type `omp_pause_resource_t`.
- 7 • The type `omp_event_t`.

8 A program that declares a new variable with the same identifier as one of the predefined

9 allocators listed in Table 2.9 on page 153 results in unspecified behavior.

10 A program that declares a new variable with the same identifier as one of the predefined memory

11 spaces listed in Table 2.7 on page 150 results in unspecified behavior.

C / C++

C++

- 12 • A class template that models the `Allocator` concept in the `omp::allocator` namespace
- 13 for each predefined memory allocator in Table 2.9 on page 153 for which the name includes
- 14 neither the `omp_` prefix nor the `_alloc` suffix.

C++

Fortran

15 The OpenMP Fortran API runtime library routines are external procedures. The return values of

16 these routines are of default kind, unless otherwise specified.

17 Interface declarations for the OpenMP Fortran runtime library routines described in this chapter

18 shall be provided in the form of a Fortran `include` file named `omp_lib.h` or a Fortran 90

19 **module** named `omp_lib`. It is implementation defined whether the `include` file or the

20 **module** file (or both) is provided.

21 These files define the following:

- 22 • The interfaces of all of the routines in this chapter.
- 23 • The `integer` parameter `omp_lock_kind`.
- 24 • The `integer` parameter `omp_nest_lock_kind`.
- 25 • The `integer` parameter `omp_sync_hint_kind`.
- 26 • The `integer` parameter `omp_lock_hint_kind` (deprecated).
- 27 • The `integer` parameter `omp_sched_kind`.
- 28 • The `integer` parameter `omp_proc_bind_kind`.
- 29 • The `integer` parameter `omp_depend_kind`.

- 1 • The **integer parameter** `omp_memspace_kind`.
- 2 • The **integer parameter** `omp_allocator_kind`.
- 3 • The **integer parameter** `omp_alloctrait_key_kind`.
- 4 • The **integer parameter** `omp_alloctrait_val_kind`.
- 5 • An **integer parameter** variable of kind `omp_memspace_kind` for each predefined
- 6 memory space in Table 2.7 on page 150.
- 7 • An **integer parameter** variable of kind `omp_allocator_kind` for each predefined
- 8 memory allocator in Table 2.9 on page 153.
- 9 • The **integer parameter** `openmp_version` with a value `yyyymm` where `yyyy` and `mm` are
- 10 the year and month designations of the version of the OpenMP Fortran API that the
- 11 implementation supports. This value matches that of the C preprocessor macro `_OPENMP`, when
- 12 a macro preprocessor is supported (see Section 2.2 on page 41).
- 13 • The **integer parameter** `omp_pause_kind`.
- 14 • The **integer parameter** `omp_event_kind`.

15 It is implementation defined whether any of the OpenMP runtime library routines that take an
16 argument are extended with a generic interface so arguments of different **KIND** type can be
17 accommodated.

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

18 3.2 Execution Environment Routines

19 This section describes routines that affect and monitor threads, processors, and the parallel
20 environment.

21 3.2.1 `omp_set_num_threads`

22 Summary

23 The `omp_set_num_threads` routine affects the number of threads to be used for subsequent
24 parallel regions that do not specify a `num_threads` clause, by setting the value of the first
25 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 *num_threads*-var ICV of the current task to the value specified in the argument.

Cross References

- *num_threads*-var ICV, see Section 2.4 on page 47.
- `parallel` construct and `num_threads` clause, see Section 2.9 on page 72.
- Determining the number of threads for a `parallel` region, see Section 2.9.1 on page 77.
- `omp_get_max_threads` routine, see Section 3.2.3 on page 334.
- `OMP_NUM_THREADS` environment variable, see Section 5.2 on page 597.

3.2.2 `omp_get_num_threads`

Summary

The `omp_get_num_threads` routine returns the number of threads in the current team.

Format

Binding

The binding region for an `omp_get_num_threads` region is the innermost enclosing `parallel` region.

Effect

The `omp_get_num_threads` routine returns the number of threads in the team executing the `parallel` region to which the routine region binds. If called from the sequential part of a program, this routine returns 1.

Cross References

- `parallel` construct, see Section 2.9 on page 72.
- Determining the number of threads for a `parallel` region, see Section 2.9.1 on page 77.
- `omp_set_num_threads` routine, see Section 3.2.1 on page 332.
- `OMP_NUM_THREADS` environment variable, see Section 5.2 on page 597.

3.2.3 `omp_get_max_threads`

Summary

The `omp_get_max_threads` routine returns an upper bound on the number of threads that could be used to form a new team if a `parallel` construct without a `num_threads` clause were encountered after execution returns from this routine.

Format

C / C++
`int omp_get_max_threads(void);`

C / C++

Fortran
`integer function omp_get_max_threads()`

Fortran

Binding

The binding task set for an `omp_get_max_threads` region is the generating task.

Effect

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.

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.

Cross References


- *nthreads-var* ICV, see Section 2.4 on page 47.
- `parallel` construct, see Section 2.9 on page 72.
- `num_threads` clause, see Section 2.9 on page 72.
- Determining the number of threads for a `parallel` region, see Section 2.9.1 on page 77.
- `omp_set_num_threads` routine, see Section 3.2.1 on page 332.
- `OMP_NUM_THREADS` environment variable, see Section 5.2 on page 597.

1 3.2.4 `omp_get_thread_num`

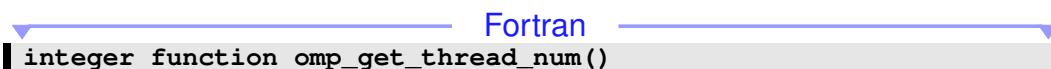
2 Summary

3 The `omp_get_thread_num` routine returns the thread number, within the current team, of the
4 calling thread.

5 Format

6  C / C++
`int omp_get_thread_num(void);`

 C / C++

7  Fortran
`integer function omp_get_thread_num()`

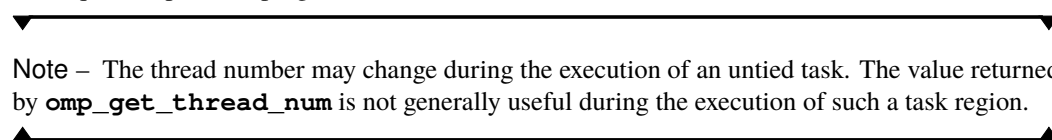
 Fortran

8 Binding

9 The binding thread set for an `omp_get_thread_num` region is the current team. The binding
10 region for an `omp_get_thread_num` region is the innermost enclosing `parallel` region.

11 Effect

12 The `omp_get_thread_num` routine returns the thread number of the calling thread, within the
13 team executing the `parallel` region to which the routine region binds. The thread number is an
14 integer between 0 and one less than the value returned by `omp_get_num_threads`, inclusive.
15 The thread number of the master thread of the team is 0. The routine returns 0 if it is called from
16 the sequential part of a program.

17  Note – The thread number may change during the execution of an untied task. The value returned
18 by `omp_get_thread_num` is not generally useful during the execution of such a task region.

19 Cross References

- 20
 - `omp_get_num_threads` routine, see Section 3.2.2 on page 333.

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 C / C++
6 Fortran

7 Binding

8 The binding thread set for an `omp_get_num_procs` region is all threads on a device. The effect
9 of executing this routine is not related to any specific region corresponding to any construct or API
10 routine.

11 Effect

12 The `omp_get_num_procs` routine returns the number of processors that are available to the
13 device at the time the routine is called. This value may change between the time that it is
14 determined by the `omp_get_num_procs` routine and the time that it is read in the calling
15 context due to system actions outside the control of the OpenMP implementation.

16 Cross References

17 None.

18 3.2.6 `omp_in_parallel`

19 Summary

20 The `omp_in_parallel` routine returns *true* if the *active-levels-var* ICV is greater than zero;
21 otherwise, it returns *false*.

1 **Format**

2 C / C++
3 Fortran
4 Fortran

4 **Binding**

5 The binding task set for an `omp_in_parallel` region is the generating task.

6 **Effect**

7 The effect of the `omp_in_parallel` routine is to return *true* if the current task is enclosed by an
8 active `parallel` region, and the `parallel` region is enclosed by the outermost initial task
9 region on the device; otherwise it returns *false*.

10 **Cross References**

- 11 • *active-levels-var*, see Section 2.4 on page 47.
- 12 • `parallel` construct, see Section 2.9 on page 72.
- 13 • `omp_get_active_level` routine, see Section 3.2.20 on page 353.

14 **3.2.7 omp_set_dynamic**

15 **Summary**

16 The `omp_set_dynamic` routine enables or disables dynamic adjustment of the number of
17 threads available for the execution of subsequent `parallel` regions by setting the value of the
18 *dyn-var* ICV.

Format

C / C++
`void omp_set_dynamic(int dynamic_threads);`

C / C++

Fortran
`subroutine omp_set_dynamic(dynamic_threads)
logical dynamic_threads`

Fortran

Binding

The binding task set for an `omp_set_dynamic` region is the generating task.

Effect

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*.

Cross References

- *dyn-var* ICV, see Section 2.4 on page 47.
- Determining the number of threads for a `parallel` region, see Section 2.9.1 on page 77.
- `omp_get_num_threads` routine, see Section 3.2.2 on page 333.
- `omp_get_dynamic` routine, see Section 3.2.8 on page 339.
- `OMP_DYNAMIC` environment variable, see Section 5.3 on page 598.

3.2.8 `omp_get_dynamic`

Summary

The `omp_get_dynamic` routine returns the value of the *dyn-var* ICV, which determines whether dynamic adjustment of the number of threads is enabled or disabled.

1

Format

2

	C / C++	
<code>int omp_get_dynamic(void);</code>		

3

	C / C++	
	Fortran	
<code>logical function omp_get_dynamic()</code>		
	Fortran	

4

Binding

5

The binding task set for an `omp_get_dynamic` region is the generating task.

6

Effect

7

8

9

This routine returns *true* if dynamic adjustment of the number of threads is enabled for the current task; it returns *false*, otherwise. If an implementation does not support dynamic adjustment of the number of threads, then this routine always returns *false*.

10

Cross References

11

12

13

14

- *dyn-var* ICV, see Section [2.4](#) on page [47](#).
- Determining the number of threads for a `parallel` region, see Section [2.9.1](#) on page [77](#).
- `omp_set_dynamic` routine, see Section [3.2.7](#) on page [338](#).
- `OMP_DYNAMIC` environment variable, see Section [5.3](#) on page [598](#).

15

3.2.9 `omp_get_cancellation`

16

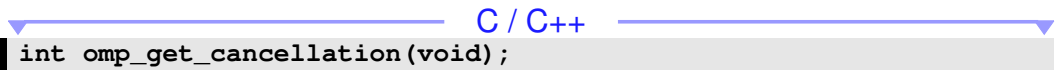
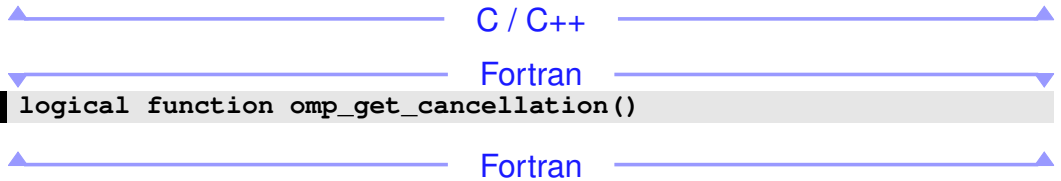
Summary

17

18

The `omp_get_cancellation` routine returns the value of the *cancel-var* ICV, which determines if cancellation is enabled or disabled.

1 **Format**

```
2         C / C++  
3         Fortran
```

4 **Binding**

5 The binding task set for an `omp_get_cancellation` region is the whole program.

6 **Effect**

7 This routine returns *true* if cancellation is enabled. It returns *false* otherwise.

8 **Cross References**

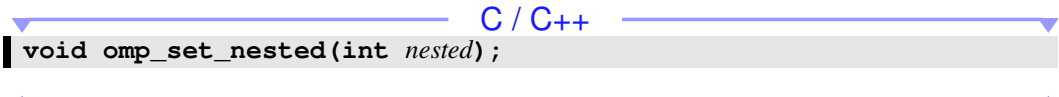
- 9 • *cancel-var* ICV, see Section [2.4.1](#) on page [47](#).
- 10 • `cancel` construct, see Section [2.21.1](#) on page [256](#).
- 11 • `OMP_CANCELLATION` environment variable, see Section [5.11](#) on page [604](#)

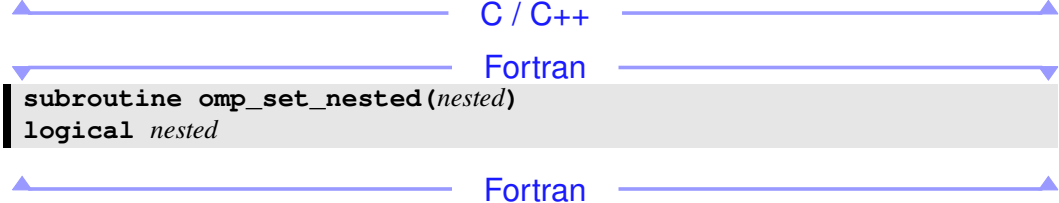
12 **3.2.10 omp_set_nested**

13 **Summary**

14 The deprecated `omp_set_nested` routine enables or disables nested parallelism, by setting the
15 *nest-var* ICV.

1 **Format**

2  `void omp_set_nested(int nested);`

3  `subroutine omp_set_nested(nested)
4 logical nested`

5 **Binding**

6 The binding task set for an `omp_set_nested` region is the generating task.

7 **Effect**

8 For implementations that support nested parallelism, if the argument to `omp_set_nested`
9 evaluates to *true*, nested parallelism is enabled for the current task; otherwise, nested parallelism is
10 disabled for the current task. For implementations that do not support nested parallelism, this
11 routine has no effect: the value of *nest-var* remains *false*. This routine has been deprecated.

12 **Cross References**

- 13 • *nest-var* ICV, see Section [2.4](#) on page [47](#).
- 14 • Determining the number of threads for a `parallel` region, see Section [2.9.1](#) on page [77](#).
- 15 • `omp_get_nested` routine, see Section [3.2.11](#) on page [342](#).
- 16 • `omp_set_max_active_levels` routine, see Section [3.2.15](#) on page [347](#).
- 17 • `omp_get_max_active_levels` routine, see Section [3.2.16](#) on page [348](#).
- 18 • `OMP_NESTED` environment variable, see Section [5.6](#) on page [601](#).

19 **3.2.11 omp_get_nested**

20 **Summary**

21 The deprecated `omp_get_nested` routine returns the value of the *nest-var* ICV, which
22 determines if nested parallelism is enabled or disabled.

1 **Format**

```
2        int omp_get_nested(void);        C / C++  
3        logical function omp_get_nested()        Fortran  
4        logical function omp_get_nested()        Fortran
```

4 **Binding**

5 The binding task set for an `omp_get_nested` region is the generating task.

6 **Effect**

7 This routine returns *true* if nested parallelism is enabled for the current task; it returns *false*,
8 otherwise. If an implementation does not support nested parallelism, this routine always returns
9 *false*. This routine has been deprecated.

10 **Cross References**

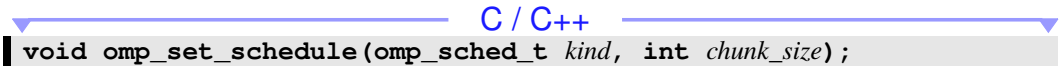
- 11 • *nest-var* ICV, see Section 2.4 on page 47.
- 12 • Determining the number of threads for a `parallel` region, see Section 2.9.1 on page 77.
- 13 • `omp_set_nested` routine, see Section 3.2.10 on page 341.
- 14 • `OMP_NESTED` environment variable, see Section 5.6 on page 601.

15 **3.2.12 omp_set_schedule**

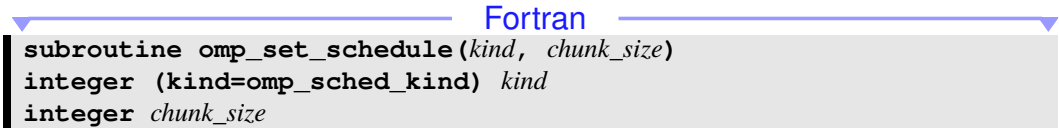
16 **Summary**

17 The `omp_set_schedule` routine affects the schedule that is applied when `runtime` is used as
18 schedule kind, by setting the value of the *run-sched-var* ICV.

1 Format

2  C / C++
`void omp_set_schedule(omp_sched_t kind, int chunk_size);`

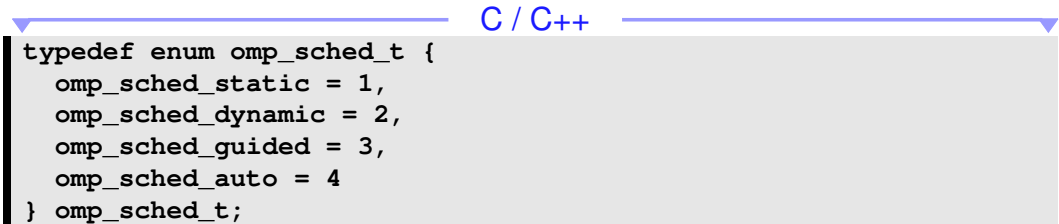
 C / C++

3  Fortran
4 `subroutine omp_set_schedule(kind, chunk_size)`
5 `integer (kind=omp_sched_kind) kind`
`integer chunk_size`

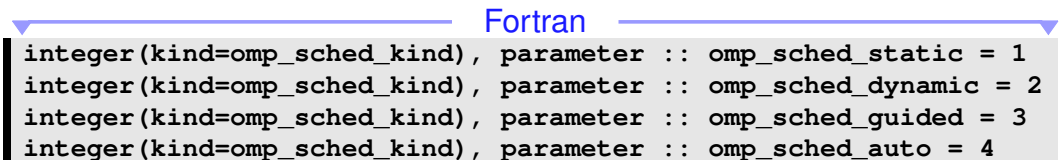
 Fortran

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:

12  C / C++
13 `typedef enum omp_sched_t {`
14 `omp_sched_static = 1,`
15 `omp_sched_dynamic = 2,`
16 `omp_sched_guided = 3,`
17 `omp_sched_auto = 4`
`} omp_sched_t;`

 C / C++

18  Fortran
19 `integer(kind=omp_sched_kind), parameter :: omp_sched_static = 1`
20 `integer(kind=omp_sched_kind), parameter :: omp_sched_dynamic = 2`
21 `integer(kind=omp_sched_kind), parameter :: omp_sched_guided = 3`
`integer(kind=omp_sched_kind), parameter :: omp_sched_auto = 4`

 Fortran

22 Binding

23 The binding task set for an **omp_set_schedule** region is the generating task.

Effect

The effect of this routine is to set the value of the *run-sched-var* ICV of the current task to the values specified in the two arguments. The schedule is set to the schedule type specified by the first argument *kind*. It can be any of the standard schedule types or any other implementation specific one. For the schedule types **static**, **dynamic**, and **guided** the *chunk_size* is set to the value of the second argument, or to the default *chunk_size* if the value of the second argument is less than 1; for the schedule type **auto** the second argument has no meaning; for implementation specific schedule types, the values and associated meanings of the second argument are implementation defined.

Cross References

- *run-sched-var* ICV, see Section 2.4 on page 47.
- Determining the schedule of a worksharing loop, see Section 2.12.2.1 on page 110.
- `omp_get_schedule` routine, see Section 3.2.13 on page 345.
- `OMP_SCHEDULE` environment variable, see Section 5.1 on page 596.

3.2.13 `omp_get_schedule`

Summary

The `omp_get_schedule` routine returns the schedule that is applied when the runtime schedule is used.

Format

C / C++
`void omp_get_schedule(omp_sched_t *kind, int *chunk_size);`

C / C++

Fortran

Fortran
`subroutine omp_get_schedule(kind, chunk_size)
integer (kind=omp_sched_kind) kind
integer chunk_size`

Fortran

1 **Binding**

2 The binding task set for an **omp_get_schedule** region is the generating task.

3 **Effect**

4 This routine returns the *run-sched-var* ICV in the task to which the routine binds. The first
5 argument *kind* returns the schedule to be used. It can be any of the standard schedule types as
6 defined in Section 3.2.12 on page 343, or any implementation specific schedule type. The second
7 argument is interpreted as in the **omp_set_schedule** call, defined in Section 3.2.12 on
8 page 343.

9 **Cross References**

- 10 • *run-sched-var* ICV, see Section 2.4 on page 47.
- 11 • Determining the schedule of a worksharing loop, see Section 2.12.2.1 on page 110.
- 12 • **omp_set_schedule** routine, see Section 3.2.12 on page 343.
- 13 • **OMP_SCHEDULE** environment variable, see Section 5.1 on page 596.

14 **3.2.14 omp_get_thread_limit**

15 **Summary**

16 The **omp_get_thread_limit** routine returns the maximum number of OpenMP threads
17 available to participate in the current contention group.

18 **Format**

```
19       C / C++  
| int omp_get_thread_limit(void);  
  
20       C / C++  
| integer function omp_get_thread_limit()  
  
      Fortran
```

1 **Binding**

2 The binding thread set for an `omp_get_thread_limit` region is all threads on the device. The
3 effect of executing this routine is not related to any specific region corresponding to any construct
4 or API routine.

5 **Effect**

6 The `omp_get_thread_limit` routine returns the value of the *thread-limit-var* ICV.

7 **Cross References**

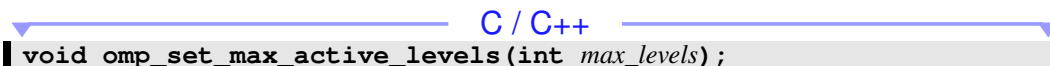
- 8 • *thread-limit-var* ICV, see Section [2.4](#) on page [47](#).
- 9 • `OMP_THREAD_LIMIT` environment variable, see Section [5.10](#) on page [604](#).

10 **3.2.15 omp_set_max_active_levels**

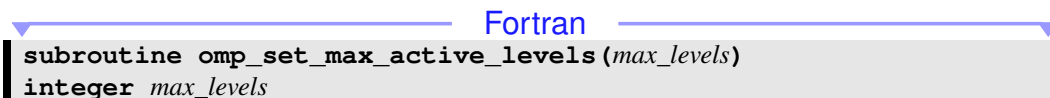
11 **Summary**

12 The `omp_set_max_active_levels` routine limits the number of nested active parallel
13 regions on the device, by setting the *max-active-levels-var* ICV

14 **Format**

15 
`void omp_set_max_active_levels(int max_levels);`

16 
`void omp_set_max_active_levels(int max_levels);`

17 
`subroutine omp_set_max_active_levels(max_levels)
integer max_levels`

18 
`subroutine omp_set_max_active_levels(max_levels)
integer max_levels`

18 **Constraints on Arguments**

19 The value of the argument passed to this routine must evaluate to a non-negative integer, otherwise
20 the behavior of this routine is implementation defined.

Binding

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.

Effect

The effect of this routine is to set the value of the *max-active-levels-var* ICV to the value specified in the argument.

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.

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.

Cross References

- *max-active-levels-var* ICV, see Section 2.4 on page 47.
- `omp_get_max_active_levels` routine, see Section 3.2.16 on page 348.
- `OMP_MAX_ACTIVE_LEVELS` environment variable, see Section 5.9 on page 603.

3.2.16 `omp_get_max_active_levels`

Summary

The `omp_get_max_active_levels` routine returns the value of the *max-active-levels-var* ICV, which determines the maximum number of nested active parallel regions on the device.

Format

```
int omp_get_max_active_levels(void);
```

C / C++

Fortran

```
1 integer function omp_get_max_active_levels()
```

Fortran

2 Binding

3 When called from a sequential part of the program, the binding thread set for an
4 `omp_get_max_active_levels` region is the encountering thread. When called from within
5 any explicit parallel region, the binding thread set (and binding region, if required) for the
6 `omp_get_max_active_levels` region is implementation defined.

7 Effect

8 The `omp_get_max_active_levels` routine returns the value of the *max-active-levels-var*
9 ICV, which determines the maximum number of nested active parallel regions on the device.

10 Cross References

- 11 • *max-active-levels-var* ICV, see Section 2.4 on page 47.
- 12 • `omp_set_max_active_levels` routine, see Section 3.2.15 on page 347.
- 13 • `OMP_MAX_ACTIVE_LEVELS` environment variable, see Section 5.9 on page 603.

14 3.2.17 `omp_get_level`

15 Summary

16 The `omp_get_level` routine returns the value of the *levels-var* ICV.

17 Format

C / C++

```
18 int omp_get_level(void);
```

C / C++

Fortran

```
integer function omp_get_level()
```

Fortran

Binding

The binding task set for an `omp_get_level` region is the generating task.

Effect

The effect of the `omp_get_level` routine is to return the number of nested `parallel` regions (whether active or inactive) enclosing the current task such that all of the `parallel` regions are enclosed by the outermost initial task region on the current device.

Cross References

- `levels-var` ICV, see Section 2.4 on page 47.
- `omp_get_active_level` routine, see Section 3.2.20 on page 353.
- `OMP_MAX_ACTIVE_LEVELS` environment variable, see Section 5.9 on page 603.

3.2.18 `omp_get_ancestor_thread_num`

Summary

The `omp_get_ancestor_thread_num` routine returns, for a given nested level of the current thread, the thread number of the ancestor of the current thread.

Format

C / C++

```
int omp_get_ancestor_thread_num(int level);
```

C / C++

Fortran

```
integer function omp_get_ancestor_thread_num(level)
integer level
```

Fortran

1 **Binding**

2 The binding thread set for an `omp_get_ancestor_thread_num` region is the encountering
3 thread. The binding region for an `omp_get_ancestor_thread_num` region is the innermost
4 enclosing `parallel` region.

5 **Effect**

6 The `omp_get_ancestor_thread_num` routine returns the thread number of the ancestor at a
7 given nest level of the current thread or the thread number of the current thread. If the requested
8 nest level is outside the range of 0 and the nest level of the current thread, as returned by the
9 `omp_get_level` routine, the routine returns -1.



10 Note – When the `omp_get_ancestor_thread_num` routine is called with a value of
11 `level=0`, the routine always returns 0. If `level=omp_get_level()`, the routine has the
12 same effect as the `omp_get_thread_num` routine.



13 **Cross References**

- 14 • `omp_get_thread_num` routine, see Section 3.2.4 on page 336.
- 15 • `omp_get_level` routine, see Section 3.2.17 on page 349.
- 16 • `omp_get_team_size` routine, see Section 3.2.19 on page 351.

17 **3.2.19 omp_get_team_size**

18 **Summary**

19 The `omp_get_team_size` routine returns, for a given nested level of the current thread, the size
20 of the thread team to which the ancestor or the current thread belongs.

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20

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 333.
- `omp_get_level` routine, see Section 3.2.17 on page 349.
- `omp_get_ancestor_thread_num` routine, see Section 3.2.18 on page 350.

1 3.2.20 `omp_get_active_level`

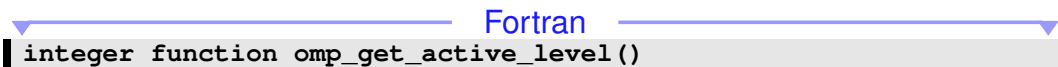
2 Summary

3 The `omp_get_active_level` routine returns the value of the *active-level-vars* ICV..

4 Format

5  C / C++
`int omp_get_active_level(void);`

 C / C++

6  Fortran
`integer function omp_get_active_level()`

 Fortran

7 Binding

8 The binding task set for the an `omp_get_active_level` region is the generating task.

9 Effect

10 The effect of the `omp_get_active_level` routine is to return the number of nested, active
11 **parallel** regions enclosing the current task such that all of the **parallel** regions are enclosed
12 by the outermost initial task region on the current device.

13 Cross References

- 14 • *active-levels-var* ICV, see Section 2.4 on page 47.
- 15 • `omp_get_level` routine, see Section 3.2.17 on page 349.

16 3.2.21 `omp_in_final`

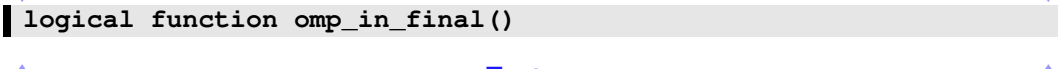
17 Summary

18 The `omp_in_final` routine returns *true* if the routine is executed in a final task region;
19 otherwise, it returns *false*.

1 **Format**

2  `int omp_in_final(void);`



3  `logical function omp_in_final()`



4 **Binding**

5 The binding task set for an `omp_in_final` region is the generating task.

6 **Effect**

7 `omp_in_final` returns *true* if the enclosing task region is final. Otherwise, it returns *false*.

8 **Cross References**

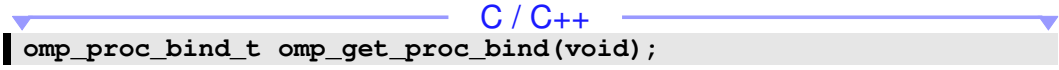
- 9 • `task` construct, see Section 2.13.1 on page 133.

10 **3.2.22 omp_get_proc_bind**

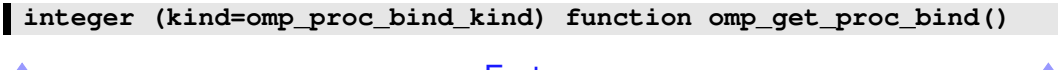
11 **Summary**

12 The `omp_get_proc_bind` routine returns the thread affinity policy to be used for the
13 subsequent nested `parallel` regions that do not specify a `proc_bind` clause.

14 **Format**

15  `omp_proc_bind_t omp_get_proc_bind(void);`



16  `integer (kind=omp_proc_bind_kind) function omp_get_proc_bind()`



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 2.9.2 on page 79 for the rules governing the thread affinity policy.

Cross References





- *bind-var* ICV, see Section 2.4 on page 47.
- Controlling OpenMP thread affinity, see Section 2.9.2 on page 79.
- `OMP_PROC_BIND` environment variable, see Section 5.4 on page 598.

3.2.23 `omp_get_num_places`

Summary

The `omp_get_num_places` routine returns the number of places available to the execution environment in the place list.

Format

	
	<code>int omp_get_num_places(void);</code>
	
	
	<code>integer function omp_get_num_places()</code>
	

Binding

The binding thread set for an `omp_get_num_places` region is all threads on a device. The effect of executing this routine is not related to any specific region corresponding to any construct or API routine.

Effect

The `omp_get_num_places` routine returns the number of places in the place list. This value is equivalent to the number of places in the *place-partition-var* ICV in the execution environment of the initial task.

Cross References

- *place-partition-var* ICV, see Section 2.4 on page 47.
- `OMP_PLACES` environment variable, see Section 5.5 on page 599.

3.2.24 `omp_get_place_num_procs`

Summary

The `omp_get_place_num_procs` routine returns the number of processors available to the execution environment in the specified place.

Format

```

C / C++
int omp_get_place_num_procs (int place_num);

C / C++

Fortran
integer function omp_get_place_num_procs (place_num)
integer place_num

Fortran
```

Binding

The binding thread set for an `omp_get_place_num_procs` region is all threads on a device. The effect of executing this routine is not related to any specific region corresponding to any construct or API routine.

Effect

The `omp_get_place_num_procs` routine returns the number of processors associated with the place numbered *place_num*. The routine returns zero when *place_num* is negative, or is equal to or larger than the value returned by `omp_get_num_places()`.

Cross References


- `OMP_PLACES` environment variable, see Section 5.5 on page 599.

1 3.2.25 omp_get_place_proc_ids


2 Summary

3 The `omp_get_place_proc_ids` routine returns the numerical identifiers of the processors
4 available to the execution environment in the specified place.

5 Format

6  C / C++
`void omp_get_place_proc_ids(int place_num, int *ids);`

 C / C++

7  Fortran
8 `subroutine omp_get_place_proc_ids(place_num, ids)`
9 `integer place_num`
`integer ids(*)`

 Fortran

10 Binding

11 The binding thread set for an `omp_get_place_proc_ids` region is all threads on a device.
12 The effect of executing this routine is not related to any specific region corresponding to any
13 construct or API routine.

14 Effect

15 The `omp_get_place_proc_ids` routine returns the numerical identifiers of each processor
16 associated with the place numbered `place_num`. The numerical identifiers are non-negative, and
17 their meaning is implementation defined. The numerical identifiers are returned in the array `ids` and
18 their order in the array is implementation defined. The array must be sufficiently large to contain
19 `omp_get_place_num_procs(place_num)` integers; otherwise, the behavior is unspecified.
20 The routine has no effect when `place_num` has a negative value, or a value equal or larger than
21 `omp_get_num_places()`.

22 Cross References


- 23 • `omp_get_place_num_procs` routine, see Section 3.2.24 on page 357.
- 24 • `omp_get_num_places` routine, see Section 3.2.23 on page 356.
- 25 • `OMP_PLACES` environment variable, see Section 5.5 on page 599.

1 3.2.26 `omp_get_place_num`

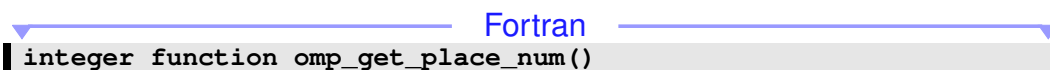
2 Summary

3 The `omp_get_place_num` routine returns the place number of the place to which the
4 encountering thread is bound.

5 Format

6  C / C++
| `int omp_get_place_num(void);`

 C / C++

7  Fortran
| `integer function omp_get_place_num()`

 Fortran

8 Binding

9 The binding thread set for an `omp_get_place_num` region is the encountering thread.

10 Effect

11 When the encountering thread is bound to a place, the `omp_get_place_num` routine returns the
12 place number associated with the thread. The returned value is between 0 and one less than the
13 value returned by `omp_get_num_places()`, inclusive. When the encountering thread is not
14 bound to a place, the routine returns -1.

15 Cross References


- 16 • Controlling OpenMP thread affinity, see Section [2.9.2](#) on page [79](#).
- 17 • `omp_get_num_places` routine, see Section [3.2.23](#) on page [356](#).
- 18 • `OMP_PLACES` environment variable, see Section [5.5](#) on page [599](#).

1 3.2.27 `omp_get_partition_num_places`

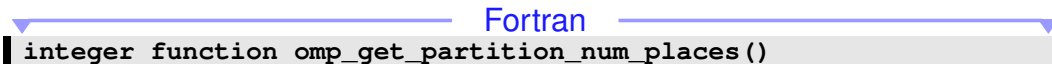
2 Summary

3 The `omp_get_partition_num_places` routine returns the number of places in the place
4 partition of the innermost implicit task.

5 Format

6  `int omp_get_partition_num_places(void);`

 C / C++

7  `integer function omp_get_partition_num_places()`

 Fortran

8 Binding

9 The binding task set for an `omp_get_partition_num_places` region is the encountering
10 implicit task.

11 Effect

12 The `omp_get_partition_num_places` routine returns the number of places in the
13 *place-partition-var* ICV.

14 Cross References

- 15 • *place-partition-var* ICV, see Section 2.4 on page 47.
- 16 • Controlling OpenMP thread affinity, see Section 2.9.2 on page 79.
- 17 • `OMP_PLACES` environment variable, see Section 5.5 on page 599.

18 3.2.28 `omp_get_partition_place_nums`

19 Summary

20 The `omp_get_partition_place_nums` routine returns the list of place numbers
21 corresponding to the places in the *place-partition-var* ICV of the innermost implicit task.

Format

C / C++
`void omp_get_partition_place_nums(int *place_nums);`

C / C++

Fortran
`subroutine omp_get_partition_place_nums(place_nums)
integer place_nums(*)`

Fortran

Binding

The binding task set for an `omp_get_partition_place_nums` region is the encountering implicit task.

Effect

The `omp_get_partition_place_nums` routine returns the list of place numbers corresponding to the places in the *place-partition-var* ICV of the innermost implicit task. The array must be sufficiently large to contain `omp_get_partition_num_places()` integers; otherwise, the behavior is unspecified.

Cross References

- *place-partition-var* ICV, see Section 2.4 on page 47.
- Controlling OpenMP thread affinity, see Section 2.9.2 on page 79.
- `omp_get_partition_num_places` routine, see Section 3.2.27 on page 360.
- `OMP_PLACES` environment variable, see Section 5.5 on page 599.

3.2.29 `omp_set_affinity_format`

Summary

The `omp_set_affinity_format` routine sets the affinity format to be used on the device by setting the value of the *affinity-format-var* ICV.

1 **Format**

```
2       | void omp_set_affinity_format(char const *format);
```

C / C++

```
3       | subroutine omp_set_affinity_format(format)
```

```
4       | character(len=*) , intent(in) : :format
```

Fortran

5 **Binding**

6 When called from a sequential part of the program, the binding thread set for an
7 **omp_set_affinity_format** region is the encountering thread. When called from within any
8 explicit parallel region, the binding thread set (and binding region, if required) for the
9 **omp_set_affinity_format** region is implementation defined.

10 **Effect**

11 The effect of **omp_set_affinity_format** routine is to copy the character string specified by
12 the *format* argument into the *affinity-format-var* ICV on the current device.

13 This routine has the described effect only when called from a sequential part of the program. When
14 called from within an explicit **parallel** region, the effect of this routine is implementation
15 defined.

16 **Cross References**

- 17 • Controlling OpenMP thread affinity, see Section 2.9.2 on page 79.
- 18 • **omp_get_affinity_format** routine, see Section 3.2.30 on page 363.
- 19 • **omp_display_affinity** routine, see Section 3.2.31 on page 364.
- 20 • **omp_capture_affinity** routine, see Section 3.2.32 on page 365.
- 21 • **OMP_DISPLAY_AFFINITY** environment variable, see Section 5.13 on page 606.
- 22 • **OMP_AFFINITY_FORMAT** environment variable, see Section 5.14 on page 607.

1 3.2.30 omp_get_affinity_format

2 Summary

3 The `omp_get_affinity_format` routine returns the value of the *affinity-format-var* ICV on
4 the device.

5 Format

6 C / C++
`size_t omp_get_affinity_format(char *buffer, size_t size);`

C / C++

7 Fortran
8 `integer function omp_get_affinity_format(buffer)`
`character(len=*) , intent (out) :: buffer`

Fortran

9 Binding

10 When called from a sequential part of the program, the binding thread set for an
11 `omp_get_affinity_format` region is the encountering thread. When called from within any
12 explicit `parallel` region, the binding thread set (and binding region, if required) for the
13 `omp_get_affinity_format` region is implementation defined.

14 Effect

15 C / C++
16 The `omp_get_affinity_format` routine returns the number of characters in the
17 *affinity-format-var* ICV on the current device excluding the terminating null byte ('`\0`') and if *size*
18 is non-zero, writes the value of the *affinity-format-var* ICV on the current device to *buffer* followed
19 by a null byte. If the return value is larger or equal to *size*, the affinity format specification is
20 truncated, with the terminating null byte stored to *buffer[size-1]*. If *size* is zero, nothing is stored
and *buffer* may be `NULL`.

C / C++

Fortran

1 The `omp_get_affinity_format` routine returns the number of characters required to hold
2 the *affinity-format-var* ICV on the current device and writes the value of the *affinity-format-var*
3 ICV on the current device to *buffer*. If the return value is larger than `len(buffer)`, the affinity
4 format specification is truncated.

Fortran

Cross References

- Controlling OpenMP thread affinity, see Section 2.9.2 on page 79.
- `omp_set_affinity_format` routine, see Section 3.2.29 on page 361.
- `omp_display_affinity` routine, see Section 3.2.31 on page 364.
- `omp_capture_affinity` routine, see Section 3.2.32 on page 365.
- `OMP_DISPLAY_AFFINITY` environment variable, see Section 5.13 on page 606.
- `OMP_AFFINITY_FORMAT` environment variable, see Section 5.14 on page 607.

3.2.31 `omp_display_affinity`

Summary

13 The `omp_display_affinity` routine prints the OpenMP thread affinity information using the
14 format specification provided.
15

Format

C / C++

```
17 void omp_display_affinity(char const *format);
```

C / C++

Fortran

```
18 subroutine omp_display_affinity(format)  
19 character(len=*) , intent(in) :: format
```

Fortran

1 **Binding**

2 The binding thread set for an `omp_display_affinity` region is the encountering thread.

3 **Effect**

4 The `omp_display_affinity` routine prints the thread affinity information of the current
5 thread in the format specified by the *format* argument, followed by a *new-line*. If the *format* is
6 **NULL** (for C/C++) or a zero-length string (for Fortran and C/C++), the value of the
7 *affinity-format-var* ICV is used.

8 **Cross References**

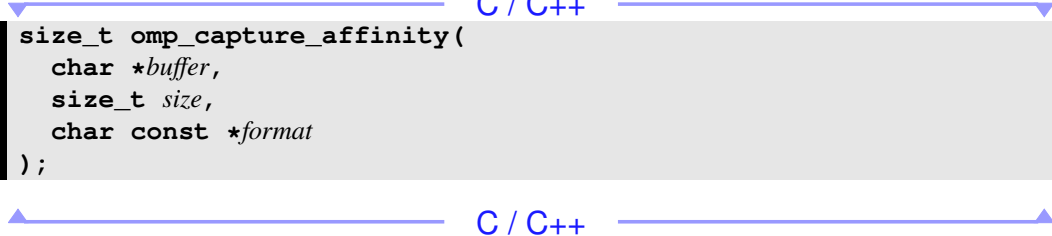
- 9 • Controlling OpenMP thread affinity, see Section [2.9.2](#) on page [79](#).
- 10 • `omp_set_affinity_format` routine, see Section [3.2.29](#) on page [361](#).
- 11 • `omp_get_affinity_format` routine, see Section [3.2.30](#) on page [363](#).
- 12 • `omp_capture_affinity` routine, see Section [3.2.32](#) on page [365](#).
- 13 • `OMP_DISPLAY_AFFINITY` environment variable, see Section [5.13](#) on page [606](#).
- 14 • `OMP_AFFINITY_FORMAT` environment variable, see Section [5.14](#) on page [607](#).

15 **3.2.32 omp_capture_affinity**

16 **Summary**

17 The `omp_capture_affinity` routine prints the OpenMP thread affinity information into a
18 buffer using the format specification provided.

19 **Format**

```
20         size_t omp_capture_affinity(  
21        char *buffer,  
22        size_t size,  
23        char const *format  
24        );
```

Fortran

```
1 integer function omp_capture_affinity(buffer,format)
2 character(len=*) ,intent(out)::buffer
3 character(len=*) ,intent(in)::format
```

Fortran

Binding

The binding thread set for an `omp_capture_affinity` region is the encountering thread.

Effect

C / C++

The `omp_capture_affinity` routine returns the number of characters in the entire thread affinity information string excluding the terminating null byte ('`\0`') and if *size* is non-zero, writes the thread affinity information of the current thread in the format specified by the *format* argument into the character string `buffer` followed by null byte. If the return value is larger or equal to *size*, the thread affinity information string is truncated, with the terminating null byte stored to `buffer[size-1]`. If *size* is zero, nothing is stored and *buffer* may be `NULL`. If the *format* is `NULL` or a zero-length string, the value of the *affinity-format-var* ICV is used.

C / C++

Fortran

The `omp_capture_affinity` routine returns the number of characters required to hold the entire thread affinity information string and prints the thread affinity information of the current thread into the character string `buffer` with the size of `len(buffer)` in the format specified by the *format* argument. If the *format* is `NULL` (for C/C++) or a zero-length string (for Fortran and C/C++), the value of the *affinity-format-var* ICV is used. If the return value is larger than `len(buffer)`, the thread affinity information string is truncated. If the *format* is a zero-length string, the value of the *affinity-format-var* ICV is used.

Fortran

Cross References

- Controlling OpenMP thread affinity, see Section 2.9.2 on page 79.
- `omp_set_affinity_format` routine, see Section 3.2.29 on page 361.
- `omp_get_affinity_format` routine, see Section 3.2.30 on page 363.
- `omp_display_affinity` routine, see Section 3.2.31 on page 364.
- `OMP_DISPLAY_AFFINITY` environment variable, see Section 5.13 on page 606.
- `OMP_AFFINITY_FORMAT` environment variable, see Section 5.14 on page 607.

3.2.33 `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 2.4 on page 47.
- `omp_get_default_device`, see Section 3.2.34 on page 368.
- `OMP_DEFAULT_DEVICE` environment variable, see Section 5.15 on page 609

3.2.34 `omp_get_default_device`

Summary

The `omp_get_default_device` routine returns the default target device.

Format

	C / C++	
	<code>int omp_get_default_device(void);</code>	
	C / C++	
	Fortran	
	<code>integer function omp_get_default_device()</code>	
	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 2.4 on page 47.
- `omp_set_default_device`, see Section 3.2.33 on page 367.
- `OMP_DEFAULT_DEVICE` environment variable, see Section 5.15 on page 609.

1 **3.2.35 omp_get_num_devices**

2 **Summary**

3 The `omp_get_num_devices` routine returns the number of target devices.

4 **Format**

```
5 | int omp_get_num_devices(void); C / C++  
6 | integer function omp_get_num_devices() Fortran
```

7 **Binding**

8 The binding task set for an `omp_get_num_devices` region is the generating task.

9 **Effect**

10 The `omp_get_num_devices` routine returns the number of available target devices. When
11 called from within a `target` region the effect of this routine is unspecified.

12 **Cross References**

13 None.

14 **3.2.36 omp_get_device_num**

15 **Summary**

16 The `omp_get_device_num` routine returns the device number of the device on which the
17 calling thread is executing.

1

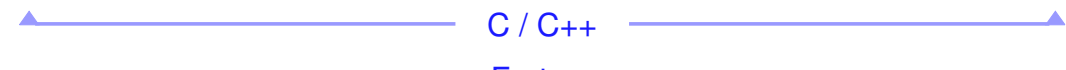
Format

2



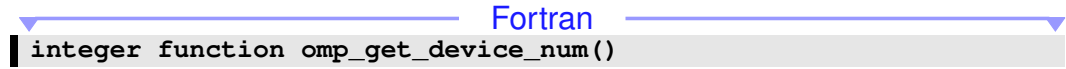
```
int omp_get_device_num(void);
```

3



```
int omp_get_device_num(void);
```

3



```
integer function omp_get_device_num()
```

4



```
integer function omp_get_device_num()
```

4

Binding

5

The binding task set for an `omp_get_devices_num` region is the generating task.

6

Effect

7

The `omp_get_device_num` routine returns the device number of the device on which the calling thread is executing. When called on the host device, it will return the same value as the `omp_get_initial_device` routine.

8

9

10

Cross References

11

- `omp_get_initial_device` routine, see Section [3.2.40](#) on page [373](#).

12

3.2.37 `omp_get_num_teams`

13


Summary


14


The `omp_get_num_teams` routine returns the number of initial teams in the current `teams` region.


15

Format


`int omp_get_num_teams(void);`


`int omp_get_num_teams(void);`


`integer function omp_get_num_teams()`


`integer function omp_get_num_teams()`

Binding

The binding task set for an `omp_get_num_teams` region is the generating task

Effect

The effect of this routine is to return the number of initial teams in the current `teams` region. The routine returns 1 if it is called from outside of a `teams` region.

Cross References


- `teams` construct, see Section [2.10](#) on page 81.


3.2.38 `omp_get_team_num`


Summary

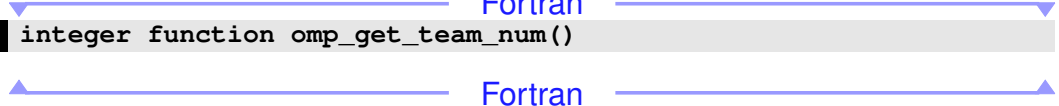
The `omp_get_team_num` routine returns the initial team number of the calling thread.

Format


`int omp_get_team_num(void);`


`int omp_get_team_num(void);`


`integer function omp_get_team_num()`


`integer function omp_get_team_num()`

1 **Binding**

2 The binding task set for an `omp_get_team_num` region is the generating task.

3 **Effect**

4 The `omp_get_team_num` routine returns the initial team number of the calling thread. The
5 initial team number is an integer between 0 and one less than the value returned by
6 `omp_get_num_teams()`, inclusive. The routine returns 0 if it is called outside of a `teams`
7 region.

8 **Cross References**

- 9 • `teams` construct, see Section 2.10 on page 81.
10 • `omp_get_num_teams` routine, see Section 3.2.37 on page 370.

11 **3.2.39** `omp_is_initial_device`

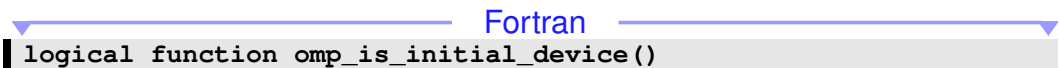
12 **Summary**

13 The `omp_is_initial_device` routine returns *true* if the current task is executing on the host
14 device; otherwise, it returns *false*.

15 **Format**

16  C / C++
`int omp_is_initial_device(void);`

17  C / C++

17  Fortran
`logical function omp_is_initial_device()`

18  Fortran

18 **Binding**

19 The binding task set for an `omp_is_initial_device` region is the generating task.

1 **Effect**

2 The effect of this routine is to return *true* if the current task is executing on the host device;
3 otherwise, it returns *false*.

4 **Cross References**

- 5 • **target** construct, see Section [2.15.5](#) on page 168

6 **3.2.40 omp_get_initial_device**

7 **Summary**

8 The **omp_get_initial_device** routine returns a device number representing the host device.

9 **Format**

10 C / C++
| `int omp_get_initial_device(void);`

11 C / C++

11 Fortran
| `integer function omp_get_initial_device()`

12 Fortran

12 **Binding**

13 The binding task set for an **omp_get_initial_device** region is the generating task.

14 **Effect**

15 The effect of this routine is to return the device number of the host device. The value of the device
16 number is implementation defined. If it is between 0 and one less than
17 **omp_get_num_devices()** then it is valid for use with all device constructs and routines; if it is
18 outside that range, then it is only valid for use with the device memory routines and not in the
19 **device** clause. When called from within a **target** region the effect of this routine is unspecified.

Cross References

- **target** construct, see Section 2.15.5 on page 168
- Device memory routines, see Section 3.6 on page 393.

3.2.41 `omp_get_max_task_priority`

Summary

The `omp_get_max_task_priority` routine returns the maximum value that can be specified in the `priority` clause.

Format

C / C++
`int omp_get_max_task_priority(void);`

C / C++

Fortran
`integer function omp_get_max_task_priority()`

Fortran

Binding

The binding thread set for an `omp_get_max_task_priority` region is all threads on the device. The effect of executing this routine is not related to any specific region corresponding to any construct or API routine.

Effect

The `omp_get_max_task_priority` routine returns the value of the `max-task-priority-var` ICV, which determines the maximum value that can be specified in the `priority` clause.

Cross References

- `max-task-priority-var`, see Section 2.4 on page 47.
- **task** construct, see Section 2.13.1 on page 133.

1 3.2.42 omp_pause_resource

2 Summary

3 The `omp_pause_resource` routine allows the runtime to relinquish resources used by OpenMP
4 on the specified device.

5 Format

C / C++

```
6 int omp_pause_resource(omp_pause_resource_t kind, int device_num  
7 );
```

C / C++

Fortran

```
8 integer function omp_pause_resource(kind, device_num)  
9 integer (kind=omp_pause_resource_kind) kind  
10 integer device_num
```

Fortran

11 Constraints on Arguments

12 The first argument passed to this routine can be one of the valid OpenMP pause kind, or any
13 implementation specific pause kind. The C/C++ header file (`omp.h`) and the Fortran include file
14 (`omp_lib.h`) and/or Fortran 90 module file (`omp_lib`) define the valid constants. The valid
15 constants must include the following, which can be extended with implementation specific values:

16 Format

C / C++

```
17 typedef enum omp_pause_resource_t {  
18     omp_pause_soft = 1,  
19     omp_pause_hard = 2  
20 } omp_pause_resource_t;
```

C / C++

Fortran

```
1 integer (kind=omp_pause_resource_kind), parameter :: &  
2   omp_pause_soft = 1  
3 integer (kind=omp_pause_resource_kind), parameter :: &  
4   omp_pause_hard = 2
```

Fortran

5 The second argument passed to this routine indicates which device is paused. The **device_num**
6 parameter must be greater than or equal to zero and less than the result of
7 **omp_get_num_devices()** or equal to the result of a call to
8 **omp_get_initial_device()**.

Binding

9 The binding task set for an **omp_pause_resource** region is the whole program.

Effect

10 The **omp_pause_resource** routine allows the runtime to relinquish resources used by OpenMP
11 on the specified device.

12 If successful, the **omp_pause_hard** value results in a hard pause for which the OpenMP state is
13 not guaranteed to persist across the **omp_pause_resource** call. Hard pause may relinquish any
14 data allocated by OpenMP on a given device, including data allocated by memory routines for that
15 device as well as data present on the device as a result of a **declare target** or **target data**
16 construct. Hard pause may also relinquish any data associated with a **threadprivate** directive.
17 When relinquished and when applicable, base language appropriate deallocation/finalization is
18 performed. When relinquished and when applicable, mapped data on a device will not be copied
19 back from the device to the host.

20 If successful, the **omp_pause_soft** value results in a soft pause for which the OpenMP state is
21 guaranteed to persist across the call, with the exception of any data associated with a
22 **threadprivate** directive which may be relinquished across the call. When relinquished and
23 when applicable, base language appropriate deallocation/finalization is performed.

24 **Note** – Hard pause may relinquish more resources, but may resume processing OpenMP regions
25 more slowly. Soft pause allows OpenMP regions to restart more quickly, but may relinquish fewer
26 resources. An OpenMP implementation will reclaim resources as needed for OpenMP regions
27 encountered after the pause region. Since a hard pause may unmap data on the specified device,
28 appropriate data mapping is required before using data on the specified device after the pause
29 region.

30 The routine returns zero in case of success, and nonzero otherwise.

1 **Tool Callbacks**

2 If the tool is not allowed to interact with the specified device after encountering this call, then the
3 runtime must call the tool finalizer for that device.

4 **Restrictions**

5 The `omp_pause_resource` routine has the following restriction:

- 6 • The routine may only be called in the sequential part and when there are no pending task waiting
7 for execution. Calling in any other circumstances may result in unspecified behavior.

8 **Cross References**

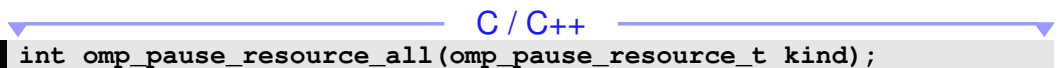
- 9 • `threadprivate` directives, see Section [2.22.2](#) on page [268](#).
- 10 • To pause resources on all devices at once, see Section [3.2.43](#) on page [377](#).

11 **3.2.43 omp_pause_resource_all**

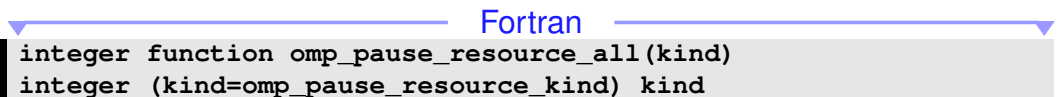
12 **Summary**

13 The `omp_pause_resource_all` routine allows the runtime to relinquish resources used by
14 OpenMP on all devices.

15 **Format**

16  `int omp_pause_resource_all(omp_pause_resource_t kind);`



17  `integer function omp_pause_resource_all(kind)
18 integer (kind=omp_pause_resource_kind) kind`



19 **Binding**

20 The binding task set for an `omp_pause_resource_all` region is the whole program.

Effect

The `omp_pause_resource_all` routine allows the runtime to relinquish resources used by OpenMP on all devices. It is equivalent to repetitively calling the `omp_pause_resource` for all of the available devices, including the host device.

The argument `kind` passed to this routine can be one of the valid OpenMP pause kind as defined in Section 3.2.42 on page 375, or any implementation specific pause kind.

Tool Callbacks

If the tool is not allowed to interact with a given device after encountering this call, then the runtime must call the tool finalizer for that device.

Restrictions

The `omp_pause_resource_all` routine has the following restriction:

- The routine may only be called in the sequential part and there are no pending task waiting for execution. Calling in any other circumstances may result in unspecified behavior.

Cross References

- To pause resources on a specific device only, see Section 3.2.42 on page 375.

3.3 Lock Routines

The OpenMP runtime library includes a set of general-purpose lock routines that can be used for synchronization. These general-purpose lock routines operate on OpenMP locks that are represented by OpenMP lock variables. OpenMP lock variables must be accessed only through the routines described in this section; programs that otherwise access OpenMP lock variables are non-conforming.

An OpenMP lock can be in one of the following states: *uninitialized*, *unlocked*, or *locked*. If a lock is in the *unlocked* state, a task can *set* the lock, which changes its state to *locked*. The task that sets the lock is then said to *own* the lock. A task that owns a lock can *unset* that lock, returning it to the *unlocked* state. A program in which a task unsets a lock that is owned by another task is non-conforming.

Two types of locks are supported: *simple locks* and *nestable locks*. A *nestable lock* can be set multiple times by the same task before being unset; a *simple lock* cannot be set if it is already owned by the task trying to set it. *Simple lock* variables are associated with *simple locks* and can

1 only be passed to *simple lock* routines. *Nestable lock* variables are associated with *nestable locks*
2 and can only be passed to *nestable lock* routines.

3 Each type of lock can also have a *synchronization hint* that contains information about the intended
4 usage of the lock by the application code. The effect of the hint is implementation defined. An
5 OpenMP implementation can use this hint to select a usage-specific lock, but hints do not change
6 the mutual exclusion semantics of locks. A conforming implementation can safely ignore the hint.

7 Constraints on the state and ownership of the lock accessed by each of the lock routines are
8 described with the routine. If these constraints are not met, the behavior of the routine is
9 unspecified.

10 The OpenMP lock routines access a lock variable such that they always read and update the most
11 current value of the lock variable. It is not necessary for an OpenMP program to include explicit
12 **flush** directives to ensure that the lock variable's value is consistent among different tasks.

13 **Binding**

14 The binding thread set for all lock routine regions is all threads in the contention group. As a
15 consequence, for each OpenMP lock, the lock routine effects relate to all tasks that call the routines,
16 without regard to which teams the threads in the contention group executing the tasks belong.

17 **Simple Lock Routines**

▼ C / C++ ▼

18 The type **omp_lock_t** represents a simple lock. For the following routines, a simple lock variable
19 must be of **omp_lock_t** type. All simple lock routines require an argument that is a pointer to a
20 variable of type **omp_lock_t**.

▲ C / C++ ▲

▼ Fortran ▼

21 For the following routines, a simple lock variable must be an integer variable of
22 **kind=omp_lock_kind**.

▲ Fortran ▲

23 The simple lock routines are as follows:

- 24 ● The **omp_init_lock** routine initializes a simple lock.
- 25 ● The **omp_init_lock_with_hint** routine initializes a simple lock and attaches a hint to it.
- 26 ● The **omp_destroy_lock** routine uninitialized a simple lock.
- 27 ● The **omp_set_lock** routine waits until a simple lock is available, and then sets it.
- 28 ● The **omp_unset_lock** routine unsets a simple lock.
- 29 ● The **omp_test_lock** routine tests a simple lock, and sets it if it is available.

Nestable Lock Routines

C / C++

The type `omp_nest_lock_t` represents a nestable lock. For the following routines, a nestable lock variable must be of `omp_nest_lock_t` type. All nestable lock routines require an argument that is a pointer to a variable of type `omp_nest_lock_t`.

C / C++

Fortran

For the following routines, a nestable lock variable must be an integer variable of `kind=omp_nest_lock_kind`.

Fortran

The nestable lock routines are as follows:

- The `omp_init_nest_lock` routine initializes a nestable lock.
- The `omp_init_nest_lock_with_hint` routine initializes a nestable lock and attaches a hint to it.
- The `omp_destroy_nest_lock` routine uninitialized a nestable lock.
- The `omp_set_nest_lock` routine waits until a nestable lock is available, and then sets it.
- The `omp_unset_nest_lock` routine unsets a nestable lock.
- The `omp_test_nest_lock` routine tests a nestable lock, and sets it if it is available

Restrictions

OpenMP lock routines have the following restrictions:

- The use of the same OpenMP lock in different contention groups results in unspecified behavior.

3.3.1 `omp_init_lock` and `omp_init_nest_lock`

Summary

These routines initialize an OpenMP lock without a hint.

1 **Format**

C / C++

```
2        void omp_init_lock(omp_lock_t *lock);  
3        void omp_init_nest_lock(omp_nest_lock_t *lock);
```

C / C++

Fortran

```
4        subroutine omp_init_lock(svar)  
5        integer (kind=omp_lock_kind) svar  
6  
7        subroutine omp_init_nest_lock(nvar)  
8        integer (kind=omp_nest_lock_kind) nvar
```

Fortran

9 **Constraints on Arguments**

10 A program that accesses a lock that is not in the uninitialized state through either routine is
11 non-conforming.

12 **Effect**

13 The effect of these routines is to initialize the lock to the unlocked state; that is, no task owns the
14 lock. In addition, the nesting count for a nestable lock is set to zero.

15 **Execution Model Events**

16 The *lock-init* or *nest-lock-init* event occurs in the thread executing a **omp_init_lock** or
17 **omp_init_nest_lock** region after initialization of the lock, but before finishing the region.

18 **Tool Callbacks**

19 A thread dispatches a registered **ompt_callback_lock_init** callback for each occurrence of
20 a *lock-init* or *nest-lock-init* event in that thread. This callback has the type signature
21 **ompt_callback_mutex_acquire_t**. The callbacks occur in the task encountering the
22 routine. The callback receives **omp_sync_hint_none** as *hint* argument and
23 **ompt_mutex_lock** or **ompt_mutex_nest_lock** as *kind* argument as appropriate.

24 **Cross References**

25 • **ompt_callback_mutex_acquire_t**, see Section [4.2.4.2.12](#) on page [458](#).

1 3.3.2 `omp_init_lock_with_hint` and 2 `omp_init_nest_lock_with_hint`

3 Summary

4 These routines initialize an OpenMP lock with a hint. The effect of the hint is
5 implementation-defined. The OpenMP implementation can ignore the hint without changing
6 program semantics.

7 Format

C / C++

```
8 void omp_init_lock_with_hint(  
9     omp_lock_t *lock,  
10    omp_sync_hint_t hint  
11 );  
12 void omp_init_nest_lock_with_hint(  
13     omp_nest_lock_t *lock,  
14     omp_sync_hint_t hint  
15 );
```

C / C++

Fortran

```
16 subroutine omp_init_lock_with_hint(svar, hint)  
17 integer (kind=omp_lock_kind) svar  
18 integer (kind=omp_sync_hint_kind) hint  
19  
20 subroutine omp_init_nest_lock_with_hint(nvar, hint)  
21 integer (kind=omp_nest_lock_kind) nvar  
22 integer (kind=omp_sync_hint_kind) hint
```

Fortran

23 Constraints on Arguments

24 A program that accesses a lock that is not in the uninitialized state through either routine is
25 non-conforming.

26 The second argument passed to these routines (*hint*) is a hint as described in Section [2.20.12](#) on
27 page [253](#).

Effect

The effect of these routines is to initialize the lock to the unlocked state and, optionally, to choose a specific lock implementation based on the hint. After initialization no task owns the lock. In addition, the nesting count for a nestable lock is set to zero.

Execution Model Events

The *lock-init* or *nest-lock-init* event occurs in the thread executing a `omp_init_lock_with_hint` or `omp_init_nest_lock_with_hint` region after initialization of the lock, but before finishing the region.

Tool Callbacks

A thread dispatches a registered `ompt_callback_lock_init` callback for each occurrence of a *lock-init* or *nest-lock-init* event in that thread. This callback has the type signature `ompt_callback_mutex_acquire_t`. The callbacks occur in the task encountering the routine. The callback receives the function's *hint* argument as *hint* argument and `ompt_mutex_lock` or `ompt_mutex_nest_lock` as *kind* argument as appropriate.

Cross References

- Synchronization Hints, see Section [2.20.12](#) on page [253](#).
- `ompt_callback_mutex_acquire_t`, see Section [4.2.4.2.12](#) on page [458](#).

3.3.3 `omp_destroy_lock` and `omp_destroy_nest_lock`

Summary

These routines ensure that the OpenMP lock is uninitialized.

Format

```
void omp_destroy_lock(omp_lock_t *lock);  
void omp_destroy_nest_lock(omp_nest_lock_t *lock);
```

C / C++

Fortran

```
1  subroutine omp_destroy_lock(svar)
2  integer (kind=omp_lock_kind) svar
3
4  subroutine omp_destroy_nest_lock(nvar)
5  integer (kind=omp_nest_lock_kind) nvar
```

Fortran

6 Constraints on Arguments

7 A program that accesses a lock that is not in the unlocked state through either routine is
8 non-conforming.

9 Effect

10 The effect of these routines is to change the state of the lock to uninitialized.

11 Execution Model Events

12 The *lock-destroy* or *nest-lock-destroy* event occurs in the thread executing a **omp_destroy_lock**
13 or **omp_destroy_nest_lock** region before finishing the region.

14 Tool Callbacks

15 A thread dispatches a registered **ompt_callback_lock_destroy** callback for each
16 occurrence of a *lock-destroy* or *nest-lock-destroy* event in that thread. This callback has the type
17 signature **ompt_callback_mutex_t**. The callbacks occur in the task encountering the routine.
18 The callbacks receive **ompt_mutex_lock** or **ompt_mutex_nest_lock** as their *kind*
19 argument as appropriate.

20 Cross References

- 21 • **ompt_callback_mutex_t**, see Section [4.2.4.2.13](#) on page [459](#).

22 3.3.4 omp_set_lock and omp_set_nest_lock

23 Summary

24 These routines provide a means of setting an OpenMP lock. The calling task region behaves as if it
25 was suspended until the lock can be set by this task.

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

Constraints on Arguments

A program that accesses a lock that is in the uninitialized state through either routine is non-conforming. A simple lock accessed by `omp_set_lock` that is in the locked state must not be owned by the task that contains the call or deadlock will result.

Effect

Each of these routines has an effect equivalent to suspension of the task executing the routine until the specified lock is available.

Note – The semantics of these routines is specified *as if* they serialize execution of the region guarded by the lock. However, implementations may implement them in other ways provided that the isolation properties are respected so that the actual execution delivers a result that could arise from some serialization.

A simple lock is available if it is unlocked. Ownership of the lock is granted to the task executing the routine.

A nestable lock is available if it is unlocked or if it is already owned by the task executing the routine. The task executing the routine is granted, or retains, ownership of the lock, and the nesting count for the lock is incremented.

1 Execution Model Events

2 The *lock-acquire* or *nest-lock-acquire* event occurs in the thread executing a `omp_set_lock` or
3 `omp_set_nest_lock` region before the associated lock is requested.

4 The *lock-acquired* or *nest-lock-acquired* event occurs in the thread executing a `omp_set_lock` or
5 `omp_set_nest_lock` region after acquiring the associated lock, if the thread did not already
6 own the lock, but before finishing the region.

7 The *nest-lock-owned* event occurs in the thread executing a `omp_set_nest_lock` region when
8 the thread already owned the lock, before finishing the region.

9 Tool Callbacks

10 A thread dispatches a registered `ompt_callback_mutex_acquire` callback for each
11 occurrence of a *lock-acquire* or *nest-lock-acquire* event in that thread. This callback has the type
12 signature `ompt_callback_mutex_acquire_t`.

13 A thread dispatches a registered `ompt_callback_mutex_acquired` callback for each
14 occurrence of a *lock-acquired* or *nest-lock-acquired* event in that thread. This callback has the type
15 signature `ompt_callback_mutex_t`.

16 A thread dispatches a registered `ompt_callback_nest_lock` callback for each occurrence of
17 a *nest-lock-owned* event in that thread. This callback has the type signature
18 `ompt_callback_nest_lock_t`. The callback receives `ompt_scope_begin` as its
19 *endpoint* argument.

20 The callbacks occur in the task encountering the lock function. The callbacks receive
21 `ompt_mutex_lock` or `ompt_mutex_nest_lock` as their *kind* argument, as appropriate.

22 Cross References

- 23 • `ompt_callback_mutex_acquire_t`, see Section [4.2.4.2.12](#) on page [458](#).
- 24 • `ompt_callback_mutex_t`, see Section [4.2.4.2.13](#) on page [459](#).
- 25 • `ompt_callback_nest_lock_t`, see Section [4.2.4.2.14](#) on page [460](#).

26 3.3.5 `omp_unset_lock` and `omp_unset_nest_lock`

27 Summary

28 These routines provide the means of unsetting an OpenMP lock.

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

Constraints on Arguments

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.

Effect

For a simple lock, the `omp_unset_lock` routine causes the lock to become unlocked.

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.

For either routine, if the lock becomes unlocked, and if one or more task regions were effectively suspended because the lock was unavailable, the effect is that one task is chosen and given ownership of the lock.

Execution Model Events

The *lock-release* or *nest-lock-release* event occurs in the thread executing a `omp_unset_lock` or `omp_unset_nest_lock` region after releasing the associated lock, but before finishing the region.

The *nest-lock-held* event occurs in the thread executing a `omp_unset_nest_lock` region when the thread still owns the lock, before finishing the region.

1 Tool Callbacks

2 A thread dispatches a registered `ompt_callback_mutex_released` callback for each
3 occurrence of a *lock-release* or *nest-lock-release* event in that thread. This callback has the type
4 signature `ompt_callback_mutex_t`. The callback occurs in the task encountering the routine.
5 The callback receives `ompt_mutex_lock` or `ompt_mutex_nest_lock` as *kind* argument as
6 appropriate.

7 A thread dispatches a registered `ompt_callback_nest_lock` callback for each occurrence of
8 a *nest-lock-held* event in that thread. This callback has the type signature
9 `ompt_callback_nest_lock_t`. The callback receives `ompt_scope_end` as its *endpoint*
10 argument.

11 Cross References

- 12 • `ompt_callback_mutex_t`, see Section 4.2.4.2.13 on page 459.
- 13 • `ompt_callback_nest_lock_t`, see Section 4.2.4.2.14 on page 460.

14 3.3.6 omp_test_lock and omp_test_nest_lock

15 Summary

16 These routines attempt to set an OpenMP lock but do not suspend execution of the task executing
17 the routine.

18 Format

19  C / C++

```
20 int omp_test_lock(omp_lock_t *lock);  
int omp_test_nest_lock(omp_nest_lock_t *lock);
```

21  C / C++

22  Fortran

```
23 logical function omp_test_lock(svar)  
24 integer (kind=omp_lock_kind) svar  
integer function omp_test_nest_lock(nvar)  
integer (kind=omp_nest_lock_kind) nvar
```

25  Fortran

1 **Constraints on Arguments**

2 A program that accesses a lock that is in the uninitialized state through either routine is
3 non-conforming. The behavior is unspecified if a simple lock accessed by **omp_test_lock** is in
4 the locked state and is owned by the task that contains the call.

5 **Effect**

6 These routines attempt to set a lock in the same manner as **omp_set_lock** and
7 **omp_set_nest_lock**, except that they do not suspend execution of the task executing the
8 routine.

9 For a simple lock, the **omp_test_lock** routine returns *true* if the lock is successfully set;
10 otherwise, it returns *false*.

11 For a nestable lock, the **omp_test_nest_lock** routine returns the new nesting count if the lock
12 is successfully set; otherwise, it returns zero.

13 **Execution Model Events**

14 The *lock-test* or *nest-lock-test* event occurs in the thread executing a **omp_test_lock** or
15 **omp_test_nest_lock** region before the associated lock is tested.

16 The *lock-test-acquired* or *nest-lock-test-acquired* event occurs in the thread executing a
17 **omp_test_lock** or **omp_test_nest_lock** region before finishing the region if the
18 associated lock was acquired and the thread did not already own the lock.

19 The *nest-lock-owned* event occurs in the thread executing a **omp_test_nest_lock** region if the
20 thread already owned the lock, before finishing the region.

21 **Tool Callbacks**

22 A thread dispatches a registered **ompt_callback_mutex_acquire** callback for each
23 occurrence of a *lock-test* or *nest-lock-test* event in that thread. This callback has the type signature
24 **ompt_callback_mutex_acquire_t**.

25 A thread dispatches a registered **ompt_callback_mutex_acquired** callback for each
26 occurrence of a *lock-test-acquired* or *nest-lock-test-acquired* event in that thread. This callback has
27 the type signature **ompt_callback_mutex_t**.

28 A thread dispatches a registered **ompt_callback_nest_lock** callback for each occurrence of
29 a *nest-lock-owned* event in that thread. This callback has the type signature
30 **ompt_callback_nest_lock_t**. The callback receives **ompt_scope_begin** as its
31 *endpoint* argument.

32 The callbacks occur in the task encountering the lock function. The callbacks receive
33 **ompt_mutex_lock** or **ompt_mutex_nest_lock** as their *kind* argument, as appropriate.

Cross References

- `ompt_callback_mutex_acquire_t`, see Section 4.2.4.2.12 on page 458.
- `ompt_callback_mutex_t`, see Section 4.2.4.2.13 on page 459.
- `ompt_callback_nest_lock_t`, see Section 4.2.4.2.14 on page 460.

3.4 Timing Routines

This section describes routines that support a portable wall clock timer.

3.4.1 `omp_get_wtime`

Summary

The `omp_get_wtime` routine returns elapsed wall clock time in seconds.

Format

C / C++
`double omp_get_wtime(void);`

C / C++

Fortran
`double precision function omp_get_wtime()`

Fortran

Binding

The binding thread set for an `omp_get_wtime` region is the encountering thread. The routine's 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 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++

```
8 double start;  
9 double end;  
10 start = omp_get_wtime();  
11 ... work to be timed ...  
12 end = omp_get_wtime();  
13 printf("Work took %f seconds\n", end - start);
```

C / C++

Fortran

```
14 DOUBLE PRECISION START, END  
15 START = omp_get_wtime()  
16 ... work to be timed ...  
17 END = omp_get_wtime()  
18 PRINT *, "Work took", END - START, "seconds"
```

Fortran

19 3.4.2 `omp_get_wtick`

20 Summary

21 The `omp_get_wtick` routine returns the precision of the timer used by `omp_get_wtime`.

1 **Format**

2 `double omp_get_wtick(void);`

3 `double precision function omp_get_wtick()`

4 **Binding**

5 The binding thread set for an `omp_get_wtick` region is the encountering thread. The routine's
6 return value is not guaranteed to be consistent across any set of threads.

7 **Effect**

8 The `omp_get_wtick` routine returns a value equal to the number of seconds between successive
9 clock ticks of the timer used by `omp_get_wtime`.

10 **3.5 Event Routines**

11 This section describes routines that support OpenMP event objects. OpenMP event objects must be
12 accessed only through routines described in this section or through the **detach** clause of the
13 **task** construct; programs that otherwise access OpenMP event objects are non-conforming.

14 **Binding**

15 The binding thread set for all event routine regions is the encountering thread.

16 **3.5.1 omp_fulfill_event**

17 **Summary**

18 This routine fulfills and destroys an OpenMP event.

Format

C / C++
`void omp_fulfill_event(omp_event_t *event_handler);`

C / C++
Fortran
Fortran
`subroutine omp_fulfill_event(event_handler)`
`integer (kind=omp_event_kind) event_handler`

Constraints on Arguments

A program that calls this routine on an event that was already fulfilled is non-conforming.

Effect

The effect of this routine is to fulfill the event associated with the event handler argument. The effect of fulfilling the event will depend on how the event was created. The event is destroyed and cannot be accessed after calling this routine, and the event handler becomes unassociated with any event.

Cross References

- `detach` clause, see Section [2.13.1](#) on page [133](#).

C / C++

3.6 Device Memory Routines

This section describes routines that support allocation of memory and management of pointers in the data environments of target devices.

3.6.1 `omp_target_alloc`

Summary

The `omp_target_alloc` routine allocates memory in a device data environment.

Format

```
void* omp_target_alloc(size_t size, int device_num);
```

Effect

The `omp_target_alloc` routine returns the device address of a storage location of *size* bytes. The storage location is dynamically allocated in the device data environment of the device specified by *device_num*, which must be greater than or equal to zero and less than the result of `omp_get_num_devices()` or the result of a call to `omp_get_initial_device()`. When called from within a **target** region the effect of this routine is unspecified.

The `omp_target_alloc` routine returns **NULL** if it cannot dynamically allocate the memory in the device data environment.

The device address returned by `omp_target_alloc` can be used in an `is_device_ptr` clause, Section 2.15.5 on page 168.

Pointer arithmetic is not supported on the device address returned by `omp_target_alloc`.

Freeing the storage returned by `omp_target_alloc` with any routine other than `omp_target_free` results in unspecified behavior.

Execution Model Events

The *target-data-allocation* event occurs when a thread allocates data on a target device.

Tool Callbacks

A thread invokes a registered `ompt_callback_target_data_op` callback for each occurrence of a *target-data-allocation* event in that thread. The callback occurs in the context of the target task. The callback has type signature `ompt_callback_target_data_op_t`.

Cross References

- **target** construct, see Section 2.15.5 on page 168
- `omp_get_num_devices` routine, see Section 3.2.35 on page 369
- `omp_get_initial_device` routine, see Section 3.2.40 on page 373
- `omp_target_free` routine, see Section 3.6.2 on page 395
- `ompt_callback_target_data_op_t`, see Section 4.2.4.2.21 on page 468.

1 3.6.2 `omp_target_free`

2 Summary

3 The `omp_target_free` routine frees the device memory allocated by the
4 `omp_target_alloc` routine.

5 Format

```
6 void omp_target_free(void *device_ptr, int device_num);
```

7 Constraints on Arguments

8 A program that calls `omp_target_free` with a non-NULL pointer that does not have a value
9 returned from `omp_target_alloc` is non-conforming. The `device_num` must be greater than or
10 equal to zero and less than the result of `omp_get_num_devices()` or the result of a call to
11 `omp_get_initial_device()`.

12 Effect

13 The `omp_target_free` routine frees the memory in the device data environment associated
14 with `device_ptr`. If `device_ptr` is NULL, the operation is ignored.

15 Synchronization must be inserted to ensure that all accesses to `device_ptr` are completed before the
16 call to `omp_target_free`.

17 When called from within a `target` region the effect of this routine is unspecified.

18 Execution Model Events

19 The *target-data-free* event occurs when a thread frees data on a target device.

20 Tool Callbacks

21 A thread invokes a registered `ompt_callback_target_data_op` callback for each
22 occurrence of a *target-data-free* event in that thread. The callback occurs in the context of the target
23 task. The callback has type signature `ompt_callback_target_data_op_t`.

24 Cross References

- 25 • `target` construct, see Section [2.15.5](#) on page 168
- 26 • `omp_get_num_devices` routine, see Section [3.2.35](#) on page 369
- 27 • `omp_get_initial_device` routine, see Section [3.2.40](#) on page 373
- 28 • `omp_target_alloc` routine, see Section [3.6.1](#) on page 393
- 29 • `ompt_callback_target_data_op_t`, see Section [4.2.4.2.21](#) on page 468.

1 3.6.3 `omp_target_is_present`

2 Summary

3 The `omp_target_is_present` routine tests whether a host pointer has corresponding storage
4 on a given device.

5 Format

```
6 int omp_target_is_present(const void *ptr, int device_num);
```

7 Constraints on Arguments

8 The value of *ptr* must be a valid host pointer or **NULL**. The *device_num* must be greater than or
9 equal to zero and less than the result of `omp_get_num_devices()` or the result of a call to
10 `omp_get_initial_device()`.

11 Effect

12 This routine returns non-zero if the specified pointer would be found present on device *device_num*
13 by a **map** clause; otherwise, it returns zero.

14 When called from within a **target** region the effect of this routine is unspecified.

15 Cross References

- 16 • **target** construct, see Section [2.15.5](#) on page [168](#).
- 17 • **map** clause, see Section [2.22.7.1](#) on page [307](#).
- 18 • `omp_get_num_devices` routine, see Section [3.2.35](#) on page [369](#)
- 19 • `omp_get_initial_device` routine, see Section [3.2.40](#) on page [373](#)

20 3.6.4 `omp_target_memcpy`

21 Summary

22 The `omp_target_memcpy` routine copies memory between any combination of host and device
23 pointers.

Format

```

1  int omp_target_memcpy (
2      void *dst,
3      const void *src,
4      size_t length,
5      size_t dst_offset,
6      size_t src_offset,
7      int dst_device_num,
8      int src_device_num
9  );
10

```

Constraints on Arguments

Each device must be compatible with the device pointer specified on the same side of the copy. The *dst_device_num* and *src_device_num* must be greater than or equal to zero and less than the result of `omp_get_num_devices()` or equal to the result of a call to `omp_get_initial_device()`.

Effect

length bytes of memory at offset *src_offset* from *src* in the device data environment of device *src_device_num* are copied to *dst* starting at offset *dst_offset* in the device data environment of device *dst_device_num*. The return value is zero on success and non-zero on failure. The host device and host device data environment can be referenced with the device number returned by `omp_get_initial_device`. This routine contains a task scheduling point.

When called from within a **target** region the effect of this routine is unspecified.

Execution Model Events

The *target-data-op* event occurs when a thread transfers data on a target device.

Tool Callbacks

A thread invokes a registered `ompt_callback_target_data_op` callback for each occurrence of a *target-data-op* event in that thread. The callback occurs in the context of the target task. The callback has type signature `ompt_callback_target_data_op_t`.

Cross References

- **target** construct, see Section 2.15.5 on page 168.
- `omp_get_initial_device` routine, see Section 3.2.40 on page 373
- `omp_target_alloc` routine, see Section 3.6.1 on page 393.
- `ompt_callback_target_data_op_t`, see Section 4.2.4.2.21 on page 468.

1 3.6.5 `omp_target_memcpy_rect`

2 Summary

3 The `omp_target_memcpy_rect` routine copies a rectangular subvolume from a
4 multi-dimensional array to another multi-dimensional array. The copies can use any combination of
5 host and device pointers.

6 Format

```
7 int omp_target_memcpy_rect (
8     void *dst,
9     const void *src,
10    size_t element_size,
11    int num_dims,
12    const size_t *volume,
13    const size_t *dst_offsets,
14    const size_t *src_offsets,
15    const size_t *dst_dimensions,
16    const size_t *src_dimensions,
17    int dst_device_num,
18    int src_device_num
19 );
```

20 Constraints on Arguments

21 The length of the offset and dimension arrays must be at least the value of `num_dims`. The
22 `dst_device_num` and `src_device_num` must be greater than or equal to zero and less than
23 the result of `omp_get_num_devices()` or equal to the result of a call to
24 `omp_get_initial_device()`.

25 The value of `num_dims` must be between 1 and the implementation-defined limit, which must be at
26 least three.

27 Effect

28 This routine copies a rectangular subvolume of `src`, in the device data environment of device
29 `src_device_num`, to `dst`, in the device data environment of device `dst_device_num`. The volume is
30 specified in terms of the size of an element, number of dimensions, and constant arrays of length
31 `num_dims`. The maximum number of dimensions supported is at least three, support for higher
32 dimensionality is implementation defined. The volume array specifies the length, in number of
33 elements, to copy in each dimension from `src` to `dst`. The `dst_offsets` (`src_offsets`) parameter
34 specifies number of elements from the origin of `dst` (`src`) in elements. The `dst_dimensions`
35 (`src_dimensions`) parameter specifies the length of each dimension of `dst` (`src`)

1 The routine returns zero if successful. If both *dst* and *src* are **NULL** pointers, the routine returns the
 2 number of dimensions supported by the implementation for the specified device numbers. The host
 3 device and host device data environment can be referenced with the device number returned by
 4 **omp_get_initial_device**. Otherwise, it returns a non-zero value. The routine contains a
 5 task scheduling point.

6 When called from within a **target** region the effect of this routine is unspecified.

7 Execution Model Events

8 The *target-data-op* event occurs when a thread transfers data on a target device.

9 Tool Callbacks

10 A thread invokes a registered **ompt_callback_target_data_op** callback for each
 11 occurrence of a *target-data-op* event in that thread. The callback occurs in the context of the target
 12 task. The callback has type signature **ompt_callback_target_data_op_t**.

13 Cross References

- 14 • **target** construct, see Section 2.15.5 on page 168.
- 15 • **omp_get_initial_device** routine, see Section 3.2.40 on page 373
- 16 • **omp_target_alloc** routine, see Section 3.6.1 on page 393.
- 17 • **ompt_callback_target_data_op_t**, see Section 4.2.4.2.21 on page 468.

18 3.6.6 omp_target_associate_ptr

19 Summary

20 The **omp_target_associate_ptr** routine maps a device pointer, which may be returned
 21 from **omp_target_alloc** or implementation-defined runtime routines, to a host pointer.

22 Format

```
23 int omp_target_associate_ptr(  
24     const void *host_ptr,  
25     const void *device_ptr,  
26     size_t size,  
27     size_t device_offset,  
28     int device_num  
29 );
```

1 Constraints on Arguments

2 The value of *device_ptr* value must be a valid pointer to device memory for the device denoted by
3 the value of *device_num*. The *device_num* argument must be greater than or equal to zero and less
4 than the result of `omp_get_num_devices()` or equal to the result of a call to
5 `omp_get_initial_device()`.

6 Effect

7 The `omp_target_associate_ptr` routine associates a device pointer in the device data
8 environment of device *device_num* with a host pointer such that when the host pointer appears in a
9 subsequent `map` clause, the associated device pointer is used as the target for data motion
10 associated with that host pointer. The *device_offset* parameter specifies what offset into *device_ptr*
11 will be used as the base address for the device side of the mapping. The reference count of the
12 resulting mapping will be infinite. After being successfully associated, the buffer pointed to by the
13 device pointer is invalidated and accessing data directly through the device pointer results in
14 unspecified behavior. The pointer can be retrieved for other uses by disassociating it. When called
15 from within a `target` region the effect of this routine is unspecified.

16 The routine returns zero if successful. Otherwise it returns a non-zero value.

17 Only one device buffer can be associated with a given host pointer value and device number pair.
18 Attempting to associate a second buffer will return non-zero. Associating the same pair of pointers
19 on the same device with the same offset has no effect and returns zero. Associating pointers that
20 share underlying storage will result in unspecified behavior. The `omp_target_is_present`
21 region can be used to test whether a given host pointer has a corresponding variable in the device
22 data environment.

23 Execution Model Events

24 The *target-data-associate* event occurs when a thread associates data on a target device.

25 Tool Callbacks

26 A thread invokes a registered `ompt_callback_target_data_op` callback for each
27 occurrence of a *target-data-associate* event in that thread. The callback occurs in the context of the
28 target task. The callback has type signature `ompt_callback_target_data_op_t`.

Cross References

- `target` construct, see Section 2.15.5 on page 168.
- `map` clause, see Section 2.22.7.1 on page 307.
- `omp_target_alloc` routine, see Section 3.6.1 on page 393.
- `omp_target_disassociate_ptr` routine, see Section 3.6.6 on page 399
- `ompt_callback_target_data_op_t`, see Section 4.2.4.2.21 on page 468.

3.6.7 `omp_target_disassociate_ptr`

Summary

The `omp_target_disassociate_ptr` removes the associated pointer for a given device from a host pointer.

Format

```
int omp_target_disassociate_ptr(const void *ptr, int device_num);
```

Constraints on Arguments

The `device_num` must be greater than or equal to zero and less than the result of `omp_get_num_devices()` or equal to the result of a call to `omp_get_initial_device()`.

Effect

The `omp_target_disassociate_ptr` removes the associated device data on device `device_num` from the presence table for host pointer `ptr`. A call to this routine on a pointer that is not `NULL` and does not have associated data on the given device results in unspecified behavior. The reference count of the mapping is reduced to zero, regardless of its current value.

When called from within a `target` region the effect of this routine is unspecified.

The routine returns zero if successful. Otherwise it returns a non-zero value.

After a call to `omp_target_disassociate_ptr`, the contents of the device buffer are invalidated.

Execution Model Events

The `target-data-disassociate` event occurs when a thread disassociates data on a target device.

1
2
3
4

5
6
7
8

Tool Callbacks

A thread invokes a registered `ompt_callback_target_data_op` callback for each occurrence of a *target-data-disassociate* event in that thread. The callback occurs in the context of the target task. The callback has type signature `ompt_callback_target_data_op_t`.

Cross References

- `target` construct, see Section [2.15.5](#) on page [168](#)
- `omp_target_associate_ptr` routine, see Section [3.6.6](#) on page [399](#)
- `ompt_callback_target_data_op_t`, see Section [4.2.4.2.21](#) on page [468](#).



1 3.7 Memory Management Routines

2 This section describes routines that support memory management on the current device.

3 Instances of memory management types must be accessed only through the routines described in
4 this section; programs that otherwise access instances of these types are non-conforming.

5 3.7.1 Memory Management Types

6 The following type definitions are used by the memory management routines:

C / C++

```
7 typedef enum {  
8     OMP_ATK_THREADMODEL = 1,  
9     OMP_ATK_ALIGNMENT = 2,  
10    OMP_ATK_ACCESS = 3,  
11    OMP_ATK_POOL_SIZE = 4,  
12    OMP_ATK_FALLBACK = 5,  
13    OMP_ATK_FB_DATA = 6,  
14    OMP_ATK_PINNED = 7,  
15    OMP_ATK_PARTITION = 8  
16 } omp_alloctrail_key_t;  
17  
18 typedef enum {  
19     OMP_ATV_FALSE = 0,  
20     OMP_ATV_TRUE = 1,  
21     OMP_ATV_DEFAULT = 2,  
22     OMP_ATV_CONTENTENDED = 3,  
23     OMP_ATV_UNCONTENTENDED = 4,  
24     OMP_ATV_SEQUENTIAL = 5,  
25     OMP_ATV_PRIVATE = 6,  
26     OMP_ATV_ALL = 7,  
27     OMP_ATV_THREAD = 8,  
28     OMP_ATV_PTEAM = 9,  
29     OMP_ATV_CGROUP = 10,  
30     OMP_ATV_DEFAULT_MEM_FB = 11,  
31     OMP_ATV_NULL_FB = 12,  
32     OMP_ATV_ABORT_FB = 13,  
33     OMP_ATV_ALLOCATOR_FB = 14,  
34     OMP_ATV_ENVIRONMENT = 15,  
35     OMP_ATV_NEAREST = 16,
```

```

1   OMP_ATV_BLOCKED = 17,
2   OMP_ATV_INTERLEAVED = 18
3   } omp_alloctrail_value_t;
4
5   typedef struct {
6       omp_alloctrail_key_t key;
7       omp_uintptr_t value;
8   } omp_alloctrail_t;
9
10  enum { OMP_NULL_ALLOCATOR = NULL };

```

← C / C++ →

← Fortran →

```

11
12  integer(kind=omp_alloctrail_key_kind), &
13      parameter :: omp_atk_threadmodel = 1
14  integer(kind=omp_alloctrail_key_kind), &
15      parameter :: omp_atk_alignment = 2
16  integer(kind=omp_alloctrail_key_kind), &
17      parameter :: omp_atk_access = 3
18  integer(kind=omp_alloctrail_key_kind), &
19      parameter :: omp_atk_pool_size = 4
20  integer(kind=omp_alloctrail_key_kind), &
21      parameter :: omp_atk_fallback = 5
22  integer(kind=omp_alloctrail_key_kind), &
23      parameter :: omp_atk_fb_data = 6
24  integer(kind=omp_alloctrail_key_kind), &
25      parameter :: omp_atk_pinned = 7
26  integer(kind=omp_alloctrail_key_kind), &
27      parameter :: omp_atk_partition = 8
28
29  integer(kind=omp_alloctrail_val_kind), &
30      parameter :: omp_atv_false = 0           ! Reserved for
31      future use
32  integer(kind=omp_alloctrail_val_kind), &
33      parameter :: omp_atv_true = 1           ! Reserved for
34      future use
35  integer(kind=omp_alloctrail_val_kind), &
36      parameter :: omp_atv_default = 2
37  integer(kind=omp_alloctrail_val_kind), &
38      parameter :: omp_atv_contended = 3
39  integer(kind=omp_alloctrail_val_kind), &
40      parameter :: omp_atv_uncontended = 4

```

```

1  integer(kind=omp_alloctrait_val_kind), &
2     parameter :: omp_atv_sequential = 5
3  integer(kind=omp_alloctrait_val_kind), &
4     parameter :: omp_atv_private = 6
5  integer(kind=omp_alloctrait_val_kind), &
6     parameter :: omp_atv_all = 7
7  integer(kind=omp_alloctrait_val_kind), &
8     parameter :: omp_atv_thread = 8
9  integer(kind=omp_alloctrait_val_kind), &
10     parameter :: omp_atv_pteam = 9
11 integer(kind=omp_alloctrait_val_kind), &
12     parameter :: omp_atv_cgroup = 10
13 integer(kind=omp_alloctrait_val_kind), &
14     parameter :: omp_atv_default_mem_fb = 11
15 integer(kind=omp_alloctrait_val_kind), &
16     parameter :: omp_atv_null_fb = 12
17 integer(kind=omp_alloctrait_val_kind), &
18     parameter :: omp_atv_abort_fb = 13
19 integer(kind=omp_alloctrait_val_kind), &
20     parameter :: omp_atv_allocator_fb = 14
21 integer(kind=omp_alloctrait_val_kind), &
22     parameter :: omp_atv_environment = 15
23 integer(kind=omp_alloctrait_val_kind), &
24     parameter :: omp_atv_nearest = 16
25 integer(kind=omp_alloctrait_val_kind), &
26     parameter :: omp_atv_blocked = 17
27 integer(kind=omp_alloctrait_val_kind), &
28     parameter :: omp_atv_interleaved = 18
29
30 type omp_alloctrait
31     integer(kind=omp_alloctrait_key_kind) key
32     integer(kind=omp_alloctrait_val_kind) value
33 end type omp_alloctrait
34
35 integer(kind=omp_allocator_kind), &
36     parameter :: omp_null_allocator = 0

```

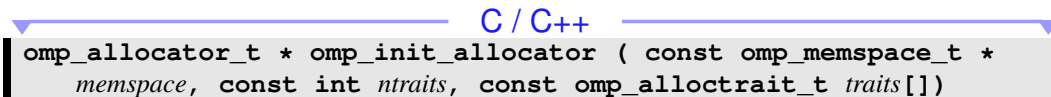
Fortran

1 3.7.2 omp_init_allocator

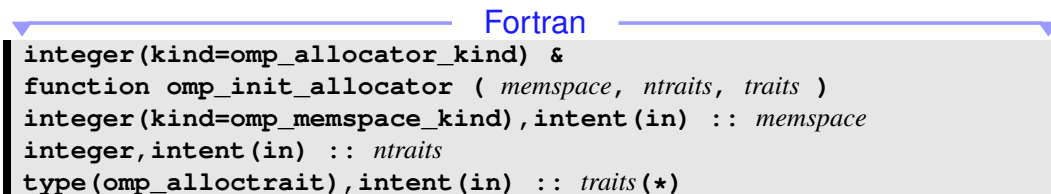
2 Summary

3 The `omp_init_allocator` routine initializes an allocator and associates it with a memory
4 space.

5 Format

6  C / C++
7 `omp_allocator_t * omp_init_allocator (const omp_memspace_t *
 memspace, const int ntraits, const omp_alloctrail_t traits[])`

8  C / C++

9  Fortran
10 `integer(kind=omp_allocator_kind) &
11 function omp_init_allocator (memspace, ntraits, traits)
12 integer(kind=omp_memspace_kind),intent(in) :: memspace
integer,intent(in) :: ntraits
type(omp_alloctrail),intent(in) :: traits(*)`

13  Fortran

13 Constraints on Arguments

14 The *memspace* argument must be one of the predefined memory spaces defined in Table 2.7.

15 If the *ntraits* argument is greater than zero, then there must be at least as many traits specified in the
16 *traits* argument. If there are fewer than *ntraits* traits the behavior is unspecified.

17 Unless a **requires** directive with the **dynamic_allocators** clause is present in the same
18 compilation unit, using this routine in a **target** region results in unspecified behavior.

19 Binding

20 The binding thread set for an `omp_init_allocator` region is all threads on a device. The
21 effect of executing this routine is not related to any specific region corresponding to any construct
22 or API routine.

Effect

The `omp_init_allocator` routine creates a new allocator that is associated with the *memspace* memory space. The allocations done through the created allocator will behave according to the allocator traits specified in the *traits* argument. The number of traits in the *traits* argument is specified by the *ntraits* argument. Specifying the same allocator trait more than once results in unspecified behavior. The routine returns a handle for the created allocator. If the special `OMP_ATV_DEFAULT` value is used for a given trait, then its value will be the default value specified in Table 2.8 for that given trait.

If *memspace* is `omp_default_mem_space` and the *traits* argument is an empty set this routine will always return a handle to an allocator. Otherwise if an allocator based on the requirements cannot be created then the special value `OMP_NULL_ALLOCATOR` is returned.

The use of an allocator returned by this routine on a device other than the one on which it was created results in unspecified behavior.

Cross References

- Memory spaces in Section 2.14.1 on page 150
- Allocators in Section 2.14.2 on page 151

3.7.3 `omp_destroy_allocator`

Summary

The `omp_destroy_allocator` routine releases all resources used by the allocator handle.

Format

C / C++
`void omp_destroy_allocator (omp_allocator_t * allocator);`

C / C++

Fortran

Fortran
`subroutine omp_destroy_allocator (allocator)
integer (kind=omp_allocator_kind), intent(in) :: allocator`

Fortran

Constraints on Arguments

The *allocator* argument must not be a predefined memory allocator.

Unless a **requires** directive with the **dynamic_allocators** clause is present in the same compilation unit, using this routine in a **target** region results in unspecified behavior.

Binding

The binding thread set for an **omp_destroy_allocator** region is all threads on a device. The effect of executing this routine is not related to any specific region corresponding to any construct or API routine.

Effect

The **omp_destroy_allocator** routine releases all resources used to implement the *allocator* handle. Accessing any memory allocated by the *allocator* after this call results in undefined behavior.

If *allocator* is **OMP_NULL_ALLOCATOR** then this routine will have no effect.

Cross References

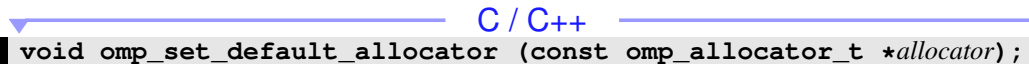
- Allocators in Section [2.14.2](#) on page [151](#)

3.7.4 **omp_set_default_allocator**

Summary

The **omp_set_default_allocator** routine sets the default memory allocator to be used by allocation calls, **allocate** directives and **allocate** clauses that do not specify an allocator.

Format


`void omp_set_default_allocator (const omp_allocator_t *allocator);`

C / C++

Fortran

```
1 subroutine omp_set_default_allocator ( allocator )  
2 integer (kind=omp_allocator_kind), intent(in) :: allocator
```

Fortran

Constraints on Arguments

The *allocator* argument must point to a valid memory allocator.

Binding

The binding task set for an **omp_set_default_allocator** region is the binding implicit task.

Effect

The effect of this routine is to set the value of the *def-allocator-var* ICV of the binding implicit task to the value specified in the *allocator* argument.

Cross References

- *def-allocator-var* ICV, see Section 2.4 on page 47.
- Memory Allocators, see Section 2.14.2 on page 151.
- **omp_alloc** routine, see Section 3.7.6 on page 410.

3.7.5 **omp_get_default_allocator**

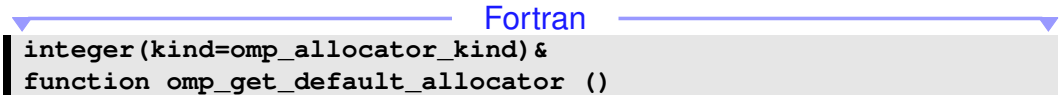
Summary

The **omp_get_default_allocator** routine returns the memory allocator to be used by allocation calls, **allocate** directives and **allocate** clauses that do not specify an allocator.

1 **Format**

2  `const omp_allocator_t * omp_get_default_allocator (void);`

 C / C++

3  `integer(kind=omp_allocator_kind) &
4 function omp_get_default_allocator ()`

 Fortran

5 **Binding**

6 The binding task set for an `omp_get_default_allocator` region is the binding implicit task.

7 **Effect**

8 The effect of this routine is to return the value of the *def-allocator-var* ICV of the binding implicit
9 task.

10 **Cross References**

- 11 • *def-allocator-var* ICV, see Section 2.4 on page 47.
- 12 • Memory Allocators, see Section 2.14.2 on page 151.
- 13 • `omp_alloc` routine, see Section 3.7.6 on page 410.

 C / C++

14 **3.7.6 omp_alloc**

15 **Summary**

16 The `omp_alloc` routine requests a memory allocation from a memory allocator.

Format

C

```
void * omp_alloc (size_t size, const omp_allocator_t *allocator);
```

C

C++

```
void * omp_alloc (
    size_t size,
    const omp_allocator_t *allocator=OMP_NULL_ALLOCATOR
);
```

C++

Constraints on Arguments

For `omp_alloc` invocations appearing in `target` regions the `allocator` argument cannot be `OMP_NULL_ALLOCATOR` and it must be an expression must evaluable by the compiler.

Effect

The `omp_alloc` routine requests a memory allocation of `size` bytes from the specified memory allocator. If the `allocator` argument is `OMP_NULL_ALLOCATOR` the memory allocator used by the routine will be the one specified by the `def-allocator-var` ICV of the binding implicit task. Upon success it returns a pointer to the allocated memory. Otherwise, the behavior specified by the `fallback` trait will be followed.

Cross References

- Memory allocators, see Section [2.14.2](#) on page [151](#).

3.7.7 `omp_free`

Summary

The `omp_free` routine deallocates previously allocated memory.

1

Format

2

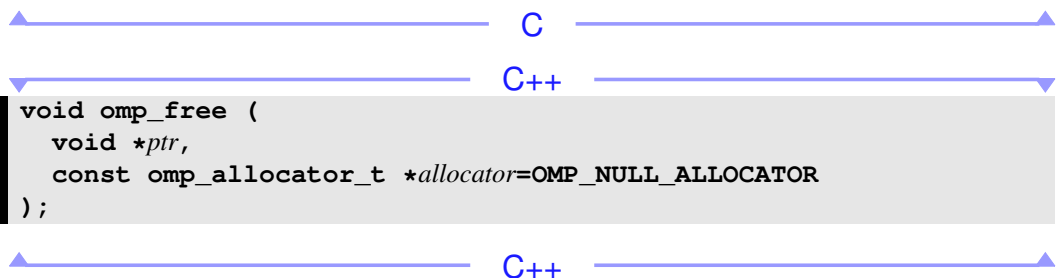

 A blue double-headed arrow spans the width of the code block. The text 'C' is centered above the arrow. The code is: `void omp_free (void *ptr, const omp_allocator_t *allocator);`

3

4

5

6


 A blue double-headed arrow spans the width of the code block. The text 'C++' is centered above the arrow. The code is: `void omp_free (
 void *ptr,
 const omp_allocator_t *allocator=OMP_NULL_ALLOCATOR
);`

7

Effect

8

9

10

11

12

13

14

The `omp_free` routine deallocates the memory to which `ptr` points. The `ptr` argument must point to memory previously allocated with a memory allocator. If the `allocator` argument is specified it must be the memory allocator to which the allocation request was made. If the `allocator` argument is `OMP_NULL_ALLOCATOR` the implementation will find the memory allocator used to allocate the memory. Using `omp_free` on memory that was already deallocated or that was allocated by an allocator that has already been destroyed with `omp_destroy_allocator` results in unspecified behavior.

15

Cross References

16

- Memory allocators, see Section [2.14.2](#) on page 151.


 A blue double-headed arrow spans the width of the code block. The text 'C / C++' is centered above the arrow.

1 3.8 Tool Control Routines

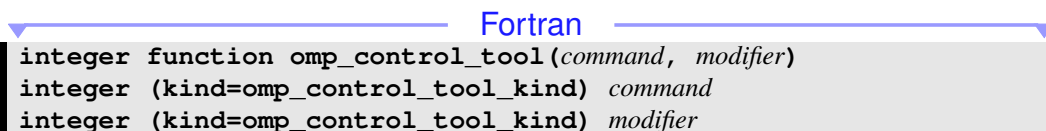
2 Summary

3 The `omp_control_tool` routine enables a program to pass commands to an active tool.

4 Format

5  `int omp_control_tool(int command, int modifier, void *arg);`

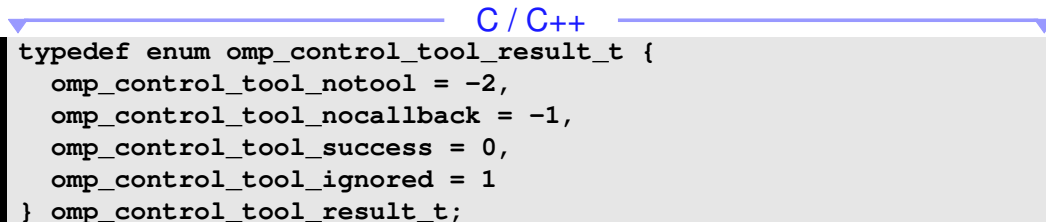


6  `integer function omp_control_tool(command, modifier)`
7 `integer (kind=omp_control_tool_kind) command`
8 `integer (kind=omp_control_tool_kind) modifier`



9 Description

10 An OpenMP program may use `omp_control_tool` to pass commands to a tool. Using
11 `omp_control_tool`, an application can request that a tool start or restart data collection when a
12 code region of interest is encountered, pause data collection when leaving the region of interest,
13 flush any data that it has collected so far, or end data collection. Additionally,
14 `omp_control_tool` can be used to pass tool-specific commands to a particular tool.

15  `typedef enum omp_control_tool_result_t {`
16 `omp_control_tool_notool = -2,`
17 `omp_control_tool_nocallback = -1,`
18 `omp_control_tool_success = 0,`
19 `omp_control_tool_ignored = 1`
20 `} omp_control_tool_result_t;`



Fortran

```
1 integer (kind=omp_control_tool_result_kind), &  
2     parameter :: omp_control_tool_notool = -2  
3 integer (kind=omp_control_tool_result_kind), &  
4     parameter :: omp_control_tool_nocallback = -1  
5 integer (kind=omp_control_tool_result_kind), &  
6     parameter :: omp_control_tool_success = 0  
7 integer (kind=omp_control_tool_result_kind), &  
8     parameter :: omp_control_tool_ignored = 1
```

Fortran

9 If no tool is active, the OpenMP implementation will return `omp_control_tool_notool`. If a
10 tool is active, but it has not registered a callback for the *tool-control* event, the OpenMP
11 implementation will return `omp_control_tool_nocallback`. An OpenMP implementation
12 may return other implementation-defined negative values < -64 ; an application may assume that
13 any negative return value indicates that a tool has not received the command. A return value of
14 `omp_control_tool_success` indicates that the tool has performed the specified command.
15 A return value of `omp_control_tool_ignored` indicates that the tool has ignored the
16 specified command. A tool may return other positive values > 64 that are tool-defined.

Constraints on Arguments

17 The following enumeration type defines four standard commands. Table 3.1 describes the actions
18 that these commands request from a tool.
19

C / C++

```
20 typedef enum omp_control_tool_t {  
21     omp_control_tool_start = 1,  
22     omp_control_tool_pause = 2,  
23     omp_control_tool_flush = 3,  
24     omp_control_tool_end = 4  
25 } omp_control_tool_t;
```

C / C++

Fortran

```
1 integer (kind=omp_control_tool_kind), &  
2     parameter :: omp_control_tool_start = 1  
3 integer (kind=omp_control_tool_kind), &  
4     parameter :: omp_control_tool_pause = 2  
5 integer (kind=omp_control_tool_kind), &  
6     parameter :: omp_control_tool_flush = 3  
7 integer (kind=omp_control_tool_kind), &  
8     parameter :: omp_control_tool_end = 4
```

Fortran

9 Tool-specific values for *command* must be ≥ 64 . Tools must ignore *command* values that they are
10 not explicitly designed to handle. Other values accepted by a tool for *command*, and any values for
11 *modifier* and *arg* are tool-defined.

TABLE 3.1: Standard tool control commands.

Command	Action
<code>omp_control_tool_start</code>	Start or restart monitoring if it is off. If monitoring is already on, this command is idempotent. If monitoring has already been turned off permanently, this command will have no effect.
<code>omp_control_tool_pause</code>	Temporarily turn monitoring off. If monitoring is already off, it is idempotent.
<code>omp_control_tool_flush</code>	Flush any data buffered by a tool. This command may be applied whether monitoring is on or off.
<code>omp_control_tool_end</code>	Turn monitoring off permanently; the tool finalizes itself and flushes all output.

12 Execution Model Events

13 The *tool-control* event occurs in the thread encountering a call to `omp_control_tool` at a point
14 inside its corresponding OpenMP region.

1 **Tool Callbacks**

2 A thread dispatches a registered `ompt_callback_control_tool` callback for each
3 occurrence of a *tool-control* event. The callback executes in the context of the call that occurs in the
4 user program. This callback has type signature `ompt_callback_control_tool_t`. The
5 callback may return any non-negative value, which will be returned to the application by the
6 OpenMP implementation as the return value of the `omp_control_tool` call that triggered the
7 callback.

8 Arguments passed to the callback are those passed by the user to `omp_control_tool`. If the
9 call is made in Fortran, the tool will be passed a **NULL** as the third argument to the callback. If any
10 of the four standard commands is presented to a tool, the tool will ignore the *modifier* and *arg*
11 argument values.

12 **Cross References**

- 13 • Tool Interface, see Chapter 4 on page 417
- 14 • `ompt_callback_control_tool_t`, see Section 4.2.4.2.26 on page 476

Tool Support

3 This chapter describes OMPT and OMPD, which are a pair of interfaces for first-party and
4 third-party tools, respectively. Section 4.2 describes OMPT—an interface for first-party tools. The
5 section begins with a description of how to initialize (Section 4.2.1) and finalize (Section 4.2.2) a
6 tool. Subsequent sections describe details of the interface, including data types shared between an
7 OpenMP implementation and a tool (Section 4.2.3), an interface that enables an OpenMP
8 implementation to determine that a tool is available (Section 4.2.1.1), type signatures for tool
9 callbacks that an OpenMP implementation may dispatch for OpenMP events (Section 4.2.4), and
10 *runtime entry points*—function interfaces provided by an OpenMP implementation for use by a tool
11 (Section 4.2.5).

12 Section 4.3 describes OMPD—an interface for third-party tools such as debuggers. Unlike OMPT,
13 a third-party tool exists in a separate process from the OpenMP program. An OpenMP
14 implementation need not maintain any extra information to support OMPD inquiries from
15 third-party tools *unless* it is explicitly instructed to do so. Section 4.3.1 discusses the mechanisms
16 for activating support for OMPD in the OpenMP runtime. Section 4.3.2 describes the data types
17 shared between the OMPD library and a third-party tool. Section 4.3 describes the API provided by
18 the OMPD library for use by a third-party tool. An OMPD library will not interact directly with the
19 OpenMP runtime for which it is designed to operate. Instead, the third-party tool must provide the
20 OMPD library with a set of callbacks that the OMPD library uses to access the OpenMP runtime.
21 This interface is given in Section 4.3. In general, a third-party’s tool’s OpenMP-related activity will
22 be conducted through the OMPD interface. However, there are a few instances where the third-party
23 tool needs to access the OpenMP runtime directly; these cases are discussed in Section 4.3.5.

1 4.1 Tool Interfaces Definitions

▼ C / C++ ▼

2 A compliant implementation must supply a set of definitions for the OMPT runtime entry points,
3 OMPT callback signatures, OMPD runtime entry points, OMPD tool callback signatures, OMPD
4 tool interface routines, and the special data types of their parameters and return values.

5 The set of definitions is provided in a header file named `omp-tools.h` and must contain a
6 declaration for each of the types defined in Sections 4.2.3 - 4.2.5 and 4.3.2 - 4.3.5.

7 In addition, the set of definitions may specify other implementation specific values.

8 The `omp_start_tool` function, the `ompd_dll_locations` function, all OMPD tool
9 interface functions, and all OMPD runtime entry points are external functions with “C” linkage.

▲ C / C++ ▲

10 4.2 OMPT

11 The OMPT interface defines mechanisms for initializing a tool, exploring the details of an OpenMP
12 implementation, examining OpenMP state associated with an OpenMP thread, interpreting an
13 OpenMP thread’s call stack, receiving notification about OpenMP *events*, tracing activity on
14 OpenMP target devices, and controlling a tool from an OpenMP application.

15 4.2.1 Activating an OMPT Tool

16 There are three steps to activating a tool. First, an OpenMP implementation determines whether a
17 tool should be initialized. If so, the OpenMP implementation invokes the tool’s initializer, enabling
18 the tool to prepare to monitor the execution on the host. Finally, a tool may arrange to monitor
19 computation that execute on target devices. This section explains how the tool and an OpenMP
20 implementation interact to accomplish these tasks.

1 4.2.1.1 Determining Whether an OMPT Tool Should be Initialized

2 A tool indicates its interest in using the OMPT interface by providing a non-**NULL** pointer to an
3 **ompt_start_tool_result_t** structure to an OpenMP implementation as a return value from
4 **ompt_start_tool**. There are three ways that a tool can provide a definition of
5 **ompt_start_tool** to an OpenMP implementation:

- 6 • statically-linking the tool's definition of **ompt_start_tool** into an OpenMP application,
- 7 • introducing a dynamically-linked library that includes the tool's definition of
8 **ompt_start_tool** into the application's address space, or
- 9 • providing the name of a dynamically-linked library appropriate for the architecture and operating
10 system used by the application in the *tool-libraries-var* ICV.

11 Immediately before an OpenMP implementation initializes itself, it determines whether it should
12 check for the presence of a tool interested in using the OMPT interface by examining the *tool-var*
13 ICV. If value of *tool-var* is *disabled*, the OpenMP implementation will initialize itself without even
14 checking whether a tool is present and the functionality of the OMPT interface will be unavailable
15 as the program executes.

16 If the value of *tool-var* is *enabled*, the OpenMP implementation will check to see if a tool has
17 provided an implementation of **ompt_start_tool**. The OpenMP implementation first checks if
18 a tool-provided implementation of **ompt_start_tool** is available in the address space, either
19 statically-linked into the application or in a dynamically-linked library loaded in the address space.
20 If multiple implementations of **ompt_start_tool** are available, the OpenMP implementation
21 will use the first tool-provided implementation of **ompt_start_tool** found.

22 If no tool-provided implementation of **ompt_start_tool** is found in the address space, the
23 OpenMP implementation will consult the *tool-libraries-var* ICV, which contains a (possibly empty)
24 list of dynamically-linked libraries. As described in detail in Section 5.19, the libraries in
25 *tool-libraries-var*, will be searched for the first usable implementation of **ompt_start_tool**
26 provided by one of the libraries in the list.

27 If a tool-provided definition of **ompt_start_tool** is found using either method, the OpenMP
28 implementation will invoke it; if it returns a non-**NULL** pointer to an
29 **ompt_start_tool_result_t** structure, the OpenMP implementation will know that a tool
30 expects to use the OMPT interface.

31 Next, the OpenMP implementation will initialize itself. If a tool provided a non-**NULL** pointer to an
32 **ompt_start_tool_result_t** structure, the OpenMP runtime will prepare itself for use of
33 the OMPT interface by a tool.

34 Cross References

- 35 • *tool-libraries-var* ICV, see Section 2.4 on page 47.
- 36 • *tool-var* ICV, see Section 2.4 on page 47.

- 1 • `ompt_start_tool_result_t`, see Section 4.2.3.1 on page 429.
- 2 • `ompt_start_tool`, see Section 4.4.2.1 on page 592.

3 4.2.1.2 Initializing an OMPT Tool

4 If a tool-provided implementation of `ompt_start_tool` returns a non-NULL pointer to an
5 `ompt_start_tool_result_t` structure, the OpenMP implementation will invoke the tool
6 initializer specified in this structure prior to the occurrence of any OpenMP *event*.

7 A tool's initializer, described in Section 4.2.4.1.1 on page 444 uses its argument *lookup* to look up
8 pointers to OMPT interface runtime entry points provided by the OpenMP implementation; this
9 process is described in Section 4.2.1.2.1 on page 421. Typically, a tool initializer will first obtain a
10 pointer to the OpenMP runtime entry point known as `ompt_set_callback` with type signature
11 `ompt_set_callback_t` and then use this runtime entry point to register tool callbacks for
12 OpenMP events, as described in Section 4.2.1.3 on page 422.

13 A tool initializer may use the OMPT interface runtime entry points known as
14 `ompt_enumerate_states` and `ompt_enumerate_mutex_impls`, which have type
15 signatures `ompt_enumerate_states_t` and `ompt_enumerate_mutex_impls_t`, to
16 determine what thread states and implementations of mutual exclusion a particular OpenMP
17 implementation employs.

18 If a tool initializer returns a non-zero value, the tool will be *activated* for the execution; otherwise,
19 the tool will be inactive.

20 Cross References

- 21 • `ompt_start_tool_result_t`, see Section 4.2.3.1 on page 429.
- 22 • `ompt_start_tool`, see Section 4.4.2.1 on page 592.
- 23 • `ompt_initialize_t`, see Section 4.2.4.1.1 on page 444.
- 24 • `ompt_callback_thread_begin_t`, see Section 4.2.4.2.1 on page 446.
- 25 • `ompt_enumerate_states_t`, see Section 4.2.5.1.1 on page 481.
- 26 • `ompt_enumerate_mutex_impls_t`, see Section 4.2.5.1.2 on page 482.
- 27 • `ompt_set_callback_t`, see Section 4.2.5.1.3 on page 483.
- 28 • `ompt_function_lookup_t`, see Section 4.2.5.3.1 on page 516.

1 4.2.1.2.1 Binding Entry Points in the OMPT Callback Interface

2 Functions that an OpenMP implementation provides to support the OMPT interface are not defined
3 as global function symbols. Instead, they are defined as runtime entry points that a tool can only
4 identify using the *lookup* function provided as an argument to the tool’s initializer. This design
5 avoids tool implementations that will fail in certain circumstances when functions defined as part of
6 the OpenMP runtime are not visible to a tool, even though the tool and the OpenMP runtime are
7 both present in the same address space. It also prevents inadvertent use of a tool support routine by
8 applications.

9 A tool’s initializer receives a function pointer to a *lookup* runtime entry point with type signature
10 **ompt_function_lookup_t** as its first argument. Using this function, a tool initializer may
11 obtain a pointer to each of the runtime entry points that an OpenMP implementation provides to
12 support the OMPT interface. Once a tool has obtained a *lookup* function, it may employ it at any
13 point in the future.

14 For each runtime entry point in the OMPT interface for the host device, Table 4.1 provides the
15 string name by which it is known and its associated type signature. Implementations can provide
16 additional, implementation-specific names and corresponding entry points as long as they don’t use
17 names that start with the prefix “**ompt_**”. These are reserved for future extensions in the OpenMP
18 specification.

19 During initialization, a tool should look up each runtime entry point in the OMPT interface by
20 name and bind a pointer maintained by the tool that it can use later to invoke the entry point as
21 needed. The entry points described in Table 4.1 enable a tool to assess what thread states and
22 mutual exclusion implementations that an OpenMP runtime supports, register tool callbacks,
23 inspect callbacks registered, introspect OpenMP state associated with threads, and use tracing to
24 monitor computations that execute on target devices.

25 Detailed information about each runtime entry point listed in Table 4.1 is included as part of the
26 description of its type signature.

27 Cross References

- 28 • **ompt_enumerate_states_t**, see Section 4.2.5.1.1 on page 481.
- 29 • **ompt_enumerate_mutex_impls_t**, see Section 4.2.5.1.2 on page 482.
- 30 • **ompt_set_callback_t**, see Section 4.2.5.1.3 on page 483.
- 31 • **ompt_get_callback_t**, see Section 4.2.5.1.4 on page 485.
- 32 • **ompt_get_thread_data_t**, see Section 4.2.5.1.5 on page 486.
- 33 • **ompt_get_num_places_t**, see Section 4.2.5.1.7 on page 488.
- 34 • **ompt_get_place_proc_ids_t**, see Section 4.2.5.1.8 on page 489.
- 35 • **ompt_get_place_num_t**, see Section 4.2.5.1.9 on page 490.

- 1 • `ompt_get_partition_place_nums_t`, see Section 4.2.5.1.10 on page 491.
- 2 • `ompt_get_proc_id_t`, see Section 4.2.5.1.11 on page 492.
- 3 • `ompt_get_state_t`, see Section 4.2.5.1.12 on page 493.
- 4 • `ompt_get_parallel_info_t`, see Section 4.2.5.1.13 on page 494.
- 5 • `ompt_get_task_info_t`, see Section 4.2.5.1.14 on page 496.
- 6 • `ompt_get_task_memory_t`, see Section 4.2.5.1.15 on page 498.
- 7 • `ompt_get_target_info_t`, see Section 4.2.5.1.16 on page 499.
- 8 • `ompt_get_num_devices_t`, see Section 4.2.5.1.17 on page 501.
- 9 • `ompt_get_num_procs_t`, see Section 4.2.5.1.6 on page 487.
- 10 • `ompt_get_unique_id_t`, see Section 4.2.5.1.18 on page 501.
- 11 • `ompt_finalize_tool_t`, see Section 4.2.5.1.19 on page 502.
- 12 • `ompt_function_lookup_t`, see Section 4.2.5.3.1 on page 516.

13 4.2.1.3 Monitoring Activity on the Host with OMPT

14 To monitor execution of an OpenMP program on the host device, a tool’s initializer must register to
15 receive notification of events that occur as an OpenMP program executes. A tool can register
16 callbacks for OpenMP events using the runtime entry point known as `ompt_set_callback`.
17 The possible return codes for `ompt_set_callback` and their meanings are shown in Table 4.4.
18 If the `ompt_set_callback` runtime entry point is called outside a tool’s initializer, registration
19 of supported callbacks may fail with a return code of `ompt_set_error`.

20 All callbacks registered with `ompt_set_callback` or returned by `ompt_get_callback` use
21 the dummy type signature `ompt_callback_t`. While this is a compromise, it is better than
22 providing unique runtime entry points with a precise type signatures to set and get the callback for
23 each unique runtime entry point type signature.

24 Table 4.2 indicates the return codes permissible when trying to register various callbacks. For
25 callbacks where the only registration return code allowed is `ompt_set_always`, an OpenMP
26 implementation must guarantee that the callback will be invoked every time a runtime event
27 associated with it occurs. Support for such callbacks is required in a minimal implementation of the
28 OMPT interface. For other callbacks where registration is allowed to return values other than
29 `ompt_set_always`, its implementation-defined whether an OpenMP implementation invokes a
30 registered callback never, sometimes, or always. If registration for a callback allows a return code
31 of `ompt_set_never`, support for invoking such a callback need not be present in a minimal
32 implementation of the OMPT interface. The return code when a callback is registered enables a tool
33 to know what to expect when the level of support for the callback can be implementation defined.

TABLE 4.1: OMPT callback interface runtime entry point names and their type signatures.

Entry Point String Name	Type signature
“ompt_enumerate_states”	ompt_enumerate_states_t
“ompt_enumerate_mutex_impls”	ompt_enumerate_mutex_impls_t
“ompt_set_callback”	ompt_set_callback_t
“ompt_get_callback”	ompt_get_callback_t
“ompt_get_thread_data”	ompt_get_thread_data_t
“ompt_get_num_places”	ompt_get_num_places_t
“ompt_get_place_proc_ids”	ompt_get_place_proc_ids_t
“ompt_get_place_num”	ompt_get_place_num_t
“ompt_get_partition_place_nums”	ompt_get_partition_place_nums_t
“ompt_get_proc_id”	ompt_get_proc_id_t
“ompt_get_state”	ompt_get_state_t
“ompt_get_parallel_info”	ompt_get_parallel_info_t
“ompt_get_task_info”	ompt_get_task_info_t
“ompt_get_task_memory”	ompt_get_task_memory_t
“ompt_get_num_devices”	ompt_get_num_devices_t
“ompt_get_num_procs”	ompt_get_num_procs_t
“ompt_get_target_info”	ompt_get_target_info_t
“ompt_get_unique_id”	ompt_get_unique_id_t
“ompt_finalize_tool”	ompt_finalize_tool_t

TABLE 4.2: Valid return codes of `ompt_set_callback` for each callback.

Return code abbreviation	N	S/P	A
<code>ompt_callback_thread_begin</code>			*
<code>ompt_callback_thread_end</code>			*
<code>ompt_callback_parallel_begin</code>			*
<code>ompt_callback_parallel_end</code>			*
<code>ompt_callback_task_create</code>			*
<code>ompt_callback_task_schedule</code>			*
<code>ompt_callback_implicit_task</code>			*
<code>ompt_callback_target</code>			*
<code>ompt_callback_target_data_op</code>			*
<code>ompt_callback_target_submit</code>			*
<code>ompt_callback_control_tool</code>			*
<code>ompt_callback_device_initialize</code>			*
<code>ompt_callback_device_finalize</code>			*
<code>ompt_callback_device_load</code>			*
<code>ompt_callback_device_unload</code>			*
<code>ompt_callback_sync_region_wait</code>	*	*	*
<code>ompt_callback_mutex_released</code>	*	*	*
<code>ompt_callback_task_dependences</code>	*	*	*
<code>ompt_callback_task_dependence</code>	*	*	*
<code>ompt_callback_work</code>	*	*	*
<code>ompt_callback_master</code>	*	*	*
<code>ompt_callback_target_map</code>	*	*	*
<code>ompt_callback_sync_region</code>	*	*	*
<code>ompt_callback_reduction</code>	*	*	*
<code>ompt_callback_lock_init</code>	*	*	*
<code>ompt_callback_lock_destroy</code>	*	*	*
<code>ompt_callback_mutex_acquire</code>	*	*	*
<code>ompt_callback_mutex_acquired</code>	*	*	*
<code>ompt_callback_nest_lock</code>	*	*	*
<code>ompt_callback_flush</code>	*	*	*
<code>ompt_callback_cancel</code>	*	*	*
<code>ompt_callback_dispatch</code>	*	*	*
N = <code>ompt_set_never</code>	S = <code>ompt_set_sometimes</code>		
P = <code>ompt_set_sometimes_paired</code>	A = <code>ompt_set_always</code>		

1 To avoid a tool interface specification that enables a tool to register unique callbacks for an
2 overwhelming number of events, the interface was collapsed in several ways. First, in cases where
3 events are naturally paired, e.g., the beginning and end of a region, and the arguments needed by the
4 callback at each endpoint were identical, the pair of events was collapsed so that a tool registers a
5 single callback that will be invoked at both endpoints with `ompt_scope_begin` or
6 `ompt_scope_end` provided as an argument to identify which endpoint the callback invocation
7 reflects. Second, when a whole class of events is amenable to uniform treatment, only a single
8 callback is provided for a family of events, e.g., a `ompt_callback_sync_region_wait`
9 callback is used for multiple kinds of synchronization regions, i.e., barrier, taskwait, and taskgroup
10 regions. Some events involve both kinds of collapsing: the aforementioned
11 `ompt_callback_sync_region_wait` represents a callback that will be invoked at each
12 endpoint for different kinds of synchronization regions.

13 Cross References

- 14 • `ompt_set_callback_t`, see Section 4.2.5.1.3 on page 483.
- 15 • `ompt_get_callback_t`, see Section 4.2.5.1.4 on page 485.

16 4.2.1.4 Tracing Activity on Target Devices with OMPT

17 A target device may or may not initialize a full OpenMP runtime system. Unless it does, it may not
18 be possible to monitor activity on a device using a tool interface based on callbacks. To
19 accommodate such cases, the OMPT interface defines a monitoring interface for tracing activity on
20 target devices. Tracing activity on a target device involves the following steps:

- 21 • To prepare to trace activity on a target device, a tool must register for an
22 `ompt_callback_device_initialize` callback. A tool may also register for an
23 `ompt_callback_device_load` callback to be notified when code is loaded onto a target
24 device or an `ompt_callback_device_unload` callback to be notified when code is
25 unloaded from a target device. A tool may also optionally register an
26 `ompt_callback_device_finalize` callback.
- 27 • When an OpenMP implementation initializes a target device, the OpenMP implementation will
28 dispatch the tool's device initialization callback on the host device. If the OpenMP
29 implementation or target device does not support tracing, the OpenMP implementation will pass
30 a `NULL` to the tool's device initializer for its *lookup* argument; otherwise, the OpenMP
31 implementation will pass a pointer to a device-specific runtime entry point with type signature
32 `ompt_function_lookup_t` to the tool's device initializer.
- 33 • If the device initializer for the tool receives a non-`NULL` *lookup* pointer, the tool may use it to
34 query which runtime entry points in the tracing interface are available for a target device and
35 bind the function pointers returned to tool variables. Table 4.3 indicates the names of the runtime
36 entry points that a target device may provide for use by a tool. Implementations can provide

TABLE 4.3: OMPT tracing interface runtime entry point names and their type signatures.

Entry Point String Name	Type Signature
"ompt_get_device_num_procs"	ompt_get_device_num_procs_t
"ompt_get_device_time"	ompt_get_device_time_t
"ompt_translate_time"	ompt_translate_time_t
"ompt_set_trace_ompt"	ompt_set_trace_ompt_t
"ompt_set_trace_native"	ompt_set_trace_native_t
"ompt_start_trace"	ompt_start_trace_t
"ompt_pause_trace"	ompt_pause_trace_t
"ompt_flush_trace"	ompt_flush_trace_t
"ompt_stop_trace"	ompt_stop_trace_t
"ompt_advance_buffer_cursor"	ompt_advance_buffer_cursor_t
"ompt_get_record_type"	ompt_get_record_type_t
"ompt_get_record_ompt"	ompt_get_record_ompt_t
"ompt_get_record_native"	ompt_get_record_native_t
"ompt_get_record_abstract"	ompt_get_record_abstract_t

1 additional, implementation-specific names and corresponding entry points as long as they don't
2 use names that start with the prefix "**ompt_**". These are reserved for future extensions in the
3 OpenMP specification.

4 If *lookup* is non-NULL, the driver for a device will provide runtime entry points that enable a tool
5 to control the device's interface for collecting traces in its *native* trace format, which may be
6 device specific. The kinds of trace records available for a device will typically be
7 implementation-defined. Some devices may also allow a tool to collect traces of records in a
8 standard format known as OMPT format, described in this document. If so, the *lookup* function
9 will return values for the runtime entry points **ompt_set_trace_ompt** and
10 **ompt_get_record_ompt**, which support collecting and decoding OMPT traces. These
11 runtime entry points are not required for all devices and will only be available for target devices
12 that support collection of standard traces in OMPT format. For some devices, their native tracing
13 format may be OMPT format. In that case, tracing can be controlled using either the runtime
14 entry points for native or OMPT tracing.

- 15 • The tool will use the **ompt_set_trace_native** and/or the **ompt_set_trace_ompt**
16 runtime entry point to specify what types of events or activities to monitor on the target device.
17 If the **ompt_set_trace_native** and/or the **ompt_set_trace_ompt** runtime entry
18 point is called outside a device initializer, registration of supported callbacks may fail with a

- 1 return code of `ompt_set_error`.
- 2 • The tool will initiate tracing on the target device by invoking `ompt_start_trace`.
3 Arguments to `ompt_start_trace` include two tool callbacks for use by the OpenMP
4 implementation to manage traces associated with the target device: one to allocate a buffer where
5 the target device can deposit trace events and a second to process a buffer of trace events from the
6 target device.
- 7 • When the target device needs a trace buffer, the OpenMP implementation will invoke the
8 tool-supplied callback function on the host device to request a new buffer.
- 9 • The OpenMP implementation will monitor execution of OpenMP constructs on the target device
10 as directed and record a trace of events or activities into a trace buffer. If the device is capable,
11 device trace records will be marked with a *host_op_id*—an identifier used to associate device
12 activities with the target operation initiated on the host that caused these activities. To correlate
13 activities on the host with activities on a device, a tool can register a
14 **`ompt_callback_target_submit`** callback. Before the host initiates each distinct activity
15 associated with a structured block for a **`target`** construct on a target device, the OpenMP
16 implementation will dispatch the **`ompt_callback_target_submit`** callback on the host in
17 the thread executing the task that encounters the **`target`** construct. Examples of activities that
18 could cause an **`ompt_callback_target_submit`** callback to be dispatched include an
19 explicit data copy between a host and target device or execution of a computation. This callback
20 provides the tool with a pair of identifiers: one that identifies the target region and a second that
21 uniquely identifies an activity associated with that region. These identifiers help the tool
22 correlate activities on the target device with their target region.
- 23 • When appropriate, e.g., when a trace buffer fills or needs to be flushed, the OpenMP
24 implementation will invoke the tool-supplied buffer completion callback to process a non-empty
25 sequence of records in a trace buffer associated with the target device.
- 26 • The tool-supplied buffer completion callback may return immediately, ignoring records in the
27 trace buffer, or it may iterate through them using the **`ompt_advance_buffer_cursor`**
28 entry point and inspect each one. A tool may inspect the type of the record at the current cursor
29 position using the **`ompt_get_record_type`** runtime entry point. A tool may choose to
30 inspect the contents of some or all records in a trace buffer using the
31 **`ompt_get_record_ompt`**, **`ompt_get_record_native`**, or
32 **`ompt_get_record_abstract`** runtime entry point. Presumably, a tool that chooses to use
33 the **`ompt_get_record_native`** runtime entry point to inspect records will have some
34 knowledge about a device’s native trace format. A tool may always use the
35 **`ompt_get_record_abstract`** runtime entry point to inspect a trace record; this runtime
36 entry point will decode the contents of a native trace record and summarize them in a standard
37 format, namely, a **`ompt_record_abstract_t`** record. Only a record in OMPT format can
38 be retrieved using the **`ompt_get_record_ompt`** runtime entry point.
- 39 • Once tracing has been started on a device, a tool may pause or resume tracing on the device at
40 any time by invoking **`ompt_pause_trace`** with an appropriate flag value as an argument.

- 1 • A tool may request that a device flush any pending trace records at any time between device
2 initialization and finalization by invoking the **ompt_flush_trace** runtime entry point for the
3 device.
- 4 • A tool may start or stop tracing on a device at any time using the **ompt_start_trace** or
5 **ompt_stop_trace** runtime entry points, respectively. When tracing is stopped on a device,
6 the OpenMP implementation will eventually gather all trace records already collected on the
7 device and present to the tool using the buffer completion callback provided by the tool.
- 8 • It is legal to shut down an OpenMP implementation while device tracing is in progress.
- 9 • When an OpenMP implementation begins to shut down, the OpenMP implementation will
10 finalize each target device. Device finalization occurs in three steps. First, the OpenMP
11 implementation halts any tracing in progress for the device. Second, the OpenMP
12 implementation flushes all trace records collected for the device and presents them to the tool
13 using the buffer completion callback associated with that device. Finally, the OpenMP
14 implementation dispatches any **ompt_callback_device_finalize** callback that was
15 previously registered by the tool.

16 Cross References

- 17 • **ompt_callback_device_initialize_t**, see Section [4.2.4.2.28](#) on page 478.
- 18 • **ompt_callback_device_finalize_t**, see Section [4.2.4.2.29](#) on page 479.
- 19 • **ompt_get_device_num_procs_t**, see Section [4.2.5.2.1](#) on page 502.
- 20 • **ompt_get_device_time**, see Section [4.2.5.2.2](#) on page 503.
- 21 • **ompt_translate_time**, see Section [4.2.5.2.3](#) on page 504.
- 22 • **ompt_set_trace_ompt**, see Section [4.2.5.2.4](#) on page 505.
- 23 • **ompt_set_trace_native**, see Section [4.2.5.2.5](#) on page 506.
- 24 • **ompt_start_trace**, see Section [4.2.5.2.6](#) on page 508.
- 25 • **ompt_pause_trace**, see Section [4.2.5.2.7](#) on page 509.
- 26 • **ompt_flush_trace**, see Section [4.2.5.2.8](#) on page 510.
- 27 • **ompt_stop_trace**, see Section [4.2.5.2.9](#) on page 510.
- 28 • **ompt_advance_buffer_cursor**, see Section [4.2.5.2.10](#) on page 511.
- 29 • **ompt_get_record_type**, see Section [4.2.5.2.11](#) on page 512.
- 30 • **ompt_get_record_ompt**, see Section [4.2.5.2.12](#) on page 513.
- 31 • **ompt_get_record_native**, see Section [4.2.5.2.13](#) on page 514.
- 32 • **ompt_get_record_abstract**, see Section [4.2.5.2.14](#) on page 515.

1 4.2.2 Finalizing an OMPT Tool

2 If `ompt_start_tool` returned a non-NULL pointer when an OpenMP implementation was
3 initialized, the tool finalizer, of type signature `ompt_finalize_t`, specified by the *finalize* field
4 in this structure will be called as the OpenMP implementation shuts down.

5 Cross References

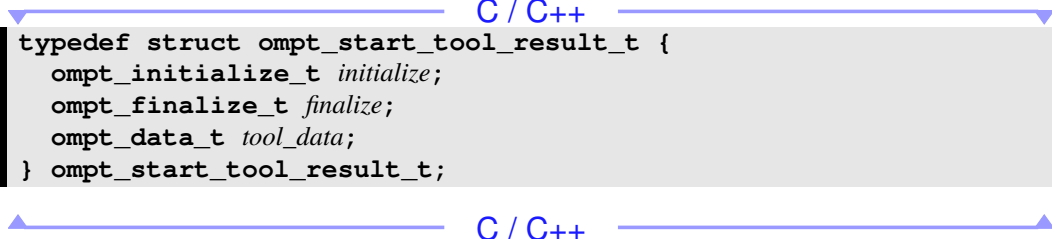
- 6 • `ompt_finalize_t`, Section [4.2.4.1.2](#) on page [445](#)

7 4.2.3 OMPT Data Types

8 4.2.3.1 Tool Initialization and Finalization

9 Summary

10 A tool's implementation of `ompt_start_tool` returns a pointer to an
11 `ompt_start_tool_result_t` structure, which contains pointers to the tool's initialization
12 and finalization callbacks as well as an `ompt_data_t` object for use by the tool.

13  C / C++
14 `typedef struct ompt_start_tool_result_t {`
15 `ompt_initialize_t initialize;`
16 `ompt_finalize_t finalize;`
17 `ompt_data_t tool_data;`
`} ompt_start_tool_result_t;`

18 Restrictions

19 The *initialize* and *finalize* callback pointer values in an `ompt_start_tool_result_t`
20 structure returned by `ompt_start_tool` must be non-NULL.

21 Cross References

- 22 • `ompt_data_t`, see Section [4.2.3.4.3](#) on page [434](#).
- 23 • `ompt_finalize_t`, see Section [4.2.4.1.2](#) on page [445](#).
- 24 • `ompt_initialize_t`, see Section [4.2.4.1.1](#) on page [444](#).
- 25 • `ompt_start_tool`, see Section [4.4.2.1](#) on page [592](#).

1 4.2.3.2 Callbacks

2 The following enumeration type indicates the integer codes used to identify OpenMP callbacks
3 when registering or querying them.

C / C++

```
4 typedef enum ompt_callbacks_t {  
5     ompt_callback_thread_begin           = 1,  
6     ompt_callback_thread_end           = 2,  
7     ompt_callback_parallel_begin       = 3,  
8     ompt_callback_parallel_end        = 4,  
9     ompt_callback_task_create          = 5,  
10    ompt_callback_task_schedule        = 6,  
11    ompt_callback_implicit_task        = 7,  
12    ompt_callback_target               = 8,  
13    ompt_callback_target_data_op       = 9,  
14    ompt_callback_target_submit        = 10,  
15    ompt_callback_control_tool         = 11,  
16    ompt_callback_device_initialize    = 12,  
17    ompt_callback_device_finalize      = 13,  
18    ompt_callback_device_load          = 14,  
19    ompt_callback_device_unload        = 15,  
20    ompt_callback_sync_region_wait     = 16,  
21    ompt_callback_mutex_released       = 17,  
22    ompt_callback_task_dependences     = 18,  
23    ompt_callback_task_dependence      = 19,  
24    ompt_callback_work                  = 20,  
25    ompt_callback_master                = 21,  
26    ompt_callback_target_map           = 22,  
27    ompt_callback_sync_region          = 23,  
28    ompt_callback_lock_init            = 24,  
29    ompt_callback_lock_destroy         = 25,  
30    ompt_callback_mutex_acquire        = 26,  
31    ompt_callback_mutex_acquired       = 27,  
32    ompt_callback_nest_lock            = 28,  
33    ompt_callback_flush                 = 29,  
34    ompt_callback_cancel                = 30,  
35    ompt_callback_reduction            = 31,  
36    ompt_callback_dispatch              = 32  
37 } ompt_callbacks_t;
```

C / C++

1 4.2.3.3 Tracing

2 4.2.3.3.1 Record Type

```
3 typedef enum ompt_record_t {
4     ompt_record_ompt           = 1,
5     ompt_record_native        = 2,
6     ompt_record_invalid       = 3
7 } ompt_record_t;
```

C / C++

8 4.2.3.3.2 Native Record Kind

```
9 typedef enum ompt_record_native_t {
10     ompt_record_native_info = 1,
11     ompt_record_native_event = 2
12 } ompt_record_native_t;
```

C / C++

13 4.2.3.3.3 Native Record Abstract Type

```
14 typedef struct ompt_record_abstract_t {
15     ompt_record_native_t rclass;
16     const char *type;
17     ompt_device_time_t start_time;
18     ompt_device_time_t end_time;
19     ompt_hwid_t hwid;
20 } ompt_record_abstract_t;
```

C / C++

Description

A `ompt_record_abstract_t` record contains several pieces of information that a tool can use to process a native record that it may not fully understand. The `rclass` field indicates whether the record is informational or represents an event; knowing this can help a tool determine how to present the record. The record `type` field points to a statically-allocated, immutable character string that provides a meaningful name that a tool might want to use to describe the event to a user. The `start_time` and `end_time` fields are used to place an event in time. The times are relative to the device clock. If an event has no associated `start_time` and/or `end_time`, its value will be `ompt_time_none`. The hardware id field, `hwid`, is used to indicate the location on the device where the event occurred. A `hwid` may represent a hardware abstraction such as a core or a hardware thread id. The meaning of a `hwid` value for a device is defined by the implementer of the software stack for the device. If there is no hardware abstraction associated with the record, the value of `hwid` will be `ompt_hwid_none`.

4.2.3.3.4 Record Type

C / C++

```
typedef struct ompt_record_ompt_t {
    ompt_callbacks_t type;
    ompt_device_time_t time;
    ompt_id_t thread_id;
    ompt_id_t target_id;
    union {
        ompt_record_thread_begin_t thread_begin;
        ompt_record_idle_t idle;
        ompt_record_parallel_begin_t parallel_begin;
        ompt_record_parallel_end_t parallel_end;
        ompt_record_task_create_t task_create;
        ompt_record_task_dependences_t task_deps;
        ompt_record_task_dependence_t task_dep;
        ompt_record_task_schedule_t task_sched;
        ompt_record_implicit_t implicit;
        ompt_record_sync_region_t sync_region;
        ompt_record_target_t target_record;
        ompt_record_target_data_op_t target_data_op;
        ompt_record_target_map_t target_map;
        ompt_record_target_kernel_t kernel;
        ompt_record_lock_init_t lock_init;
        ompt_record_lock_destroy_t lock_destroy;
        ompt_record_mutex_acquire_t mutex_acquire;
        ompt_record_mutex_t mutex;
    };
};
```



```
1      ompt_record_nest_lock_t nest_lock;
2      ompt_record_master_t master;
3      ompt_record_work_t work;
4      ompt_record_flush_t flush;
5      } record;
6  } ompt_record_ompt_t;
```

C / C++

7 Description

8 The field *type* specifies the type of record provided by this structure. According to the type, event
9 specific information is stored in the matching *record* entry.

10 Restrictions

11 If the *type* is set to `ompt_callback_thread_end_t`, the value of *record* is undefined.

12 4.2.3.4 Miscellaneous Type Definitions

13 This section describes miscellaneous types and enumerations used by the tool interface.

14 4.2.3.4.1 ompt_callback_t

15 Pointers to tool callback functions with many different type signatures are passed to the
16 `ompt_set_callback` runtime entry point and returned by the `ompt_get_callback`
17 runtime entry point. For convenience, these runtime entry points expect all type signatures to be
18 cast to a dummy type `ompt_callback_t`.

```
19 typedef void (*ompt_callback_t) (void);
```

C / C++

1 4.2.3.4.2 `ompt_id_t`

2 When tracing asynchronous activity on OpenMP devices, tools need identifiers to correlate target
3 regions and operations initiated by the host with associated activities on a target device. In addition,
4 tools need identifiers to refer to parallel regions and tasks that execute on a device. OpenMP
5 implementations use identifiers of type `ompt_id_t` type for each of these purposes.

```
6 | typedef uint64_t ompt_id_t;
```

C / C++

7 `ompt_id_none` is defined as an instance of type `ompt_id_t` with the value 0.

8 Identifiers created on each device must be unique from the time an OpenMP implementation is
9 initialized until it is shut down. Specifically, this means that (1) identifiers for each target region
10 and target operation instance initiated by the host device must be unique over time on the host, and
11 (2) identifiers for parallel and task region instances that execute on a device must be unique over
12 time within that device.

13 Tools should not assume that `ompt_id_t` values are small or densely allocated.

14 4.2.3.4.3 `ompt_data_t`

15 Threads, parallel regions, and task regions each have an associated data object of type
16 `ompt_data_t` reserved for use by a tool. When an OpenMP implementation creates a thread or
17 an instance of a parallel or task region, it will initialize its associated `ompt_data_t` object with
18 the value `ompt_data_none`.

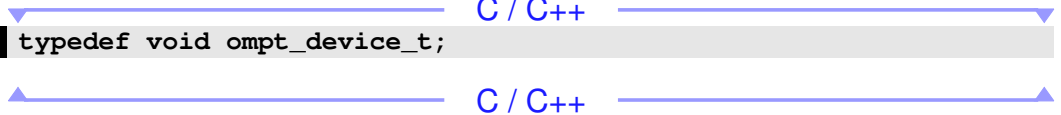
```
19 | typedef union ompt_data_t {  
20 |     uint64_t value;  
21 |     void *ptr;  
22 | } ompt_data_t;
```

C / C++

23 `ompt_data_none` is defined as an instance of type `ompt_data_t` with the data and pointer
24 fields equal to 0.

1 4.2.3.4.4 `ompt_device_t`

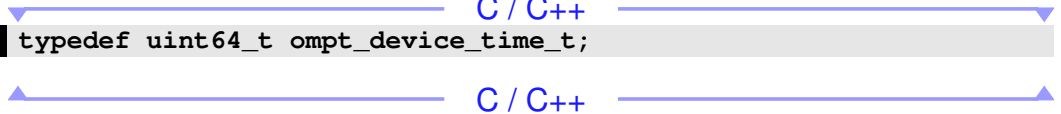
2 `ompt_device_t` is an opaque object representing a device.

3  `typedef void ompt_device_t;`

The code block is enclosed in a grey rectangular box. Above the box, a blue horizontal line with downward-pointing triangles at both ends spans the width of the box, with the text "C / C++" centered above it. Below the box, a blue horizontal line with upward-pointing triangles at both ends spans the width of the box, with the text "C / C++" centered below it.

4 4.2.3.4.5 `ompt_device_time_t`

5 `ompt_device_time_t` is an opaque object representing a raw time value from a device.
6 `ompt_time_none` refers to an unknown or unspecified time.

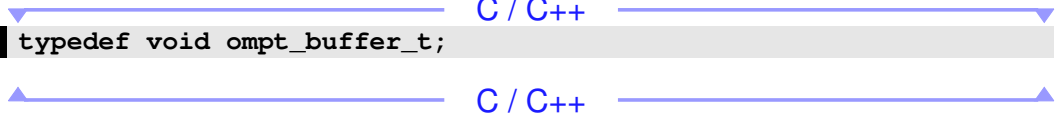
7  `typedef uint64_t ompt_device_time_t;`

The code block is enclosed in a grey rectangular box. Above the box, a blue horizontal line with downward-pointing triangles at both ends spans the width of the box, with the text "C / C++" centered above it. Below the box, a blue horizontal line with upward-pointing triangles at both ends spans the width of the box, with the text "C / C++" centered below it.

8 `ompt_time_none` is defined as an instance of type `ompt_device_time_t` with the value 0.

9 4.2.3.4.6 `ompt_buffer_t`

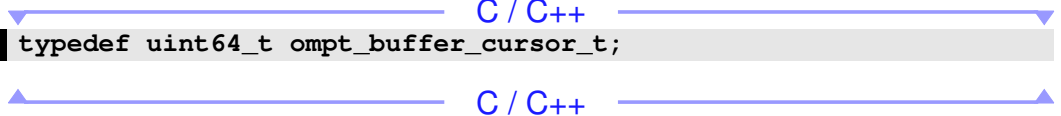
10 `ompt_buffer_t` is an opaque object handle for a target buffer.

11  `typedef void ompt_buffer_t;`

The code block is enclosed in a grey rectangular box. Above the box, a blue horizontal line with downward-pointing triangles at both ends spans the width of the box, with the text "C / C++" centered above it. Below the box, a blue horizontal line with upward-pointing triangles at both ends spans the width of the box, with the text "C / C++" centered below it.

12 4.2.3.4.7 `ompt_buffer_cursor_t`

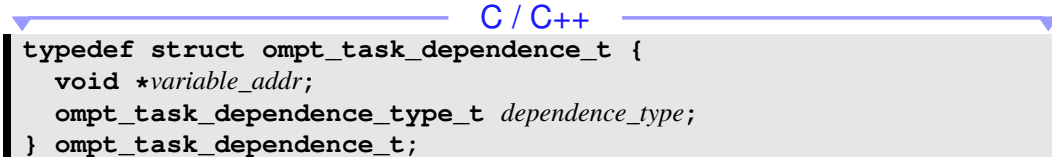
13 `ompt_buffer_cursor_t` is an opaque handle for a position in a target buffer.

14  `typedef uint64_t ompt_buffer_cursor_t;`

The code block is enclosed in a grey rectangular box. Above the box, a blue horizontal line with downward-pointing triangles at both ends spans the width of the box, with the text "C / C++" centered above it. Below the box, a blue horizontal line with upward-pointing triangles at both ends spans the width of the box, with the text "C / C++" centered below it.

1 4.2.3.4.8 `ompt_task_dependence_t`

2 `ompt_task_dependence_t` is a task dependence.

3  C / C++
4 `typedef struct ompt_task_dependence_t {`
5 `void *variable_addr;`
6 `ompt_task_dependence_type_t dependence_type;`
7 `} ompt_task_dependence_t;`

7 Description

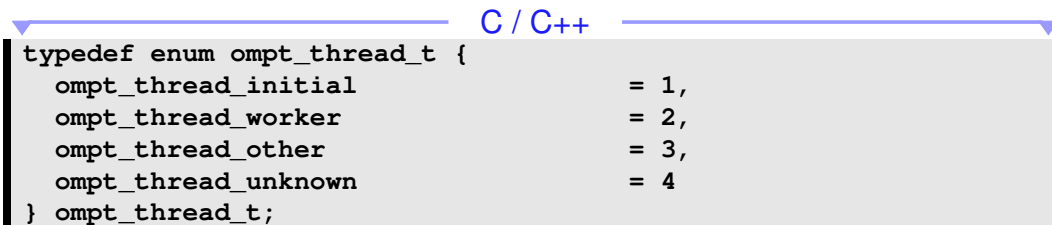
8 `ompt_task_dependence_t` is a structure to hold information about a depend clause. The
9 element `variable_addr` points to the storage location of the dependence. The element
10 `dependence_type` indicates the type of dependence described.

11 Cross References

- 12 • `ompt_task_dependence_type_t`, see Section [4.2.3.4.22](#) on page [442](#).

13 4.2.3.4.9 `ompt_thread_t`

14 `ompt_thread_t` is an enumeration that defines the valid thread type values.

15  C / C++
16 `typedef enum ompt_thread_t {`
17 `ompt_thread_initial = 1,`
18 `ompt_thread_worker = 2,`
19 `ompt_thread_other = 3,`
20 `ompt_thread_unknown = 4`
21 `} ompt_thread_t;`

22 Any *initial thread* has thread type `ompt_thread_initial`. All *OpenMP threads* that are not
23 initial threads have thread type `ompt_thread_worker`. A thread employed by an OpenMP
24 implementation that does not execute user code has thread type `ompt_thread_other`. Any
25 thread created outside an OpenMP implementation that is not an *initial thread* has thread type
`ompt_thread_unknown`.

1 4.2.3.4.10 ompt_scope_endpoint_t

2 ompt_scope_endpoint_t is an enumeration that defines valid scope endpoint values.

```
3                                     C / C++
4 typedef enum ompt_scope_endpoint_t {
5     ompt_scope_begin                = 1,
6     ompt_scope_end                  = 2
7 } ompt_scope_endpoint_t;
```

7 4.2.3.4.11 ompt_dispatch_t

8 ompt_dispatch_t is an enumeration that defines the valid dispatch kind values.

```
9                                     C / C++
10 typedef enum ompt_dispatch_t {
11     ompt_dispatch_iteration          = 1,
12     ompt_dispatch_section            = 2
13 } ompt_dispatch_t;
```

13 4.2.3.4.12 ompt_sync_region_t

14 ompt_sync_region_t is an enumeration that defines the valid sync region kind values.

```
15                                     C / C++
16 typedef enum ompt_sync_region_t {
17     ompt_sync_region_barrier        = 1,
18     ompt_sync_region_barrier_implicit = 2,
19     ompt_sync_region_barrier_explicit = 3,
20     ompt_sync_region_barrier_implementation = 4,
21     ompt_sync_region_taskwait      = 5,
22     ompt_sync_region_taskgroup     = 6,
23     ompt_sync_region_reduction     = 7
24 } ompt_sync_region_kind_t;
```

1 4.2.3.4.13 ompt_target_data_op_t

2 ompt_target_data_op_t is an enumeration that defines the valid target data operation values.

```
3                                     C / C++
4 typedef enum ompt_target_data_op_t {
5     ompt_target_data_alloc           = 1,
6     ompt_target_data_transfer_to_device = 2,
7     ompt_target_data_transfer_from_device = 3,
8     ompt_target_data_delete         = 4,
9     ompt_target_data_associate      = 5,
10    ompt_target_data_disassociate    = 6
11 } ompt_target_data_op_t;
```

▲ C / C++ ▼

11 4.2.3.4.14 ompt_work_t

12 ompt_work_t is an enumeration that defines the valid work type values.

```
13                                     C / C++
14 typedef enum ompt_work_t {
15     ompt_work_loop                   = 1,
16     ompt_work_sections               = 2,
17     ompt_work_single_executor       = 3,
18     ompt_work_single_other          = 4,
19     ompt_work_workshare              = 5,
20     ompt_work_distribute             = 6,
21     ompt_work_taskloop               = 7
22 } ompt_work_t;
```

▲ C / C++ ▼

1 4.2.3.4.15 ompt_mutex_t

2 ompt_mutex_t is an enumeration that defines the valid mutex kind values.

```
3                                     C / C++
4 typedef enum ompt_mutex_t {
5     ompt_mutex_lock                 = 1,
6     ompt_mutex_nest_lock            = 2,
7     ompt_mutex_critical              = 3,
8     ompt_mutex_atomic               = 4,
9     ompt_mutex_ordered               = 5
10 } ompt_mutex_t;
```

10 4.2.3.4.16 ompt_native_mon_flag_t

11 ompt_native_mon_flag_t is an enumeration that defines the valid native monitoring flag
12 values.

```
13                                     C / C++
14 typedef enum ompt_native_mon_flag_t {
15     ompt_native_data_motion_explicit = 0x01,
16     ompt_native_data_motion_implicit = 0x02,
17     ompt_native_kernel_invocation    = 0x04,
18     ompt_native_kernel_execution     = 0x08,
19     ompt_native_driver                = 0x10,
20     ompt_native_runtime               = 0x20,
21     ompt_native_overhead              = 0x40,
22     ompt_native_idleness              = 0x80
23 } ompt_native_mon_flag_t;
```

1 4.2.3.4.17 `ompt_task_flag_t`

2 `ompt_task_flag_t` is an enumeration that defines the valid task type values. The least
3 significant byte provides information about the general classification of the task. The other bits
4 represent properties of the task.

```
5                                     C / C++  
6 typedef enum ompt_task_flag_t {  
7     ompt_task_initial                = 0x00000001,  
8     ompt_task_implicit              = 0x00000002,  
9     ompt_task_explicit              = 0x00000004,  
10    ompt_task_target                = 0x00000008,  
11    ompt_task_undeferred            = 0x08000000,  
12    ompt_task_untied                = 0x10000000,  
13    ompt_task_final                 = 0x20000000,  
14    ompt_task_mergeable             = 0x40000000,  
15    ompt_task_merged                = 0x80000000  
} ompt_task_flag_t;  
                                     C / C++
```

16 4.2.3.4.18 `ompt_task_status_t`

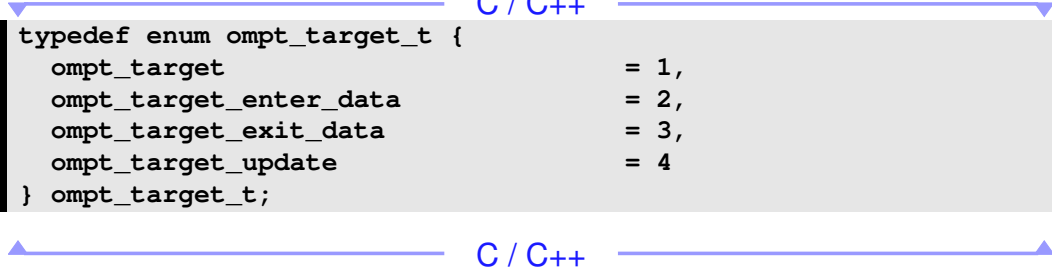
17 `ompt_task_status_t` is an enumeration that explains the reasons for switching a task that
18 reached a task scheduling point.

```
19                                     C / C++  
20 typedef enum ompt_task_status_t {  
21     ompt_task_complete = 1,  
22     ompt_task_yield   = 2,  
23     ompt_task_cancel  = 3,  
24     ompt_task_switch  = 4  
} ompt_task_status_t;  
                                     C / C++
```

25 The value `ompt_task_complete` indicates the completion of task that encountered the task
26 scheduling point. The value `ompt_task_yield` indicates that the task encountered a
27 `taskyield` construct. The value `ompt_task_cancel` indicates that the task is canceled due
28 to the encountering of an active cancellation point resulting in the cancellation of that task. The
29 value `ompt_task_switch` is used in the remaining cases of task switches.

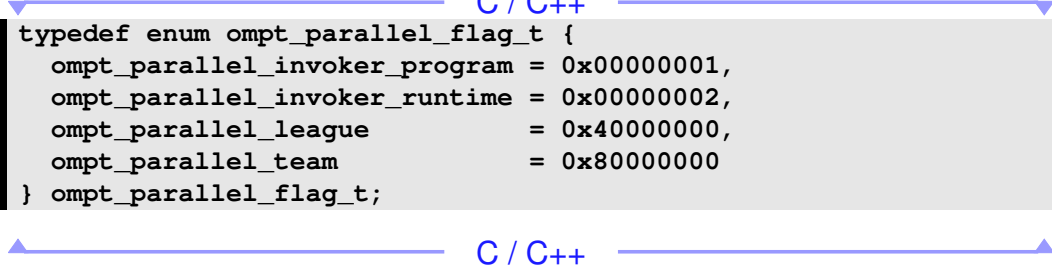
1 4.2.3.4.19 `ompt_target_t`

2 `ompt_target_t` is an enumeration that defines the valid target type values.

3  C / C++
4 `typedef enum ompt_target_t {`
5 `ompt_target = 1,`
6 `ompt_target_enter_data = 2,`
7 `ompt_target_exit_data = 3,`
8 `ompt_target_update = 4`
`} ompt_target_t;`

9 4.2.3.4.20 `ompt_parallel_flag_t`

10 `ompt_parallel_flag_t` is an enumeration that defines the valid invoker values.

11  C / C++
12 `typedef enum ompt_parallel_flag_t {`
13 `ompt_parallel_invoker_program = 0x00000001,`
14 `ompt_parallel_invoker_runtime = 0x00000002,`
15 `ompt_parallel_league = 0x40000000,`
16 `ompt_parallel_team = 0x80000000`
`} ompt_parallel_flag_t;`

17 **Description**

18 The value `ompt_parallel_invoker_program` indicates that on the master thread for a
19 parallel region, the outlined function associated with implicit tasks for the region is invoked directly
20 by the application.

21 The value `ompt_parallel_invoker_runtime` indicates that on the master thread for a
22 parallel region, the outlined function associated with implicit tasks for the region is invoked by the
23 runtime.

24 The value `ompt_parallel_league` indicates that the callback indicates the creation of a
25 league of teams by a `teams` construct.

26 The value `ompt_parallel_team` indicates that the callback indicates the creation of a team of
27 threads by a `parallel` construct.

1 4.2.3.4.21 ompt_target_map_flag_t

2 ompt_target_map_flag_t is an enumeration that defines the valid target map flag values.

```
3                                     C / C++
4 typedef enum ompt_target_map_flag_t {
5     ompt_target_map_flag_to           = 0x01,
6     ompt_target_map_flag_from        = 0x02,
7     ompt_target_map_flag_alloc       = 0x04,
8     ompt_target_map_flag_release     = 0x08,
9     ompt_target_map_flag_delete      = 0x10,
10    ompt_target_map_flag_implicit     = 0x20
11 } ompt_target_map_flag_t;
```

C / C++

11 4.2.3.4.22 ompt_task_dependence_type_t

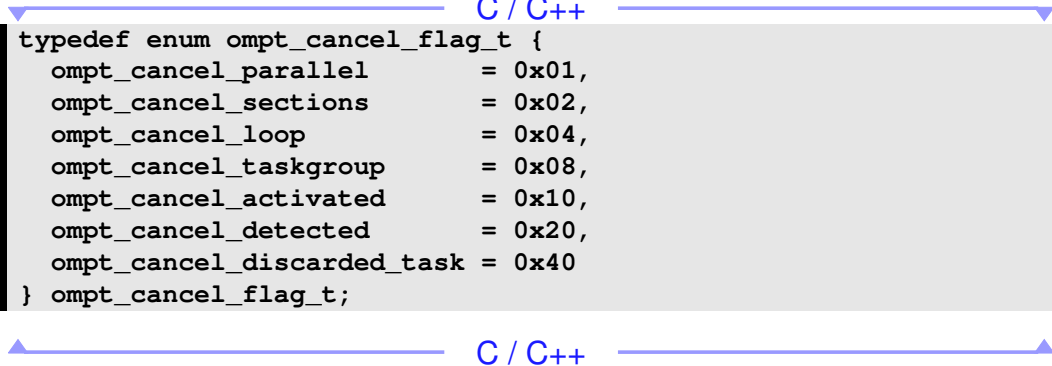
12 ompt_task_dependence_type_t is an enumeration that defines the valid task dependence
13 type values.

```
14                                     C / C++
15 typedef enum ompt_task_dependence_type_t {
16     ompt_task_dependence_type_in     = 1,
17     ompt_task_dependence_type_out    = 2,
18     ompt_task_dependence_type_inout = 3,
19     ompt_task_dependence_type_mutexinoutset = 4
20 } ompt_task_dependence_type_t;
```

C / C++

1 4.2.3.4.23 `ompt_cancel_flag_t`

2 `ompt_cancel_flag_t` is an enumeration that defines the valid cancel flag values.

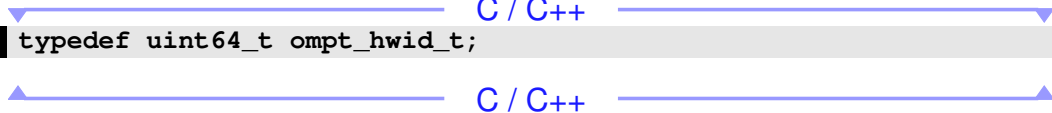
```
3  C / C++  
4 typedef enum ompt_cancel_flag_t {  
5     ompt_cancel_parallel      = 0x01,  
6     ompt_cancel_sections     = 0x02,  
7     ompt_cancel_loop         = 0x04,  
8     ompt_cancel_taskgroup    = 0x08,  
9     ompt_cancel_activated    = 0x10,  
10    ompt_cancel_detected      = 0x20,  
11    ompt_cancel_discarded_task = 0x40  
    } ompt_cancel_flag_t;
```

12 Cross References

- 13 • `ompt_cancel_t` data type, see Section [4.2.4.2.27](#) on page [477](#).

14 4.2.3.4.24 `ompt_hwid_t`

15 `ompt_hwid_t` is an opaque object representing a hardware identifier for a target device.
16 `ompt_hwid_none` refers to an unknown or unspecified hardware id. If there is no `hwid`
17 associated with a `ompt_record_abstract_t`, the value of `hwid` shall be
18 `ompt_hwid_none`.

```
19  C / C++  
20 typedef uint64_t ompt_hwid_t;
```

20 `ompt_hwid_none` is defined as an instance of type `ompt_hwid_t` with the value 0.

21 4.2.4 OMPT Tool Callback Signatures and Trace Records

22 Restrictions

23 Tool callbacks may not use OpenMP directives or call any runtime library routines described in
24 Section [3](#).

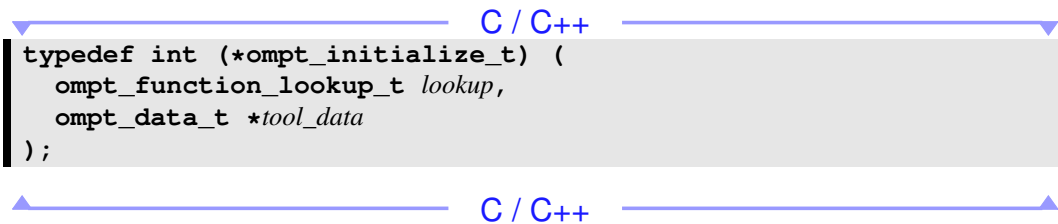
1 4.2.4.1 Initialization and Finalization Callback Signature

2 4.2.4.1.1 `ompt_initialize_t`

3 Summary

4 A tool implements an initializer with the type signature `ompt_initialize_t` to initialize the
5 tool's use of the OMPT interface.

6 Format

```
7  C / C++  
8 typedef int (*ompt_initialize_t) (  
9     ompt_function_lookup_t lookup,  
10    ompt_data_t *tool_data  
11 );
```

11 Description

12 For a tool to use the OMPT interface of an OpenMP implementation, the tool's implementation of
13 `ompt_start_tool` must return a non-NULL pointer to an `ompt_start_tool_result_t`
14 structure that contains a non-NULL pointer to a tool initializer with type signature
15 `ompt_initialize_t`. An OpenMP implementation will call the tool initializer after fully
16 initializing itself but before beginning execution of any OpenMP construct or completing execution
17 of any environment routine invocation.

18 The initializer returns a non-zero value if it succeeds.

19 Description of Arguments

20 The argument `lookup` is a callback to an OpenMP runtime routine that a tool must use to obtain a
21 pointer to each runtime entry point in the OMPT interface. The argument `tool_data` is a pointer to
22 the `tool_data` field in the `ompt_start_tool_result_t` structure returned by
23 `ompt_start_tool`. The expected actions of a tool initializer are described in Section 4.2.1.2 on
24 page 420.

Cross References

- `ompt_start_tool_result_t`, see Section 4.2.3.1 on page 429.
- `ompt_data_t`, see Section 4.2.3.4.3 on page 434.
- `ompt_start_tool`, see Section 4.4.2.1 on page 592.
- `ompt_function_lookup_t`, see Section 4.2.5.3.1 on page 516.

4.2.4.1.2 `ompt_finalize_t`

Summary

A tool implements a finalizer with the type signature `ompt_finalize_t` to finalize the tool's use of the OMPT interface.

Format

```

C / C++
typedef void (*ompt_finalize_t) (
    ompt_data_t *tool_data
);
C / C++
```

Description

For a tool to use the OMPT interface of an OpenMP implementation, the tool's implementation of `ompt_start_tool` must return a non-**NULL** pointer to an `ompt_start_tool_result_t` structure that contains a non-**NULL** pointer to a tool finalizer with type signature `ompt_finalize_t`. An OpenMP implementation will call the tool finalizer after the last OMPT *event* as the OpenMP implementation shuts down.

Description of Arguments

The argument `tool_data` is a pointer to the `tool_data` field in the `ompt_start_tool_result_t` structure returned by `ompt_start_tool`.

Cross References

- `ompt_start_tool_result_t`, see Section 4.2.3.1 on page 429.
- `ompt_data_t`, see Section 4.2.3.4.3 on page 434.
- `ompt_start_tool`, see Section 4.4.2.1 on page 592.

4.2.4.2 Event Callback Signatures and Trace Records

This section describes the signatures of tool callback functions that an OMPT tool might register and that are called during runtime of an OpenMP program.

4.2.4.2.1 `ompt_callback_thread_begin_t`

Format

C / C++

```
typedef void (*ompt_callback_thread_begin_t) (  
    ompt_thread_t thread_type,  
    ompt_data_t *thread_data  
);
```

C / C++

Trace Record

C / C++

```
typedef struct ompt_record_thread_begin_t {  
    ompt_thread_t thread_type;  
} ompt_record_thread_begin_t;
```

C / C++

Description of Arguments

The argument `thread_type` indicates the type of the new thread: initial, worker, or other.

The binding of argument `thread_data` is the new thread.

Cross References

- `ompt_data_t` type, see Section [4.2.3.4.3](#) on page [434](#).
- `ompt_thread_t` type, see Section [4.2.3.4.9](#) on page [436](#).

4.2.4.2.2 `ompt_callback_thread_end_t`

Format

```
typedef void (*ompt_callback_thread_end_t) (  
    ompt_data_t *thread_data  
);
```

Description of Arguments

The binding of argument *thread_data* is the thread that is terminating.

Cross References

- `ompt_data_t` type, see Section [4.2.3.4.3](#) on page [434](#).
- `ompt_record_ompt_t` type, see Section [4.2.3.3.4](#) on page [432](#).

4.2.4.2.3 `ompt_callback_parallel_begin_t`

Format

```
typedef void (*ompt_callback_parallel_begin_t) (  
    ompt_data_t *encountering_task_data,  
    const omp_frame_t *encountering_task_frame,  
    ompt_data_t *parallel_data,  
    unsigned int requested_parallelism,  
    int flag,  
    const void *codeptr_ra  
);
```

Trace Record

C / C++

```
typedef struct omp_t_record_parallel_begin_t {
    omp_id_t encountering_task_id;
    omp_id_t parallel_id;
    unsigned int requested_parallelism;
    int flag;
    const void *codeptr_ra;
} omp_t_record_parallel_begin_t;
```

C / C++

Description of Arguments

The binding of argument *encountering_task_data* is the encountering task.

The argument *encountering_task_frame* points to the frame object associated with the encountering task.

The binding of argument *parallel_data* is the parallel or teams region that is beginning.

The argument *requested_parallelism* indicates the number of threads or teams requested by the user.

The argument *flag* indicates whether the code for the parallel region is inlined into the application or invoked by the runtime and also whether the region is a parallel or teams region.

The argument *codeptr_ra* is used to relate the implementation of an OpenMP region back to its source code. In cases where a runtime routine implements the region associated with this callback, *codeptr_ra* is expected to contain the return address of the call to the runtime routine. In cases where the implementation of this feature is inlined, *codeptr_ra* is expected to contain the return address of the invocation of this callback. In cases where attribution to source code is impossible or inappropriate, *codeptr_ra* may be **NULL**.

Cross References

- **omp_data_t** type, see Section [4.2.3.4.3](#) on page [434](#).
- **omp_frame_t** type, see Section [4.4.1.2](#) on page [589](#).
- **omp_parallel_flag_t** type, see Section [4.2.3.4.20](#) on page [441](#).

1 4.2.4.2.4 `ompt_callback_parallel_end_t`

2 **Format**

```
3 typedef void (*ompt_callback_parallel_end_t) (  
4     ompt_data_t *parallel_data,  
5     ompt_data_t *encountering_task_data,  
6     int flag,  
7     const void *codeptr_ra  
8 );
```

C / C++

9 **Trace Record**

```
10 typedef struct ompt_record_parallel_end_t {  
11     ompt_id_t parallel_id;  
12     ompt_id_t encountering_task_id;  
13     int flag;  
14     const void *codeptr_ra;  
15 } ompt_record_parallel_end_t;
```

C / C++

16 **Description of Arguments**

17 The binding of argument *parallel_data* is the parallel or teams region that is ending.

18 The binding of argument *encountering_task_data* is the encountering task.

19 The argument *flag* indicates whether the execution of the parallel region is inlined into the
20 application or invoked by the runtime and also whether the region is a parallel or teams region.

21 The argument *codeptr_ra* is used to relate the implementation of an OpenMP region back to its
22 source code. In cases where a runtime routine implements the region associated with this callback,
23 *codeptr_ra* is expected to contain the return address of the call to the runtime routine. In cases
24 where the implementation of this feature is inlined, *codeptr_ra* is expected to contain the return
25 address of the invocation of this callback. In cases where attribution to source code is impossible or
26 inappropriate, *codeptr_ra* may be **NULL**.

Cross References

- `ompt_data_t` type signature, see Section [4.2.3.4.3](#) on page [434](#).
- `ompt_parallel_flag_t` type signature, see Section [4.2.3.4.20](#) on page [441](#).

4.2.4.2.5 `ompt_callback_master_t`

Format

C / C++

```
typedef void (*ompt_callback_master_t) (  
    ompt_scope_endpoint_t endpoint,  
    ompt_data_t *parallel_data,  
    ompt_data_t *task_data,  
    const void *codeptr_ra  
);
```

C / C++

Trace Record

C / C++

```
typedef struct ompt_record_master_t {  
    ompt_scope_endpoint_t endpoint;  
    ompt_id_t parallel_id;  
    ompt_id_t task_id;  
    const void *codeptr_ra;  
} ompt_record_master_t;
```

C / C++

Description of Arguments

The argument *endpoint* indicates whether the callback is signalling the beginning or the end of a scope.

The binding of argument *parallel_data* is the current parallel region.

The binding of argument *task_data* is the encountering task.

The argument *codeptr_ra* is used to relate the implementation of an OpenMP region back to its source code. In cases where a runtime routine implements the region associated with this callback, *codeptr_ra* is expected to contain the return address of the call to the runtime routine. In cases where the implementation of this feature is inlined, *codeptr_ra* is expected to contain the return address of the invocation of this callback. In cases where attribution to source code is impossible or inappropriate, *codeptr_ra* may be **NULL**.

Cross References

- `ompt_data_t` type signature, see Section 4.2.3.4.3 on page 434.
- `ompt_scope_endpoint_t` type, see Section 4.2.3.4.10 on page 437.

4.2.4.2.6 `ompt_callback_task_create_t`

Format

```
C / C++
typedef void (*ompt_callback_task_create_t) (
    ompt_data_t *encountering_task_data,
    const omp_frame_t *encountering_task_frame,
    ompt_data_t *new_task_data,
    int flag,
    int has_dependences,
    const void *codeptr_ra
);
C / C++
```

1

Trace Record

C / C++

2

```

typedef struct omp_t_record_task_create_t {
  omp_id_t  encountering_task_id;
  omp_id_t  new_task_id;
  int  flag;
  int  has_dependences;
  const void *codeptr_ra;
} omp_t_record_task_create_t;

```

3

4

5

6

7

8

C / C++

9

Description of Arguments

10

The binding of argument *encountering_task_data* is the encountering task. This parameter is **NULL** for an initial task.

11

12

The argument *encountering_task_frame* points to the frame object associated with the encountering task. This parameter is **NULL** for an initial task.

13

14

The binding of argument *new_task_data* is the created task.

15

The argument *flag* indicates the kind of the task: initial, explicit or target. Values for *flag* are composed by or-ing elements of enum **omp_task_flag_t**.

16

17

The argument *has_dependences* indicates whether created task has dependences.

18

The argument *codeptr_ra* is used to relate the implementation of an OpenMP region back to its source code. In cases where a runtime routine implements the region associated with this callback, *codeptr_ra* is expected to contain the return address of the call to the runtime routine. In cases where the implementation of this feature is inlined, *codeptr_ra* is expected to contain the return address of the invocation of this callback. In cases where attribution to source code is impossible or inappropriate, *codeptr_ra* may be **NULL**.

19

20

21

22

23

24

Cross References

25

- **omp_data_t** type, see Section [4.2.3.4.3](#) on page [434](#).

26



- **omp_frame_t** type, see Section [4.4.1.2](#) on page [589](#).

27



- **omp_task_flag_t** type, see Section [4.2.3.4.17](#) on page [440](#).

1 4.2.4.2.7 `ompt_callback_task_dependences_t`

2 **Format**

3  C / C++ 
4 `typedef void (*ompt_callback_task_dependences_t) (
5 ompt_data_t *task_data,
6 const ompt_task_dependence_t *deps,
7 int ndeps
8);`

8 **Trace Record**

9  C / C++ 
10 `typedef struct ompt_record_task_dependences_t {
11 ompt_id_t task_id;
12 ompt_task_dependence_t dep;
13 int ndeps;
14 } ompt_record_task_dependences_t;`

14 **Description of Arguments**

15 The binding of argument `task_data` is the task being created.

16 The argument `deps` lists all dependences of a new task.

17 The argument `ndeps` specifies the length of the list. The memory for `deps` is owned by the caller;
18 the tool cannot rely on the data after the callback returns.


19 The performance monitor interface for tracing activity on target devices will provide one record per
20 dependence.

21 **Cross References**


- 22 • `ompt_data_t` type, see Section [4.2.3.4.3](#) on page [434](#).
- 23 • `ompt_task_dependence_t` type, see Section [4.2.3.4.8](#) on page [436](#).

1 4.2.4.2.8 ompt_callback_task_dependence_t

2 Format

```
3  C / C++  
4 typedef void (*ompt_callback_task_dependence_t) (  
5     ompt_data_t *src_task_data,  
6     ompt_data_t *sink_task_data  
7 );
```

7 Trace Record

```
8  C / C++  
9 typedef struct ompt_record_task_dependence_t {  
10     ompt_id_t src_task_id;  
11     ompt_id_t sink_task_id;  
12 } ompt_record_task_dependence_t;
```

12 Description of Arguments

13 The binding of argument *src_task_data* is a running task with an outgoing dependence.


14 The binding of argument *sink_task_data* is a task with an unsatisfied incoming dependence.

15 Cross References

- 16 • `ompt_data_t` type signature, see Section 4.2.3.4.3 on page 434.

17 4.2.4.2.9 ompt_callback_task_schedule_t

18 Format

```
19  C / C++  
20 typedef void (*ompt_callback_task_schedule_t) (  
21     ompt_data_t *prior_task_data,  
22     ompt_task_status_t prior_task_status,  
23     ompt_data_t *next_task_data  
24 );
```

 C / C++

1 Trace Record

C / C++

```
2 typedef struct ompt_record_task_schedule_t {  
3     ompt_id_t prior_task_id;  
4     ompt_task_status_t prior_task_status;  
5     ompt_id_t next_task_id;  
6 } ompt_record_task_schedule_t;
```

C / C++

7 Description of Arguments

8 The argument *prior_task_status* indicates the status of the task that arrived at a task scheduling
9 point.

10 The binding of argument *prior_task_data* is the task that arrived at the scheduling point.

11 The binding of argument *next_task_data* is the task that will resume at the scheduling point.

12 Cross References

- 13 • `ompt_data_t` type, see Section [4.2.3.4.3](#) on page [434](#).
- 14 • `ompt_task_status_t` type, see Section [4.2.3.4.18](#) on page [440](#).

15 4.2.4.2.10 `ompt_callback_implicit_task_t`

16 Format

C / C++

```
17 typedef void (*ompt_callback_implicit_task_t) (  
18     ompt_scope_endpoint_t endpoint,  
19     ompt_data_t *parallel_data,  
20     ompt_data_t *task_data,  
21     unsigned int actual_parallelism,  
22     unsigned int index  
23 );
```

C / C++

1

Trace Record

C / C++

2

```

typedef struct ompt_record_implicit_t {
  ompt_scope_endpoint_t endpoint;
  ompt_id_t parallel_id;
  ompt_id_t task_id;
  unsigned int actual_parallelism;
  unsigned int index;
} ompt_record_implicit_t;

```

3

4

5

6

7

8

C / C++

9

Description of Arguments

10

The argument *endpoint* indicates whether the callback is signalling the beginning or the end of a scope.

11

12

The binding of argument *parallel_data* is the current parallel region. For the *implicit-task-end* event, this argument is **NULL**.

13

14

The binding of argument *task_data* is the implicit task executing the parallel region's structured block.

15

16

The argument *actual_parallelism* indicates the number of threads in the **parallel** region, respectively the number of teams in the **teams** region. For the *implicit-task-end* and the *initial-task-end* events, this argument is **0**.

17

18

19

The argument *index* indicates the thread number or team number of the calling thread, within the team or league executing the parallel or teams region to which the implicit region binds.

20

21

Cross References

22

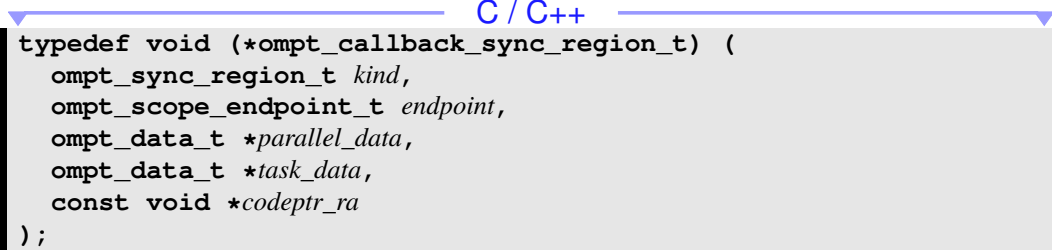
- **ompt_data_t** type, see Section [4.2.3.4.3](#) on page [434](#).

23

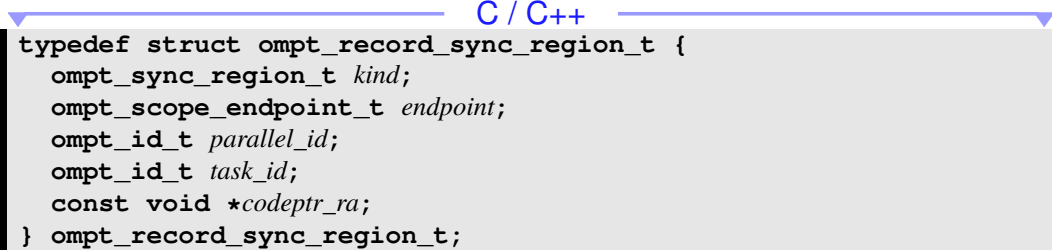
- **ompt_scope_endpoint_t** enumeration type, see Section [4.2.3.4.10](#) on page [437](#).

1 4.2.4.2.11 `ompt_callback_sync_region_t`

2 Format

3  C / C++
4 `typedef void (*ompt_callback_sync_region_t) (
5 ompt_sync_region_t kind,
6 ompt_scope_endpoint_t endpoint,
7 ompt_data_t *parallel_data,
8 ompt_data_t *task_data,
9 const void *codeptr_ra
10);`

10 Trace Record

11  C / C++
12 `typedef struct ompt_record_sync_region_t {
13 ompt_sync_region_t kind;
14 ompt_scope_endpoint_t endpoint;
15 ompt_id_t parallel_id;
16 ompt_id_t task_id;
17 const void *codeptr_ra;
18 } ompt_record_sync_region_t;`

18 Description of Arguments

19 The argument *kind* indicates the kind of synchronization.

20 The argument *endpoint* indicates whether the callback is signalling the beginning or the end of a
21 scope.

22 The binding of argument *parallel_data* is the current parallel region. For the *barrier-end* event at
23 the end of a parallel region, this argument is **NULL**.

24 The binding of argument *task_data* is the current task.

25 The argument *codeptr_ra* is used to relate the implementation of an OpenMP region back to its
26 source code. In cases where a runtime routine implements the region associated with this callback,
27 *codeptr_ra* is expected to contain the return address of the call to the runtime routine. In cases
28 where the implementation of this feature is inlined, *codeptr_ra* is expected to contain the return
29 address of the invocation of this callback. In cases where attribution to source code is impossible or
30 inappropriate, *codeptr_ra* may be **NULL**.

Cross References

- `ompt_data_t` type, see Section 4.2.3.4.3 on page 434.
- `ompt_sync_region_t` type, see Section 4.2.3.4.12 on page 437.
- `ompt_scope_endpoint_t` type, see Section 4.2.3.4.10 on page 437.

4.2.4.2.12 `ompt_callback_mutex_acquire_t`

Format

C / C++

```
typedef void (*ompt_callback_mutex_acquire_t) (  
    ompt_mutex_t kind,  
    unsigned int hint,  
    unsigned int impl,  
    omp_wait_id_t wait_id,  
    const void *codeptr_ra  
);
```

C / C++

Trace Record

C / C++

```
typedef struct ompt_record_mutex_acquire_t {  
    ompt_mutex_t kind;  
    unsigned int hint;  
    unsigned int impl;  
    omp_wait_id_t wait_id;  
    const void *codeptr_ra;  
} ompt_record_mutex_acquire_t;
```

C / C++

Description of Arguments

The argument *kind* indicates the kind of the lock.

The argument *hint* indicates the hint provided when initializing an implementation of mutual exclusion. If no hint is available when a thread initiates acquisition of mutual exclusion, the runtime may supply `omp_sync_hint_none` as the value for *hint*.

The argument *impl* indicates the mechanism chosen by the runtime to implement the mutual exclusion.

The argument *wait_id* indicates the object being awaited.

The argument *codeptr_ra* is used to relate the implementation of an OpenMP region back to its source code. In cases where a runtime routine implements the region associated with this callback, *codeptr_ra* is expected to contain the return address of the call to the runtime routine. In cases where the implementation of this feature is inlined, *codeptr_ra* is expected to contain the return address of the invocation of this callback. In cases where attribution to source code is impossible or inappropriate, *codeptr_ra* may be `NULL`.

Cross References

- `omp_wait_id_t` type, see Section 4.4.1.3 on page 591.
- `ompt_mutex_t` type, see Section 4.2.3.4.15 on page 439.

4.2.4.2.13 `ompt_callback_mutex_t`

Format

```

C / C++
typedef void (*ompt_callback_mutex_t) (
    ompt_mutex_t kind,
    ompt_wait_id_t wait_id,
    const void *codeptr_ra
);
```

C / C++

Trace Record

```
typedef struct omp_t_record_mutex_t {  
    omp_mutex_t kind;  
    omp_wait_id_t wait_id;  
    const void *codeptr_ra;  
} omp_t_record_mutex_t;
```

Description of Arguments

The argument *kind* indicates the kind of mutual exclusion event.

The argument *wait_id* indicates the object being awaited.

The argument *codeptr_ra* is used to relate the implementation of an OpenMP region back to its source code. In cases where a runtime routine implements the region associated with this callback, *codeptr_ra* is expected to contain the return address of the call to the runtime routine. In cases where the implementation of this feature is inlined, *codeptr_ra* is expected to contain the return address of the invocation of this callback. In cases where attribution to source code is impossible or inappropriate, *codeptr_ra* may be **NULL**.

Cross References

- `omp_wait_id_t` type signature, see Section 4.4.1.3 on page 591.
- `omp_mutex_t` type signature, see Section 4.2.3.4.15 on page 439.

4.2.4.2.14 `omp_callback_nest_lock_t`

Format

```
typedef void (*omp_callback_nest_lock_t) (  
    omp_scope_endpoint_t endpoint,  
    omp_wait_id_t wait_id,  
    const void *codeptr_ra  
);
```

Trace Record

C / C++

```
typedef struct ompt_record_nest_lock_t {
    ompt_scope_endpoint_t endpoint;
    ompt_wait_id_t wait_id;
    const void *codeptr_ra;
} ompt_record_nest_lock_t;
```

C / C++

Description of Arguments

The argument *endpoint* indicates whether the callback is signalling the beginning or the end of a scope.

The argument *wait_id* indicates the object being awaited.

The argument *codeptr_ra* is used to relate the implementation of an OpenMP region back to its source code. In cases where a runtime routine implements the region associated with this callback, *codeptr_ra* is expected to contain the return address of the call to the runtime routine. In cases where the implementation of this feature is inlined, *codeptr_ra* is expected to contain the return address of the invocation of this callback. In cases where attribution to source code is impossible or inappropriate, *codeptr_ra* may be **NULL**.

Cross References

- `ompt_wait_id_t` type signature, see Section 4.4.1.3 on page 591.
- `ompt_scope_endpoint_t` type signature, see Section 4.2.3.4.10 on page 437.

4.2.4.2.15 `ompt_callback_work_t`

Format

C / C++

```
typedef void (*ompt_callback_work_t) (
    ompt_work_t wstype,
    ompt_scope_endpoint_t endpoint,
    ompt_data_t *parallel_data,
    ompt_data_t *task_data,
    uint64_t count,
    const void *codeptr_ra
);
```

C / C++

1

Trace Record

C / C++

2

```

typedef struct ompt_record_work_t {
  ompt_work_t wstype;
  ompt_scope_endpoint_t endpoint;
  ompt_id_t parallel_id;
  ompt_id_t task_id;
  uint64_t count;
  const void *codeptr_ra;
} ompt_record_work_t;

```

3

4

5

6

7

8

9

C / C++

10

Description of Arguments

11

The argument *wstype* indicates the kind of worksharing region.

12

The argument *endpoint* indicates whether the callback is signalling the beginning or the end of a scope.

13

14

The binding of argument *parallel_data* is the current parallel region.

15

The binding of argument *task_data* is the current task.

16

The argument *count* is a measure of the quantity of work involved in the worksharing construct. For a worksharing-loop construct, *count* represents the number of iterations of the loop. For a **taskloop** construct, *count* represents the number of iterations in the iteration space, which may be the result of collapsing several associated loops. For a **sections** construct, *count* represents the number of sections. For a **workshare** construct, *count* represents the units of work, as defined by the **workshare** construct. For a **single** construct, *count* is always 1. When the *endpoint* argument is signaling the end of a scope, a *count* value of 0 indicates that the actual *count* value is not available.

17

18

19

20

21

22

23

24

The argument *codeptr_ra* is used to relate the implementation of an OpenMP region back to its source code. In cases where a runtime routine implements the region associated with this callback, *codeptr_ra* is expected to contain the return address of the call to the runtime routine. In cases where the implementation of this feature is inlined, *codeptr_ra* is expected to contain the return address of the invocation of this callback. In cases where attribution to source code is impossible or inappropriate, *codeptr_ra* may be **NULL**.

25

26

27

28

29

Cross References

- worksharing constructs, see Section 2.11 on page 85 and Section 2.12.2 on page 100.
- `ompt_data_t` type signature, see Section 4.2.3.4.3 on page 434.
- `ompt_scope_endpoint_t` type signature, see Section 4.2.3.4.10 on page 437.
- `ompt_work_t` type signature, see Section 4.2.3.4.14 on page 438.

4.2.4.2.16 `ompt_callback_flush_t`

Format

C / C++

```
typedef void (*ompt_callback_flush_t) (  
    ompt_data_t *thread_data,  
    const void *codeptr_ra  
);
```

C / C++

Trace Record

C / C++

```
typedef struct ompt_record_flush_t {  
    const void *codeptr_ra;  
} ompt_record_flush_t;
```

C / C++

Description of Arguments

The argument `codeptr_ra` is used to relate the implementation of an OpenMP region back to its source code. In cases where a runtime routine implements the region associated with this callback, `codeptr_ra` is expected to contain the return address of the call to the runtime routine. In cases where the implementation of this feature is inlined, `codeptr_ra` is expected to contain the return address of the invocation of this callback. In cases where attribution to source code is impossible or inappropriate, `codeptr_ra` may be **NULL**.

Cross References

- `ompt_data_t` type signature, see Section 4.2.3.4.3 on page 434.

1 4.2.4.2.17 `ompt_callback_dispatch_t`

2 Format

```
3 typedef void (*ompt_callback_dispatch_t) (  
4     ompt_data_t *parallel_data,  
5     ompt_data_t *task_data,  
6     ompt_dispatch_t kind,  
7     ompt_data_t instance  
8 );
```

C / C++

9 Trace Record

```
10 typedef struct ompt_record_dispatch_t {  
11     ompt_id_t parallel_id;  
12     ompt_id_t task_id;  
13     ompt_dispatch_t kind;  
14     ompt_data_t instance;  
15 } ompt_record_dispatch_t;
```

C / C++

16 Description of Arguments

17 The argument *kind* indicates whether a loop iteration or a section is being dispatched.

18 For a loop iteration, the argument *instance.value* contains the iteration variable value. For a
19 structured block in the **sections** construct, *instance.ptr* contains a code address identifying the
20 structured block. In cases where a runtime routine implements the structured block associated with
21 this callback, *instance.ptr* is expected to contain the return address of the call to the runtime
22 routine. In cases where the implementation of the structured block is inlined, *instance.ptr* is
23 expected to contain the return address of the invocation of this callback.

24 Cross References

- 25 • `ompt_data_t` type signature, see Section [4.2.3.4.3](#) on page [434](#).
- 26 • `ompt_dispatch_t` type, see Section [4.2.3.4.11](#) on page [437](#).

1 4.2.4.2.18 ompt_callback_target_t

2 Format

C / C++

```
3 typedef void (*ompt_callback_target_t) (  
4     ompt_target_t kind,  
5     ompt_scope_endpoint_t endpoint,  
6     uint64_t device_num,  
7     ompt_data_t *task_data,  
8     ompt_id_t target_id,  
9     const void *codeptr_ra  
10 );
```

C / C++

11 Trace Record

C / C++

```
12 typedef struct ompt_record_target_t {  
13     ompt_target_t kind;  
14     ompt_scope_endpoint_t endpoint;  
15     uint64_t device_num;  
16     ompt_data_t *task_data;  
17     ompt_id_t target_id;  
18     const void *codeptr_ra;  
19 } ompt_record_target_t;
```

C / C++

Description of Arguments

The argument *kind* indicates the kind of target region.

The argument *endpoint* indicates whether the callback is signalling the beginning or the end of a scope.

The argument *device_num* indicates the id of the device which will execute the target region.

The binding of argument *task_data* is the generating task.

The binding of argument *target_id* is the target region.

The argument *codeptr_ra* is used to relate the implementation of an OpenMP region back to its source code. In cases where a runtime routine implements the region associated with this callback, *codeptr_ra* is expected to contain the return address of the call to the runtime routine. In cases where the implementation of this feature is inlined, *codeptr_ra* is expected to contain the return address of the invocation of this callback. In cases where attribution to source code is impossible or inappropriate, *codeptr_ra* may be **NULL**.

Cross References

- `ompt_id_t` type, see Section [4.2.3.4.2](#) on page [434](#).
- `ompt_data_t` type signature, see Section [4.2.3.4.3](#) on page [434](#).
- `ompt_scope_endpoint_t` type signature, see Section [4.2.3.4.10](#) on page [437](#).
- `ompt_target_t` type signature, see Section [4.2.3.4.19](#) on page [441](#).

4.2.4.2.19 `ompt_callback_device_load_t`

Summary

The OpenMP runtime invokes this callback to notify a tool immediately after loading code onto the specified device.

Format

```

1  typedef void (*ompt_callback_device_load_t) (
2      uint64_t device_num,
3      const char *filename,
4      int64_t offset_in_file,
5      void *vma_in_file,
6      size_t bytes
7      void *host_addr,
8      void *device_addr,
9      uint64_t module_id
10 );

```

11 **ompt_addr_none** is defined as a pointer with the value ~0.

12 Description of Arguments

13 The argument *device_num* specifies the device.

14 The argument *filename* indicates the name of a file in which the device code can be found. A NULL
15 *filename* indicates that the code is not available in a file in the file system.

16 The argument *offset_in_file* indicates an offset into *filename* at which the code can be found. A
17 value of -1 indicates that no offset is provided.

18 The argument *vma_in_file* indicates an virtual address in *filename* at which the code can be found.
19 A value of *ompt_addr_none* indicates that a virtual address in the file is not available.

20 The argument *bytes* indicates the size of the device code object in bytes.

21 The argument *host_addr* indicates where a copy of the device code is available in host memory. A
22 value of *ompt_addr_none* indicates that a host code address is not available.

23 The argument *device_addr* indicates where the device code has been loaded in device memory. A
24 value of *ompt_addr_none* indicates that a device code address is not available.

25 The argument *module_id* is an identifier that is associated with the device code object.

26 4.2.4.2.20 ompt_callback_device_unload_t

27 Summary

28 The OpenMP runtime invokes this callback to notify a tool immediately prior to unloading code
29 from the specified device.

1

Format

C / C++

2

```

typedef void (*ompt_callback_device_unload_t) (
    uint64_t device_num,
    uint64_t module_id
);

```

3

4

5

C / C++

6

Description of Arguments

7

The argument *device_num* specifies the device.

8

The argument *module_id* is an identifier that is associated with the device code object.

9 4.2.4.2.21 ompt_callback_target_data_op_t

10

Format

C / C++

11

```

typedef void (*ompt_callback_target_data_op_t) (
    ompt_id_t target_id,
    ompt_id_t host_op_id,
    ompt_target_data_op_t optype,
    void *src_addr,
    int src_device_num,
    void *dest_addr,
    int dest_device_num,
    size_t bytes,
    const void *codeptr_ra
);

```

12

13

14

15

16

17

18

19

20

21

C / C++

1

Trace Record

C / C++

2

```
typedef struct ompt_record_target_data_op_t {  
3     ompt_id_t host_op_id;  
4     ompt_target_data_op_t optype;  
5     void *src_addr;  
6     int src_device_num;  
7     void *dest_addr;  
8     int dest_device_num;  
9     size_t bytes;  
10    ompt_device_time_t end_time;  
11    const void *codeptr_ra;  
12 } ompt_record_target_data_op_t;
```

3

4

5

6

7

8

9

10

11

12

C / C++

13

Description

14

An OpenMP implementation will dispatch a registered `ompt_callback_target_data_op` callback when device memory is allocated or freed, as well as when data is copied to or from a device.

15

16

17

Note – An OpenMP implementation may aggregate program variables and data operations upon them. For instance, an OpenMP implementation may synthesize a composite to represent multiple scalars and then allocate, free, or copy this composite as a whole rather than performing data operations on each scalar individually. For that reason, a tool should not expect to see separate data operations on each variable.

18

19

20

21

Description of Arguments

The argument *host_op_id* is a unique identifier for a data operations on a target device.

The argument *optype* indicates the kind of data mapping.

The argument *src_addr* indicates the address of data before the operation, where applicable.

The argument *src_device_num* indicates the source device number for the data operation, where applicable.

The argument *dest_addr* indicates the address of data after the operation.

The argument *dest_device_num* indicates the destination device number for the data operation.

It is implementation defined whether in some operations *src_addr* or *dest_addr* might point to an intermediate buffer.

The argument *bytes* indicates the size of data.

The argument *codeptr_ra* is used to relate the implementation of an OpenMP region back to its source code. In cases where a runtime routine implements the region associated with this callback, *codeptr_ra* is expected to contain the return address of the call to the runtime routine. In cases where the implementation of this feature is inlined, *codeptr_ra* is expected to contain the return address of the invocation of this callback. In cases where attribution to source code is impossible or inappropriate, *codeptr_ra* may be **NULL**.

Cross References

- `ompt_id_t` type, see Section [4.2.3.4.2](#) on page [434](#).
- `ompt_target_data_op_t` type signature, see Section [4.2.3.4.13](#) on page [438](#).

1 4.2.4.2.22 `ompt_callback_target_map_t`

2 **Format**

C / C++

```
3 typedef void (*ompt_callback_target_map_t) (  
4     ompt_id_t target_id,  
5     unsigned int nitems,  
6     void **host_addr,  
7     void **device_addr,  
8     size_t *bytes,  
9     unsigned int *mapping_flags,  
10    const void *codeptr_ra  
11 );
```

C / C++

12 **Trace Record**

C / C++

```
13 typedef struct ompt_record_target_map_t {  
14     ompt_id_t target_id;  
15     unsigned int nitems;  
16     void **host_addr;  
17     void **device_addr;  
18     size_t *bytes;  
19     unsigned int *mapping_flags;  
20     const void *codeptr_ra;  
21 } ompt_record_target_map_t;
```

C / C++

22 **Description**

23 An instance of a **target**, **target data**, or **target enter data** construct may contain
24 one or more **map** clauses. An OpenMP implementation may report the set of mappings associated
25 with **map** clauses for a construct with a single **ompt_callback_target_map** callback to
26 report the effect of all mappings or multiple **ompt_callback_target_map** callbacks with
27 each reporting a subset of the mappings. Furthermore, an OpenMP implementation may omit
28 mappings that it determines are unnecessary. If an OpenMP implementation issues multiple
29 **ompt_callback_target_map** callbacks, these callbacks may be interleaved with
30 **ompt_callback_target_data_op** callbacks used to report data operations associated with
31 the mappings.

Description of Arguments

The binding of argument *target_id* is the target region.

The argument *nitems* indicates the number of data mappings being reported by this callback.

The argument *host_addr* indicates an array of addresses of data on host side.

The argument *device_addr* indicates an array of addresses of data on device side.

The argument *bytes* indicates an array of size of data.

The argument *mapping_flags* indicates the kind of data mapping. Flags for a mapping include one or more values specified by the type `ompt_target_map_flag_t`.

The argument *codeptr_ra* is used to relate the implementation of an OpenMP region back to its source code. In cases where a runtime routine implements the region associated with this callback, *codeptr_ra* is expected to contain the return address of the call to the runtime routine. In cases where the implementation of this feature is inlined, *codeptr_ra* is expected to contain the return address of the invocation of this callback. In cases where attribution to source code is impossible or inappropriate, *codeptr_ra* may be `NULL`.

Cross References

- `ompt_id_t` type, see Section 4.2.3.4.2 on page 434.
- `ompt_callback_target_data_op_t`, see Section 4.2.4.2.21 on page 468.
- `ompt_target_map_flag_t` type, see Section 4.2.3.4.21 on page 442.

4.2.4.2.23 `ompt_callback_target_submit_t`

Format

```
typedef void (*ompt_callback_target_submit_t) (  
    ompt_id_t target_id,  
    ompt_id_t host_op_id,  
    unsigned int requested_num_teams  
);
```

Description

A thread dispatches a registered `ompt_callback_target_submit` callback on the host when a target task creates an initial task on a target device.

Description of Arguments

The argument *target_id* is a unique identifier for the associated target region.

The argument *host_op_id* is a unique identifier for the initial task on the target device.

The argument *requested_num_teams* is the number of teams that the host is requesting to execute the kernel. The actual number of teams that execute the kernel may be smaller and generally won't be known until the kernel begins to execute on the device.

Constraints on Arguments

The argument *target_id* indicates the instance of the target construct to which the computation belongs.

The argument *host_op_id* provides a unique host-side identifier that represents the computation on the device.

Trace Record

```
C / C++
typedef struct ompt_record_target_kernel_t {
    ompt_id_t host_op_id;
    unsigned int requested_num_teams;
    unsigned int granted_num_teams;
    ompt_device_time_t end_time;
} ompt_record_target_kernel_t;
```

If a tool has configured a device to trace kernel execution using `ompt_set_trace_ompt`, the device will log a `ompt_record_target_kernel_t` record in a trace. The fields in the record are as follows:

- The *host_op_id* field contains a unique identifier that a tool can use to correlate a `ompt_record_target_kernel_t` record with its associated `ompt_callback_target_submit` callback on the host.
- The *requested_num_teams* field contains the number of teams that the host requested to execute the kernel.
- The *granted_num_teams* field contains the number of teams that the device actually used to execute the kernel.
- The time when the initial task began execution on the device is recorded in the *time* field of an enclosing `ompt_record_t` structure; the time when the initial task completed execution on the device is recorded in the *end_time* field.

Cross References

- `ompt_id_t` type, see Section [4.2.3.4.2](#) on page 434.

4.2.4.2.24 `ompt_callback_buffer_request_t`

Summary

The OpenMP runtime will invoke a callback with type signature `ompt_callback_buffer_request_t` to request a buffer to store event records for a device.

Format

```
C / C++
typedef void (*ompt_callback_buffer_request_t) (
    uint64_t device_num,
    ompt_buffer_t **buffer,
    size_t *bytes
);
```

C / C++

Description

This callback requests a buffer to store trace records for the specified device.

A buffer request callback may set `*bytes` to 0 if it does not want to provide a buffer for any reason. If a callback sets `*bytes` to 0, further recording of events for the device will be disabled until the next invocation of `ompt_start_trace`. This will cause the device to drop future trace records until recording is restarted.

The buffer request callback is not required to be *async signal safe*.

Description of Arguments

The argument `device_num` specifies the device.

A tool should set `*buffer` to point to a buffer where device events may be recorded and `*bytes` to the length of that buffer.

Cross References

- `ompt_buffer_t` type, see Section [4.2.3.4.6](#) on page 435.

1 4.2.4.2.25 `ompt_callback_buffer_complete_t`

2 **Summary**

3 A device triggers a call to `ompt_callback_buffer_complete_t` when no further records
4 will be recorded in an event buffer and all records written to the buffer are valid.

5 **Format**

```
6 C / C++  
7 typedef void (*ompt_callback_buffer_complete_t) (  
8     uint64_t device_num,  
9     ompt_buffer_t *buffer,  
10    size_t bytes,  
11    ompt_buffer_cursor_t begin,  
12    int buffer_owned  
    );
```

13 **Description**

14 This callback provides a tool with a buffer containing trace records for the specified device.
15 Typically, a tool will iterate through the records in the buffer and process them.

16 The OpenMP implementation will make these callbacks on a thread that is not an OpenMP master
17 or worker.

18 The callee may delete the buffer if the argument `buffer_owned=0`.

19 The buffer completion callback is not required to be *async signal safe*.

20 **Description of Arguments**

21 The argument `device_num` indicates the device whose events the buffer contains.

22 The argument `buffer` is the address of a buffer previously allocated by a *buffer request* callback.

23 The argument `bytes` indicates the full size of the buffer.

24 The argument `begin` is an opaque cursor that indicates the position at the beginning of the first
25 record in the buffer.

26 The argument `buffer_owned` is 1 if the data pointed to by `buffer` can be deleted by the callback and
27 0 otherwise. If multiple devices accumulate trace events into a single buffer, this callback might be
28 invoked with a pointer to one or more trace records in a shared buffer with `buffer_owned = 0`. In
29 this case, the callback may not delete the buffer.

Cross References

- `ompt_buffer_t` type, see Section 4.2.3.4.6 on page 435.
- `ompt_buffer_cursor_t` type, see Section 4.2.3.4.7 on page 435.

4.2.4.2.26 `ompt_callback_control_tool_t`

Format

```
C / C++
typedef int (*ompt_callback_control_tool_t) (
    uint64_t command,
    uint64_t modifier,
    void *arg,
    const void *codeptr_ra
);
C / C++
```

Description

The tool control callback may return any non-negative value, which will be returned to the application by the OpenMP implementation as the return value of the `omp_control_tool` call that triggered the callback.

Description of Arguments

The argument *command* passes a command from an application to a tool. Standard values for *command* are defined by `omp_control_tool_t`, defined in Section 3.8 on page 413.

The argument *modifier* passes a command modifier from an application to a tool.

This callback allows tool-specific values for *command* and *modifier*. Tools must ignore *command* values that they are not explicitly designed to handle.

The argument *arg* is a void pointer that enables a tool and an application to pass arbitrary state back and forth. The argument *arg* may be **NULL**.

The argument *codeptr_ra* is used to relate the implementation of an OpenMP region back to its source code. In cases where a runtime routine implements the region associated with this callback, *codeptr_ra* is expected to contain the return address of the call to the runtime routine. In cases where the implementation of this feature is inlined, *codeptr_ra* is expected to contain the return address of the invocation of this callback. In cases where attribution to source code is impossible or inappropriate, *codeptr_ra* may be **NULL**.

Constraints on Arguments

Tool-specific values for *command* must be ≥ 64 .

Cross References

- `omp_control_tool_t` enumeration type, see Section 3.8 on page 413.

4.2.4.2.27 `ompt_callback_cancel_t`

Format

```
                                     C / C++
typedef void (*ompt_callback_cancel_t) (
    ompt_data_t *task_data,
    int flags,
    const void *codeptr_ra
);
```

C / C++

Description of Arguments

The argument *task_data* corresponds to the task encountering a **cancel** construct, a **cancellation point** construct, or a construct defined as having an implicit cancellation point.

The argument *flags*, defined by the enumeration `ompt_cancel_flag_t`, indicates whether the cancel is activated by the current task, or detected as being activated by another task. The construct being canceled is also described in the *flags*. When several constructs are detected as being concurrently canceled, each corresponding bit in the flags will be set.

The argument *codeptr_ra* is used to relate the implementation of an OpenMP region back to its source code. In cases where a runtime routine implements the region associated with this callback, *codeptr_ra* is expected to contain the return address of the call to the runtime routine. In cases where the implementation of this feature is inlined, *codeptr_ra* is expected to contain the return address of the invocation of this callback. In cases where attribution to source code is impossible or inappropriate, *codeptr_ra* may be **NULL**.

Cross References

- `omp_cancel_flag_t` enumeration type, see Section 4.2.3.4.23 on page 443.

1 4.2.4.2.28 `ompt_callback_device_initialize_t`

2 Summary

3 The tool callback with type signature `ompt_callback_device_initialize_t` initializes a
4 tool's tracing interface for a device.

5 Format

```
6 typedef void (*ompt_callback_device_initialize_t) (  
7     uint64_t device_num,  
8     const char *type,  
9     ompt_device_t *device,  
10    ompt_function_lookup_t lookup,  
11    const char *documentation  
12 );
```

C / C++

13 Description

14 A tool that wants to asynchronously collect a trace of activities on a device should register a
15 callback with type signature `ompt_callback_device_initialize_t` for the
16 `ompt_callback_device_initialize` OpenMP event. An OpenMP implementation will
17 invoke this callback for a device after OpenMP is initialized for the device but before beginning
18 execution of any OpenMP construct on the device.

19 Description of Arguments

20 The argument *device_num* identifies the logical device being initialized.

21 The argument *type* is a character string indicating the type of the device. A device type string is a
22 semicolon separated character string that includes at a minimum the vendor and model name of the
23 device. This may be followed by a semicolon-separated sequence of properties that describe a
24 device's hardware or software.

25 The argument *device* is a pointer to an opaque object that represents the target device instance. The
26 pointer to the device instance object is used by functions in the device tracing interface to identify
27 the device being addressed.

28 The argument *lookup* is a pointer to a runtime callback that a tool must use to obtain pointers to
29 runtime entry points in the device's OMPT tracing interface. If a device does not support tracing, it
30 should provide **NULL** for *lookup*.

1 The argument *documentation* is a string that describes how to use any device-specific runtime entry
2 points that can be obtained using *lookup*. This documentation string could simply be a pointer to
3 external documentation, or it could be inline descriptions that includes names and type signatures
4 for any device-specific interfaces that are available through *lookup* along with descriptions of how
5 to use these interface functions to control monitoring and analysis of device traces.

6 Constraints on Arguments

7 The arguments *type* and *documentation* must be immutable strings that are defined for the lifetime
8 of a program execution.

9 Effect

10 A tool's device initializer has several duties. First, it should use *type* to determine whether the tool
11 has any special knowledge about a device's hardware and/or software. Second, it should use *lookup*
12 to look up pointers to runtime entry points in the OMPT tracing interface for the device. Finally,
13 using these runtime entry points, it can then set up tracing for a device.

14 Initializing tracing for a target device is described in section Section [4.2.1.4](#) on page [425](#).

15 Cross References

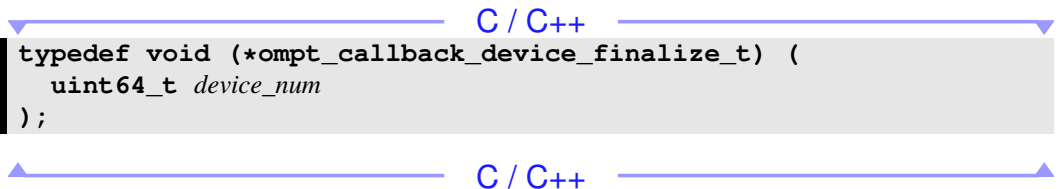
- 16 • `ompt_function_lookup_t`, see Section [4.2.5.3.1](#) on page [516](#).

17 4.2.4.2.29 `ompt_callback_device_finalize_t`

18 Summary

19 The tool callback with type signature `ompt_callback_device_finalize_t` finalizes a
20 tool's tracing interface for a device.

21 Format

```
22  C / C++  
23 typedef void (*ompt_callback_device_finalize_t) (  
24     uint64_t device_num  
25 );
```

1 **Description of Arguments**

2 The argument *device_num* identifies the logical device being finalized.

3 **Description**

4 An OpenMP implementation dispatches a finalization callback for a device immediately prior to
5 finalizing the device. Prior to dispatching a finalization callback for a device on which tracing is
6 active, the OpenMP implementation will stop tracing on the device and synchronously flush all
7 trace records for the device that have not yet been reported to the tool. If any trace records for the
8 device need to be flushed, the OpenMP implementation will issue one or more buffer completion
9 callbacks with type signature `ompt_callback_buffer_complete_t` as needed.

10 **Cross References**

- 11 • `ompt_callback_buffer_complete_t`, see Section [4.2.4.2.25](#) on page [475](#).

12 **4.2.5 OMPT Runtime Entry Points for Tools**

13 The OMPT interface supports two principal sets of runtime entry points for tools. One set of
14 runtime entry points enables a tool to register callbacks for OpenMP events and to inspect the state
15 of an OpenMP thread while executing in a tool callback or a signal handler. The second set of
16 runtime entry points enables a tool to trace activities on a device. When directed by the tracing
17 interface, an OpenMP implementation will trace activities on a device, collect buffers full of trace
18 records, and invoke callbacks on the host to process these records. Runtime entry points for tools in
19 an OpenMP implementation should not be global symbols since tools cannot rely on the visibility
20 of such symbols in general.

21 In addition, the OMPT interface supports runtime entry points for two classes of lookup routines.
22 The first class of lookup routines contains a single member: a routine that returns runtime entry
23 points in the OMPT callback interface. The second class of lookup routines includes a unique
24 lookup routine for each kind of device that can return runtime entry points in a device's OMPT
25 tracing interface.

26 **Restrictions**

27 Calling an OMPT runtime entry point from a signal handler before a *thread-begin* or after a
28 *thread-end* event on a thread results in unspecified behavior.

29 Calling an OMPT device runtime entry point after a *device-finalize* event for that device results in
30 unspecified behavior.

1 4.2.5.1 Entry Points in the OMPT Callback Interface

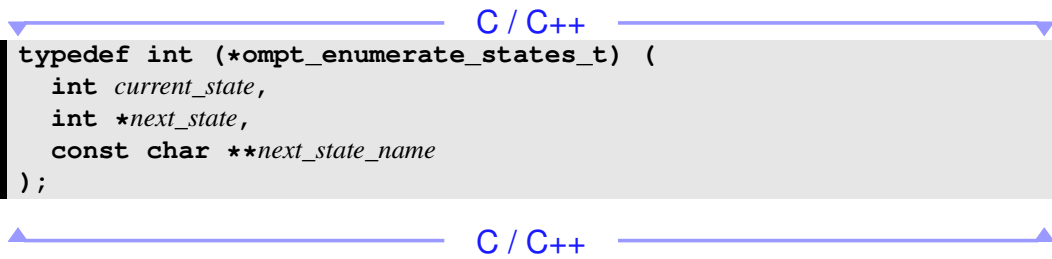
2 Entry points in the OMPT callback interface enable a tool to register callbacks for OpenMP events
3 and to inspect the state of an OpenMP thread while executing in a tool callback or a signal handler.
4 A tool obtains pointers to these runtime entry points using the lookup function passed to the tool's
5 initializer for the callback interface.

6 4.2.5.1.1 `ompt_enumerate_states_t`

7 Summary

8 A runtime entry point known as `ompt_enumerate_states` with type signature
9 `ompt_enumerate_states_t` enumerates the thread states supported by an OpenMP
10 implementation.

11 Format

12 
13 `typedef int (*ompt_enumerate_states_t) (
14 int current_state,
15 int *next_state,
16 const char **next_state_name
17);`

17 Description

18 An OpenMP implementation may support only a subset of the states defined by the
19 `omp_state_t` enumeration type. In addition, an OpenMP implementation may support
20 implementation-specific states. The `ompt_enumerate_states` runtime entry point enables a
21 tool to enumerate the thread states supported by an OpenMP implementation.

22 When a thread state supported by an OpenMP implementation is passed as the first argument to the
23 runtime entry point, the runtime entry point will assign the next thread state in the enumeration to
24 the variable passed by reference as the runtime entry point's second argument and assign the name
25 associated with the next thread state to the character pointer passed by reference as the third
26 argument.

27 Whenever one or more states are left in the enumeration, the enumerate states runtime entry point
28 will return 1. When the last state in the enumeration is passed as the first argument, the runtime
29 entry point will return 0 indicating that the enumeration is complete.

Description of Arguments

The argument *current_state* must be a thread state supported by the OpenMP implementation. To begin enumerating the states that an OpenMP implementation supports, a tool should pass **omp_state_undefined** as *current_state*. Subsequent invocations of the runtime entry point by the tool should pass the value assigned to the variable passed by reference as the second argument to the previous call.

The argument *next_state* is a pointer to an integer where the entry point will return the value of the next state in the enumeration.

The argument *next_state_name* is a pointer to a character string pointer, where the entry point will return a string describing the next state.

Constraints on Arguments

Any string returned through the argument *next_state_name* must be immutable and defined for the lifetime of a program execution.

Cross References

- **omp_state_t**, see Section 4.4.1.1 on page 584.

4.2.5.1.2 **ompt_enumerate_mutex_impls_t**

Summary

A runtime entry point known as **ompt_enumerate_mutex_impls** with type signature **ompt_enumerate_mutex_impls_t** enumerates the kinds of mutual exclusion implementations that an OpenMP implementation employs.

Format

```
C / C++
typedef int (*ompt_enumerate_mutex_impls_t) (
    int current_impl,
    int *next_impl,
    const char **next_impl_name
);
```

ompt_mutex_impl_none is defined as an integer with the value 0.

Description

An OpenMP implementation may implement mutual exclusion for locks, nest locks, critical sections, and atomic regions in several different ways. The `ompt_enumerate_mutex_impls` runtime entry point enables a tool to enumerate the kinds of mutual exclusion implementations that an OpenMP implementation employs. The value `ompt_mutex_impl_none` is reserved to indicate an invalid implementation.

When a mutex kind supported by an OpenMP implementation is passed as the first argument to the runtime entry point, the runtime entry point will assign the next mutex kind in the enumeration to the variable passed by reference as the runtime entry point's second argument and assign the name associated with the next mutex kind to the character pointer passed by reference as the third argument.

Whenever one or more mutex kinds are left in the enumeration, the runtime entry point to enumerate mutex implementations will return 1. When the last mutex kind in the enumeration is passed as the first argument, the runtime entry point will return 0 indicating that the enumeration is complete.

Description of Arguments

The argument `current_impl` must be a mutex implementation kind supported by an OpenMP implementation. To begin enumerating the mutex implementation kinds that an OpenMP implementation supports, a tool should pass `ompt_mutex_impl_none` as the first argument of the enumerate mutex kinds runtime entry point. Subsequent invocations of the runtime entry point by the tool should pass the value assigned to the variable passed by reference as the second argument to the previous call.

The argument `next_impl` is a pointer to an integer where the entry point will return the value of the next mutex implementation in the enumeration.

The argument `next_impl_name` is a pointer to a character string pointer, where the entry point will return a string describing the next mutex implementation.

Constraints on Arguments

Any string returned through the argument `next_impl_name` must be immutable and defined for the lifetime of a program execution.

4.2.5.1.3 `ompt_set_callback_t`

Summary

A runtime entry point known as `ompt_set_callback` with type signature `ompt_set_callback_t` registers a pointer to a tool callback that an OpenMP implementation will invoke when a host OpenMP event occurs.

Format

```
typedef int (*ompt_set_callback_t) (  
    ompt_callbacks_t which,  
    ompt_callback_t callback  
);
```

Description

OpenMP implementations can inform tools about events that occur during the execution of an OpenMP program using callbacks. To register a tool callback for an OpenMP event on the current device, a tool uses the runtime entry point known as `ompt_set_callback` with type signature `ompt_set_callback_t`.

The return value of the `ompt_set_callback` runtime entry point may indicate several possible outcomes. Callback registration may fail if it is called outside the initializer for the callback interface, returning `ompt_set_error`. Otherwise, the return value of `ompt_set_callback` indicates whether *dispatching* a callback leads to its invocation. A return value of `ompt_set_never` indicates that the callback will never be invoked at runtime. A return value of `ompt_set_sometimes` indicates that the callback will be invoked at runtime for an implementation-defined subset of associated event occurrences. A return value of `ompt_set_sometimes_paired` is similar to `ompt_set_sometimes`, but provides an additional guarantee for callbacks with an *endpoint* parameter. Namely, it guarantees that a callback with an *endpoint* value of `ompt_scope_begin` is invoked if and only if the same callback with *endpoint* value of `ompt_scope_end` will also be invoked sometime in the future. A return value of `ompt_set_always` indicates that the callback will be always invoked at runtime for associated event occurrences.

Description of Arguments

The argument *which* indicates the callback being registered.

The argument *callback* is a tool callback function.

A tool may pass a **NULL** value for *callback* to disable any callback associated with *which*. If disabling was successful, `ompt_set_always` is returned.

Constraints on Arguments

When a tool registers a callback for an event, the type signature for the callback must match the type signature appropriate for the event.

TABLE 4.4: Return codes for `ompt_set_callback` and `ompt_set_trace_ompt`.

```
typedef enum ompt_set_result_t {
    ompt_set_error          = 0,
    ompt_set_never         = 1,
    ompt_set_sometimes     = 2,
    ompt_set_sometimes_paired = 3,
    ompt_set_always        = 4
} ompt_set_result_t;
```

Cross References

- `ompt_callbacks_t` enumeration type, see Section 4.2.3.2 on page 430.
- `ompt_callback_t` type, see Section 4.2.3.4.1 on page 433.
- `ompt_get_callback_t` host callback type signature, see Section 4.2.5.1.4 on page 485.

4.2.5.1.4 `ompt_get_callback_t`

Summary

A runtime entry point known as `ompt_get_callback` with type signature `ompt_get_callback_t` retrieves a pointer to a tool callback routine (if any) that an OpenMP implementation will invoke when an OpenMP event occurs.

Format

```
C / C++
typedef int (*ompt_get_callback_t) (
    ompt_callbacks_t which,
    ompt_callback_t *callback
);
```

C / C++

1
2
3
4
5
6
7
8

Description

A tool uses the runtime entry point known as **ompt_get_callback** with type signature **ompt_get_callback_t** to obtain a pointer to the tool callback that an OpenMP implementation will invoke when a host OpenMP event occurs. If a non-**NULL** tool callback is registered for the specified event, the pointer to the tool callback will be assigned to the variable passed by reference as the second argument and the entry point will return 1; otherwise, it will return 0. If the entry point returns 0, the value of the variable passed by reference as the second argument is undefined.

9
10
11

Description of Arguments

The argument *which* indicates the callback being inspected.
The argument *callback* returns a pointer to the callback being inspected.

12
13

Constraints on Arguments

The second argument passed to the entry point must be a reference to a variable of specified type.

14
15
16
17

Cross References

- **ompt_callbacks_t** enumeration type, see Section 4.2.3.2 on page 430.
- **ompt_callback_t** type, see Section 4.2.3.4.1 on page 433.
- **ompt_set_callback_t** type signature, see Section 4.2.5.1.3 on page 483.

18 4.2.5.1.5 ompt_get_thread_data_t

19
20
21

Summary

A runtime entry point known as **ompt_get_thread_data** with type signature **ompt_get_thread_data_t** returns the address of the thread data object for the current thread.

22

Format

23

```

| _____ C / C++ _____
| typedef ompt_data_t>(*ompt_get_thread_data_t) (void);
| _____ C / C++ _____

```

1 **Binding**

2 The binding thread for runtime entry point known as `ompt_get_thread_data` is the current
3 thread.

4 **Description**

5 Each OpenMP thread has an associated thread data object of type `ompt_data_t`. A tool uses the
6 runtime entry point known as `ompt_get_thread_data` with type signature
7 **`ompt_get_thread_data_t`** to obtain a pointer to the thread data object, if any, associated
8 with the current thread.

9 A tool may use a pointer to an OpenMP thread's data object obtained from this runtime entry point
10 to inspect or modify the value of the data object. When an OpenMP thread is created, its data
11 object will be initialized with value `ompt_data_none`.

12 This runtime entry point is *async signal safe*.

13 **Cross References**

- 14 • `ompt_data_t` type, see Section [4.2.3.4.3](#) on page [434](#).

15 **4.2.5.1.6 ompt_get_num_procs_t**

16 **Summary**

17 A runtime entry point known as `ompt_get_num_procs` with type signature
18 **`ompt_get_num_procs_t`** returns the number of processors currently available to the execution
19 environment on the host device.

20 **Format**

21

```

C / C++
| typedef int (*ompt_get_num_procs_t) (void);
C / C++

```

22 **Binding**

23 The binding thread set for runtime entry point known as `ompt_get_num_procs` is all threads
24 on a device.

Description

The `ompt_get_num_procs` runtime entry point returns the number of processors that are available on the host device at the time the routine is called. This value may change between the time that it is determined and the time that it is read in the calling context due to system actions outside the control of the OpenMP implementation.

This runtime entry point is *async signal safe*.

4.2.5.1.7 `ompt_get_num_places_t`

Summary

A runtime entry point known as `ompt_get_num_places` with type signature `ompt_get_num_places_t` returns the number of places available to the execution environment in the place list.

Format

```
▼────────────────────────────────── C / C++ ───────────────────────────────────▼  
| typedef int (*ompt_get_num_places_t) (void); |  
▲────────────────────────────────── C / C++ ───────────────────────────────────▲
```

Binding

The binding thread set for the runtime entry point known as `ompt_get_num_places` is all threads on a device.

Description

The runtime entry point known as `ompt_get_num_places` returns the number of places in the place list. This value is equivalent to the number of places in the *place-partition-var* ICV in the execution environment of the initial task.

This runtime entry point is *async signal safe*.

Cross References

- *place-partition-var* ICV, see Section 2.4 on page 47.
- `OMP_PLACES` environment variable, see Section 5.5 on page 599.

1 4.2.5.1.8 `ompt_get_place_proc_ids_t`

2 **Summary**

3 A runtime entry point known as `ompt_get_place_proc_ids` with type signature
4 `ompt_get_place_proc_ids_t` returns the numerical identifiers of the processors available
5 to the execution environment in the specified place.

6 **Format**

```
7 C / C++  
8 typedef int (*ompt_get_place_proc_ids_t) (  
9     int place_num,  
10    int ids_size,  
11    int *ids  
    );  
12 C / C++
```

12 **Binding**

13 The binding thread set for the runtime entry point known as `ompt_get_place_proc_ids` is
14 all threads on a device.

15 **Description**

16 The runtime entry point known as `ompt_get_place_proc_ids` with type signature
17 `ompt_get_place_proc_ids_t` returns the numerical identifiers of each processor associated
18 with the specified place. The numerical identifiers returned are non-negative, and their meaning is
19 implementation defined.

20 **Description of Arguments**

21 The argument `place_num` specifies the place being queried.

22 The argument `ids_size` indicates the size of the result array specified by argument `ids`.

23 The argument `ids` is an array where the routine can return a vector of processor identifiers in the
24 specified place.

1 **Effect**

2 If the array *ids* of size *ids_size* is large enough to contain all identifiers, they are returned in *ids* and
3 their order in the array is implementation defined.

4 Otherwise, if the *ids* array is too small, the values in *ids* when the function returns are unspecified.

5 In both cases, the routine returns the number of numerical identifiers available to the execution
6 environment in the specified place.

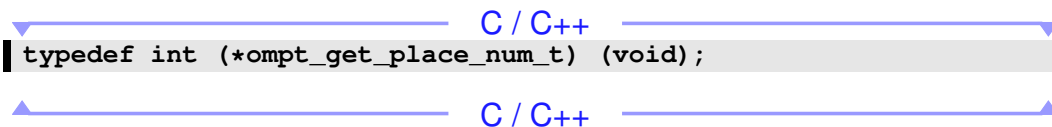
7 **4.2.5.1.9 ompt_get_place_num_t**

8 **Summary**

9 A runtime entry point known as **ompt_get_place_num** with type signature
10 **ompt_get_place_num_t** returns the place number of the place to which the current thread is
11 bound.

12 **Format**

```
13       | typedef int (*ompt_get_place_num_t) (void);
```



14 **Binding**

15 The binding thread set for the runtime entry point known as **ompt_get_place_num** is the
16 current thread.

17 **Description**

18 When the current thread is bound to a place, the runtime entry point known as
19 **ompt_get_place_num** returns the place number associated with the thread. The returned value
20 is between 0 and one less than the value returned by runtime entry point known as
21 **ompt_get_num_places**, inclusive. When the current thread is not bound to a place, the routine
22 returns -1.

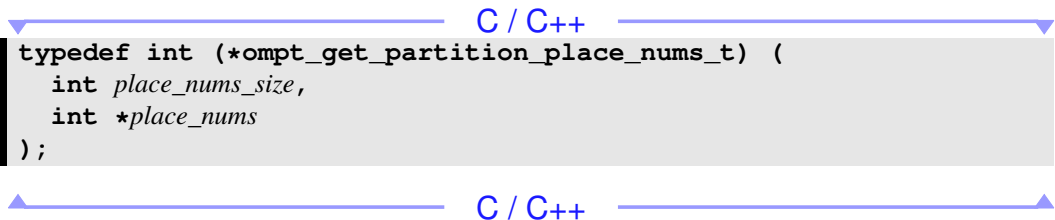
23 This runtime entry point is *async signal safe*.

1 4.2.5.1.10 `ompt_get_partition_place_nums_t`

2 **Summary**

3 A runtime entry point known as `ompt_get_partition_place_nums` with type signature
4 `ompt_get_partition_place_nums_t` returns the list of place numbers corresponding to
5 the places in the *place-partition-var* ICV of the innermost implicit task.

6 **Format**

```
7  C / C++  
8 typedef int (*ompt_get_partition_place_nums_t) (  
9     int place_nums_size,  
10    int *place_nums  
11 );
```

11 **Binding**

12 The binding task set for the runtime entry point known as
13 `ompt_get_partition_place_nums` is the current implicit task.

14 **Description**

15 The runtime entry point known as `ompt_get_partition_place_nums` with type signature
16 `ompt_get_partition_place_nums_t` returns the list of place numbers corresponding to
17 the places in the *place-partition-var* ICV of the innermost implicit task.

18 This runtime entry point is *async signal safe*.

19 **Description of Arguments**

20 The argument *place_nums_size* indicates the size of the result array specified by argument
21 *place_nums*.

22 The argument *place_nums* is an array where the routine can return a vector of place identifiers.

Effect

If the array *place_nums* of size *place_nums_size* is large enough to contain all identifiers, they are returned in *place_nums* and their order in the array is implementation defined.

Otherwise, if the *place_nums* array is too small, the values in *place_nums* when the function returns are unspecified.

In both cases, the routine returns the number of places in the *place-partition-var* ICV of the innermost implicit task.

Cross References

- *place-partition-var* ICV, see Section 2.4 on page 47.
- `OMP_PLACES` environment variable, see Section 5.5 on page 599.

4.2.5.1.11 `ompt_get_proc_id_t`

Summary

A runtime entry point known as `ompt_get_proc_id` with type signature `ompt_get_proc_id_t` returns the numerical identifier of the processor of the current thread.

Format

```
typedef int (*ompt_get_proc_id_t) (void);
```

Binding

The binding thread set for the runtime entry point known as `ompt_get_proc_id` is the current thread.

Description

The runtime entry point known as `ompt_get_proc_id` returns the numerical identifier of the processor of the current thread. A defined numerical identifier is non-negative, and its meaning is implementation defined. A negative number indicates a failure to retrieve the numerical identifier.

This runtime entry point is *async signal safe*.

1 4.2.5.1.12 `ompt_get_state_t`

2 Summary

3 A runtime entry point known as `ompt_get_state` with type signature `ompt_get_state_t`
4 returns the state and the wait identifier of the current thread.

5 Format

```
6 typedef int (*ompt_get_state_t) (  
7     omp_wait_id_t *wait_id  
8 );
```

C / C++

9 Binding

10 The binding thread for runtime entry point known as `ompt_get_state` is the current thread.

11 Description

12 Each OpenMP thread has an associated state and a wait identifier. If a thread's state indicates that
13 the thread is waiting for mutual exclusion, the thread's wait identifier will contain an opaque handle
14 that indicates the data object upon which the thread is waiting.

15 To retrieve the state and wait identifier for the current thread, a tool uses the runtime entry point
16 known as `ompt_get_state` with type signature `ompt_get_state_t`.

17 The returned value may be any one of the states predefined by `omp_state_t` or a value that
18 represents any implementation specific state. The tool may obtain a string representation for each
19 state with the function known as `ompt_enumerate_states`.

20 If the returned state indicates that the thread is waiting for a lock, nest lock, critical section, atomic
21 region, or ordered region the value of the thread's wait identifier will be assigned to a non-**NULL**
22 wait identifier passed as an argument.

23 This runtime entry point is *async signal safe*.

24 Description of Arguments

25 The argument `wait_id` is a pointer to an opaque handle available to receive the value of the thread's
26 wait identifier. If the `wait_id` pointer is not **NULL**, the entry point will assign the value of the
27 thread's wait identifier `*wait_id`. If the returned state is not one of the specified wait states, the
28 value of `*wait_id` is undefined after the call.

1 **Constraints on Arguments**

2 The argument passed to the entry point must be a reference to a variable of the specified type or
3 **NULL**.

4 **Cross References**

- 5 • `omp_wait_id_t` type, see Section 4.4.1.3 on page 591.
- 6 • `omp_state_t` type, see Section 4.4.1.1 on page 584.
- 7 • `ompt_enumerate_states_t` type, see Section 4.2.5.1.1 on page 481.

8 **4.2.5.1.13 ompt_get_parallel_info_t**

9 **Summary**

10 A runtime entry point known as `ompt_get_parallel_info` with type signature
11 `ompt_get_parallel_info_t` returns information about the parallel region, if any, at the
12 specified ancestor level for the current execution context.

13 **Format**

C / C++

```

14 typedef int (*ompt_get_parallel_info_t) (
15     int ancestor_level,
16     ompt_data_t **parallel_data,
17     int *team_size
18 );

```

C / C++

Description

During execution, an OpenMP program may employ nested parallel regions. To obtain information about a parallel region, a tool uses the runtime entry point known as `ompt_get_parallel_info` with type signature `ompt_get_parallel_info_t`. This runtime entry point can be used to obtain information about the current parallel region, if any, and any enclosing parallel regions for the current execution context.

The entry point returns 2 if there is a parallel region at the specified ancestor level and the information is available, 1 if there is a parallel region at the specified ancestor level but the information is currently unavailable, and 0 otherwise.

A tool may use the pointer to a parallel region's data object that it obtains from this runtime entry point to inspect or modify the value of the data object. When a parallel region is created, its data object will be initialized with the value `ompt_data_none`.

This runtime entry point is *async signal safe*.

Description of Arguments

The argument *ancestor_level* specifies the parallel region of interest to a tool by its ancestor level. Ancestor level 0 refers to the innermost parallel region; information about enclosing parallel regions may be obtained using larger ancestor levels.

The argument *parallel_data* returns the parallel data if the argument is not **NULL**.

The argument *team_size* returns the team size if the argument is not **NULL**.

Effect

If the runtime entry point returns 0 or 1, no argument is modified. Otherwise, the entry point has the effects described below.

If a non-**NULL** value was passed for *parallel_data*, the value returned in **parallel_data* is a pointer to a data word associated with the parallel region at the specified level.

If a non-**NULL** value was passed for *team_size*, the value returned in **team_size* is the number of threads in the team associated with the parallel region.

Constraints on Arguments

While argument *ancestor_level* is passed by value, all other arguments to the entry point must be pointers to variables of the specified types or **NULL**.

Restrictions

Between a *parallel-begin* event and an *implicit-task-begin* event, a call to `ompt_get_parallel_info(0, ...)` may return information about the outer parallel team, the new parallel team or an inconsistent state.

If a thread is in the state `omp_state_wait_barrier_implicit_parallel`, a call to `ompt_get_parallel_info` may return a pointer to a copy of the specified parallel region's *parallel_data* rather than a pointer to the data word for the region itself. This convention enables the master thread for a parallel region to free storage for the region immediately after the region ends, yet avoid having some other thread in the region's team potentially reference the region's *parallel_data* object after it has been freed.

Cross References

- `ompt_data_t` type, see Section 4.2.3.4.3 on page 434.

4.2.5.1.14 `ompt_get_task_info_t`

Summary

A runtime entry point known as `ompt_get_task_info` with type signature `ompt_get_task_info_t` provides information about the task, if any, at the specified ancestor level in the current execution context.

Format

```
                                     C / C++
typedef int (*ompt_get_task_info_t) (
    int ancestor_level,
    int *flag,
    ompt_data_t **task_data,
    omp_frame_t **task_frame,
    ompt_data_t **parallel_data,
    int *thread_num
);
```

C / C++

Description

During execution, an OpenMP thread may be executing an OpenMP task. Additionally, the thread's stack may contain procedure frames associated with suspended OpenMP tasks or OpenMP runtime system routines. To obtain information about any task on the current thread's stack, a tool uses the runtime entry point known as `ompt_get_task_info` with type signature `ompt_get_task_info_t`.

Ancestor level 0 refers to the active task; information about other tasks with associated frames present on the stack in the current execution context may be queried at higher ancestor levels.

The `ompt_get_task_info` runtime entry point returns 2 if there is a task region at the specified ancestor level and the information is available, 1 if there is a task region at the specified ancestor level but the information is currently unavailable, and 0 otherwise.

If a task exists at the specified ancestor level and the information is available, information will be returned in the variables passed by reference to the entry point. If no task region exists at the specified ancestor level or the information is unavailable, the values of variables passed by reference to the entry point will be undefined when the entry point returns.

A tool may use a pointer to a data object for a task or parallel region that it obtains from this runtime entry point to inspect or modify the value of the data object. When either a parallel region or a task region is created, its data object will be initialized with the value `ompt_data_none`.

This runtime entry point is *async signal safe*.

Description of Arguments

The argument `ancestor_level` specifies the task region of interest to a tool by its ancestor level. Ancestor level 0 refers to the active task; information about ancestor tasks found in the current execution context may be queried at higher ancestor levels.

The argument `flag` returns the task type if the argument is not `NULL`.

The argument `task_data` returns the task data if the argument is not `NULL`.

The argument `task_frame` returns the task frame pointer if the argument is not `NULL`.

The argument `parallel_data` returns the parallel data if the argument is not `NULL`.

The argument `thread_num` returns the thread number if the argument is not `NULL`.

Effect

If the runtime entry point returns 0 or 1, no argument is modified. Otherwise, the entry point has the effects described below.

If a non-**NULL** value was passed for *flag*, the value returned in **flag* represents the type of the task at the specified level. Task types that a tool may observe on a thread's stack include initial, implicit, explicit, and target tasks.

If a non-**NULL** value was passed for *task_data*, the value returned in **task_data* is a pointer to a data word associated with the task at the specified level.

If a non-**NULL** value was passed for *task_frame*, the value returned in **task_frame* is a pointer to the **omp_frame_t** structure associated with the task at the specified level. Appendix B discusses an example that illustrates the use of **omp_frame_t** structures with multiple threads and nested parallelism.

If a non-**NULL** value was passed for *parallel_data*, the value returned in **parallel_data* is a pointer to a data word associated with the parallel region containing the task at the specified level. If the task at the specified level is an initial task, the value of **parallel_data* will be **NULL**.

If a non-**NULL** value was passed for *thread_num*, the value returned in **thread_num* indicates the number of the thread in the parallel region executing the task.

Constraints on Arguments

While argument *ancestor_level* is passed by value, all other arguments to the entry point must be pointers to variables of the specified types or **NULL**.

Cross References

- **ompt_data_t** type, see Section 4.2.3.4.3 on page 434.
- **omp_frame_t** type, see Section 4.4.1.2 on page 589.
- **ompt_task_flag_t** type, see Section 4.2.3.4.17 on page 440.

4.2.5.1.15 **ompt_get_task_memory_t**

Summary

A runtime entry point known as **ompt_get_task_memory** with type signature **ompt_get_task_memory_t** provides information about memory ranges that are associated with the task at task creation to store the data environment of the task for the execution.

Format

C / C++

```
typedef int (*ompt_get_task_memory_t) (  
    void **addr,  
    size_t *size,  
    int block  
);
```

C / C++

Description

During execution, an OpenMP thread may be executing an OpenMP task. The OpenMP implementation needs to preserve the data environment from the creation of the task to the execution of the task. To obtain information about the memory ranges used to store the data environment for the current task, a tool uses the runtime entry point known as `ompt_get_task_memory` with type signature `ompt_get_task_memory_t`.

There might be multiple memory ranges used to store these data. The `block` argument allows the tool to iterate over these memory ranges.

The `ompt_get_task_memory` runtime entry point returns 1 if there are more memory ranges available, and 0 otherwise.

If there is no memory used for a task, `size` is set to 0. In this case, `addr` is unspecified.

This runtime entry point is *async signal safe*.

Description of Arguments

The argument `addr` is a pointer to a void pointer return value to provide the start address of a memory block.

The argument `size` is a pointer to a size type return value to provide the size of the memory block.

The argument `block` is an integer value to specify the memory block of interest.

4.2.5.1.16 `ompt_get_target_info_t`

Summary

A runtime entry point known as `ompt_get_target_info` with type signature `ompt_get_target_info_t` returns identifiers that specify a thread's current target region and target operation id, if any.

Format

C / C++

```
typedef int (*ompt_get_target_info_t) (  
    uint64_t *device_num,  
    ompt_id_t *target_id,  
    ompt_id_t *host_op_id  
);
```

C / C++

Description

A tool can query whether an OpenMP thread is in a target region by invoking the entry point known as `ompt_get_target_info` with type signature `ompt_get_target_info_t`. This runtime entry point returns 1 if the current thread is in a target region and 0 otherwise. If the entry point returns 0, the values of the variables passed by reference as its arguments are undefined.

If the current thread is in a target region, the entry point will return information about the current device, active target region, and active host operation, if any.

This runtime entry point is *async signal safe*.

Description of Arguments

If the host is in a **target** region, `device_num` returns the target device.

If the host is in a **target** region, `target_id` returns the **target** region identifier.

If the current thread is in the process of initiating an operation on a target device (e.g., copying data to or from an accelerator or launching a kernel) `host_op_id` returns the identifier for the operation; otherwise, `host_op_id` returns `ompt_id_none`.

Constraints on Arguments

Arguments passed to the entry point must be valid references to variables of the specified types.

Cross References

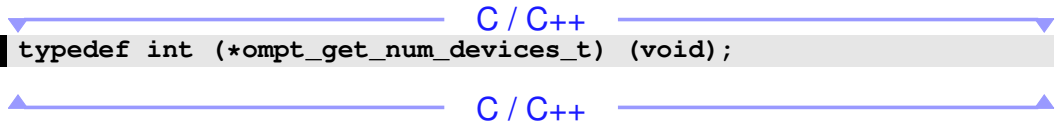
- `ompt_id_t` type, see Section [4.2.3.4.2](#) on page [434](#).

1 4.2.5.1.17 `ompt_get_num_devices_t`

2 **Summary**

3 A runtime entry point known as `ompt_get_num_devices` with type signature
4 `ompt_get_num_devices_t` returns the number of available devices.

5 **Format**

6  `typedef int (*ompt_get_num_devices_t) (void);`

The diagram shows a blue arrow pointing down from the text "C / C++" above the code, and a blue arrow pointing up from the text "C / C++" below the code, both pointing to the code line.

7 **Description**

8 An OpenMP program may execute on one or more devices. A tool may determine the number of
9 devices available to an OpenMP program by invoking a runtime entry point known as
10 `ompt_get_num_devices` with type signature `ompt_get_num_devices_t`.

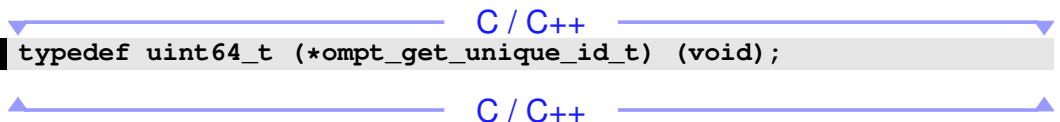
11 This runtime entry point is *async signal safe*.

12 4.2.5.1.18 `ompt_get_unique_id_t`

13 **Summary**

14 A runtime entry point known as `ompt_get_unique_id` with type signature
15 `ompt_get_unique_id_t` returns a unique number.

16 **Format**

17  `typedef uint64_t (*ompt_get_unique_id_t) (void);`

The diagram shows a blue arrow pointing down from the text "C / C++" above the code, and a blue arrow pointing up from the text "C / C++" below the code, both pointing to the code line.

18 **Description**

19 A tool may obtain a number that is unique for the duration of an OpenMP program by invoking a
20 runtime entry point known as `ompt_get_unique_id` with type signature
21 `ompt_get_unique_id_t`. Successive invocations may not result in consecutive or even
22 increasing numbers.

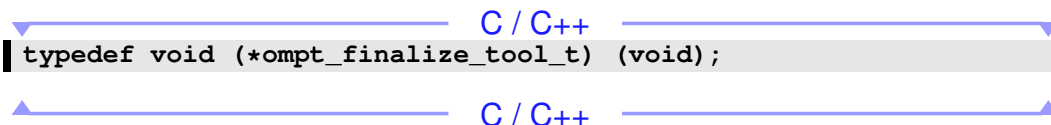
23 This runtime entry point is *async signal safe*.

1 4.2.5.1.19 `ompt_finalize_tool_t`

2 **Summary**

3 A runtime entry point known as `ompt_finalize_tool` with type signature
4 `ompt_finalize_tool_t` enables the tool to finalize itself.

5 **Format**

6 The code snippet is enclosed in a light gray box. Above the box, a blue double-headed arrow spans the width of the code, with "C / C++" centered above it. Below the box, another blue double-headed arrow spans the width, with "C / C++" centered below it.

```
typedef void (*ompt_finalize_tool_t) (void);
```

7 **Description**

8 A tool may detect that the execution of an OpenMP program is ending before the OpenMP
9 implementation does. To facilitate clean termination of the tool, the tool may invoke a runtime
10 entry point known as `ompt_finalize_tool` with type signature `ompt_finalize_tool_t`.
11 Upon completion of the `ompt_finalize_tool` routine, no OMPT callbacks are dispatched.

12 **Effect**

13 The `ompt_finalize_tool` routine detaches the tool from the runtime and unregisters all
14 callbacks and invalidates all OMPT entry points passed to the tool in the lookup-function. Upon
15 completion of the `ompt_finalize_tool` routine, no further callbacks on any thread will be
16 issued.

17 Before the callbacks get unregistered, the OpenMP runtime should make all efforts to dispatch all
18 outstanding registered callbacks as well as dispatch callbacks that would be encountered during
19 shutdown of the runtime, if possible in the current execution context.

20 4.2.5.2 **Entry Points in the OMPT Device Tracing Interface**

21 4.2.5.2.1 `ompt_get_device_num_procs_t`

22 **Summary**

23 A runtime entry point for a device known as `ompt_get_device_num_procs` with type
24 signature `ompt_get_device_num_procs_t` returns the number of processors currently
25 available to the execution environment on the specified device.

1 **Format**

```
2        C / C++  
3        typedef int (*ompt_get_device_num_procs_t) (  
4            ompt_device_t *device  
5        );
```

C / C++

5 **Description**

6 A runtime entry point for a device known as `ompt_get_device_num_procs` with type
7 signature `ompt_get_device_num_procs_t` returns the number of processors that are
8 available on the device at the time the routine is called. This value may change between the time
9 that it is determined and the time that it is read in the calling context due to system actions outside
10 the control of the OpenMP implementation.

11 **Description of Arguments**

12 The argument *device* is a pointer to an opaque object that represents the target device instance. The
13 pointer to the device instance object is used by functions in the device tracing interface to identify
14 the device being addressed.

15 **Cross References**

- 16 • `ompt_device_t`, see Section [4.2.3.4.4](#) on page [435](#).

17 **4.2.5.2.2 ompt_get_device_time_t**

18 **Summary**

19 A runtime entry point for a device known as `ompt_get_device_time` with type signature
20 `ompt_get_device_time_t` returns the current time on the specified device.

21 **Format**

```
22        C / C++  
23        typedef ompt_device_time_t (*ompt_get_device_time_t) (  
24            ompt_device_t *device  
25        );
```

C / C++

Description

Host and target devices are typically distinct and run independently. If host and target devices are different hardware components, they may use different clock generators. For this reason, there may be no common time base for ordering host-side and device-side events.

A runtime entry point for a device known as `ompt_get_device_time` with type signature `ompt_get_device_time_t` returns the current time on the specified device. A tool can use this information to align time stamps from different devices.

Description of Arguments

The argument *device* is a pointer to an opaque object that represents the target device instance. The pointer to the device instance object is used by functions in the device tracing interface to identify the device being addressed.

Cross References

- `ompt_device_t`, see Section [4.2.3.4.4](#) on page [435](#).
- `ompt_device_time_t`, see Section [4.2.3.4.5](#) on page [435](#).

4.2.5.2.3 `ompt_translate_time_t`

Summary

A runtime entry point for a device known as `ompt_translate_time` with type signature `ompt_translate_time_t` translates a time value obtained from the specified device to a corresponding time value on the host device.

Format

```

C / C++
typedef double (*ompt_translate_time_t) (
    ompt_device_t *device,
    ompt_device_time_t time
);
C / C++
```


Description

A runtime entry point for a device known as `ompt_translate_time` with type signature `ompt_translate_time_t` translates a time value obtained from the specified device to a corresponding time value on the host device. The returned value for the host time has the same meaning as the value returned from `omp_get_wtime`.

Note – The accuracy of time translations may degrade if they are not performed promptly after a device time value is received if either the host or device vary their clock speeds. Prompt translation of device times to host times is recommended.

Description of Arguments

The argument *device* is a pointer to an opaque object that represents the target device instance. The pointer to the device instance object is used by functions in the device tracing interface to identify the device being addressed.

The argument *time* is a time from the specified device.

Cross References

- `ompt_device_t`, see Section 4.2.3.4.4 on page 435.
- `ompt_device_time_t`, see Section 4.2.3.4.5 on page 435.

4.2.5.2.4 `ompt_set_trace_ompt_t`

Summary

A runtime entry point for a device known as `ompt_set_trace_ompt` with type signature `ompt_set_trace_ompt_t` enables or disables the recording of trace records for one or more types of OMPT events.

Format

```
C / C++
typedef int (*ompt_set_trace_ompt_t) (
    ompt_device_t *device,
    unsigned int enable,
    unsigned int etype
);
```

C / C++

Description of Arguments

The argument *device* is a pointer to an opaque object that represents the target device instance. The pointer to the device instance object is used by functions in the device tracing interface to identify the device being addressed.

The argument *enable* indicates whether tracing should be enabled or disabled for the event or events specified by argument *etype*. A positive value for *enable* indicates that recording of one or more events specified by *etype* should be enabled; a value of 0 for *enable* indicates that recording of events should be disabled by this invocation.

An argument *etype* value 0 indicates that traces for all event types will be enabled or disabled. Passing a positive value for *etype* indicates that recording should be enabled or disabled for the event in `ompt_callbacks_t` that matches *etype*.

Effect

Table 4.5 shows the possible return codes for `ompt_set_trace_ompt`. If a single invocation of `ompt_set_trace_ompt` is used to enable or disable more than one event (i.e., `etype=0`), the return code will be 3 if tracing is possible for one or more events but not for others.

TABLE 4.5: Meaning of return codes for `ompt_set_trace_ompt` and `ompt_set_trace_native`.

Return Code	Meaning
0	error
1	event will never occur
2	event may occur but no tracing is possible
3	event may occur and will be traced when convenient
4	event may occur and will always be traced if event occurs

Cross References

- `ompt_callbacks_t`, see Section 4.2.3.2 on page 430.
- `ompt_device_t`, see Section 4.2.3.4.4 on page 435.

4.2.5.2.5 `ompt_set_trace_native_t`

Summary

A runtime entry point for a device known as `ompt_set_trace_native` with type signature `ompt_set_trace_native_t` enables or disables the recording of native trace records for a device.

Format

```
C / C++
typedef int (*ompt_set_trace_native_t) (
    ompt_device_t *device,
    int enable,
    int flags
);
```

Description

This interface is designed for use by a tool with no knowledge about an attached device. If a tool knows how to program a particular attached device, it may opt to invoke native control functions directly using pointers obtained through the `lookup` function associated with the device and described in the `documentation` string that is provided to the device initializer callback.

Description of Arguments

The argument `device` is a pointer to an opaque object that represents the target device instance. The pointer to the device instance object is used by functions in the device tracing interface to identify the device being addressed.

The argument `enable` indicates whether recording of events should be enabled or disabled by this invocation.

The argument `flags` specifies the kinds of native device monitoring to enable or disable. Each kind of monitoring is specified by a flag bit. Flags can be composed by using logical `OR` to combine enumeration values from type `ompt_native_mon_flag_t`. Table 4.5 shows the possible return codes for `ompt_set_trace_native`. If a single invocation of `ompt_set_trace_ompt` is used to enable/disable more than one kind of monitoring, the return code will be 3 if tracing is possible for one or more kinds of monitoring but not for others.

To start, pause, flush, or stop tracing for a specific target device associated with the handle `device`, a tool invokes the `ompt_start_trace`, `ompt_pause_trace`, `ompt_flush_trace`, or `ompt_stop_trace` runtime entry point for the device.

Cross References

- `ompt_device_t`, see Section 4.2.3.4.4 on page 435.

4.2.5.2.6 `ompt_start_trace_t`

Summary

A runtime entry point for a device known as `ompt_start_trace` with type signature `ompt_start_trace_t` starts tracing of activity on a specific device.

Format

```
C / C++
typedef int (*ompt_start_trace_t) (
    ompt_device_t *device,
    ompt_callback_buffer_request_t request,
    ompt_callback_buffer_complete_t complete
);
```

Description

A tool may initiate tracing on a device by invoking the device's `ompt_start_trace` runtime entry point.

Under normal operating conditions, every event buffer provided to a device by a tool callback will be returned to the tool before the OpenMP runtime shuts down. If an exceptional condition terminates execution of an OpenMP program, the OpenMP runtime may not return buffers provided to the device.

An invocation of `ompt_start_trace` returns 1 if the command succeeded and 0 otherwise.

Description of Arguments

The argument *device* is a pointer to an opaque object that represents the target device instance. The pointer to the device instance object is used by functions in the device tracing interface to identify the device being addressed.

The argument *buffer request* specifies a tool callback that will supply a device with a buffer to deposit events.

The argument *buffer complete* specifies a tool callback that will be invoked by the OpenMP implementation to empty a buffer containing event records.

Cross References

- `ompt_device_t`, see Section 4.2.3.4.4 on page 435.
- `ompt_callback_buffer_request_t`, see Section 4.2.4.2.24 on page 474.
- `ompt_callback_buffer_complete_t`, see Section 4.2.4.2.25 on page 475.

4.2.5.2.7 `ompt_pause_trace_t`

Summary

A runtime entry point for a device known as `ompt_pause_trace` with type signature `ompt_pause_trace_t` pauses or restarts activity tracing on a specific device.

```

C / C++
typedef int (*ompt_pause_trace_t) (
    ompt_device_t *device,
    int begin_pause
);
```

C / C++

Description

A tool may pause or resume tracing on a device by invoking the device's `ompt_pause_trace` runtime entry point. An invocation of `ompt_pause_trace` returns 1 if the command succeeded and 0 otherwise.

Redundant pause or resume commands are idempotent and will return 1 indicating success.

Description of Arguments

The argument `device` is a pointer to an opaque object that represents the target device instance. The pointer to the device instance object is used by functions in the device tracing interface to identify the device being addressed.

The argument `begin_pause` indicates whether to pause or resume tracing. To resume tracing, zero should be supplied for `begin_pause`.

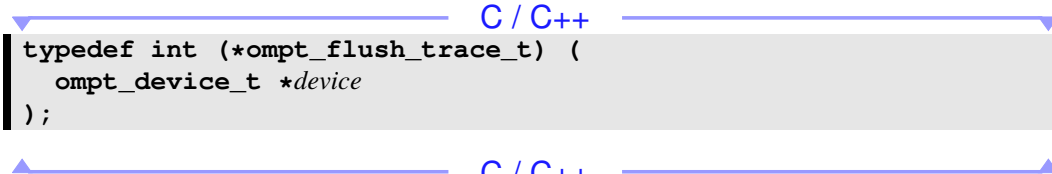
Cross References

- `ompt_device_t`, see Section 4.2.3.4.4 on page 435.

1 4.2.5.2.8 `ompt_flush_trace_t`

2 Summary

3 A runtime entry point for a device known as `ompt_flush_trace` with type signature
4 `ompt_flush_trace_t` causes all pending trace records for the specified device to be delivered
5 to the tool.

```
6  C / C++  
7 typedef int (*ompt_flush_trace_t) (  
8     ompt_device_t *device  
9 );
```

9 Description

10 A tool may request that a device flush any pending trace records by invoking the
11 `ompt_flush_trace` runtime entry point for the device. Invoking `ompt_flush_trace`
12 causes the OpenMP implementation to issue a sequence of zero or more buffer completion
13 callbacks to deliver to the tool all trace records that have been collected prior to the flush. An
14 invocation of `ompt_flush_trace` returns 1 if the command succeeded and 0 otherwise.

15 Description of Arguments

16 The argument *device* is a pointer to an opaque object that represents the target device instance. The
17 pointer to the device instance object is used by functions in the device tracing interface to identify
18 the device being addressed.

19 Cross References

- 20 • `ompt_device_t`, see Section [4.2.3.4.4](#) on page 435.

21 4.2.5.2.9 `ompt_stop_trace_t`

22 Summary

23 A runtime entry point for a device known as `ompt_stop_trace` with type signature
24 `ompt_stop_trace_t` stops tracing for a device.

C / C++

```
1 typedef int (*ompt_stop_trace_t) (  
2     ompt_device_t *device  
3 );
```

C / C++

4 Description

5 A tool may halt tracing on a device and request that the device flush any pending trace records by
6 invoking the **ompt_stop_trace** runtime entry point for the device. An invocation of
7 **ompt_stop_trace** returns 1 if the command succeeded and 0 otherwise.

8 Description of Arguments

9 The argument *device* is a pointer to an opaque object that represents the target device instance. The
10 pointer to the device instance object is used by functions in the device tracing interface to identify
11 the device being addressed.

12 Cross References

- 13 • **ompt_device_t**, see Section [4.2.3.4.4](#) on page [435](#).

14 4.2.5.2.10 ompt_advance_buffer_cursor_t

15 Summary

16 A runtime entry point for a device known as **ompt_advance_buffer_cursor** with type
17 signature **ompt_advance_buffer_cursor_t** advances a trace buffer cursor to the next
18 record.

19 Format

C / C++

```
20 typedef int (*ompt_advance_buffer_cursor_t) (  
21     ompt_buffer_t *buffer,  
22     size_t size,  
23     ompt_buffer_cursor_t current,  
24     ompt_buffer_cursor_t *next  
25 );
```

C / C++

1 **Description**

2 It returns *true* if the advance is successful and the next position in the buffer is valid.

3 **Description of Arguments**

4 The argument *device* is a pointer to an opaque object that represents the target device instance. The
5 pointer to the device instance object is used by functions in the device tracing interface to identify
6 the device being addressed.

7 The argument *buffer* indicates a trace buffer associated with the cursors.

8 The argument *size* indicates the size of *buffer* in bytes.

9 The argument *current* is an opaque buffer cursor.

10 The argument *next* returns the next value of a opaque buffer cursor.

11 **Cross References**

12 • `ompt_device_t`, see Section [4.2.3.4.4](#) on page [435](#).

13 • `ompt_buffer_cursor_t`, see Section [4.2.3.4.7](#) on page [435](#).

14 **4.2.5.2.11 ompt_get_record_type_t**

15 **Summary**

16 A runtime entry point for a device known as `ompt_get_record_type` with type signature
17 `ompt_get_record_type_t` inspects the type of a trace record for a device.

18 **Format**

```
19      C / C++  
20      typedef ompt_record_t (*ompt_get_record_type_t) (  
21          ompt_buffer_t *buffer,  
22          ompt_buffer_cursor_t current  
23      );  
24      C / C++
```


Description

Trace records for a device may be in one of two forms: a *native* record format, which may be device-specific, or an *OMPT* record format, where each trace record corresponds to an OpenMP *event* and fields in the record structure are mostly the arguments that would be passed to the OMPT callback for the event.

A runtime entry point for a device known as `ompt_get_record_type` with type signature `ompt_get_record_type_t` inspects the type of a trace record and indicates whether the record at the current position in the provided trace buffer is an OMPT record, a native record, or an invalid record. An invalid record type is returned if the cursor is out of bounds.

Description of Arguments

The argument *buffer* indicates a trace buffer.

The argument *current* is an opaque buffer cursor.

Cross References

- `ompt_buffer_t`, see Section 4.2.3.4.6 on page 435.
- `ompt_buffer_cursor_t`, see Section 4.2.3.4.7 on page 435.

4.2.5.2.12 `ompt_get_record_ompt_t`

Summary

A runtime entry point for a device known as `ompt_get_record_ompt` with type signature `ompt_get_record_ompt_t` obtains a pointer to an OMPT trace record from a trace buffer associated with a device.

Format

```
C / C++
typedef ompt_record_ompt_t *(*ompt_get_record_ompt_t) (
    ompt_buffer_t *buffer,
    ompt_buffer_cursor_t current
);
```

C / C++

1 **Description**

2 This function returns a pointer that may point to a record in the trace buffer, or it may point to a
3 record in thread local storage where the information extracted from a record was assembled. The
4 information available for an event depends upon its type.

5 The return value of type `ompt_record_ompt_t` defines a union type that can represent
6 information for any OMPT event record type. Another call to the runtime entry point may
7 overwrite the contents of the fields in a record returned by a prior invocation.

8 **Description of Arguments**

9 The argument *buffer* indicates a trace buffer.

10 The argument *current* is an opaque buffer cursor.

11 **Cross References**

12 • `ompt_record_ompt_t`, see Section 4.2.3.3.4 on page 432.

13 • `ompt_device_t`, see Section 4.2.3.4.4 on page 435.

14 • `ompt_buffer_cursor_t`, see Section 4.2.3.4.7 on page 435.

15 **4.2.5.2.13 ompt_get_record_native_t**

16 **Summary**

17 A runtime entry point for a device known as `ompt_get_record_native` with type signature
18 `ompt_get_record_native_t` obtains a pointer to a native trace record from a trace buffer
19 associated with a device.

20 **Format**

21 C / C++

```
22      typedef void *(*ompt_get_record_native_t) (  
23          ompt_buffer_t *buffer,  
24          ompt_buffer_cursor_t current,  
25          ompt_id_t *host_op_id  
26      );
```

C / C++

Description

The pointer returned may point into the specified trace buffer, or into thread local storage where the information extracted from a trace record was assembled. The information available for a native event depends upon its type. If the function returns a non-NULL result, it will also set ***host_op_id** to identify host-side identifier for the operation associated with the record. A subsequent call to **ompt_get_record_native** may overwrite the contents of the fields in a record returned by a prior invocation.

Description of Arguments

The argument *buffer* indicates a trace buffer.

The argument *current* is an opaque buffer cursor.

The argument *host_op_id* is a pointer to an identifier that will be returned by the function. The entry point will set **host_op_id* to the value of a host-side identifier for an operation on a target device that was created when the operation was initiated by the host.

Cross References

- **ompt_id_t**, see Section [4.2.3.4.2](#) on page [434](#).
- **ompt_buffer_t**, see Section [4.2.3.4.6](#) on page [435](#).
- **ompt_buffer_cursor_t**, see Section [4.2.3.4.7](#) on page [435](#).

4.2.5.2.14 ompt_get_record_abstract_t

Summary

A runtime entry point for a device known as **ompt_get_record_abstract** with type signature **ompt_get_record_abstract_t** summarizes the context of a native (device-specific) trace record.

Format

```
C / C++  
typedef ompt_record_abstract_t *  
(*ompt_get_record_abstract_t) (  
    void *native_record  
) ;  
C / C++
```

1 **Description**

2 An OpenMP implementation may execute on a device that logs trace records in a native
3 (device-specific) format unknown to a tool. A tool can use the `ompt_get_record_abstract`
4 runtime entry point for the device with type signature `ompt_get_record_abstract_t` to
5 decode a native trace record that it does not understand into a standard form that it can interpret.

6 **Description of Arguments**

7 The argument *native_record* is a pointer to a native trace record.

8 **Cross References**

- 9 • `ompt_record_abstract_t`, see Section 4.2.3.3.3 on page 431.

10 **4.2.5.3 Lookup Entry Point**

11 **4.2.5.3.1 ompt_function_lookup_t**

12 **Summary**

13 A tool uses a lookup routine with type signature `ompt_function_lookup_t` to obtain
14 pointers to runtime entry points that are part of the OMPT interface.

15 **Format**

```
16           C / C++  
17           typedef void (*ompt_interface_fn_t) (void);  
18           typedef ompt_interface_fn_t (*ompt_function_lookup_t) (  
19           const char *interface_function_name  
20           );  
          C / C++
```

Description

An OpenMP implementation provides a pointer to a lookup routine as an argument to tool callbacks used to initialize tool support for monitoring an OpenMP device using either tracing or callbacks.

When an OpenMP implementation invokes a tool initializer to configure the OMPT callback interface, the OpenMP implementation will pass the initializer a lookup function that the tool can use to obtain pointers to runtime entry points that implement routines that are part of the OMPT callback interface.

When an OpenMP implementation invokes a tool initializer to configure the OMPT tracing interface for a device, the Open implementation will pass the device tracing initializer a lookup function that the tool can use to obtain pointers to runtime entry points that implement tracing control routines appropriate for that device.

A tool can call the lookup function to obtain a pointer to a runtime entry point.

Description of Arguments

The argument *interface_function_name* is a C string that represents the name of a runtime entry point.

Cross References

- Entry points in the OMPT callback interface, see Table 4.1 on page 423 for a list and Section 4.2.5.1 on page 481 for detailed definitions.
- Tool initializer for a device's OMPT tracing interface, Section 4.2.1.4 on page 425.
- Entry points in the OMPT tracing interface, see Table 4.3 on page 426 for a list and Section 4.2.5.2 on page 502 for detailed definitions.
- Tool initializer for the OMPT callback interface, Section 4.2.4.1.1 on page 444

1 4.3 OMPD

2 The OMPD interface is designed to allow a *third-party tool* such as a debugger to inspect the
3 OpenMP state of a live program or core file in an implementation agnostic manner. That is, a tool
4 that uses OMPD should work with any conforming OpenMP implementation. The model for
5 OMPD is that an OpenMP implementor provides a library that a third-party tool can dynamically
6 load. Using the interface exported by the OMPD library, the external tool can inspect the OpenMP
7 state of a program using that implementation of OpenMP. In order to satisfy requests from the
8 third-party tool, the OMPD library may need to read data from, or find the addresses of symbols in,
9 the OpenMP program. The OMPD library does this by using the callback interface the third-party
10 tool must make available to the OMPD library.

11 The diagram shown in Section C on page 623 shows how the different components fits together.
12 The third-party tool loads the OMPD library that matches the OpenMP runtime being used by the
13 OpenMP program. The library exports the API defined later in this document, which the tool uses
14 to get OpenMP information about the OpenMP program. The OMPD library will need to look up
15 the symbols, or read data out of the program. It does not do this directly, but instead asks the tool to
16 perform these operations for it using a callback interface exported by the tool.

17 This architectural layout insulates the tool from the details of the internal structure of the OpenMP
18 runtime. Similarly, the OMPD library does not need to be concerned about how to access the
19 OpenMP program. Decoupling the library and tool in this way allows for great flexibility in how
20 the OpenMP program and tool are deployed, so that, for example, there is no requirement that tool
21 and OpenMP program execute on the same machine. Generally the tool does not interact directly
22 with the OpenMP runtime in the OpenMP program, and instead uses the OMPD library for this
23 purpose. However, there are a few instances where the tool does need to access the OpenMP
24 runtime directly. These cases fall into two broad categories. The first is during initialization, where
25 the tool needs to be able to look up symbols and read variables in OpenMP runtime in order to
26 identify the OMPD library it should use. This is discussed in Sections 4.4.2.2 and 4.4.2.3.

27 The second category relates to arranging for the tool to be notified when certain events occur
28 during the execution of the OpenMP program. The model used for this purpose is that the OpenMP
29 implementation is required to define certain symbols in the runtime code. This is discussed in
30 Section 4.3.5. Each of these symbols corresponds to an event type. The runtime must ensure that
31 control passes through the appropriate named location when events occur. If the tool wants to get
32 notification of an event, it can plant a breakpoint at the matching location.

33 The code locations can, but do not need to, be functions. They can, for example, simply be labels.
34 However, the names must have external C linkage.

1 4.3.1 Activating an OMPD Tool

2 The tool and the OpenMP program the tool controls exist as separate processes.; thus coordination
3 is required between the OpenMP runtime and the external tool for OMPD to be used successfully.

4 4.3.1.1 Enabling the Runtime for OMPD

5 In order to support external tools, the OpenMP runtime may need to collect and maintain
6 information that it might otherwise not do, perhaps for performance reasons, or because it is not
7 otherwise needed. The OpenMP runtime collects whatever information is necessary to support
8 OMPD if the environment variable **OMP_DEBUG** is set to *enabled*.

9 Cross References

- 10 • **OMP_DEBUG**, Section 5.20 on page 612
- 11 • Activating an OMPT Tool, Section 4.2.1 on page 418

12 4.3.1.2 Finding the OMPD library

13 An OpenMP runtime may have more than one matching OMPD library for tools to use. The tool
14 must be able to locate the right library to use for the OpenMP program it is examining.

15 As part of the OpenMP interface, OMPD requires that the OpenMP runtime system provides a
16 public variable **ompd_dll_locations**, which is an **argv**-style vector of filename string
17 pointers that provides the name(s) of any compatible OMPD library. **ompd_dll_locations**
18 must have **C** linkage. The tool uses the name of the variable verbatim, and in particular, will not
19 apply any name mangling before performing the look up.

20 **ompd_dll_locations** points to a NULL-terminated vector of zero or more NULL-terminated
21 pathname strings. There are no filename conventions for pathname strings. The last entry in the
22 vector is NULL. The vector of string pointers must be fully initialized *before*
23 **ompd_dll_locations** is set to a non-NULL value, such that if the tool, such as a debugger,
24 stops execution of the OpenMP program at any point where **ompd_dll_locations** is
25 non-NULL, then the vector of strings it points to is valid and complete.

26 The programming model or architecture of the tool (and hence that of the required OMPD) does
27 not have to match that of the OpenMP program being examined. It is the responsibility of the tool
28 to interpret the contents of **ompd_dll_locations** to find a suitable OMPD that matches its
29 own architectural characteristics. On platforms that support different programming models (*e.g.*,
30 32-bit vs 64-bit), OpenMP implementers are encouraged to provide OMPD library for all models,
31 and which can handle OpenMP programs of any model. Thus, for example, a 32-bit debugger using

1 OMPD should be able to debug a 64-bit OpenMP program by loading a 32-bit OMPD that can
2 manage a 64-bit OpenMP runtime.

3 **Cross References**

- 4 • Identifying the Matching OMPD, Section [4.4.2.2](#) on page [593](#)

5 **4.3.2 OMPD Data Types**

6 In this section, we define the types, structures, and functions for the OMPD API.

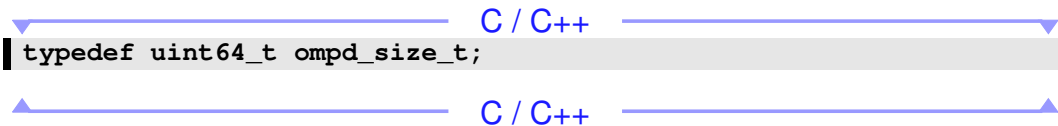
7 **4.3.2.1 Basic Types**

8 The following describes the basic OMPD API types.

9 **4.3.2.1.1 Size Type**

10 This type is used to specify the number of bytes in opaque data objects passed across the OMPD
11 API.

12 **Format**

13 The code snippet is `typedef uint64_t ompd_size_t;` enclosed in a light gray box. Above the box, a blue double-headed arrow spans the width of the code, with the text "C / C++" centered above it. Below the box, another blue double-headed arrow spans the width of the code, with the text "C / C++" centered below it.

```
13 | typedef uint64_t ompd_size_t;
```

14 **4.3.2.1.2 Wait ID Type**

15 This type identifies what a thread is waiting for.

1 Format

```
2 | typedef uint64_t ompd_wait_id_t;
```

3 C / C++

3 4.3.2.1.3 Basic Value Types

4 These definitions represent a word, address, and segment value types.

5 Format

```
6 | typedef uint64_t ompd_addr_t;
```

```
7 | typedef int64_t ompd_word_t;
```

```
8 | typedef uint64_t ompd_seg_t;
```

9 C / C++

9 The *ompd_addr_t* type represents an unsigned integer address in an OpenMP process. The
10 *ompd_word_t* type represents a signed version of *ompd_addr_t* to hold a signed integer of the
11 OpenMP process. The *ompd_seg_t* type represents an unsigned integer segment value.

12 4.3.2.1.4 Address Type

13 This type is a structure that OMPD uses to specify device addresses, which may or may not be
14 segmented.

15 Format

```
16 | typedef struct {
```

```
17 |     ompd_seg_t segment;
```

```
18 |     ompd_addr_t address;
```

```
19 | } ompd_address_t;
```

20 C / C++

20 *ompd_segment_none* is defined as an instance of type *ompd_seg_t* with the value 0.

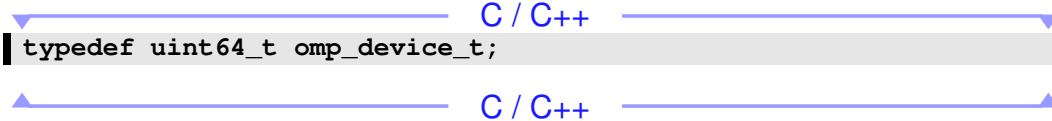
21 For non-segmented architectures, *ompd_segment_none* is used in the *segment* field of
22 *ompd_address_t*.

1 4.3.2.2 System Device Identifiers

2 Different OpenMP runtimes may utilize different underlying devices. The type used to hold an
3 device identifier can vary in size and format, and therefore is not explicitly represented in the
4 OMPD API. Device identifiers are passed across the interface using a device-identifier ‘kind’, a
5 pointer to where the device identifier is stored, and the size of the device identifier in bytes. The
6 OMPD library and tool using it must agree on the format of what is being passed. Each different
7 kind of device identifier uses a unique unsigned 64-bit integer value.

8 Recommended values of `omp_device_t` are defined in the `ompd_types.h` header file, which
9 is available on <http://www.openmp.org/>.

10 Format

11 

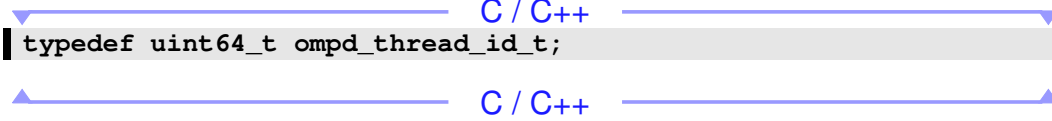
```
typedef uint64_t omp_device_t;
```

12 4.3.2.3 Thread Identifiers

13 Different OpenMP runtimes may use different underlying native thread implementations. The type
14 used to hold a thread identifier can vary in size and format, and therefore is not explicitly
15 represented in the OMPD API. Thread identifiers are passed across the interface using a
16 thread-identifier ‘kind’, a pointer to where the thread identifier is stored, and the size of the thread
17 identifier in bytes. The OMPD library and tool using it must agree on the format of what is being
18 passed. Each different kind of thread identifier uses a unique unsigned 64-bit integer value.

19 Recommended values of `ompd_thread_id_t` are defined in the `ompd_types.h` header file,
20 which is available on <http://www.openmp.org/>.

21 Format

22 

```
typedef uint64_t ompd_thread_id_t;
```

1 4.3.2.4 OMPD Handle Types

2 Each operation of the OMPD interface that applies to a particular address space, thread, parallel
3 region, or task must explicitly specify a *handle* for the operation. OMPD employs handles for
4 address spaces (for a host or target device), threads, parallel regions, and tasks. A handle for an
5 entity is constant while the entity itself is alive. Handles are defined by the OMPD library, and are
6 opaque to the tool. The following defines the OMPD handle types:

7 Format

```
8                                     C / C++  
9 typedef struct _ompd_aspace_handle_s ompd_address_space_handle_t  
10 ;  
11 typedef struct _ompd_thread_handle_s ompd_thread_handle_t;  
12 typedef struct _ompd_parallel_handle_s ompd_parallel_handle_t;  
13 typedef struct _ompd_task_handle_s ompd_task_handle_t;
```

13 Defining externally visible type names in this way introduces type safety to the interface, and helps
14 to catch instances where incorrect handles are passed by the tool to the OMPD library. The
15 **structs** do not need to be defined at all. The OMPD library must cast incoming (pointers to)
16 handles to the appropriate internal, private types.

17 4.3.2.5 OMPD Scope Types

18 Summary

19 The `ompd_scope_t` type describes OMPD scope types for OMPD tool interface routines.

20 Format

```
21                                     C / C++  
22 typedef enum ompd_scope_t {  
23     ompd_scope_global = 1,  
24     ompd_scope_address_space = 2,  
25     ompd_scope_thread = 3,  
26     ompd_scope_parallel = 4,  
27     ompd_scope_implicit_task = 5,  
28     ompd_scope_task = 6  
29 } ompd_scope_t;
```

1 **Description**

2 When used in an interface function call, the scope type and the ompd handle must match according
3 to Table 4.6.

TABLE 4.6: Mapping of scope type and OMPD handles

Scope types	Handles
<i>ompd_scope_global</i>	address space handle for the host device
<i>ompd_scope_address_space</i>	any address space handle
<i>ompd_scope_thread</i>	any thread handle
<i>ompd_scope_parallel</i>	any parallel handle
<i>ompd_scope_implicit_task</i>	task handle for an implicit task
<i>ompd_scope_task</i>	any task handle

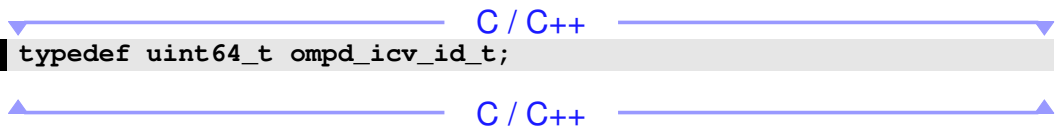
4 **4.3.2.6 ICV ID Type**

5 **Summary**

6 This type identifies what a thread is waiting for.

7 **Format**

```
8 | typedef uint64_t ompd_icv_id_t;
```



9 `ompd_icv_undefined` is defined as an instance of type `ompd_icv_id_t` with the value 0.

10 **4.3.2.7 Tool Context Types**

11 A third-party tool uses contexts to uniquely identifies abstractions. These contexts are opaque to the
12 OMPD library and are defined as follows:

1 **Format**

C / C++

```
2 typedef struct _ompd_aspace_cont_s ompd_address_space_context_t;  
3 typedef struct _ompd_thread_cont_s ompd_thread_context_t;
```

C / C++

4 **4.3.2.8 Return Code Types**

5 Each OMPD operation has a return code. The return code types and their semantics are defined as
6 follows:

7 **Format**

C / C++

```
8 typedef enum {  
9     ompd_rc_ok = 0,  
10    ompd_rc_unavailable = 1,  
11    ompd_rc_stale_handle = 2,  
12    ompd_rc_bad_input = 3,  
13    ompd_rc_error = 4,  
14    ompd_rc_unsupported = 5,  
15    ompd_rc_needs_state_tracking = 6,  
16    ompd_rc_incompatible = 7,  
17    ompd_rc_device_read_error = 8,  
18    ompd_rc_device_write_error = 9,  
19    ompd_rc_nomem = 10,  
20 } ompd_rc_t;
```

C / C++

Description

`ompd_rc_ok` is returned when the operation is successful.

`ompd_rc_unavailable` is returned when information is not available for the specified context.

`ompd_rc_stale_handle` is returned when the specified handle is no longer valid.

`ompd_rc_bad_input` is returned when the input parameters (other than handle) are invalid.

`ompd_rc_error` is returned when a fatal error occurred.

`ompd_rc_unsupported` is returned when the requested operation is not supported.

`ompd_rc_needs_state_tracking` is returned when the state tracking operation failed because state tracking is not currently enabled.

`ompd_rc_incompatible` is returned when this OMPD is incompatible with the OpenMP program.

`ompd_rc_device_read_error` is returned when a read operation failed on the device.

`ompd_rc_device_write_error` is returned when a write operation failed to the device.

`ompd_rc_nomem` is returned when unable to allocate memory.

4.3.2.9 Primitive Types

The following structure contains members that the OMPD library can use to interrogate the tool about the “sizeof” of primitive types in the OpenMP architecture address space.

Format

```
typedef struct {  
    uint8_t  sizeof_char;  
    uint8_t  sizeof_short;  
    uint8_t  sizeof_int;  
    uint8_t  sizeof_long;  
    uint8_t  sizeof_long_long;  
    uint8_t  sizeof_pointer;  
} ompd_device_type_sizes_t;
```

C / C++

C / C++

Description

The fields of `ompd_device_type_sizes_t` give the sizes of the eponymous basic types used by the OpenMP runtime. As the tool and the OMPD library, by definition, have the same architecture and programming model, the size of the fields can be given as `int`.

Cross References

- `ompd_callback_sizeof_fn_t`, Section [4.3.3.2.2](#) on page [531](#)

4.3.3 OMPD Tool Callback Interface

For the OMPD library to provide information about the internal state of the OpenMP runtime system in an OpenMP process or core file, it must have a means to extract information from the OpenMP process that the tool is debugging. The OpenMP process that the tool is operating on may be either a “live” process or a core file, and a thread may be either a “live” thread in an OpenMP process, or a thread in a core file. To enable the OMPD library to extract state information from an OpenMP process or core file, the tool must supply the OMPD library with callback functions to inquire about the size of primitive types in the device of the OpenMP process, look up the addresses of symbols, as well as read and write memory in the device. The OMPD library then uses these callbacks to implement its interface operations. The OMPD library will only call the callback functions in direct response to calls made by the tool to the OMPD library. Signatures for the tool callbacks used by the OMPD library are given below.

4.3.3.1 Memory Management of OMPD Library

The OMPD library must not access the heap manager directly. Instead, if it needs heap memory it must use the memory allocation and deallocation callback functions that are described in this section, `ompd_callback_memory_alloc_fn_t` (see Section [4.3.3.1.1](#) on page [528](#)) and `ompd_callback_memory_free_fn_t` (see Section [4.3.3.1.2](#) on page [529](#)), which are provided by the tool to obtain and release heap memory. This will ensure that the library does not interfere with any custom memory management scheme that the tool may use.

If the OMPD library is implemented in **C++**, memory management operators like `new` and `delete` in all their variants, *must all* be overloaded and implemented in terms of the callbacks provided by the tool. The OMPD library must be coded so that any of its definitions of `new` or `delete` do not interfere with any defined by the tool.

In some cases, the OMPD library will need to allocate memory to return results to the tool. This memory will then be ‘owned’ by the tool, which will be responsible for releasing it. It is therefore vital that the OMPD library and the tool use the same memory manager.

1 OMPD handles are created by the OMPD library. These are opaque to the tool, and depending on
2 the specific implementation of OMPD may have complex internal structure. The tool cannot know
3 whether the handle pointers returned by the API correspond to discrete heap allocations.
4 Consequently, the tool must not simply deallocate a handle by passing an address it receives from
5 the OMPD library to its own memory manager. Instead, the API includes functions that the tool
6 must use when it no longer needs a handle.

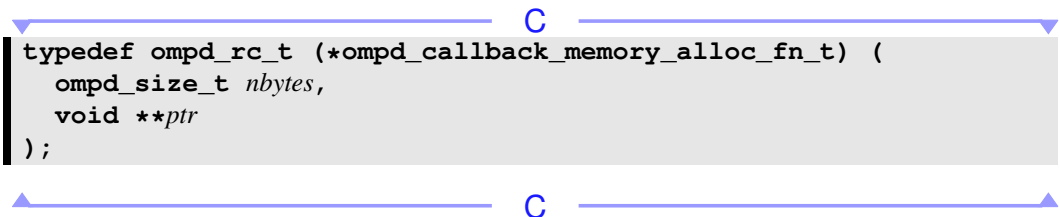
7 Contexts are created by a tool and passed to the OMPD library. The OMPD library does not need to
8 release contexts; instead this will be done by the tool after it releases any handles that may be
9 referencing the contexts.

10 4.3.3.1.1 `ompd_callback_memory_alloc_fn_t`

11 Summary

12 The type signature of the callback routine provided by the tool to be used by the OMPD library to
13 allocate memory.

```
14 typedef ompd_rc_t (*ompd_callback_memory_alloc_fn_t) (  
15     ompd_size_t nbytes,  
16     void **ptr  
17 );
```

A code block containing the C type signature for `ompd_callback_memory_alloc_fn_t`. The code is enclosed in a light gray box. Above and below the box are blue double-headed arrows with a 'C' in the center, indicating the C language context.

18 Description

19 The OMPD library may call the `ompd_callback_memory_alloc_fn_t` callback function to
20 allocate memory.

21 Description of Arguments

22 The argument `nbytes` gives the size in bytes of the block of memory the OMPD library wants
23 allocated.

24 The address of the newly allocated block of memory is returned in `*ptr`. The newly allocated block
25 is suitably aligned for any type of variable, and is not guaranteed to be zeroed.

26 Cross References

- 27 • `ompd_size_t`, Section [4.3.2.1.1](#) on page [520](#)
- 28 • `ompd_rc_t`, Section [4.3.2.8](#) on page [525](#)

1 4.3.3.1.2 `ompd_callback_memory_free_fn_t`

2 **Summary**

3 The type signature of the callback routine provided by the tool to be used by the OMPD library to
4 deallocate memory.

```
5 typedef ompd_rc_t (*ompd_callback_memory_free_fn_t) (  
6     void *ptr  
7 );
```

8 **Description**

9 The OMPD library calls the `ompd_callback_memory_free_fn_t` callback function to
10 deallocate memory obtained from a prior call to the `ompd_callback_memory_alloc_fn_t`
11 callback function.

12 **Description of Arguments**

13 *ptr* is the address of the block to be deallocated.

14 **Cross References**

- 15 • `ompd_callback_memory_alloc_fn_t`, Section [4.3.3.1.1](#) on page [528](#)
- 16 • `ompd_rc_t`, Section [4.3.2.8](#) on page [525](#)
- 17 • `ompd_callbacks_t`, Section [4.3.3.6](#) on page [538](#)



18 4.3.3.2 **Context Management and Navigation**

19 The tool provides the OMPD library with callbacks to manage and navigate context relationships.

1 4.3.3.2.1 ompd_callback_get_thread_context_for_thread_id_fn_t

2 Summary

3 The type signature of the callback routine provided by the third party tool the OMPD library can
4 use to map a thread identifier to a tool *thread context*.

5  

```
6 typedef ompd_rc_t  
7 (*ompd_callback_get_thread_context_for_thread_id_fn_t) (  
8     ompd_address_space_context_t *address_space_context,  
9     ompd_thread_id_t kind,  
10    ompd_size_t sizeof_thread_id,  
11    const void *thread_id,  
12    ompd_thread_context_t **thread_context  
13 );
```

13 Description

14 This callback maps a thread identifier within the address space identified by *address_space_context*
15 to a tool thread context. The OMPD library can use the thread context, for example, to access
16 thread local storage (TLS).

17 Description of Arguments

18 The input argument *address_space_context* is an opaque handle provided by the tool to reference
19 an address space. The input arguments *kind*, *sizeof_thread_id*, and *thread_id* represent a thread
20 identifier. On return the output argument *thread_context* provides an opaque handle to the OMPD
21 library that maps a thread identifier to a tool thread context.

22 Restrictions

23 The *thread_context* provided by this function is valid until the OMPD library returns from the
24 OMPD tool interface routine.

Cross References

- `ompd_rc_t`, Section 4.3.2.8 on page 525
- `ompd_address_space_context_t`, Section 4.3.2.7 on page 524
- `ompd_thread_id_t`, Section 4.3.2.3 on page 522
- `ompd_size_t`, Section 4.3.2.1.1 on page 520
- `ompd_thread_context_t`, Section 4.3.2.7 on page 524

4.3.3.2.2 `ompd_callback_sizeof_fn_t`

Summary

The type signature of the callback routine provided by the tool the OMPD library can use to find the sizes of the primitive types in an address space.

```
typedef ompd_rc_t (*ompd_callback_sizeof_fn_t) (  
    ompd_address_space_context_t *address_space_context,  
    ompd_device_type_sizes_t *sizes  
);
```

Description

This callback may be called by the OMPD library to obtain the sizes of the basic primitive types for a given address space.

Description of Arguments

The callback returns the sizes of the basic primitive types used by the `address_space_context` in `*sizes`.

Cross References

- `ompd_address_space_context_t`, Section 4.3.2.7 on page 524
- `ompd_device_type_sizes_t`, Section 4.3.2.9 on page 526
- `ompd_rc_t`, Section 4.3.2.8 on page 525
- `ompd_callbacks_t`, Section 4.3.3.6 on page 538

1 4.3.3.3 Accessing Memory in the OpenMP Program or Runtime

2 The OMPD library may need to read from, or write to, the OpenMP program. It cannot do this
3 directly, but instead must use the callbacks provided to it by the tool, which will perform the
4 operation on its behalf.

5 4.3.3.3.1 ompd_callback_symbol_addr_fn_t

6 Summary

7 The type signature of the callback provided by the tool the OMPD library can use to look up the
8 addresses of symbols in an OpenMP program.

```
9 typedef ompd_rc_t (*ompd_callback_symbol_addr_fn_t) (  
10     ompd_address_space_context_t *address_space_context,  
11     ompd_thread_context_t *thread_context,  
12     const char *symbol_name,  
13     ompd_address_t *symbol_addr,  
14     const char *file_name  
15 );
```

16 Description

17 This callback function may be called by the OMPD library to look up addresses of symbols within
18 an specified address space of the tool.

19 Description of Arguments

20 This callback looks up the symbol provided in *symbol_name*.

21 The *address_space_context* input parameter is the tool's representation of the address space of the
22 process, core file, or device. The use of a NULL *address_space_context* results in unspecified
23 behavior.

24 The *thread_context* is an optional input parameter which should be NULL for global memory
25 access. If *thread_context* is not NULL, it gives the thread specific context for the symbol lookup,
26 for the purpose of calculating thread local storage (TLS) addresses. If the *thread_context* parameter
27 is not NULL, the thread that the *thread_context* argument refers to must be associated either to the
28 process or to the device that corresponds to the *address_space_context* argument.

1 The *symbol_name* supplied by the OMPD library is used verbatim by the tool, and in particular, no
2 name mangling, demangling or other transformations are performed prior to the lookup. The
3 *symbol_name* parameter must correspond to a statically allocated symbol within the specified
4 address space. The symbol can correspond to any type of object, such as a variable, thread local
5 storage variable, function, or untyped label. The symbol must be defined and can have a local,
6 global, or weak binding.

7 The *file_name* parameter is an optional input parameter that indicates the name of the shared library
8 where the symbol is defined, and is intended to help the third party tool disambiguate symbols that
9 are defined multiple times across the executable or shared library files. The shared library name
10 may not be an exact match for the name seen by the tool. If the *file_name* parameter is NULL, the
11 tool will try first finding the symbol in the executable file, and, if the symbol is not found, the tool
12 will try finding the symbol in the shared libraries in the order in which the shared libraries are
13 loaded into the address space. If the *file_name* parameter is not NULL, the tool will try first finding
14 the symbol in the libraries that match the name in the *file_name* parameter, and, if the symbol is not
15 found, the tool will find the symbol following the procedure as if the *file_name* parameter is NULL.

16 The callback does not support finding symbols that are dynamically allocated on the call stack, or
17 statically allocated symbols defined within the scope of a function or subroutine.

18 The callback returns the address of the symbol in **symbol_addr*.

19 **Cross References**

- 20 • `ompd_address_space_context_t`, Section [4.3.2.7](#) on page [524](#)
- 21 • `ompd_thread_context_t`, Section [4.3.2.7](#) on page [524](#)
- 22 • `ompd_address_t`, Section [4.3.2.1.4](#) on page [521](#)
- 23 • `ompd_rc_t`, Section [4.3.2.8](#) on page [525](#)
- 24 • `ompd_callbacks_t`, Section [4.3.3.6](#) on page [538](#)

25 **4.3.3.3.2 `ompd_callback_memory_read_fn_t`**

26 **Summary**

27 The type signature of the callback provided by the tool the OMPD library can use to read data out
28 of an OpenMP program.

```

1  typedef ompd_rc_t (*ompd_callback_memory_read_fn_t) (
2      ompd_address_space_context_t *address_space_context,
3      ompd_thread_context_t *thread_context,
4      const ompd_address_t *addr,
5      ompd_size_t nbytes,
6      void *buffer
7  );

```

8 Description

9 The function **read_memory** of this type copies a block of data from *addr* within the address
10 space to the tool *buffer*.

11 The function **read_string** of this type copies a string pointed to by *addr*, including the
12 terminating null byte (' \0 '), to the tool *buffer*. At most *nbytes* bytes are copied. If there is no null
13 byte among the first *nbytes* bytes, the string placed in *buffer* will not be null-terminated.

14 Description of Arguments

15 The address from which the data are to be read out of the OpenMP program specified by
16 *address_space_context* is given by *addr*. *nbytes* gives the number of bytes to be transferred. The
17 *thread_context* argument is optional for global memory access, and in this case should be NULL. If
18 it is not NULL, *thread_context* identifies the thread specific context for the memory access for the
19 purpose of accessing thread local storage (TLS).

20 The data are returned through *buffer*, which is allocated and owned by the OMPD library. The
21 contents of the buffer are unstructured, raw bytes. It is the responsibility of the OMPD library to
22 arrange for any transformations such as byte-swapping that may be necessary (see Section 4.3.3.4.1
23 on page 536) to interpret the data returned.

24 Cross References

- 25 • **ompd_address_space_context_t**, Section 4.3.2.7 on page 524
- 26 • **ompd_thread_context_t**, Section 4.3.2.7 on page 524
- 27 • **ompd_address_t**, Section 4.3.2.1.4 on page 521
- 28 • **ompd_size_t**, Section 4.3.2.1.1 on page 520
- 29 • **ompd_rc_t**, Section 4.3.2.8 on page 525
- 30 • **ompd_callback_device_host_fn_t**, Section 4.3.3.4.1 on page 536
- 31 • **ompd_callbacks_t**, Section 4.3.3.6 on page 538

1 4.3.3.3 ompd_callback_memory_write_fn_t

2 Summary

3 The type signature of the callback provided by the tool the OMPD library can use to write data to
4 an OpenMP program.

C

```
5 typedef ompd_rc_t (*ompd_callback_memory_write_fn_t) (  
6     ompd_address_space_context_t *address_space_context,  
7     ompd_thread_context_t *thread_context,  
8     const ompd_address_t *addr,  
9     ompd_size_t nbytes,  
10    const void *buffer  
11 );
```

C

12 Description

13 The OMPD library may call this function callback to have the tool write a block of data to a
14 location within an address space from a provided buffer.

15 Description of Arguments

16 The address to which the data are to be written in the OpenMP program specified by
17 *address_space_context* is given by *addr*. *nbytes* gives the number of bytes to be transferred. The
18 *thread_context* argument is optional for global memory access, and in this case should be NULL. If
19 it is not NULL, *thread_context* identifies the thread specific context for the memory access for the
20 purpose of accessing thread local storage (TLS).

21 The data to be written are passed through *buffer*, which is allocated and owned by the OMPD
22 library. The contents of the buffer are unstructured, raw bytes. It is the responsibility of the OMPD
23 library to arrange for any transformations such as byte-swapping that may be necessary
24 (see Section 4.3.3.4.1 on page 536) to render the data into a form compatible with the OpenMP
25 runtime.

Cross References

- `ompd_address_space_context_t` Section 4.3.2.7 on page 524
- `ompd_thread_context_t`, Section 4.3.2.7 on page 524
- `ompd_address_t`, Section 4.3.2.1.4 on page 521
- `ompd_size_t`, Section 4.3.2.1.1 on page 520
- `ompd_rc_t`, Section 4.3.2.8 on page 525
- `ompd_callback_device_host_fn_t`, Section 4.3.3.4.1 on page 536
- `ompd_callbacks_t`, Section 4.3.3.6 on page 538

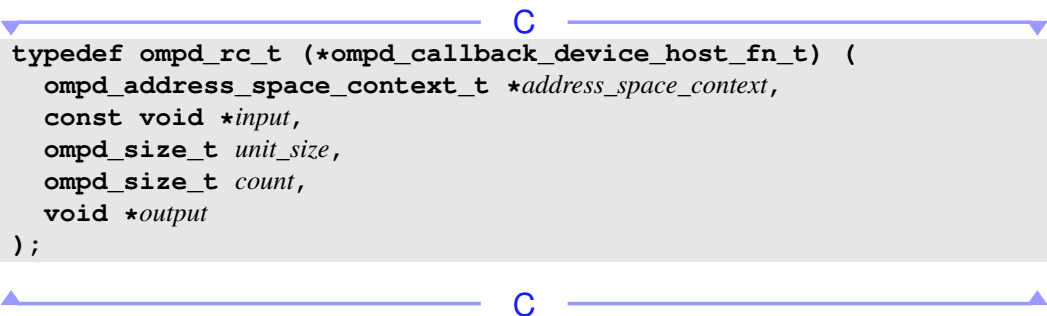
4.3.3.4 Data Format Conversion

The architecture and/or programming-model of tool and OMPD library may be different from that of the OpenMP program being examined. Consequently, the conventions for representing data will differ. The callback interface includes operations for converting between the conventions, such as byte order ('endianness'), used by the tool and OMPD library on the one hand, and the OpenMP program on the other.

4.3.3.4.1 `ompd_callback_device_host_fn_t`

Summary

The type signature of the callback provided by the tool the OMPD library can use to convert data between the formats used by the tool and OMPD library, and the OpenMP program.



```
typedef ompd_rc_t (*ompd_callback_device_host_fn_t) (  
    ompd_address_space_context_t *address_space_context,  
    const void *input,  
    ompd_size_t unit_size,  
    ompd_size_t count,  
    void *output  
);
```


Description

The OMPD library may call the `ompd_callback_print_string_fn_t` callback function to emit output, such as logging or debug information. If the tool does not want to allow the OMPD library to emit output, the tool can provide to the OMPD library a **NULL** value for the `ompd_callback_print_string_fn_t` callback function. Note that the OMPD library is prohibited from writing to file descriptors that it did not open.

Description of Arguments

The input *string* parameter is the null-terminated string to be printed. No conversion or formatting is performed on the string.

The input *category* parameter is the category of the string to be printed. The value of *category* is implementation defined.

Cross References

- `ompd_rc_t`, Section [4.3.2.8](#) on page [525](#)
- `ompd_callbacks_t`, Section [4.3.3.6](#) on page [538](#)

4.3.3.6 The Callback Interface

Summary

All the OMPD library's interactions with the OpenMP program must be through a set of callbacks provided to it by the tool which loaded it. These callbacks must also be used for allocating or releasing resources, such as memory, that the library needs.

```
typedef struct {
    ompd_callback_memory_alloc_fn_t alloc_memory;
    ompd_callback_memory_free_fn_t free_memory;
    ompd_callback_print_string_fn_t print_string;
    ompd_callback_sizeof_fn_t sizeof_type;
    ompd_callback_symbol_addr_fn_t symbol_addr_lookup;
    ompd_callback_memory_read_fn_t read_memory;
    ompd_callback_memory_write_fn_t write_memory;
    ompd_callback_memory_read_fn_t read_string;
    ompd_callback_device_host_fn_t device_to_host;
    ompd_callback_device_host_fn_t host_to_device;
    ompd_callback_get_thread_context_for_thread_id_fn_t
        get_thread_context_for_thread_id;
} ompd_callbacks_t;
```

Description

The set of callbacks the OMPD library should use is collected in the `ompd_callbacks_t` record structure. An instance of this type is passed to the OMPD library as a parameter to `ompd_initialize` (see Section 4.3.4.1.1 on page 540). Each field points to a function that the OMPD library should use for interacting with the OpenMP program, or getting memory from the tool.

The `alloc_memory` and `free_memory` fields are pointers to functions the OMPD library uses to allocate and release dynamic memory.

`print_string` points to a function that prints a string.

The architectures or programming models of the OMPD library and party tool may be different from that of the OpenMP program being examined. `sizeof_type` points to function that allows the OMPD library to determine the sizes of the basic integer and pointer types used by the OpenMP program. Because of the differences in architecture or programming model, the conventions for representing data in the OMPD library and the OpenMP program may be different. The `device_to_host` field points to a function which translates data from the conventions used by the OpenMP program to that used by the tool and OMPD library. The reverse operation is performed by the function pointed to by the `host_to_device` field.

The OMPD library may need to access memory in the OpenMP program. The `symbol_addr_lookup` field points to a callback the OMPD library can use to find the address of a global or thread local storage (TLS) symbol. The `read_memory`, `write_memory` and `read_string` fields are pointers to functions for reading from, and writing to, global or TLS memory in the OpenMP program, respectively.

`get_thread_context_for_thread_id` is a pointer to a function the OMPD library can use to obtain a thread context that corresponds to a thread identifier.

Cross References

- `ompd_callback_memory_alloc_fn_t`, Section 4.3.3.1.1 on page 528
- `ompd_callback_memory_free_fn_t`, Section 4.3.3.1.2 on page 529
- `ompd_callback_print_string_fn_t`, Section 4.3.3.5.1 on page 537
- `ompd_callback_sizeof_fn_t`, Section 4.3.3.2.2 on page 531
- `ompd_callback_symbol_addr_fn_t`, Section 4.3.3.3.1 on page 532
- `ompd_callback_memory_read_fn_t`, Section 4.3.3.3.2 on page 533
- `ompd_callback_memory_write_fn_t`, Section 4.3.3.3.3 on page 535
- `ompd_callback_device_host_fn_t`, Section 4.3.3.4.1 on page 536

- `ompd_callback_get_thread_context_for_thread_id_fn_t`, Section 4.3.3.2.1 on page 530

4.3.4 OMPD Tool Interface Routines

4.3.4.1 Per OMPD Library Initialization and Finalization

The OMPD library must be initialized exactly once after it is loaded, and finalized exactly once before it is unloaded. Per OpenMP process or core file initialization and finalization are also required.

Once loaded, the tool can determine the version of the OMPD API supported by the library by calling `ompd_get_api_version` (see Section 4.3.4.1.2 on page 541). If the tool supports the version returned by `ompd_get_api_version`, the tool starts the initialization by calling `ompd_initialize` (see Section 4.3.4.1.1 on page 540) using the version of the OMPD API supported by the library. If the tool does not support the version returned by `ompd_get_api_version`, it may attempt to call `ompd_initialize` with a different version.

4.3.4.1.1 `ompd_initialize`

Summary

The `ompd_initialize` function initializes the OMPD library.

Format

```
ompd_rc_t ompd_initialize(  
    ompd_word_t api_version,  
    const ompd_callbacks_t *callbacks  
);
```

Description

The above initialization is performed for each OMPD library that is loaded by an OMPD using tool. There may be more than one library present in a third-party tool, such as a debugger, because the tool may be controlling a number of devices that may be using different runtime systems which require different OMPD libraries. This initialization must be performed exactly once before the tool can begin operating on a OpenMP process or core file.

Description of Arguments

The *api_version* input argument is the OMPD API version that the tool is requesting to use. The tool may call `ompd_get_api_version` to obtain the latest version supported by the OMPD library.

The tool provides the OMPD library with a set of callback functions in the *callbacks* input argument which enables the OMPD library to allocate and deallocate memory in the tool's address space, to lookup the sizes of basic primitive types in the device, to lookup symbols in the device, as well as to read and write memory in the device.

Cross References

- `ompd_rc_t` type, see Section 4.3.2.8 on page 525.
- `ompd_callbacks_t` type, see Section 4.3.3.6 on page 538.
- `ompd_get_api_version` call, see Section 4.3.4.1.2 on page 541.

4.3.4.1.2 `ompd_get_api_version`

Summary

The `ompd_get_api_version` function returns the OMPD API version.

Format

```
ompd_rc_t ompd_get_api_version(ompd_word_t *version);
```

Description

The tool may call this function to obtain the latest OMPD API version number of the OMPD library.

Description of Arguments

The latest version number is returned in to the location pointed to by the *version* output argument

Cross References

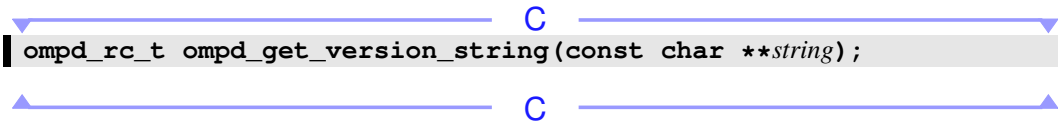
- `ompd_rc_t` type, see Section 4.3.2.8 on page 525.

1 4.3.4.1.3 `ompd_get_version_string`

2 **Summary**

3 The `ompd_get_version_string` function returns a descriptive string for the OMPD API
4 version.

5 **Format**

6 The code signature is: `ompd_rc_t ompd_get_version_string(const char **string);`. The signature is enclosed in a grey box with blue arrows pointing outwards from the top and bottom.

7 **Description**

8 The tool may call this function to obtain a pointer to a descriptive version string of the OMPD
9 library.

10 **Description of Arguments**

11 A pointer to a descriptive version string will be placed into the `*string` output argument. The string
12 returned by the OMPD library is ‘owned’ by the library, and it must not be modified or released by
13 the tool. It is guaranteed to remain valid for as long as the library is loaded.

14 `ompd_get_version_string` may be called before `ompd_initialize`
15 (see Section 4.3.4.1.1 on page 540). Accordingly, the OMPD library must not use heap or stack
16 memory for the string it returns to the tool.

17 The signatures of `ompd_get_api_version` (see Section 4.3.4.1.2 on page 541) and
18 `ompd_get_version_string` are guaranteed not to change in future versions of the API. In
19 contrast, the type definitions and prototypes in the rest of the API do not carry the same guarantee.
20 Therefore an OMPD using tool should check the version of the API of a loaded OMPD library
21 before calling any other function of the API.

22 **Cross References**

- 23 • `ompd_rc_t` type, see Section 4.3.2.8 on page 525.

24 4.3.4.1.4 `ompd_finalize`

25 **Summary**

26 When the tool is finished with the OMPD library it should call `ompd_finalize` before
27 unloading the library.

Description

On return from `ompd_process_initialize` the address space handle is owned by the tool. This function must be called before any OMPD operations are performed on the OpenMP process. `ompd_process_initialize` gives the OMPD library an opportunity to confirm that it is capable of handling the OpenMP process or core file identified by the `context`. Incompatibility is signaled by a return value of `ompd_rc_incompatible`. On return, the handle is owned by the tool, which must release it using `ompd_release_address_space_handle`.

Description of Arguments

The input argument `context` is an opaque handle provided by the tool to address an address space. On return the output argument `handle` provides an opaque handle to the tool for this address space, which the tool is responsible for releasing when it is no longer needed

Cross References

- `ompd_rc_t` type, see Section 4.3.2.8 on page 525.
- `ompd_address_space_context_t` type, see Section 4.3.2.7 on page 524.
- `ompd_address_space_handle_t` type, see Section 4.3.2.4 on page 523.
- `ompd_release_address_space_handle` type, see Section 4.3.4.2.3 on page 545.

4.3.4.2.2 `ompd_device_initialize`

Summary

A tool obtains an address space handle for a device that has at least one active target region by calling `ompd_device_initialize`.

Format

```
ompd_rc_t ompd_device_initialize(  
    ompd_address_space_handle_t *process_handle,  
    ompd_address_space_context_t *device_context,  
    omp_device_t kind,  
    ompd_size_t sizeof_id,  
    void *id,  
    ompd_address_space_handle_t **device_handle  
);
```


Description

On return from `ompd_device_initialize` the address space handle is owned by the tool.

Description of Arguments

The input argument `process_handle` is an opaque handle provided by the tool to reference the address space of the OpenMP process. The input argument `device_context` is an opaque handle provided by the tool to reference a device address space. The input arguments `kind`, `sizeof_id`, and `id` represent a device identifier. On return the output argument `device_handle` provides an opaque handle to the tool for this address space.

Cross References

- `ompd_rc_t` type, see Section 4.3.2.8 on page 525.
- `ompd_address_space_context_t` type, see Section 4.3.2.7 on page 524.
- `ompd_device_t` type, see Section 4.3.2.2 on page 522.
- `ompd_size_t` type, see Section 4.3.2.1.1 on page 520.
- `ompd_address_space_handle_t` type, see Section 4.3.2.4 on page 523.

4.3.4.2.3 `ompd_release_address_space_handle`

Summary

A tool calls `ompd_release_address_space_handle` to release an address space handle.

Format

```
      C  
ompd_rc_t ompd_release_address_space_handle(  
    ompd_address_space_handle_t *handle  
);
```

Description

When the tool is finished with the OpenMP process address space handle it should call `ompd_release_address_space_handle` to release the handle and give the OMPD library the opportunity to release any resources it may have related to the address space.

1 **Description of Arguments**

2 The input argument *handle* is an opaque handle for an address space to be released.

3 **Restrictions**

4 Using an address space context after releasing the corresponding address space handle results in
5 undefined behavior.

6 **Cross References**

- 7 • `ompd_rc_t` type, see Section 4.3.2.8 on page 525.
8 • `ompd_address_space_handle_t` type, see Section 4.3.2.4 on page 523.

9 **4.3.4.3 Thread and Signal Safety**

10 The OMPD library does not need to be reentrant. It is the responsibility of the tool to ensure that
11 only one thread enters the OMPD library at a time. The OMPD library must not install signal
12 handlers or otherwise interfere with the tool’s signal configuration.

13 **4.3.4.4 Address Space Information**

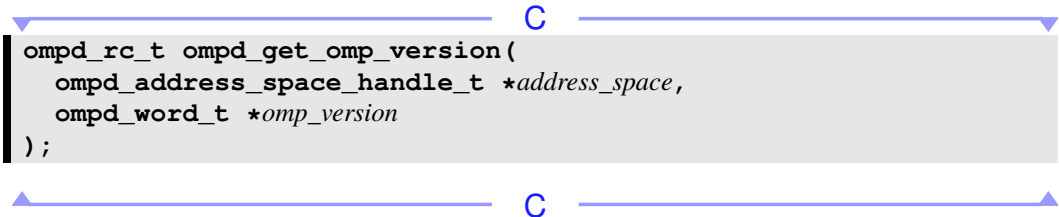
14 **4.3.4.4.1 `ompd_get_omp_version`**

15 **Summary**

16 This function may be called by the tool to obtain the version of the OpenMP API associated with an
17 address space.

18 **Format**

```
19 ompd_rc_t ompd_get_omp_version(  
20     ompd_address_space_handle_t *address_space,  
21     ompd_word_t *omp_version  
22 );
```

A diagram showing the C language signature of the function `ompd_get_omp_version`. The signature is enclosed in a light gray box with a vertical bar on the left. Above the box, a blue arrow points from the text 'C' to the right. Below the box, another blue arrow points from the text 'C' to the right.

1 Description

2 The tool may call this function to obtain the version of the OpenMP API associated with the
3 address space.

4 Description of Arguments

5 The input argument *address_space* is an opaque handle provided by the tool to reference the
6 address space of the OpenMP process or device.

7 Upon return, the output argument *omp_version* will contain the version of the OpenMP runtime in
8 the *_OPENMP* version macro format.

9 Cross References

- 10 • `ompd_rc_t` type, see Section [4.3.2.8](#) on page [525](#).
- 11 • `ompd_address_space_handle_t` type, see Section [4.3.2.4](#) on page [523](#).

12 4.3.4.4.2 `ompd_get_omp_version_string`

13 Summary

14 The `ompd_get_omp_version_string` function returns a descriptive string for the OpenMP
15 API version associated with an address space.

16 Format

```
17 ompd_rc_t ompd_get_omp_version_string(  
18     ompd_address_space_handle_t *address_space,  
19     const char **string  
20 );
```

21 Description

22 After initialization, the tool may call this function to obtain the version of the OpenMP API
23 associated with an address space.

Description of Arguments

The input argument `address_space` is an opaque handle provided by the tool to reference the address space of the OpenMP process or device. A pointer to a descriptive version string will be placed into the `*string` output argument.

After returning from the call, the string is ‘owned’ by the third-party tool. The string storage must be allocated by the OMPD library using the memory allocation callback provided by the tool. The third-party tool is responsible for releasing the memory.

Cross References

- `ompd_rc_t` type, see Section 4.3.2.8 on page 525.
- `ompd_address_space_handle_t` type, see Section 4.3.2.4 on page 523.

4.3.4.5 Thread Handles

4.3.4.5.1 `ompd_get_thread_in_parallel`

Summary

The `ompd_get_thread_in_parallel` operation enables a tool to obtain handles for OpenMP threads associated with a parallel region.

Format

```
ompd_rc_t ompd_get_thread_in_parallel(  
    ompd_parallel_handle_t *parallel_handle,  
    int thread_num,  
    ompd_thread_handle_t **thread_handle  
);
```

Description

A successful invocation of `ompd_get_thread_in_parallel` returns a pointer to a thread handle in `*thread_handle`. This call yields meaningful results only if all OpenMP threads in the parallel region are stopped.

Description of Arguments

The input argument *parallel_handle* is an opaque handle for a parallel region and selects the parallel region to operate on. The input argument **thread_num** selects the thread of the team to be returned. On return the output argument *thread_handle* is an opaque handle for the selected thread.

Restrictions

The value of **thread_num** must be a non-negative integer smaller than the team size provided as the *ompd-team-size-var* from **ompd_get_icv_from_scope**.

Cross References

- **ompd_rc_t** type, see Section 4.3.2.8 on page 525.
- **ompd_parallel_handle_t** type, see Section 4.3.2.4 on page 523.
- **ompd_thread_handle_t** type, see Section 4.3.2.4 on page 523.
- **ompd_get_icv_from_scope** call, see Section 4.3.4.9.2 on page 573.

4.3.4.5.2 **ompd_get_thread_handle**

Summary

Mapping a native thread to an OMPD thread handle.

Format

```
ompd_rc_t ompd_get_thread_handle(  
    ompd_address_space_handle_t *handle,  
    ompd_thread_id_t kind,  
    ompd_size_t sizeof_thread_id,  
    const void *thread_id,  
    ompd_thread_handle_t **thread_handle  
);
```

Description

OMP provides the function `ompd_get_thread_handle` to inquire whether a native thread is an OpenMP thread or not. On success, the thread identifier is an OpenMP thread and `*thread_handle` is initialized to a pointer to the thread handle for the OpenMP thread.

Description of Arguments

The input argument *handle* is an opaque handle provided by the tool to reference to an address space. The input arguments *kind*, *sizeof_thread_id*, and *thread_id* represent a thread identifier. On return the output argument *thread_handle* provides an opaque handle to the tool for thread within the provided address space.

The thread identifier **thread_id* is guaranteed to be valid for the duration of the call. If the OMPD library needs to retain the thread identifier it must copy it.

Cross References

- `ompd_rc_t` type, see Section 4.3.2.8 on page 525.
- `ompd_address_space_handle_t` type, see Section 4.3.2.4 on page 523.
- `ompd_thread_id_t` type, see Section 4.3.2.3 on page 522.
- `ompd_size_t` type, see Section 4.3.2.1.1 on page 520.
- `ompd_thread_handle_t` type, see Section 4.3.2.4 on page 523.

4.3.4.5.3 `ompd_release_thread_handle`

Summary

This operation releases a thread handle.

Format

```
ompd_rc_t ompd_release_thread_handle(  
    ompd_thread_handle_t *thread_handle  
);
```

Description

Thread handles are opaque to tools, which therefore cannot release them directly. Instead, when the tool is finished with a thread handle it must pass it to the OMPD `ompd_release_thread_handle` routine for disposal.

Description of Arguments

The input argument *thread_handle* is an opaque handle for a thread to be released.

Cross References

- `ompd_rc_t` type, see Section 4.3.2.8 on page 525.
- `ompd_thread_handle_t` type, see Section 4.3.2.4 on page 523.

4.3.4.5.4 `ompd_thread_handle_compare`

Summary

The `ompd_thread_handle_compare` operation allows tools to compare two thread handles.

Format

```
C
ompd_rc_t ompd_thread_handle_compare(
    ompd_thread_handle_t *thread_handle_1,
    ompd_thread_handle_t *thread_handle_2,
    int *cmp_value
);
C
```

Description

The internal structure of thread handles is opaque to a tool. While the tool can easily compare pointers to thread handles, it cannot determine whether handles of two different addresses refer to the same underlying thread. This function can be used to compare thread handles.

On success, `ompd_thread_handle_compare` returns in `*cmp_value` a signed integer value that indicates how the underlying threads compare: a value less than, equal to, or greater than 0 indicates that the thread corresponding to `thread_handle_1` is, respectively, less than, equal to, or greater than that corresponding to `thread_handle_2`.

Description of Arguments

The input arguments *thread_handle_1* and *thread_handle_2* are opaque handles for threads. On return the output argument *cmp_value* is set to a signed integer value.

Cross References

- `ompd_rc_t` type, see Section 4.3.2.8 on page 525.
- `ompd_thread_handle_t` type, see Section 4.3.2.4 on page 523.

4.3.4.5.5 `ompd_get_thread_id`

Summary

Mapping an OMPD thread handle to a native thread.

Format

```
ompd_rc_t ompd_get_thread_id(  
    ompd_thread_handle_t *thread_handle,  
    ompd_thread_id_t kind,  
    ompd_size_t sizeof_thread_id,  
    void *thread_id  
);
```

Description

`ompd_get_thread_id` performs the mapping between an OMPD thread handle and a thread identifier.

Description of Arguments

The input argument *thread_handle* is an opaque thread handle. The input argument *kind* represents the thread identifier. The input argument *sizeof_thread_id* represents the size of the thread identifier. The output argument *thread_id* is a buffer that represents a thread identifier.

Cross References

- `ompd_rc_t` type, see Section 4.3.2.8 on page 525.
- `ompd_thread_handle_t` type, see Section 4.3.2.4 on page 523.
- `ompd_thread_id_t` type, see Section 4.3.2.3 on page 522.
- `ompd_size_t` type, see Section 4.3.2.1.1 on page 520.

4.3.4.6 Parallel Region Handles

4.3.4.6.1 `ompd_get_current_parallel_handle`

Summary

The `ompd_get_current_parallel_handle` operation enables the tool to obtain a pointer to the parallel handle for the current parallel region associated with an OpenMP thread.

Format

```
C
ompd_rc_t ompd_get_current_parallel_handle(
    ompd_thread_handle_t *thread_handle,
    ompd_parallel_handle_t **parallel_handle
);
```

Description

This call is meaningful only if the thread whose handle is provided is stopped. The parallel handle must be released by calling `ompd_release_parallel_handle`.

Description of Arguments

The input argument `thread_handle` is an opaque handle for a thread and selects the thread to operate on. On return the output argument `parallel_handle` is set to a handle for the parallel region currently executing on this thread if there is any.

Cross References

- `ompd_rc_t` type, see Section 4.3.2.8 on page 525.
- `ompd_thread_handle_t` type, see Section 4.3.2.4 on page 523.
- `ompd_parallel_handle_t` type, see Section 4.3.2.4 on page 523.
- `ompd_release_parallel_handle` call, see Section 4.3.4.6.4 on page 556.

4.3.4.6.2 `ompd_get_enclosing_parallel_handle`

Summary

The `ompd_get_enclosing_parallel_handle` operation enables a tool to obtain a pointer to the parallel handle for the parallel region enclosing the parallel region specified by `parallel_handle`.

Format

```
ompd_rc_t ompd_get_enclosing_parallel_handle(  
    ompd_parallel_handle_t *parallel_handle,  
    ompd_parallel_handle_t **enclosing_parallel_handle  
);
```

Description

This call is meaningful only if at least one thread in the parallel region is stopped. A pointer to the parallel handle for the enclosing region is returned in `*enclosing_parallel_handle`. After the call the handle is owned by the tool, which must release it when it is no longer required by calling `ompd_release_parallel_handle`.

Description of Arguments

The input argument `parallel_handle` is an opaque handle for a parallel region and selects the parallel region to operate on. On return the output argument `parallel_handle` is set to a handle for the parallel region enclosing the selected parallel region.

Cross References

- `ompd_rc_t` type, see Section [4.3.2.8](#) on page [525](#).
- `ompd_parallel_handle_t` type, see Section [4.3.2.4](#) on page [523](#).
- `ompd_release_parallel_handle` call, see Section [4.3.4.6.4](#) on page [556](#).

4.3.4.6.3 `ompd_get_task_parallel_handle`

Summary

The `ompd_get_task_parallel_handle` operation enables a tool to obtain a pointer to the parallel handle for the parallel region enclosing the task region specified by `task_handle`.

Format

```
ompd_rc_t ompd_get_task_parallel_handle(  
    ompd_task_handle_t *task_handle,  
    ompd_parallel_handle_t **task_parallel_handle  
);
```

Description

This call is meaningful only if at least one thread in the parallel region is stopped. A pointer to the parallel regions handle is returned in `*task_parallel_handle`. The parallel handle is owned by the tool, which must release it by calling `ompd_release_parallel_handle`.

Description of Arguments

The input argument `task_handle` is an opaque handle for a task and selects the task to operate on. On return the output argument `parallel_handle` is set to a handle for the parallel region enclosing the selected task.

Cross References

- `ompd_rc_t` type, see Section [4.3.2.8](#) on page [525](#).
- `ompd_task_handle_t` type, see Section [4.3.2.4](#) on page [523](#).
- `ompd_parallel_handle_t` type, see Section [4.3.2.4](#) on page [523](#).
- `ompd_release_parallel_handle` call, see Section [4.3.4.6.4](#) on page [556](#).

1 4.3.4.6.4 ompd_release_parallel_handle

2 Summary

3 This operation allows releasing a parallel region handle.

4 Format

```
5 ompd_rc_t ompd_release_parallel_handle(  
6     ompd_parallel_handle_t *parallel_handle  
7 );
```

8 Description

9 Parallel region handles are opaque to the tool, which therefore cannot release them directly. Instead,
10 when the tool is finished with a parallel region handle it must pass it to the OMPD
11 **ompd_release_parallel_handle** routine for disposal.

12 Description of Arguments

13 The input argument *parallel_handle* is an opaque handle for a parallel region to be released.

14 Cross References

- 15 • **ompd_rc_t** type, see Section [4.3.2.8](#) on page [525](#).
- 16 • **ompd_parallel_handle_t** type, see Section [4.3.2.4](#) on page [523](#).

17 4.3.4.6.5 ompd_parallel_handle_compare

18 Summary

19 The **ompd_parallel_handle_compare** operation allows a tool to compare two parallel
20 region handles.

Format

```
ompd_rc_t ompd_parallel_handle_compare(  
    ompd_parallel_handle_t *parallel_handle_1,  
    ompd_parallel_handle_t *parallel_handle_2,  
    int *cmp_value  
);
```

Description

The internal structure of parallel region handles is opaque to the tool. While the tool can easily compare pointers to parallel region handles, it cannot determine whether handles at two different addresses refer to the same underlying parallel region.

On success, `ompd_parallel_handle_compare` returns in `*cmp_value` a signed integer value that indicates how the underlying parallel regions compare: a value less than, equal to, or greater than 0 indicates that the region corresponding to `parallel_handle_1` is, respectively, less than, equal to, or greater than that corresponding to `parallel_handle_2`.

For OMPD libraries that always have a single, unique, underlying parallel region handle for a given parallel region, this operation reduces to a simple comparison of the pointers. However, other implementations may take a different approach, and therefore the only reliable way of determining whether two different pointers to parallel regions handles refer the same or distinct parallel regions is to use `ompd_parallel_handle_compare`.

Allowing parallel region handles to be compared allows the tool to hold them in ordered collections. The means by which parallel region handles are ordered is implementation-defined.

Description of Arguments

The input arguments `parallel_handle_1` and `parallel_handle_2` are opaque handles corresponding to parallel regions. On return the output argument `cmp_value` returns a signed integer value that indicates how the underlying parallel regions compare.

Cross References

- `ompd_rc_t` type, see Section [4.3.2.8](#) on page [525](#).
- `ompd_parallel_handle_t` type, see Section [4.3.2.4](#) on page [523](#).

1 4.3.4.7 Task Handles

2 4.3.4.7.1 ompd_get_current_task_handle

3 Summary

4 A tool uses the **ompd_get_current_task_handle** operation to obtain a pointer to the task
5 handle for the current task region associated with an OpenMP thread.

6 Format

```
7 ompd_rc_t ompd_get_current_task_handle(  
8     ompd_thread_handle_t *thread_handle,  
9     ompd_task_handle_t **task_handle  
10 );
```

C

11 Description

12 This call is meaningful only if the thread whose handle is provided is stopped. The task handle
13 must be released by calling **ompd_release_task_handle**.

14 Description of Arguments

15 The input argument *thread_handle* is an opaque handle for a thread and selects the thread to operate
16 on. On return the output argument *task_handle* is set to a handle for the task currently executing on
17 the thread.

18 Cross References

- 19 • **ompd_rc_t** type, see Section [4.3.2.8](#) on page [525](#).
- 20 • **ompd_thread_handle_t** type, see Section [4.3.2.4](#) on page [523](#).
- 21 • **ompd_task_handle_t** type, see Section [4.3.2.4](#) on page [523](#).
- 22 • **ompd_release_task_handle** call, see Section [4.3.4.7.5](#) on page [562](#).

1 4.3.4.7.2 ompd_get_generating_task_handle

2 Summary

3 The OMPD tool interface routine **ompd_get_generating_task_handle** provides access to
4 the task that encountered the OpenMP task construct which caused the task represented by
5 **task_handle** to be created.

6 Format

```
7 ompd_rc_t ompd_get_generating_task_handle(  
8     ompd_task_handle_t *task_handle,  
9     ompd_task_handle_t **generating_task_handle  
10 );
```

11 Description

12 In this operation, the generating task is the OpenMP task that was active when the task specified by
13 **task_handle** was created. This call is meaningful only if the thread executing the task specified
14 by **task_handle** is stopped. The generating task handle must be released by calling
15 **ompd_release_task_handle**.

16 Description of Arguments

17 The input argument *task_handle* is an opaque handle for a task and selects the task to operate on.
18 On return the output argument *generating_task_handle* is set to a handle for the task that created
19 the selected task.

20 Cross References

- 21 • **ompd_rc_t** type, see Section [4.3.2.8](#) on page [525](#).
- 22 • **ompd_task_handle_t** type, see Section [4.3.2.4](#) on page [523](#).
- 23 • **ompd_release_task_handle** call, see Section [4.3.4.7.5](#) on page [562](#).

1 4.3.4.7.3 ompd_get_scheduling_task_handle

2 Summary

3 The OMPD tool interface routine **ompd_get_generating_task_handle** provides access to
4 the task handle for the task that scheduled the task represented by **task_handle** on a task
5 scheduling point.

6 Format

```
7 ompd_rc_t ompd_get_scheduling_task_handle(  
8     ompd_task_handle_t *task_handle,  
9     ompd_task_handle_t **scheduling_task_handle  
10 );
```

11 Description

12 The scheduling task in this routine is the OpenMP task that was active when the task specified by
13 **task_handle** was scheduled. This call is meaningful only if the thread executing the task
14 specified by **task_handle** is stopped. The scheduling task handle must be released by calling
15 **ompd_release_task_handle**.

16 Description of Arguments

17 The input argument *task_handle* is an opaque handle for a task and selects the task to operate on.
18 On return the output argument *scheduling_task_handle* is set to a handle for the task that is still on
19 the stack of execution on the same thread and was deferred in favor of executing the selected task.

20 Cross References

- 21 • **ompd_rc_t** type, see Section [4.3.2.8](#) on page [525](#).
- 22 • **ompd_task_handle_t** type, see Section [4.3.2.4](#) on page [523](#).
- 23 • **ompd_release_task_handle** call, see Section [4.3.4.7.5](#) on page [562](#).

1 4.3.4.7.4 `ompd_get_task_in_parallel`

2 **Summary**

3 The `ompd_get_task_in_parallel` operation enables a tool to obtain handles for the implicit
4 tasks associated with a parallel region.

5 **Format**

```
6 ompd_rc_t ompd_get_task_in_parallel(  
7     ompd_parallel_handle_t *parallel_handle,  
8     int thread_num,  
9     ompd_task_handle_t **task_handle  
10 );
```

11 **Description**

12 A successful invocation of `ompd_get_task_in_parallel` returns a pointer to a task handle
13 in `*task_handle`. This call yields meaningful results only if all OpenMP threads in the parallel
14 region are stopped.

15 **Description of Arguments**

16 The input argument `parallel_handle` is an opaque handle for a parallel region and selects the
17 parallel region to operate on. The input argument `thread_num` selects the implicit task of the
18 team to be returned. The selected implicit task would return `thread_num` from a call of the
19 `ompd_get_thread_num()` routine. On return the output argument `task_handle` is an opaque
20 handle for the selected implicit task.

21 **Restrictions**

22 The value of `thread_num` must be a non-negative integer smaller than the team size provided as
23 the `ompd-team-size-var` from `ompd_get_icv_from_scope`.

24 **Cross References**

- 25 • `ompd_rc_t` type, see Section [4.3.2.8](#) on page [525](#).
- 26 • `ompd_parallel_handle_t` type, see Section [4.3.2.4](#) on page [523](#).
- 27 • `ompd_task_handle_t` type, see Section [4.3.2.4](#) on page [523](#).
- 28 • `ompd_get_icv_from_scope` call, see Section [4.3.4.9.2](#) on page [573](#).

1 4.3.4.7.5 ompd_release_task_handle

2 Summary

3 This operation releases a task handle.

4 Format

```
5 ompd_rc_t ompd_release_task_handle(  
6     ompd_task_handle_t *task_handle  
7 );
```

8 Description

9 Task handles are opaque to the tool, which therefore cannot release them directly. Instead, when the
10 tool is finished with a task handle it must pass it to the OMPD **ompd_release_task_handle**
11 routine for disposal.

12 Description of Arguments

13 The input argument *task_handle* is an opaque handle for a task to be released.

14 Cross References

- 15 • **ompd_rc_t** type, see Section [4.3.2.8](#) on page [525](#).
- 16 • **ompd_task_handle_t** type, see Section [4.3.2.4](#) on page [523](#).

17 4.3.4.7.6 ompd_task_handle_compare

18 Summary

19 The **ompd_task_handle_compare** operations allows a tool to compare task handles.

Format

```
ompd_rc_t ompd_task_handle_compare(  
    ompd_task_handle_t *task_handle_1,  
    ompd_task_handle_t *task_handle_2,  
    int *cmp_value  
);
```

Description

The internal structure of task handles is opaque to the tool. While the tool can easily compare pointers to task handles, it cannot determine whether handles at two different addresses refer to the same underlying task.

On success, `ompd_task_handle_compare` returns in `*cmp_value` a signed integer value that indicates how the underlying tasks compare: a value less than, equal to, or greater than 0 indicates that the task corresponding to `task_handle_1` is, respectively, less than, equal to, or greater than that corresponding to `task_handle_2`.

For OMPD libraries that always have a single, unique, underlying task handle for a given task, this operation reduces to a simple comparison of the pointers. However, other implementations may take a different approach, and therefore the only reliable way of determining whether two different pointers to task handles refer the same or distinct task is to use `ompd_task_handle_compare`.

Allowing task handles to be compared allows the tool to hold them in ordered collections. The means by which task handles are ordered is implementation-defined.

Description of Arguments

The input arguments `task_handle_1` and `task_handle_2` are opaque handles corresponding to tasks. On return the output argument `cmp_value` returns a signed integer value that indicates how the underlying tasks compare.

Cross References

- `ompd_rc_t` type, see Section [4.3.2.8](#) on page [525](#).
- `ompd_task_handle_t` type, see Section [4.3.2.4](#) on page [523](#).

1 4.3.4.7.7 `ompd_get_task_function`

2 **Summary**

3 Task Function Entry Point

4 **Format**

```
5 ompd_rc_t ompd_get_task_function (  
6     ompd_task_handle_t *task_handle,  
7     ompd_address_t *entry_point  
8 );
```

9 **Description**

10 The `ompd_get_task_function` returns the entry point of the code that corresponds to the
11 body of code executed by the task:

12 **Description of Arguments**

13 The input argument `task_handle` is an opaque handle for a task and selects the task to operate on.
14 On return the output argument `entry_point` is set an address that describes the begin of application
15 code which executes the task region.

16 **Cross References**

- 17 • `ompd_rc_t` type, see Section [4.3.2.8](#) on page [525](#).
- 18 • `ompd_task_handle_t` type, see Section [4.3.2.4](#) on page [523](#).
- 19 • `ompd_address_t` type, see Section [4.3.2.1.4](#) on page [521](#).

20 4.3.4.7.8 `ompd_get_task_frame`

21 **Summary**

22 For the specified task, extract the task's frame pointers maintained by an OpenMP implementation.

Format

```
ompd_rc_t ompd_get_task_frame (  
    ompd_task_handle_t *task_handle,  
    ompd_address_t *exit_frame,  
    ompd_address_t *enter_frame  
);
```

Description

An OpenMP implementation maintains an `omp_frame_t` object for every implicit or explicit task. For the task identified by `task_handle`, `ompd_get_task_frame` extracts the `enter_frame` and `exit_frame` fields of the task's `omp_frame_t` object.

Description of Arguments

The argument `task_handle` specifies an OpenMP task.

The argument `exit_frame` is a pointer to an `ompd_address_t` object that the OMPD library will modify to return the segment and address that represent the value of the `exit_frame` field of the `omp_frame_t` object associated with the specified task.

The argument `enter_frame` is a pointer to an `ompd_address_t` object that the OMPD library will modify to return the segment and address that represent the value of the `enter_frame` field of the `omp_frame_t` object associated with the specified task.

Cross References

- `ompd_rc_t` type, see Section [4.3.2.8](#) on page [525](#).
- `ompd_task_handle_t` type, see Section [4.3.2.4](#) on page [523](#).
- `ompd_address_t` type, see Section [4.3.2.1.4](#) on page [521](#).
- `omp_frame_t` type, see Section [4.4.1.2](#) on page [589](#).

4.3.4.7.9 `ompd_enumerate_states`

Summary

Enumerate thread states supported by an OpenMP implementation.

Format

```
ompd_rc_t ompd_enumerate_states (  
    ompd_address_space_handle_t *address_space_handle,  
    ompd_word_t current_state,  
    ompd_word_t *next_state,  
    const char **next_state_name,  
    ompd_word_t *more_enums  
);
```

Description

An OpenMP implementation may support only a subset of the states defined by the **omp_states_t** enumeration type. In addition, an OpenMP implementation may support implementation-specific states. The **ompd_enumerate_states** call enables a tool to enumerate the thread states supported by an OpenMP implementation.

When the *current_state* input argument is set to a thread state supported by an OpenMP implementation, the call will assign the value and string name of the next thread state in the enumeration to the locations pointed to by the *next_state* and *next_state_name* output arguments, respectively.

After returning from the call, the string *next_state_name* is ‘owned’ by the third-party tool. The string storage must be allocated by the OMPD library using the memory allocation callback provided by the tool. The third-party tool is responsible for releasing the memory.

Whenever one or more states are left in the enumeration, the call will set the location pointed to by the *more_enums* output argument to 1. When the last state in the enumeration is passed in *current_state*, the call will set the location pointed to by the *more_enums* output argument to 0

Description of Arguments

The address space is identified by the input argument *address_space_handle*.

The input argument *current_state* must be a thread state supported by the OpenMP implementation. To begin enumerating the states that an OpenMP implementation supports, a tool should pass **omp_state_undefined** as the value of the input argument *current_state*. Subsequent calls to **ompd_enumerate_states** by the tool should pass the value returned by the call in the *next_state* output argument.

The output argument *next_state* is a pointer to an integer where the call will return the value of the next state in the enumeration.

1 The output argument *next_state_name* is a pointer to a character string pointer, where the call will
2 return a string describing the next state.

3 The output argument *more_enums* is a pointer to an integer where the call will return a value of 1
4 when there are more states left to enumerate or a value of 0 when there are not.

5 **Constraints on Arguments**

6 Any string returned through the argument *next_state_name* must be immutable and defined for the
7 lifetime of a program execution.

8 **Cross References**

- 9 • `ompd_address_space_handle_t` type, see Section [4.3.2.4](#) on page [523](#).
- 10 • `ompd_state_t` type, see Section [4.4.1.1](#) on page [584](#).
- 11 • `ompd_rc_t` type, see Section [4.3.2.8](#) on page [525](#).

12 **4.3.4.7.10 ompd_get_state**

13 **Summary**

14 This function allows a third party tool to interrogate the OMPD library about the state of a thread.

15 **Format**

```
16 ompd_rc_t ompd_get_state (  
17     ompd_thread_handle_t *thread_handle,  
18     ompd_word_t *state,  
19     omp_wait_id_t *wait_id  
20 );
```

C

C

21 **Description**

22 This function returns the state of an OpenMP thread.

Description of Arguments

The input argument *thread_handle* is a thread handle. The output argument *state* represents the state of the thread that is represented by the thread handle. The thread states are represented by values returned by `ompd_enumerate_states`.

The output argument *wait_id* is a pointer to an opaque handle available to receive the value of the thread's wait identifier. If the *wait_id* pointer is not **NULL**, it will contain the value of the thread's wait identifier **wait_id*. If the thread state is not one of the specified wait states, the value of **wait_id* is undefined.

Cross References

- `ompd_rc_t` type, see Section 4.3.2.8 on page 525.
- `ompd_thread_handle_t` type, see Section 4.3.2.4 on page 523.
- `ompd_wait_id_t` type, see Section 4.3.2.1.2 on page 520.
- `ompd_enumerate_states` call, see Section 4.3.4.7.9 on page 565.

4.3.4.8 Display Control variables

4.3.4.8.1 `ompd_get_display_control_vars`

Summary

Returns a list of name/value pairs for the OpenMP control variables that are user-controllable and important to the operation or performance of OpenMP programs.

Format

```
ompd_rc_t ompd_get_display_control_vars (  
    ompd_address_space_handle_t *address_space_handle,  
    const char * const * *control_vars  
);
```


Description

The function `ompd_get_display_control_vars` returns a NULL-terminated vector of strings of name/value pairs of control variables whose settings are (a) user controllable, and (b) important to the operation or performance of an OpenMP runtime system. The control variables exposed through this interface include all of the OpenMP environment variables, settings that may come from vendor or platform-specific environment variables, and other settings that affect the operation or functioning of an OpenMP runtime.

The format of the strings is:

name=a string

The third-party tool must not modify the vector or the strings (i.e., they are both `const`). The strings are NULL terminated. The vector is NULL terminated.

After returning from the call, the vector and strings are ‘owned’ by the third third-party tool. Providing the termination constraints are satisfied, the OMPD library is free to use static or dynamic memory for the vector and/or the strings, and to arrange them in memory as it pleases. If dynamic memory is used, then the OMPD library must use the `allocate` callback it received in the call to `ompd_initialize`. As the third-party tool cannot make any assumptions about the origin or layout of the memory used for the vector or strings, it cannot release the display control variables directly when they are no longer needed; instead it must use `ompd_release_display_control_vars~()`.

Description of Arguments

The address space is identified by the input argument `address_space_handle`. The vector of display control variables is returned through the output argument `control_vars`.

Cross References

- `ompd_rc_t` type, see Section [4.3.2.8](#) on page [525](#).
- `ompd_address_space_handle_t` type, see Section [4.3.2.4](#) on page [523](#).
- `ompd_release_display_control_vars` type, see Section [4.3.4.8.2](#) on page [569](#).
- `ompd_initialize` call, see Section [4.3.4.1.1](#) on page [540](#).

4.3.4.8.2 `ompd_release_display_control_vars`

Summary

Releases a list of name/value pairs of OpenMP control variables previously acquired using `ompd_get_display_control_vars`.

Format

```
ompd_rc_t ompd_release_display_control_vars (  
    const char * const **control_vars  
);
```

Description

The vector and strings returned from `ompd_get_display_control_vars` are ‘owned’ by the third-party tool, but allocated by the OMPD library. Because the third-party tool doesn’t know how the memory for the vector and strings was allocated, it cannot deallocate the memory itself. Instead, the third-party tool must call `ompd_release_display_control_vars` to release the vector and strings.

Description of Arguments

The input parameter `control_vars` is the vector of display control variables to be released.

Cross References

- `ompd_rc_t` type, see Section [4.3.2.8](#) on page [525](#).
- `ompd_get_display_control_vars` call, see Section [4.3.4.8.1](#) on page [568](#).

4.3.4.9 Accessing Scope Specific Information

4.3.4.9.1 `ompd_enumerate_icvs`

Summary

Enumerate ICVs supported by an OpenMP implementation.

Format

```
ompd_rc_t ompd_enumerate_icvs (  
    ompd_address_space_handle *handle,  
    ompd_icv_id_t current,  
    ompd_icv_id_t *next_id,  
    const char **next_icv_name,  
    ompd_scope_t *next_scope,  
    int *more  
);
```

Description

In addition to the ICVs listed in Table 2.1, an OpenMP implementation must support the OMPD specific ICVs listed in Table 4.7 in the OMPD interface. An OpenMP implementation might support additional implementation specific variables.

Also an implementation might decide to store ICVs in a different scope than suggested in Table 2.3. The `ompd_enumerate_icvs` call enables a tool to enumerate the ICVs supported by an OpenMP implementation and the related scopes.

When the `current` input argument is set to a value supported by an OpenMP implementation, the call will assign the value, string name, and scope of the next ICV in the enumeration to the locations pointed to by the `next_id`, `next_icv_name`, and `next_scope` output arguments, respectively.

After returning from the call, the string `next_icv_name` is ‘owned’ by the third-party tool. The string storage must be allocated by the OMPD library using the memory allocation callback provided by the tool. The third-party tool is responsible for releasing the memory.

Whenever one or more ICV are left in the enumeration, the call will set the location pointed to by the `more` output argument to 1. When the last ICV in the enumeration is passed in `current`, the call will set the location pointed to by the `more` output argument to 0

Description of Arguments

The address space is identified by the input argument *address_space_handle*.

The input argument *current* must be an ICV supported by the OpenMP implementation. To begin enumerating the ICVs that an OpenMP implementation supports, a tool should pass **ompd_icv_undefined** as the value of the input argument *current*. Subsequent calls to **ompd_enumerate_icvs** by the tool should pass the value returned by the call in the *next_id* output argument.

The output argument *next_id* is a pointer to an integer where the call will return the id of the next ICV in the enumeration.

The output argument *next_icv* is a pointer to a character string pointer, where the call will return a string providing the name of the next ICV.

The output argument *next_scope* is a pointer to a scope enum value, where the call will return the scope for the next ICV.

The output argument *more_enums* is a pointer to an integer where the call will return a value of 1 when there are more ICV left to enumerate or a value of 0 when there are not.

Constraints on Arguments

Any string returned through the argument *next_icv* must be immutable and defined for the lifetime of a program execution.

TABLE 4.7: OMPD specific ICVs

Variable	Scope	Meaning
<i>ompd-num-procs-var</i>	device	return value of omp_get_num_procs() when executed on this device
<i>ompd-thread-num-var</i>	task	return value of omp_get_thread_num() when executed in this task
<i>ompd-final-var</i>	task	return value of omp_in_final() when executed in this task
<i>ompd-implicit-var</i>	task	the task is an implicit task
<i>ompd-team-size-var</i>	team	return value of omp_get_num_threads() when executed in this team

Cross References

- `ompd_address_space_handle_t` type, see Section 4.3.2.4 on page 523.
- `ompd_scope_t` type, see Section 4.3.2.5 on page 523.
- `ompd_icv_id_t` type, see Section 4.3.2.6 on page 524.
- `ompd_rc_t` type, see Section 4.3.2.8 on page 525.

4.3.4.9.2 `ompd_get_icv_from_scope`

Summary

Returns the value of an ICV as present in the provided scope.

Format

```
ompd_rc_t ompd_get_icv_from_scope (  
    void *handle,  
    ompd_scope_t scope,  
    ompd_icv_id_t icv_id,  
    ompd_word_t *icv_value  
);
```

Description

The function `ompd_get_icv_from_scope` provides access to the internal control variables as defined in Tables 2.1 and 4.7.

Description of Arguments

The argument *handle* provides an OpenMP scope handle. The argument *scope* specifies the kind of scope provided in *handle*. The argument *icv_name* specifies the name of the requested ICV. On successful return, the output argument *icv_value* is set to the value of the requested ICV.

Constraints on Arguments

If the ICV cannot be represented by an integer type value, the function returns `ompd_rc_incompatible`.

The provided *handle* must match the *scope* as defined in Section 4.3.2.6 on page 524.

The provided *scope* must match the scope for *icv_id* as requested by `ompd_enumerate_icvs`.

Cross References

- `ompd_address_space_handle_t` type, see Section 4.3.2.4 on page 523.
- `ompd_thread_handle_t` type, see Section 4.3.2.4 on page 523.
- `ompd_parallel_handle_t` type, see Section 4.3.2.4 on page 523.
- `ompd_task_handle_t` type, see Section 4.3.2.4 on page 523.
- `ompd_scope_t` type, see Section 4.3.2.5 on page 523.
- `ompd_icv_id_t` type, see Section 4.3.2.6 on page 524.
- `ompd_rc_t` type, see Section 4.3.2.8 on page 525.

4.3.4.9.3 `ompd_get_icv_string_from_scope`

Summary

Returns the value of an ICV as present in the provided scope.

Format

```
ompd_rc_t ompd_get_icv_string_from_scope (  
    void *handle,  
    ompd_scope_t scope,  
    ompd_icv_id_t icv_id,  
    const char **icv_string  
);
```

Description

The function `ompd_get_icv_string_from_scope` provides access to the internal control variables as defined in Table 2.1.

Description of Arguments

The argument *handle* provides an OpenMP scope handle. The argument *scope* specifies the kind of scope provided in *handle*. The argument *icv_id* specifies the id of the requested ICV. On successful return, the output argument *icv_string* points to a string representation of the requested ICV.

After returning from the call, the string *icv_string* is ‘owned’ by the third-party tool. The string storage must be allocated by the OMPD library using the memory allocation callback provided by the tool. The third-party tool is responsible for releasing the memory.

Constraints on Arguments

Any string passed through the argument *icv_string* must be allocated by the OMPD library with the memory alloc callback `ompd_callback_memory_alloc_fn_t` and freed by the tool.

The provided *handle* must match the *scope* as defined in Section 4.3.2.6 on page 524.

The provided *scope* must match the scope for *icv_id* as requested by `ompd_enumerate_icvs`.

Cross References

- `ompd_address_space_handle_t` type, see Section 4.3.2.4 on page 523.
- `ompd_thread_handle_t` type, see Section 4.3.2.4 on page 523.
- `ompd_parallel_handle_t` type, see Section 4.3.2.4 on page 523.
- `ompd_task_handle_t` type, see Section 4.3.2.4 on page 523.
- `ompd_scope_t` type, see Section 4.3.2.5 on page 523.
- `ompd_icv_id_t` type, see Section 4.3.2.6 on page 524.
- `ompd_rc_t` type, see Section 4.3.2.8 on page 525.

4.3.4.9.4 `ompd_get_tool_data`

Summary

The `ompd_get_tool_data` function provides access to the OMPT data variable stored for each OpenMP scope.

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24

Format

```
ompd_rc_t ompd_get_tool_data(  
    void* handle,  
    ompd_scope_t scope,  
    ompd_word_t *value,  
    ompd_address_t *ptr  
);
```

Description

The function `ompd_get_tool_data` provides access to the OMPT tool data stored for each scope.

If the runtime library has no support for OMPT, the function returns `ompd_rc_unsupported`.

Description of Arguments

The argument *handle* provides an OpenMP scope handle. The argument *scope* specifies the kind of scope provided in *handle*. On return the output argument *value* is set to the *value* field of the `ompt_data_t` union stored for the selected scope. On return the output argument *ptr* is set to the *ptr* field of the `ompt_data_t` union stored for the selected scope.

Cross References

- `ompd_address_space_handle_t` type, see Section [4.3.2.4](#) on page [523](#).
- `ompd_thread_handle_t` type, see Section [4.3.2.4](#) on page [523](#).
- `ompd_parallel_handle_t` type, see Section [4.3.2.4](#) on page [523](#).
- `ompd_task_handle_t` type, see Section [4.3.2.4](#) on page [523](#).
- `ompd_scope_t` type, see Section [4.3.2.5](#) on page [523](#).
- `ompd_rc_t` type, see Section [4.3.2.8](#) on page [525](#).
- `ompt_data_t` type, see Section [4.2.3.4.3](#) on page [434](#).

1 4.3.5 Runtime Entry Points for OMPD

2 Most of the tool's OpenMP-related activity on an OpenMP program will be performed through the
3 OMPD interface. However, supporting OMPD introduced some requirements of the OpenMP
4 runtime. These fall into three categories: entrypoints the user's code in OpenMP program can call;
5 locations in the OpenMP runtime through which control must pass when specific events occur; and
6 data that must be accessible to the tool.

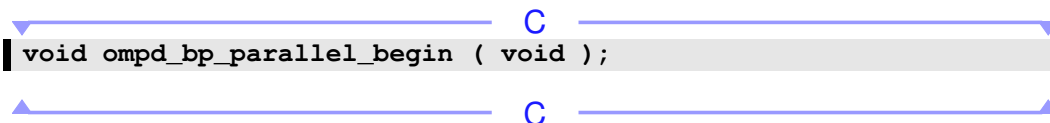
7 Neither a tool nor an OpenMP runtime system know what application code a program will launch
8 as parallel regions or tasks until the program invokes the runtime system and provides a code
9 address as an argument. To help a tool control the execution of an OpenMP program launching
10 parallel regions or tasks, the OpenMP runtime must define a number of symbols through which
11 execution must pass when particular events occur *and* data collection for OMPD is enabled. These
12 locations may, but do not have to, be subroutines. They may, for example, be labeled locations. The
13 symbols must all have external, **C**, linkage.

14 A tool can gain notification of the event by planting a breakpoint at the corresponding named
15 location.

16 4.3.5.1 Beginning Parallel Regions

17 Summary

18 The OpenMP runtime must execute through `ompd_bp_parallel_begin` when a new parallel
19 region is launched.

```
20  void ompd_bp_parallel_begin ( void );
```

21 Description

22 When starting a new parallel region, the runtime must allow execution to flow through
23 `ompd_bp_parallel_begin`. This should occur after a task encounters a parallel construct, but
24 before any implicit task starts to execute the parallel region's work.

25 Control passes through `ompd_bp_parallel_begin` once per region, and not once for each
26 thread per region.

27 At the point where the runtime reaches `ompd_bp_parallel_begin`, a tool can map the
28 encountering native thread to an OpenMP thread handle using `ompd_get_thread_handle`. At
29 this point the handle returned by `ompd_get_current_parallel_handle` is that of the new

1 parallel region. The tool can find the entry point of the user code that the new parallel region will
2 execute by passing the parallel handle region to **ompd_get_parallel_function**.

3 The number of threads participating in the parallel region is provided by the internal variable
4 *ompd-team-size-var* from **ompd_get_icv_from_scope**.

5 The task handle returned by **ompd_get_current_task_handle** will be that of the task
6 encountering the parallel construct.

7 The ‘reenter runtime’ address in the information returned by **ompd_get_task_frame** will be
8 that of the stack frame where the thread called the OpenMP runtime to handle the parallel
9 construct. The ‘exit runtime’ address will be for the stack frame where the thread left the OpenMP
10 runtime to execute the user code that encountered the parallel construct.

11 Restrictions

12 **ompd_bp_parallel_begin** has external C linkage, and no demangling or other
13 transformations are required by a tool to look up its address in the OpenMP program.

14 Conceptually **ompd_bp_parallel_begin** has the type signature given above. However, it
15 does not need to be a function, but can be a labeled location in the runtime code.

16 Cross References

- 17 • **ompd_get_thread_handle**, Section [4.3.4.5.2](#) on page [549](#)
- 18 • **ompd_get_current_parallel_handle**, Section [4.3.4.6.1](#) on page [553](#)
- 19 • **ompd_get_task_function**, Section [4.3.4.7.7](#) on page [564](#)
- 20 • **ompd_get_icv_from_scope** call, see Section [4.3.4.9.2](#) on page [573](#).
- 21 • **ompd_get_current_task_handle**, Section [4.3.4.7.1](#) on page [558](#)
- 22 • **ompd_get_task_frame**, Section [4.3.4.7.8](#) on page [564](#)

23 4.3.5.2 Ending Parallel Regions

24 The OpenMP runtime must execute through **ompd_bp_parallel_end** when a parallel region
25 ends.

```
1 | void ompd_bp_parallel_end ( void );
```

2 Description

3 When a parallel region finishes, the OpenMP runtime must allow execution to flow through
4 **ompd_bp_parallel_end**.

5 Control passes through **ompd_bp_parallel_end** once per region, and not once for each thread
6 per region.

7 At the point the runtime reaches **ompd_bp_parallel_end** the tool can map the encountering
8 native thread to an OpenMP thread handle using **ompd_get_thread_handle**.
9 **ompd_get_current_parallel_handle** returns the handle of the terminating parallel
10 region.

11 **ompd_get_current_task_handle** returns the handle of the task that encountered the
12 parallel construct that initiated the parallel region just terminating. The ‘reenter runtime’ address in
13 the frame information returned by **ompd_get_task_frame** will be that for the stack frame in
14 which the thread called the OpenMP runtime to start the parallel construct just terminating. The
15 ‘exit runtime’ address will refer to the stack frame where the thread left the OpenMP runtime to
16 execute the user code that invoked the parallel construct just terminating.

17 After passing **ompd_bp_parallel_end**, any *parallel_handle* acquired for this parallel region is
18 invalid and should be released.

19 Restrictions

20 **ompd_bp_parallel_end** has external **C** linkage, and no demangling or other transformations
21 are required by a tool to look up its address in the OpenMP program.

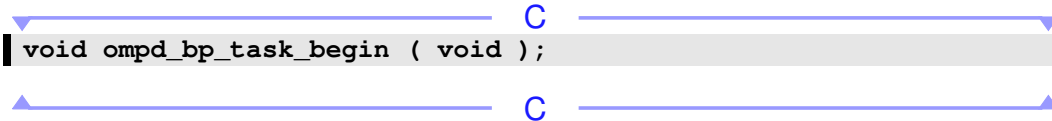
22 Conceptually **ompd_bp_parallel_end** has the type signature given above. However, it does
23 not need to be a function, but can be a labeled location in the runtime code.

24 Cross References

- 25 • **ompd_get_thread_handle**, Section [4.3.4.5.2](#) on page [549](#)
- 26 • **ompd_get_current_parallel_handle**, Section [4.3.4.6.1](#) on page [553](#)
- 27 • **ompd_get_current_task_handle**, Section [4.3.4.7.1](#) on page [558](#)
- 28 • **ompd_get_task_frame**, Section [4.3.4.7.8](#) on page [564](#)

1 4.3.5.3 Beginning Task Regions

2 The OpenMP runtime must execute through `ompd_bp_task_begin` when a new task is started.

3 
`void ompd_bp_task_begin (void);`

4 Description

5 When starting a new task region, the OpenMP runtime system must allow control to pass through
6 `ompd_bp_task_begin`.

7 The OpenMP runtime system will execute through this location after the task construct is
8 encountered, but before the new explicit task starts. At the point where the runtime reaches
9 `ompd_bp_task_begin` the tool can map the native thread to an OpenMP handle using
10 `ompd_get_thread_handle`.

11 `ompd_get_current_task_handle` returns the handle of the new task region. The entry
12 point of the user code to be executed by the new task is returned from
13 `ompd_get_task_function`.

14 Restrictions

15 `ompd_bp_task_begin` has external C linkage, and no demangling or other transformations are
16 required by a tool to look up its address in the OpenMP program.

17 Conceptually `ompd_bp_task_begin` has the type signature given above. However, it does not
18 need to be a function, but can be a labeled location in the runtime code.

19 Cross References

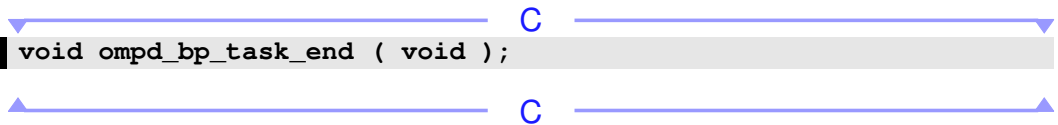
- 20 • `ompd_get_thread_handle`, Section [4.3.4.5.2](#) on page [549](#)
- 21 • `ompd_get_current_task_handle`, Section [4.3.4.7.1](#) on page [558](#)
- 22 • `ompd_get_task_function`, Section [4.3.4.7.7](#) on page [564](#)

1 4.3.5.4 Ending Task Regions

2 Summary

3 The OpenMP runtime must execute through `ompd_bp_task_end` when a task region ends.

```
4 void ompd_bp_task_end ( void );
```

A diagram consisting of two horizontal blue lines with arrowheads at both ends. The top line has a 'C' centered above it. The bottom line has a 'C' centered below it. The lines are positioned above and below the function signature in the code block above.

5 Description

6 When a task region completes, the OpenMP runtime system must allow execution to flow through
7 the location `ompd_bp_task_end`.

8 At the point where the runtime reaches `ompd_bp_task_end` the tool can use
9 `ompd_get_thread_handle` to map the encountering native thread to the corresponding
10 OpenMP thread handle. At this point `ompd_get_current_task_handle` returns the handle
11 for the terminating task.

12 After passing `ompd_bp_task_end`, any *task_handle* acquired for this task region is invalid and
13 should be released.

14 Restrictions

15 `ompd_bp_task_end` has external **C** linkage, and no demangling or other transformations are
16 required by a tool to look up its address in the OpenMP program.

17 Conceptually `ompd_bp_task_end` has the type signature given above. However, it does not
18 need to be a function, but can be a labeled location in the runtime code.

19 Cross References



- 20 • `ompd_get_thread_handle`, Section [4.3.4.5.2](#) on page [549](#)
- 21 • `ompd_get_current_task_handle`, Section [4.3.4.7.1](#) on page [558](#)

1 4.3.5.5 Beginning OpenMP Thread

2 Summary

3 The OpenMP runtime must execute through `ompd_bp_thread_begin` at the *thread-begin* and
4 the *implicit-thread-begin* event.

```
5 void ompd_bp_thread_begin ( void );
```



6 Description

7 When starting an OpenMP thread, the runtime must allow execution to flow through
8 `ompd_bp_thread_begin`. This should occur before the thread executes any OpenMP region's
9 work.

10 Restrictions

11 `ompd_bp_thread_begin` has external C linkage, and no demangling or other transformations
12 are required by a tool to look up its address in the OpenMP program.



13 Conceptually `ompd_bp_thread_begin` has the type signature given above. However, it does
14 not need to be a function, but can be a labeled location in the runtime code.

15 4.3.5.6 Ending OpenMP Thread

16 Summary

17 The OpenMP runtime must execute through `ompd_bp_thread_end` at the *thread-end* and the
18 *implicit-thread-end* event.

```
19 void ompd_bp_thread_end ( void );
```



20 Description

21 When terminating an OpenMP thread, the runtime must allow execution to flow through
22 `ompd_bp_thread_end`. This should occur after the thread executes any OpenMP region's work.

23 After passing `ompd_bp_thread_end`, any *thread_handle* acquired for this thread is invalid and
24 should be released.

Restrictions

`ompd_bp_thread_end` has external **C** linkage, and no demangling or other transformations are required by a tool to look up its address in the OpenMP program.

Conceptually `ompd_bp_thread_end` has the type signature given above. However, it does not need to be a function, but can be a labeled location in the runtime code.

Cross References

- `ompd_get_thread_handle`, Section 4.3.4.5.2 on page 549

4.3.5.7 Beginning OpenMP Device

Summary

The OpenMP runtime must execute through `ompd_bp_device_begin` at the *device-initialize* event.

```
void ompd_bp_device_begin ( void );
```

Description

When initializing a device for executing a target region, the runtime must allow execution to flow through `ompd_bp_device_begin`. This should occur before any OpenMP region's work executes on the device.

Restrictions



`ompd_bp_device_begin` has external **C** linkage, and no demangling or other transformations are required by a tool to look up its address in the OpenMP program.

Conceptually `ompd_bp_device_begin` has the type signature given above. However, it does not need to be a function, but can be a labeled location in the runtime code.

1 4.3.5.8 Ending OpenMP Device

2 Summary

3 The OpenMP runtime must execute through `ompd_bp_device_end` at the *device-finalize* event.

4  `void ompd_bp_device_end (void);` 

5 Description

6 When terminating an OpenMP thread, the runtime must allow execution to flow through
7 `ompd_bp_device_end`. This should occur after the thread executes any OpenMP region's work.

8 After passing `ompd_bp_device_end`, any *address_space_handle* acquired for this device is
9 invalid and should be released.

10 Restrictions

11 `ompd_bp_device_end` has external **C** linkage, and no demangling or other transformations are
12 required by a tool to look up its address in the OpenMP program.

13 Conceptually `ompd_bp_device_end` has the type signature given above. However, it does not
14 need to be a function, but can be a labeled location in the runtime code.

15 4.4 Tool Foundation

16 4.4.1 Data Types

17 4.4.1.1 Thread States

18 To enable a tool to understand the behavior of an executing program, an OpenMP implementation
19 maintains a state for each thread. The state maintained for a thread is an approximation of the
20 thread's instantaneous state.

1 A thread's state will be one of the values of the enumeration type `omp_state_t` or an
 2 implementation-defined state value of 512 or higher. Thread states in the enumeration fall into
 3 several classes: work, barrier wait, task wait, mutex wait, target wait, and miscellaneous.

```

4 typedef enum omp_state_t {
5     omp_state_work_serial           = 0x000,
6     omp_state_work_parallel        = 0x001,
7     omp_state_work_reduction       = 0x002,
8
9     omp_state_wait_barrier         = 0x010,
10    omp_state_wait_barrier_implicit_parallel = 0x011,
11    omp_state_wait_barrier_implicit_workshare = 0x012,
12    omp_state_wait_barrier_implicit = 0x013,
13    omp_state_wait_barrier_explicit = 0x014,
14
15    omp_state_wait_taskwait         = 0x020,
16    omp_state_wait_taskgroup        = 0x021,
17
18    omp_state_wait_mutex            = 0x040,
19    omp_state_wait_lock              = 0x041,
20    omp_state_wait_critical          = 0x042,
21    omp_state_wait_atomic           = 0x043,
22    omp_state_wait_ordered          = 0x044,
23
24    omp_state_wait_target            = 0x080,
25    omp_state_wait_target_map        = 0x081,
26    omp_state_wait_target_update    = 0x082,
27
28    omp_state_idle                   = 0x100,
29    omp_state_overhead               = 0x101,
30    omp_state_undefined              = 0x102
31 } omp_state_t;
  
```

32 A tool can query the OpenMP state of a thread at any time. If a tool queries the state of a thread that
 33 is not associated with OpenMP, the implementation reports the state as
 34 `omp_state_undefined`.

35 Some values of the enumeration type `omp_state_t` are used by all OpenMP implementations,
 36 e.g., `omp_state_work_serial`, which indicates that a thread is executing in a serial region,
 37 and `omp_state_work_parallel`, which indicates that a thread is executing in a parallel
 38 region. Other values of the enumeration type describe a thread's state at different levels of

1 specificity. For instance, an OpenMP implementation may use the state
2 **omp_state_wait_barrier** to represent all waiting at barriers. It may differentiate between
3 waiting at implicit or explicit barriers using **omp_state_wait_barrier_implicit** and
4 **omp_state_wait_barrier_explicit**. To provide full detail about the type of an implicit
5 barrier, a runtime may report **omp_state_wait_barrier_implicit_parallel** or
6 **omp_state_wait_barrier_implicit_workshare** as appropriate.

7 For states that represent waiting, an OpenMP implementation has the choice of transitioning a
8 thread to such states early or late. For instance, when an OpenMP thread is trying to acquire a lock,
9 there are several points at which an OpenMP implementation transition the thread to the
10 **omp_state_wait_lock** state. One implementation may transition the thread to the state early
11 before the thread attempts to acquire a lock. Another implementation may transition the thread to
12 the state late, only if the thread begins to spin or block to wait for an unavailable lock. A third
13 implementation may transition the thread to the state even later, e.g., only after the thread waits for
14 a significant amount of time.

15 The following sections describe the classes of states and the states in each class.

16 4.4.1.1.1 Work States

17 An OpenMP implementation reports a thread in a work state when the thread is performing serial
18 work, parallel work, or a reduction.

19 **omp_state_work_serial**

20 The thread is executing code outside all parallel regions.

21 **omp_state_work_parallel**

22 The thread is executing code within the scope of a parallel region construct.

23 **omp_state_work_reduction**

24 The thread is combining partial reduction results from threads in its team. An OpenMP
25 implementation might never report a thread in this state; a thread combining partial reduction
26 results may have its state reported as **omp_state_work_parallel** or
27 **omp_state_overhead**.

1 4.4.1.1.2 Barrier Wait States

2 An OpenMP implementation reports that a thread is in a barrier wait state when the thread is
3 awaiting completion of a barrier.

4 **omp_state_wait_barrier**

5 The thread is waiting at either an implicit or explicit barrier. A thread may enter this state early,
6 when the thread encounters a barrier, or late, when the thread begins to wait at the barrier. An
7 implementation may never report a thread in this state; instead, a thread may have its state reported
8 as **omp_state_wait_barrier_implicit** or
9 **omp_state_wait_barrier_explicit**, as appropriate.

10 **omp_state_wait_barrier_implicit**

11 The thread is waiting at an implicit barrier in a parallel region. A thread may enter this state early,
12 when the thread encounters a barrier, or late, when the thread begins to wait at the barrier. An
13 OpenMP implementation may report **omp_state_wait_barrier** for implicit barriers.

14 **omp_state_wait_barrier_implicit_parallel**

15 The description of when a thread reports a state associated with an implicit barrier is described for
16 state **omp_state_wait_barrier_implicit**. An OpenMP implementation may report
17 **omp_state_wait_barrier_implicit_parallel** for an implicit barrier that occurs at
18 the end of a parallel region. As explained in Section 4.2.4.2.11 on page 457, reporting the state
19 **omp_state_wait_barrier_implicit_parallel** permits a weaker contract between a
20 runtime and a tool that enables a simpler and faster implementation of parallel regions.

21 **omp_state_wait_barrier_implicit_workshare**

22 The description of when a thread reports a state associated with an implicit barrier is described for
23 state **omp_state_wait_barrier_implicit**. An OpenMP implementation may report
24 **omp_state_wait_barrier_implicit_parallel** for an implicit barrier that occurs at
25 the end of a worksharing construct.

26 **omp_state_wait_barrier_explicit**

27 The thread is waiting at an explicit barrier in a parallel region. A thread may enter this state early,
28 when the thread encounters a barrier, or late, when the thread begins to wait at the barrier. An
29 implementation may report **omp_state_wait_barrier** for explicit barriers.

30 4.4.1.1.3 Task Wait States

31 **omp_state_wait_taskwait**

32 The thread is waiting at a taskwait construct. A thread may enter this state early, when the thread
33 encounters a taskwait construct, or late, when the thread begins to wait for an uncompleted task.

1 **omp_state_wait_taskgroup**

2 The thread is waiting at the end of a taskgroup construct. A thread may enter this state early, when
3 the thread encounters the end of a taskgroup construct, or late, when the thread begins to wait for
4 an uncompleted task.

5 **4.4.1.1.4 Mutex Wait States**

6 OpenMP provides several mechanisms that enforce mutual exclusion: locks as well as critical,
7 atomic, and ordered sections. This grouping contains all states used to indicate that a thread is
8 awaiting exclusive access to a lock, critical section, variable, or ordered section.

9 An OpenMP implementation may report a thread waiting for any type of mutual exclusion using
10 either a state that precisely identifies the type of mutual exclusion, or a more generic state such as
11 **omp_state_wait_mutex** or **omp_state_wait_lock**. This flexibility may significantly
12 simplify the maintenance of states associated with mutual exclusion in the runtime when various
13 mechanisms for mutual exclusion rely on a common implementation, e.g., locks.

14 **omp_state_wait_mutex**

15 The thread is waiting for a mutex of an unspecified type. A thread may enter this state early, when
16 a thread encounters a lock acquisition or a region that requires mutual exclusion, or late, when the
17 thread begins to wait.

18 **omp_state_wait_lock**

19 The thread is waiting for a lock or nest lock. A thread may enter this state early, when a thread
20 encounters a lock **set** routine, or late, when the thread begins to wait for a lock.

21 **omp_state_wait_critical**

22 The thread is waiting to enter a critical region. A thread may enter this state early, when the thread
23 encounters a critical construct, or late, when the thread begins to wait to enter the critical region.

24 **omp_state_wait_atomic**

25 The thread is waiting to enter an atomic region. A thread may enter this state early, when the
26 thread encounters an atomic construct, or late, when the thread begins to wait to enter the atomic
27 region. An implementation may opt not to report this state when using atomic hardware
28 instructions that support non-blocking atomic implementations.

29 **omp_state_wait_ordered**

30 The thread is waiting to enter an ordered region. A thread may enter this state early, when the
31 thread encounters an ordered construct, or late, when the thread begins to wait to enter the ordered
32 region.

1 4.4.1.1.5 Target Wait States

2 `omp_state_wait_target`

3 The thread is waiting for a target region to complete.

4 `omp_state_wait_target_map`

5 The thread is waiting for a target data mapping operation to complete. An implementation may
6 report `omp_state_wait_target` for target data constructs.

7 `omp_state_wait_target_update`

8 The thread is waiting for a target update operation to complete. An implementation may report
9 `omp_state_wait_target` for target update constructs.

10 4.4.1.1.6 Miscellaneous States

11 `omp_state_idle`

12 The thread is idle, waiting for work.

13 `omp_state_overhead`

14 A thread may be reported as being in the overhead state at any point while executing within an
15 OpenMP runtime, except while waiting indefinitely at a synchronization point. An OpenMP
16 implementation report a thread's state as a work state for some or all of the time the thread spends
17 in executing in the OpenMP runtime.

18 `omp_state_undefined`

19 This state is reserved for threads that are not user threads, initial threads, threads currently in an
20 OpenMP team, or threads waiting to become part of an OpenMP team.

21 4.4.1.2 Frames

```
22 typedef struct omp_frame_t {  
23     void *exit_frame;  
24     void *enter_frame;  
25 } omp_frame_t;
```

C / C++

C / C++

Description

When executing an OpenMP program, at times, one or more procedure frames associated with the OpenMP runtime may appear on a thread's stack between frames associated with tasks. To help a tool determine whether a procedure frame on the call stack belongs to a task or not, for each task whose frames appear on the stack, the runtime maintains an `omp_frame_t` object that indicates a contiguous sequence of procedure frames associated with the task. Each `omp_frame_t` object is associated with the task to which the procedure frames belong. Each non-merged initial, implicit, explicit, or target task with one or more frames on a thread's stack will have an associated `omp_frame_t` object.

An `omp_frame_t` object associated with a task contains a pair of pointers: `exit_frame` and `enter_frame`. The field names were chosen, respectively, to reflect that they typically contain a pointer to a procedure frame on the stack when *exiting* the OpenMP runtime into code for a task or *entering* the OpenMP runtime from a task.

The `exit_frame` field of a task's `omp_frame_t` object contains the canonical frame address for the procedure frame that transfers control to the structured block for the task. The value of `exit_frame` is **NULL** until just prior to beginning execution of the structured block for the task. A task's `exit_frame` may point to a procedure frame that belongs to the OpenMP runtime or one that belongs to another task. The `exit_frame` for the `omp_frame_t` object associated with an *initial task* is **NULL**.

The `enter_frame` field of a task's `omp_frame_t` object contains the canonical frame address of a task procedure frame that invoked the OpenMP runtime causing the current task to suspend and another task to execute. If a task with frames on the stack has not suspended, the value of `enter_frame` for the `omp_frame_t` object associated with the task may contain **NULL**. The value of `enter_frame` in a task's `omp_frame_t` is reset to **NULL** just before a suspended task resumes execution.

An `omp_frame_t`'s lifetime begins when a task is created and ends when the task is destroyed. Tools should not assume that a frame structure remains at a constant location in memory throughout a task's lifetime. A pointer to a task's `omp_frame_t` object is passed to some callbacks; a pointer to a task's `omp_frame_t` object can also be retrieved by a tool at any time, including in a signal handler, by invoking the `ompt_get_task_info` runtime entry point (described in Section 4.2.5.1.14).

Table 4.8 describes various states in which an `omp_frame_t` object may be observed and their meaning. In the presence of nested parallelism, a tool may observe a sequence of `omp_frame_t` objects for a thread. Appendix B illustrates use of `omp_frame_t` objects with nested parallelism.

TABLE 4.8: Meaning of various states of an `omp_frame_t` object.

<i>exit_frame</i> / <i>enter_frame</i> state	<i>enter_frame</i> is NULL	<i>enter_frame</i> is non- NULL
<i>exit_frame</i> is NULL	case 1) initial task during execution case 2) task that is created but not yet scheduled or already finished	initial task suspended while another task executes
<i>exit_frame</i> is non- NULL	non-initial task that has been scheduled	non-initial task suspended while another task executes

Note – A monitoring tool using asynchronous sampling can observe values of *exit_frame* and *enter_frame* at inconvenient times. Tools must be prepared to observe and handle `omp_frame_t` objects observed just prior to when their field values will be set or cleared.

4.4.1.3 Wait Identifiers

Each thread instance maintains a *wait identifier* of type `omp_wait_id_t`. When a task executing on a thread is waiting for mutual exclusion, the thread's wait identifier indicates what the thread is awaiting. A wait identifier may represent a critical section *name*, a lock, a program variable accessed in an atomic region, or a synchronization object internal to an OpenMP implementation.

```
C / C++
typedef uint64_t omp_wait_id_t;
```

`omp_wait_id_none` is defined as an instance of type `omp_wait_id_t` with the value 0.

When a thread is not in a wait state, the value of the thread's wait identifier is undefined.

4.4.2 Global Symbols

Many of the interfaces between tools and an OpenMP implementation are invisible to users. This section describes a few global symbols used by OMPT and OMPD tools to coordinate with an OpenMP implementation.

1 4.4.2.1 `ompt_start_tool`

2 Summary

3 If a tool wants to use the OMPT interface provided by an OpenMP implementation, the tool must
4 implement the function `ompt_start_tool` to announce its interest.

5 Format

```
6 ompt_start_tool_result_t *ompt_start_tool(  
7     unsigned int omp_version,  
8     const char *runtime_version  
9 );
```

10 Description

11 For a tool to use the OMPT interface provided by an OpenMP implementation, the tool must define
12 a globally-visible implementation of the function `ompt_start_tool`.

13 A tool may indicate its intent to use the OMPT interface provided by an OpenMP implementation
14 by having `ompt_start_tool` return a non-`NULL` pointer to an
15 `ompt_start_tool_result_t` structure, which contains pointers to tool initialization and
16 finalization callbacks along with a tool data word that an OpenMP implementation must pass by
17 reference to these callbacks.

18 A tool may use its argument `omp_version` to determine whether it is compatible with the OMPT
19 interface provided by an OpenMP implementation.

20 If a tool implements `ompt_start_tool` but has no interest in using the OMPT interface in a
21 particular execution, `ompt_start_tool` should return `NULL`.

22 Description of Arguments

23 The argument `omp_version` is the value of the `_OPENMP` version macro associated with the
24 OpenMP API implementation. This value identifies the OpenMP API version supported by an
25 OpenMP implementation, which specifies the version of the OMPT interface that it supports.

26 The argument `runtime_version` is a version string that unambiguously identifies the OpenMP
27 implementation.

1 Constraints on Arguments

2 The argument *runtime_version* must be an immutable string that is defined for the lifetime of a
3 program execution.

4 Effect

5 If a tool returns a non-NULL pointer to an `ompt_start_tool_result_t` structure, an
6 OpenMP implementation will call the tool initializer specified by the *initialize* field in this structure
7 before beginning execution of any OpenMP construct or completing execution of any environment
8 routine invocation; the OpenMP implementation will call the tool finalizer specified by the *finalize*
9 field in this structure when the OpenMP implementation shuts down.

10 Cross References

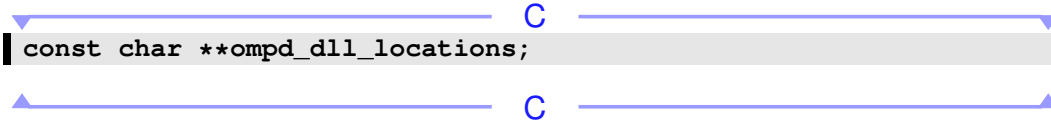
- 11 • `ompt_start_tool_result_t`, see Section 4.2.3.1 on page 429.

12 4.4.2.2 `ompd_dll_locations`

13 Summary

14 The global variable `ompd_dll_locations` indicates where a tool should look for OMPD
15 libraries that are compatible with the OpenMP implementation.

```
16 const char **ompd_dll_locations;
```



17 Description

18 `ompd_dll_locations` is an `argv`-style vector of filename strings that provide the names of
19 any OMPD libraries that are compatible with the OpenMP runtime. The vector is
20 NULL-terminated.

21 The programming model or architecture of the third-party tool, and hence that of the required
22 OMPD library, might not match that of the OpenMP program to be examined. On platforms that
23 support multiple programming models (*e.g.*, 32- v. 64-bit), or in heterogenous environments where
24 the architectures of the OpenMP program and third-party tool may be different, OpenMP
25 implementors are encouraged to provide OMPD libraries for all models. The vector, therefore, may
26 name libraries that are not compatible with the third-party tool. This is legal, and it is up to the
27 third-party tool to check that a library is compatible. (Typically, a tool might iterate over the vector
28 until a compatible library is found.)

1 **Restrictions**

2 **ompd_dll_locations** has external C linkage, no demangling or other transformations are
3 required by a third-party tool before looking up its address in the OpenMP program.

4 The vector and its members must be fully initialized before **ompd_dll_locations** is set to a
5 non-NULL value. That is, if **ompd_dll_locations** is not NULL, the vector and its contents
6 are valid.

7 **Cross References**

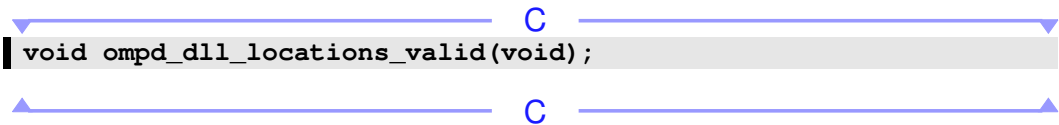
- 8 • **ompd_dll_locations_valid**, Section 4.4.2.3 on page 594
- 9 • Finding the OMPD library, Section 4.3.1.2 on page 519

10 **4.4.2.3 ompd_dll_locations_valid**

11 **Summary**

12 The OpenMP runtime notifies third-party tools that **ompd_dll_locations** is valid by allowing
13 execution to pass through a location identified by the symbol **ompd_dll_locations_valid**.

```
14       | void ompd_dll_locations_valid(void);
```



15 **Description**

16 Depending on how the OpenMP runtime is implemented, **ompd_dll_locations** might not be
17 a static variable, and therefore needs to be initialized at runtime. The OpenMP runtime notifies
18 third-party tools that **ompd_dll_locations** is valid by having execution pass through a
19 location identified by the symbol **ompd_dll_locations_valid**. If
20 **ompd_dll_locations** is NULL, a third-party tool, e.g., a debugger can place a breakpoint at
21 **ompd_dll_locations_valid** to be notified when **ompd_dll_locations** has been
22 initialized. In practice, the symbol **ompd_dll_locations_valid** need not be a function;
23 instead, it may be a labeled machine instruction through which execution passes once the vector is
24 valid.

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 2.4 on page 47). 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 following examples demonstrate how the OpenMP environment variables can be set in
12 different environments:

- 13 • csh-like shells:

```
14 | setenv OMP_SCHEDULE "dynamic"
```

- 15 • bash-like shells:

```
16 | export OMP_SCHEDULE="dynamic"
```

- 17 • Windows Command Line:

```
18 | set OMP_SCHEDULE=dynamic
```

1 5.1 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 2.12.2 on
10 page 100 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 343.

15 Examples:

```
16 setenv OMP_SCHEDULE "guided, 4"  
17 setenv OMP_SCHEDULE "dynamic"
```

18 Cross References

- 19 • *run-sched-var* ICV, see Section 2.4 on page 47.
- 20 • Worksharing-Loop construct, see Section 2.12.2 on page 100.
- 21 • Parallel worksharing-loop construct, see Section 2.16.1 on page 185.
- 22 • **omp_set_schedule** routine, see Section 3.2.12 on page 343.
- 23 • **omp_get_schedule** routine, see Section 3.2.13 on page 345.

1 5.2 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 2.4 on page 47 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 2.9.1 on page 77 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:

```
14 | setenv OMP_NUM_THREADS 4,3,2
```

15 Cross References

- 16 • *nthreads-var* ICV, see Section 2.4 on page 47.
- 17 • **num_threads** clause, Section 2.9 on page 72.
- 18 • **omp_set_num_threads** routine, see Section 3.2.1 on page 332.
- 19 • **omp_get_num_threads** routine, see Section 3.2.2 on page 333.
- 20 • **omp_get_max_threads** routine, see Section 3.2.3 on page 334.
- 21 • **omp_get_team_size** routine, see Section 3.2.19 on page 351.

1 5.3 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:

```
10 | setenv OMP_DYNAMIC true
```

11 Cross References

- 12 • *dyn-var* ICV, see Section 2.4 on page 47.
- 13 • `omp_set_dynamic` routine, see Section 3.2.7 on page 338.
- 14 • `omp_get_dynamic` routine, see Section 3.2.8 on page 339.

15 5.4 OMP_PROC_BIND

16 The `OMP_PROC_BIND` environment variable sets the initial value of the *bind-var* ICV. The value
17 of this environment variable is either `true`, `false`, or a comma separated list of `master`,
18 `close`, or `spread`. The values of the list set the thread affinity policy to be used for parallel
19 regions at the corresponding nested level.

20 If the environment variable is set to `false`, the execution environment may move OpenMP threads
21 between OpenMP places, thread affinity is disabled, and `proc_bind` clauses on `parallel`
22 constructs are ignored.

23 Otherwise, the execution environment should not move OpenMP threads between OpenMP places,
24 thread affinity is enabled, and the initial thread is bound to the first place in the OpenMP place list
25 prior to the first active parallel region.

26 The behavior of the program is implementation defined if the value in the `OMP_PROC_BIND`
27 environment variable is not `true`, `false`, or a comma separated list of `master`, `close`, or
28 `spread`. The behavior is also implementation defined if an initial thread cannot be bound to the
29 first place in the OpenMP place list.

1 Examples:

```
2 setenv OMP_PROC_BIND false  
3 setenv OMP_PROC_BIND "spread, spread, close"
```

4 Cross References

- 5 • *bind-var* ICV, see Section 2.4 on page 47.
- 6 • `proc_bind` clause, see Section 2.9.2 on page 79.
- 7 • `omp_get_proc_bind` routine, see Section 3.2.22 on page 354.

8 5.5 OMP_PLACES

9 A list of places can be specified in the **OMP_PLACES** environment variable. The
10 *place-partition-var* ICV obtains its initial value from the **OMP_PLACES** value, and makes the list
11 available to the execution environment. The value of **OMP_PLACES** can be one of two types of
12 values: either an abstract name describing a set of places or an explicit list of places described by
13 non-negative numbers.

14 The **OMP_PLACES** environment variable can be defined using an explicit ordered list of
15 comma-separated places. A place is defined by an unordered set of comma-separated non-negative
16 numbers enclosed by braces. The meaning of the numbers and how the numbering is done are
17 implementation defined. Generally, the numbers represent the smallest unit of execution exposed by
18 the execution environment, typically a hardware thread.

19 Intervals may also be used to define places. Intervals can be specified using the *<lower-bound>* :
20 *<length>* : *<stride>* notation to represent the following list of numbers: “*<lower-bound>*,
21 *<lower-bound>* + *<stride>*, ..., *<lower-bound>* + (*<length>* - 1)**<stride>*.” When *<stride>* is
22 omitted, a unit stride is assumed. Intervals can specify numbers within a place as well as sequences
23 of places.

24 An exclusion operator “!” can also be used to exclude the number or place immediately following
25 the operator.

26 Alternatively, the abstract names listed in Table 5.1 should be understood by the execution and
27 runtime environment. The precise definitions of the abstract names are implementation defined. An
28 implementation may also add abstract names as appropriate for the target platform.

29 The abstract name may be appended by a positive number in parentheses to denote the length of the
30 place list to be created, that is *abstract_name(num_places)*. When requesting fewer places than
31 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.

TABLE 5.1: Defined Abstract Names for **OMP_PLACES**

Abstract Name	Meaning
threads	Each place corresponds to a single hardware thread on the target machine.
cores	Each place corresponds to a single core (having one or more hardware threads) on the target machine.
sockets	Each place corresponds to a single socket (consisting of one or more cores) on the target machine.

3 The behavior of the program is implementation defined when the execution environment cannot
 4 map a numerical value (either explicitly defined or implicitly derived from an interval) within the
 5 **OMP_PLACES** list to a processor on the target platform, or if it maps to an unavailable processor.
 6 The behavior is also implementation defined when the **OMP_PLACES** environment variable is
 7 defined using an abstract name.

8 The following grammar describes the values accepted for the **OMP_PLACES** environment variable.

```

    <list>   ⊢ <p-list> | <aname>
    <p-list> ⊢ <p-interval> | <p-list>, <p-interval>
    <p-interval> ⊢ <place>:<len>:<stride> | <place>:<len> | <place> | !<place>
    <place>   ⊢ {<res-list>}
    <res-list> ⊢ <res-interval> | <res-list>, <res-interval>
    <res-interval> ⊢ <res>:<num-places>:<stride> | <res>:<num-places> | <res> | !<res>
    <aname>   ⊢ <word>(<num-places>) | <word>
    <word>   ⊢ sockets | cores | threads | <implementation-defined abstract name>
    <res>    ⊢ non-negative integer
    <num-places> ⊢ positive integer
    <stride>  ⊢ integer
    <len>    ⊢ positive integer
  
```


1 Examples:

```
2 setenv OMP_PLACES threads
3 setenv OMP_PLACES "threads (4) "
4 setenv OMP_PLACES
5     "{0,1,2,3},{4,5,6,7},{8,9,10,11},{12,13,14,15}"
6 setenv OMP_PLACES "{0:4},{4:4},{8:4},{12:4}"
7 setenv OMP_PLACES "{0:4}:4:4"
```

8 where each of the last three definitions corresponds to the same 4 places including the smallest
9 units of execution exposed by the execution environment numbered, in turn, 0 to 3, 4 to 7, 8 to 11,
10 and 12 to 15.

11 Cross References

- 12 • *place-partition-var*, Section 2.4 on page 47.
- 13 • Controlling OpenMP thread affinity, Section 2.9.2 on page 79.
- 14 • `omp_get_num_places` routine, see Section 3.2.23 on page 356.
- 15 • `omp_get_place_num_procs` routine, see Section 3.2.24 on page 357.
- 16 • `omp_get_place_proc_ids` routine, see Section 3.2.25 on page 358.
- 17 • `omp_get_place_num` routine, see Section 3.2.26 on page 359.
- 18 • `omp_get_partition_num_places` routine, see Section 3.2.27 on page 360.
- 19 • `omp_get_partition_place_nums` routine, see Section 3.2.28 on page 360.

20 5.6 OMP_NESTED

21 The deprecated `OMP_NESTED` environment variable controls nested parallelism by setting the
22 initial value of the *nest-var* ICV. The value of this environment variable must be **true** or **false**.
23 If the environment variable is set to **true**, nested parallelism is enabled; if set to **false**, nested
24 parallelism is disabled. The behavior of the program is implementation defined if the value of
25 `OMP_NESTED` is neither **true** nor **false**.

26 Example:

```
27 setenv OMP_NESTED false
```

Cross References

- *nest-var* ICV, see Section 2.4 on page 47.
- `omp_set_nested` routine, see Section 3.2.10 on page 341.
- `omp_get_team_size` routine, see Section 3.2.19 on page 351.

5.7 OMP_STACKSIZE

The `OMP_STACKSIZE` environment variable controls the size of the stack for threads created by the OpenMP implementation, by setting the value of the *stacksize-var* ICV. The environment variable does not control the size of the stack for an initial thread.

The value of this environment variable takes the form:

size | *size***B** | *size***K** | *size***M** | *size***G**

where:

- *size* is a positive integer that specifies the size of the stack for threads that are created by the OpenMP implementation.
- **B**, **K**, **M**, and **G** are letters that specify whether the given size is in Bytes, Kilobytes (1024 Bytes), Megabytes (1024 Kilobytes), or Gigabytes (1024 Megabytes), respectively. If one of these letters is present, there may be white space between *size* and the letter.

If only *size* is specified and none of **B**, **K**, **M**, or **G** is specified, then *size* is assumed to be in Kilobytes.

The behavior of the program is implementation defined if `OMP_STACKSIZE` does not conform to the above format, or if the implementation cannot provide a stack with the requested size.

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
```

Cross References

- *stacksize-var* ICV, see Section 2.4 on page 47.

1 5.8 OMP_WAIT_POLICY

2 The **OMP_WAIT_POLICY** environment variable provides a hint to an OpenMP implementation
3 about the desired behavior of waiting threads by setting the *wait-policy-var* ICV. A compliant
4 OpenMP implementation may or may not abide by the setting of the environment variable.

5 The value of this environment variable takes the form:

6 **ACTIVE | PASSIVE**

7 The **ACTIVE** value specifies that waiting threads should mostly be active, consuming processor
8 cycles, while waiting. An OpenMP implementation may, for example, make waiting threads spin.

9 The **PASSIVE** value specifies that waiting threads should mostly be passive, not consuming
10 processor cycles, while waiting. For example, an OpenMP implementation may make waiting
11 threads yield the processor to other threads or go to sleep.

12 The details of the **ACTIVE** and **PASSIVE** behaviors are implementation defined.

13 Examples:

```
14 setenv OMP_WAIT_POLICY ACTIVE  
15 setenv OMP_WAIT_POLICY active  
16 setenv OMP_WAIT_POLICY PASSIVE  
17 setenv OMP_WAIT_POLICY passive
```

18 Cross References

- 19
- *wait-policy-var* ICV, see Section 2.4 on page 47.

20 5.9 OMP_MAX_ACTIVE_LEVELS

21 The **OMP_MAX_ACTIVE_LEVELS** environment variable controls the maximum number of nested
22 active **parallel** regions by setting the initial value of the *max-active-levels-var* ICV.

23 The value of this environment variable must be a non-negative integer. The behavior of the
24 program is implementation defined if the requested value of **OMP_MAX_ACTIVE_LEVELS** is
25 greater than the maximum number of nested active parallel levels an implementation can support,
26 or if the value is not a non-negative integer.

Cross References

- *max-active-levels-var* ICV, see Section 2.4 on page 47.
- `omp_set_max_active_levels` routine, see Section 3.2.15 on page 347.
- `omp_get_max_active_levels` routine, see Section 3.2.16 on page 348.

5.10 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 2.4 on page 47.
- `omp_get_thread_limit` routine, see Section 3.2.14 on page 346.

5.11 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 2.4.1 on page 47.
- `cancel` construct, see Section 2.21.1 on page 256.
- `cancellation point` construct, see Section 2.21.2 on page 260.
- `omp_get_cancellation` routine, see Section 3.2.9 on page 340.

1 5.12 OMP_DISPLAY_ENV

2 The **OMP_DISPLAY_ENV** environment variable instructs the runtime to display the OpenMP
3 version number and the value of the ICVs associated with the environment variables described in
4 Chapter 5, as *name = value* pairs. The runtime displays this information once, after processing the
5 environment variables and before any user calls to change the ICV values by runtime routines
6 defined in Chapter 3.

7 The value of the **OMP_DISPLAY_ENV** environment variable may be set to one of these values:

8 **TRUE | FALSE | VERBOSE**

9 The **TRUE** value instructs the runtime to display the OpenMP version number defined by the
10 **_OPENMP** version macro (or the **openmp_version** Fortran parameter) value and the initial ICV
11 values for the environment variables listed in Chapter 5. The **VERBOSE** value indicates that the
12 runtime may also display the values of runtime variables that may be modified by vendor-specific
13 environment variables. The runtime does not display any information when the
14 **OMP_DISPLAY_ENV** environment variable is **FALSE** or undefined. For all values of the
15 environment variable other than **TRUE**, **FALSE**, and **VERBOSE**, the displayed information is
16 unspecified.

17 The display begins with "OPENMP DISPLAY ENVIRONMENT BEGIN", followed by the
18 **_OPENMP** version macro (or the **openmp_version** Fortran parameter) value and ICV values, in
19 the format *NAME '=' VALUE*. *NAME* corresponds to the macro or environment variable name,
20 optionally prepended by a bracketed *device-type*. *VALUE* corresponds to the value of the macro or
21 ICV associated with this environment variable. Values should be enclosed in single quotes. The
22 display is terminated with "OPENMP DISPLAY ENVIRONMENT END".

23 Example:

```
24 % setenv OMP_DISPLAY_ENV TRUE
```

25 The above example causes an OpenMP implementation to generate output of the following form:

```
26 OPENMP DISPLAY ENVIRONMENT BEGIN  
27   _OPENMP=' 201811'  
28   [host] OMP_SCHEDULE=' GUIDED, 4'  
29   [host] OMP_NUM_THREADS=' 4, 3, 2'  
30   [device] OMP_NUM_THREADS=' 2'  
31   [host,device] OMP_DYNAMIC=' TRUE'  
32   [host] OMP_PLACES=' {0:4}, {4:4}, {8:4}, {12:4}'  
33   ...  
34 OPENMP DISPLAY ENVIRONMENT END
```

1 5.13 OMP_DISPLAY_AFFINITY

2 The **OMP_DISPLAY_AFFINITY** environment variable instructs the runtime to display formatted
3 affinity information for all OpenMP threads in the parallel region upon entering the first parallel
4 region and when there is any change in the information accessible by the format specifiers listed in
5 table 5.2. If there is a change of affinity of any thread in a parallel region, thread affinity
6 information for all threads in that region will be displayed. There is no specific order in displaying
7 thread affinity information for all threads in the same parallel region.

8 The value of the **OMP_DISPLAY_AFFINITY** environment variable may be set to one of these
9 values:

10 **TRUE** | **FALSE**

11 The **TRUE** value instructs the runtime to display the OpenMP thread affinity information, and uses
12 the format setting defined in the *affinity-format-var* ICV.

13 The runtime does not display the OpenMP thread affinity information when the value of the
14 **OMP_DISPLAY_AFFINITY** environment variable is **FALSE** or undefined. For all values of the
15 environment variable other than **TRUE** or **FALSE**, the display action is implementation defined.

16 Example:

```
17 setenv OMP_DISPLAY_AFFINITY TRUE
```

18 The above example causes an OpenMP implementation to display OpenMP thread affinity
19 information during execution of the program, in a format given by the *affinity-format-var* ICV. The
20 following is a sample output:

```
21 thread_level= 1, thread_id= 0, thread_affinity= 0,1  
22 thread_level= 1, thread_id= 1, thread_affinity= 2,3
```

23 Cross References

- 24 • Controlling OpenMP thread affinity, see Section 2.9.2 on page 79.
- 25 • **omp_set_affinity_format** routine, see Section 3.2.29 on page 361.
- 26 • **omp_get_affinity_format** routine, see Section 3.2.30 on page 363.
- 27 • **omp_display_affinity** routine, see Section 3.2.31 on page 364.
- 28 • **omp_capture_affinity** routine, see Section 3.2.32 on page 365.
- 29 • **OMP_AFFINITY_FORMAT** environment variable, see Section 5.14 on page 607.

1 5.14 OMP_AFFINITY_FORMAT

2 The **OMP_AFFINITY_FORMAT** environment variable sets the initial value of the
3 *affinity-format-var* ICV which defines the format when displaying OpenMP thread affinity
4 information.

5 The value of this environment variable is a character string that may contain as substrings one or
6 more field specifiers, in addition to other characters. The format of each field specifier is

7 **%[[[0].] size] type**

8 where an individual field specifier must contain the percent symbol (%) and a type. The type can be
9 a single character short name or its corresponding long name delimited with curly braces, such as
10 %n or %{**thread_num**}. A literal percent is specified as %%. Field specifiers can be provided in
11 any order.

12 The 0 modifier indicates whether or not to add leading zeros to the output, following any indication
13 of sign or base. The . modifier indicates the output should be right justified when *size* is specified.
14 By default, output is left justified. The minimum field length is *size*, which is decimal digit string
15 with non-zero first digit. If no *size* is specified, the actual length needed to print the field will be
16 used. If the 0 modifier is used with *type* of **a**, **{thread_affinity}**, **h**, **{host}**, or a type that
17 is not printed as a number, the result is unspecified.

18 Any other characters in the format string that are not part of a field specifier will be included
19 literally in the output.

20 Available field types are:

TABLE 5.2: Available Field Types for Formatting OpenMP Thread Affinity Information

Short Name	Long Name	Meaning
L	thread_level	The value returned by omp_get_level() .
n	thread_num	The value returned by omp_get_thread_num() .
h	host	The name for the host machine on which the OpenMP program is running.
P	process_id	The process identifier used by the implementation.
T	thread_identifier	The thread identifier for a native thread defined by the implementation.

table continued on next page

table continued from previous page

Short Name	textbfLong Name	Meaning
N	num_threads	The value returned by <code>omp_get_num_threads()</code> .
A	ancestor_tnum	The value returned by <code>omp_get_ancestor_thread_num(level)</code> , where <i>level</i> is <code>omp_get_level()</code> minus 1.
a	thread_affinity	The list of numerical identifiers, in the format of a comma-separated list of integers or integer ranges, representing processors on which a thread may execute, subject to OpenMP thread affinity control and/or other external affinity mechanisms.

1 Implementations may define additional field types. If an implementation does not have information
2 for a field type, "undefined" is printed for this field when displaying the OpenMP thread affinity
3 information.

4 Example:

```
5 setenv OMP_AFFINITY_FORMAT  
6 "Thread Affinity: %0.3L %.8n %.15{thread_affinity} %.12h"
```

7 The above example causes an OpenMP implementation to display OpenMP thread affinity
8 information in the following form:

```
9 Thread Affinity: 001 0 0-1,16-17 nid003  
10 Thread Affinity: 001 1 2-3,18-19 nid003
```

11 Cross References

- 12 • Controlling OpenMP thread affinity, see Section 2.9.2 on page 79.
- 13 • `omp_set_affinity_format` routine, see Section 3.2.29 on page 361.
- 14 • `omp_get_affinity_format` routine, see Section 3.2.30 on page 363.
- 15 • `omp_display_affinity` routine, see Section 3.2.31 on page 364.
- 16 • `omp_capture_affinity` routine, see Section 3.2.32 on page 365.
- 17 • `OMP_DISPLAY_AFFINITY` environment variable, see Section 5.13 on page 606.

1 5.15 OMP_DEFAULT_DEVICE

2 The **OMP_DEFAULT_DEVICE** environment variable sets the device number to use in device
3 constructs by setting the initial value of the *default-device-var* ICV.

4 The value of this environment variable must be a non-negative integer value.

5 Cross References

- 6 • *default-device-var* ICV, see Section 2.4 on page 47.
- 7 • device constructs, Section 2.15 on page 158.

8 5.16 OMP_MAX_TASK_PRIORITY

9 The **OMP_MAX_TASK_PRIORITY** environment variable controls the use of task priorities by
10 setting the initial value of the *max-task-priority-var* ICV. The value of this environment variable
11 must be a non-negative integer.

12 Example:

```
13 | % setenv OMP_MAX_TASK_PRIORITY 20
```

14 Cross References

- 15 • *max-task-priority-var* ICV, see Section 2.4 on page 47.
- 16 • Tasking Constructs, see Section 2.13 on page 133.
- 17 • **omp_get_max_task_priority** routine, see Section 3.2.41 on page 374.

1 5.17 OMP_TARGET_OFFLOAD

2 The **OMP_TARGET_OFFLOAD** environment variable sets the initial value of the *target-offload-var*
3 ICV. The value of the **OMP_TARGET_OFFLOAD** environment variable may be set to one of these
4 values:

5 **MANDATORY | DISABLED | DEFAULT**

6 The **MANDATORY** value specifies that a device construct or a device memory routine must execute
7 on a target device. If the device construct cannot execute on its target device, or if a device memory
8 routine fails to execute, a warning is issued and the program execution aborts. Device constructs
9 are exempt from this behavior when an if-clause is present and the if-clause expression evaluates to
10 false.

11 The support of **DISABLED** is implementation defined. If an implementation supports it, the
12 behavior should be that a device construct must execute on the host. The behavior with this
13 environment value is equivalent to an if clause present on all device constructs, where each of these
14 if clause expressions evaluate to false. Device memory routines behave as if all device number
15 parameters are set to the value returned by `omp_get_initial_device()`. The
16 `omp_get_initial_device()` routine returns that no target device is available

17 The **DEFAULT** value specifies that when one or more target devices are available, the runtime
18 behaves as if this environment variable is set to **MANDATORY**; otherwise, the runtime behaves as if
19 this environment variable is set to **DISABLED**.

20 Example:

```
21 | % setenv OMP_TARGET_OFFLOAD MANDATORY
```

22 Cross References

- 23 • *target-offload-icv* ICV, see Section 2.4 on page 47.
- 24 • device constructs, Section 2.15 on page 158.

1 5.18 OMP_TOOL

2 The **OMP_TOOL** environment variable sets the *tool-var* ICV which controls whether an OpenMP
3 runtime will try to register a first party tool. The value of this environment variable must be
4 **enabled** or **disabled**. If **OMP_TOOL** is set to any value other than **enabled** or **disabled**,
5 the behavior is unspecified. If **OMP_TOOL** is not defined, the default value for *tool-var* is
6 **enabled**.

7 Example:

```
8 | % setenv OMP_TOOL enabled
```

9 Cross References

- 10 • *tool-var* ICV, see Section 2.4 on page 47.
- 11 • Tool Interface, see Section 4 on page 417.

12 5.19 OMP_TOOL_LIBRARIES

13 The **OMP_TOOL_LIBRARIES** environment variable sets the *tool-libraries-var* ICV to a list of tool
14 libraries that will be considered for use on a device where an OpenMP implementation is being
15 initialized. The value of this environment variable must be a colon-separated list of
16 dynamically-linked libraries, each specified by an absolute path.

17 If the *tool-var* ICV is not enabled, the value of *tool-libraries-var* will be ignored. Otherwise, if
18 **ompt_start_tool**, a global function symbol for a tool initializer, isn't visible in the address
19 space on a device where OpenMP is being initialized or if **ompt_start_tool** returns **NULL**, an
20 OpenMP implementation will consider libraries in the *tool-libraries-var* list in a left to right order.
21 The OpenMP implementation will search the list for a library that meets two criteria: it can be
22 dynamically loaded on the current device and it defines the symbol **ompt_start_tool**. If an
23 OpenMP implementation finds a suitable library, no further libraries in the list will be considered.

24 Cross References

- 25 • *tool-libraries-var* ICV, see Section 2.4 on page 47.
- 26 • Tool Interface, see Section 4 on page 417.
- 27 • **ompt_start_tool** routine, see Section 4.4.2.1 on page 592.

1 5.20 OMP_DEBUG

2 The **OMP_DEBUG** environment variable sets the *debug-var* ICV which controls whether an
3 OpenMP runtime will collect information that an OMPD library may need to support a tool. The
4 value of this environment variable must be **enabled** or **disabled**. If **OMP_DEBUG** is set to any
5 value other than **enabled** or **disabled**, the behavior is implementation defined.

6 Example:

```
7 | % setenv OMP_DEBUG enabled
```

8 Cross References

- 9 • *debug-var* ICV, see Section [2.4](#) on page [47](#).
- 10 • Tool Interface, see Section [4](#) on page [417](#).
- 11 • Enabling the Runtime for OMPD, see Section [4.3.1.1](#) on page [519](#).

12 5.21 OMP_ALLOCATOR

13 **OMP_ALLOCATOR** sets the *def-allocator-var* ICV that specifies the default allocator for allocation
14 calls, directives and clauses that do not specify an allocator. The value of this environment variable
15 is a predefined allocator from Table [2.9](#) on page [153](#). The value of this environment variable is not
16 case sensitive.

17 Cross References

- 18 • *def-allocator-var* ICV, see Section [2.4](#) on page [47](#).
- 19 • Memory allocators, see Section [2.14.2](#) on page [151](#).
- 20 • **omp_set_default_allocator** routine, see Section [3.7.4](#) on page [408](#).
- 21 • **omp_get_default_allocator** routine, see Section [3.7.5](#) on page [409](#).

2

OpenMP Implementation-Defined

3

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 • **Device address**: an address in a *device data environment* (see Section 1.2.6 on page 13).
- 10 • **Memory model**: the minimum size at which a memory update may also read and write back
11 adjacent variables that are part of another variable (as array or structure elements) is
12 implementation defined but is no larger than required by the base language (see Section 1.4.1 on
13 page 22).
- 14 • **Memory model**: Implementations are allowed to relax the ordering imposed by implicit flush
15 operations when the result is only visible to programs using non-sequentially consistent atomic
16 directives (see Section 1.4.6 on page 27).
- 17 • **Internal control variables**: the initial values of *dyn-var*, *nest-var*, *nthreads-var*, *run-sched-var*,
18 *def-sched-var*, *bind-var*, *stacksize-var*, *wait-policy-var*, *thread-limit-var*, *max-active-levels-var*,
19 *place-partition-var*, *affinity-format-var*, *default-device-var* and *def-allocator-var* are
20 implementation defined. The method for initializing a target device's internal control variable is
21 implementation defined (see Section 2.4.2 on page 49).
- 22 • **Dynamic adjustment of threads**: providing the ability to dynamically adjust the number of
23 threads is implementation defined . Implementations are allowed to deliver fewer threads (but at
24 least one) than indicated in Algorithm 2-1 even if dynamic adjustment is disabled (see
25 Section 2.9.1 on page 77).

- 1 • **Thread affinity:** For the **close** thread affinity policy, if $T > P$ and P does not divide T evenly,
2 the exact number of threads in a particular place is implementation defined. For the **spread**
3 thread affinity, if $T > P$ and P does not divide T evenly, the exact number of threads in a
4 particular subpartition is implementation defined. The determination of whether the affinity
5 request can be fulfilled is implementation defined. If not, the number of threads in the team and
6 their mapping to places become implementation defined (see Section 2.9.2 on page 79).
- 7 • **declare variant directive:** whether, for some specific OpenMP context, the prototype of
8 the variant should differ from that of the base function, and if so how it should differ, is
9 implementation defined.
- 10 • **Teams construct:** the assignment of initial threads to places and the values of the
11 *place-partition-var* and *default-device-var* ICVs for each initial thread (see Section 2.10 on
12 page 81) are implementation defined.
- 13 • **Worksharing-Loop directive:** the integer type (or kind, for Fortran) used to compute the
14 iteration count of a collapsed loop is implementation defined. The effect of the
15 **schedule(runtime)** clause when the *run-sched-var* ICV is set to **auto** is implementation
16 defined. The *simd_width* used when a **simd** schedule modifier is specified is implementation
17 defined (see Section 2.12.2 on page 100).
- 18 • **sections construct:** the method of scheduling the structured blocks among threads in the
19 team is implementation defined (see Section 2.11.1 on page 86).
- 20 • **single construct:** the method of choosing a thread to execute the structured block is
21 implementation defined (see Section 2.11.2 on page 89)
- 22 • **simd construct:** the integer type (or kind, for Fortran) used to compute the iteration count for
23 the collapsed loop is implementation defined. The number of iterations that are executed
24 concurrently at any given time is implementation defined. If the *alignment* parameter is not
25 specified in the **aligned** clause, the default alignments for the SIMD instructions are
26 implementation defined (see Section 2.12.3.1 on page 111).
- 27 • **declare simd directive:** if the parameter of the **simdlen** clause is not a constant positive
28 integer expression, the number of concurrent arguments for the function is implementation
29 defined. If the *alignment* parameter of the **aligned** clause is not specified, the default
30 alignments for SIMD instructions are implementation defined (see Section 2.23.1 on page 317).
- 31 • **taskloop construct:** The number of loop iterations assigned to a task created from a
32 **taskloop** construct is implementation defined, unless the **grainsize** or **num_tasks**
33 clauses are specified. The integer type (or kind, for Fortran) used to compute the iteration count
34 for the collapsed loop is implementation defined (see Section 2.13.2 on page 138).
- 35 • **is_device_ptr clause:** Support for pointers created outside of the OpenMP device data
36 management routines is implementation defined (see Section 2.15.5 on page 168).
- 37 • **target construct:** the effect of invoking a virtual member function of an object on a device
38 other than the device on which the object was constructed is implementation defined (see

1 Section 2.15.5 on page 168).

- 2 • **teams construct**: the number of teams that are created is implementation defined but less than
3 or equal to the value of the **num_teams** clause if specified. The maximum number of threads
4 participating in the contention group that each team initiates is implementation defined but less
5 than or equal to the value of the **thread_limit** clause if specified (see Section 2.10 on
6 page 81).
- 7 • **distribute construct**: the integer type (or kind, for Fortran) used to compute the iteration
8 count for the collapsed loop is implementation defined (see Section 2.12.4.1 on page 117).
- 9 • **distribute construct**: If no **dist_schedule** clause is specified then the schedule for the
10 **distribute** construct is implementation defined (see Section 2.12.4.1 on page 117).
- 11 • **critical construct**: the effect of using a **hint** clause is implementation defined (see
12 Section 2.20.1 on page 216 and Section 2.20.12 on page 253).
- 13 • **atomic construct**: a compliant implementation may enforce exclusive access between **atomic**
14 regions that update different storage locations. The circumstances under which this occurs are
15 implementation defined. If the storage location designated by x is not size-aligned (that is, if the
16 byte alignment of x is not a multiple of the size of x), then the behavior of the atomic region is
17 implementation defined (see Section 2.20.7 on page 227). The effect of using a **hint** clause is
18 implementation defined (see Section 2.20.7 on page 227 and Section 2.20.12 on page 253).

▼ Fortran ▲

- 19 • **Data-sharing attributes**: The data-sharing attributes of dummy arguments without the **VALUE**
20 attribute are implementation-defined if the associated actual argument is shared, except for the
21 conditions specified (see Section 2.22.1.2 on page 266).
- 22 • **threadprivate directive**: if the conditions for values of data in the threadprivate objects of
23 threads (other than an initial thread) to persist between two consecutive active parallel regions do
24 not all hold, the allocation status of an allocatable variable in the second region is
25 implementation defined (see Section 2.22.2 on page 268).
- 26 • **Runtime library definitions**: it is implementation defined whether the include file **omp_lib.h**
27 or the module **omp_lib** (or both) is provided. It is implementation defined whether any of the
28 OpenMP runtime library routines that take an argument are extended with a generic interface so
29 arguments of different **KIND** type can be accommodated (see Section 3.1 on page 330).

▲ Fortran ▼

- 30 • **omp_set_num_threads routine**: if the argument is not a positive integer the behavior is
31 implementation defined (see Section 3.2.1 on page 332).
- 32 • **omp_set_schedule routine**: for implementation specific schedule types, the values and
33 associated meanings of the second argument are implementation defined. (see Section 3.2.12 on
34 page 343).

- 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 347).
- 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 348).
- 10 • **omp_get_place_proc_ids routine:** the meaning of the nonnegative numerical identifiers
11 returned by the **omp_get_place_proc_ids** routine is implementation defined (see
12 Section 3.2.25 on page 358).
- 13 • **omp_set_affinity_format routine:** when called from within any explicit **parallel**
14 region the binding thread set (and binding region, if required) for the
15 **omp_set_affinity_format** region is implementation defined and the behavior is
16 implementation defined. If the argument does not conform to the specified format then the result
17 is implementation defined (see Section 3.2.29 on page 361).
- 18 • **omp_get_affinity_format routine:** when called from within any explicit **parallel**
19 region the binding thread set (and binding region, if required) for the
20 **omp_get_affinity_format** region is implementation defined (see Section 3.2.30 on
21 page 363).
- 22 • **omp_display_affinity routine:** if the argument does not conform to the specified format
23 then the result is implementation defined (see Section 3.2.31 on page 364).
- 24 • **omp_capture_affinity routine:** if the *format* argument does not conform to the specified
25 format then the result is implementation defined (see Section 3.2.32 on page 365).
- 26 • **omp_get_initial_device routine:** the value of the device number is implementation
27 defined (see Section 3.2.40 on page 373).
- 28 • **omp_init_lock_with_hint** and **omp_init_nest_lock_with_hint routines:** if
29 hints are stored with a lock variable, the effect of the hints on the locks are implementation
30 defined (see Section 3.3.2 on page 382).
- 31 • **omp_target_memcpy_rect routine:** the maximum number of dimensions supported is
32 implementation defined, but must be at least three (see Section 3.6.5 on page 398).
- 33 • **OMP_SCHEDULE environment variable:** if the value does not conform to the specified format
34 then the result is implementation defined (see Section 5.1 on page 596).
- 35 • **OMP_NUM_THREADS environment variable:** if any value of the list specified in the
36 **OMP_NUM_THREADS** environment variable leads to a number of threads that is greater than the
37 implementation can support, or if any value is not a positive integer, then the result is
38 implementation defined (see Section 5.2 on page 597).

- 1 ● **OMP_PROC_BIND environment variable:** if the value is not **true**, **false**, or a comma
2 separated list of **master**, **close**, or **spread**, the behavior is implementation defined. The
3 behavior is also implementation defined if an initial thread cannot be bound to the first place in
4 the OpenMP place list (see Section 5.4 on page 598).
- 5 ● **OMP_DYNAMIC environment variable:** if the value is neither **true** nor **false** the behavior is
6 implementation defined (see Section 5.3 on page 598).
- 7 ● **OMP_NESTED environment variable:** if the value is neither **true** nor **false** the behavior is
8 implementation defined (see Section 5.6 on page 601).
- 9 ● **OMP_STACKSIZE environment variable:** if the value does not conform to the specified format
10 or the implementation cannot provide a stack of the specified size then the behavior is
11 implementation defined (see Section 5.7 on page 602).
- 12 ● **OMP_WAIT_POLICY environment variable:** the details of the **ACTIVE** and **PASSIVE**
13 behaviors are implementation defined (see Section 5.8 on page 603).
- 14 ● **OMP_MAX_ACTIVE_LEVELS environment variable:** if the value is not a non-negative integer
15 or is greater than the number of parallel levels an implementation can support then the behavior
16 is implementation defined (see Section 5.9 on page 603).
- 17 ● **OMP_THREAD_LIMIT environment variable:** if the requested value is greater than the number
18 of threads an implementation can support, or if the value is not a positive integer, the behavior of
19 the program is implementation defined (see Section 5.10 on page 604).
- 20 ● **OMP_PLACES environment variable:** the meaning of the numbers specified in the environment
21 variable and how the numbering is done are implementation defined. The precise definitions of
22 the abstract names are implementation defined. An implementation may add
23 implementation-defined abstract names as appropriate for the target platform. When creating a
24 place list of n elements by appending the number n to an abstract name, the determination of
25 which resources to include in the place list is implementation defined. When requesting more
26 resources than available, the length of the place list is also implementation defined. The behavior
27 of the program is implementation defined when the execution environment cannot map a
28 numerical value (either explicitly defined or implicitly derived from an interval) within the
29 **OMP_PLACES** list to a processor on the target platform, or if it maps to an unavailable processor.
30 The behavior is also implementation defined when the **OMP_PLACES** environment variable is
31 defined using an abstract name (see Section 5.5 on page 599).
- 32 ● **OMP_AFFINITY_FORMAT environment variable:** if the value does not conform to the
33 specified format then the result is implementation defined (see Section 5.14 on page 607).
- 34 ● **OMPT thread states:** The set of OMPT thread states supported is implementation defined (see
35 Section 4.4.1.1 on page 584).
- 36 ● **ompt_callback_idle tool callback:** if a tool attempts to register a callback with this string
37 name using the runtime entry point **ompt_set_callback**, it is implementation defined

1 whether the registered callback may never or sometimes invoke this callback for the associated
2 events (see Table 4.2 on page 424)

3 • **ompt_callback_sync_region_wait tool callback:** if a tool attempts to register a
4 callback with this string name using the runtime entry point **ompt_set_callback**, it is
5 implementation defined whether the registered callback may never or sometimes invoke this
6 callback for the associated events (see Table 4.2 on page 424)

7 • **ompt_callback_mutex_released tool callback:** if a tool attempts to register a callback
8 with this string name using the runtime entry point **ompt_set_callback**, it is
9 implementation defined whether the registered callback may never or sometimes invoke this
10 callback for the associated events (see Table 4.2 on page 424)

11 • **ompt_callback_task_dependences tool callback:** if a tool attempts to register a
12 callback with this string name using the runtime entry point **ompt_set_callback**, it is
13 implementation defined whether the registered callback may never or sometimes invoke this
14 callback for the associated events (see Table 4.2 on page 424)

15 • **ompt_callback_task_dependence tool callback:** if a tool attempts to register a
16 callback with this string name using the runtime entry point **ompt_set_callback**, it is
17 implementation defined whether the registered callback may never or sometimes invoke this
18 callback for the associated events (see Table 4.2 on page 424)

19 • **ompt_callback_work tool callback:** if a tool attempts to register a callback with this string
20 name using the runtime entry point **ompt_set_callback**, it is implementation defined
21 whether the registered callback may never or sometimes invoke this callback for the associated
22 events (see Table 4.2 on page 424)

23 • **ompt_callback_master tool callback:** if a tool attempts to register a callback with this
24 string name using the runtime entry point **ompt_set_callback**, it is implementation defined
25 whether the registered callback may never or sometimes invoke this callback for the associated
26 events (see Table 4.2 on page 424)

27 • **ompt_callback_target_map tool callback:** if a tool attempts to register a callback with
28 this string name using the runtime entry point **ompt_set_callback**, it is implementation
29 defined whether the registered callback may never or sometimes invoke this callback for the
30 associated events (see Table 4.2 on page 424)

31 • **ompt_callback_sync_region tool callback:** if a tool attempts to register a callback with
32 this string name using the runtime entry point **ompt_set_callback**, it is implementation
33 defined whether the registered callback may never or sometimes invoke this callback for the
34 associated events (see Table 4.2 on page 424)

35 • **ompt_callback_lock_init tool callback:** if a tool attempts to register a callback with
36 this string name using the runtime entry point **ompt_set_callback**, it is implementation
37 defined whether the registered callback may never or sometimes invoke this callback for the
38 associated events (see Table 4.2 on page 424)

- 1 ● **ompt_callback_lock_destroy tool callback**: if a tool attempts to register a callback
2 with this string name using the runtime entry point **ompt_set_callback**, it is
3 implementation defined whether the registered callback may never or sometimes invoke this
4 callback for the associated events (see Table 4.2 on page 424)
- 5 ● **ompt_callback_mutex_acquire tool callback**: if a tool attempts to register a callback
6 with this string name using the runtime entry point **ompt_set_callback**, it is
7 implementation defined whether the registered callback may never or sometimes invoke this
8 callback for the associated events (see Table 4.2 on page 424)
- 9 ● **ompt_callback_mutex_acquired tool callback**: if a tool attempts to register a callback
10 with this string name using the runtime entry point **ompt_set_callback**, it is
11 implementation defined whether the registered callback may never or sometimes invoke this
12 callback for the associated events (see Table 4.2 on page 424)
- 13 ● **ompt_callback_nest_lock tool callback**: if a tool attempts to register a callback with
14 this string name using the runtime entry point **ompt_set_callback**, it is implementation
15 defined whether the registered callback may never or sometimes invoke this callback for the
16 associated events (see Table 4.2 on page 424)
- 17 ● **ompt_callback_flush tool callback**: if a tool attempts to register a callback with this
18 string name using the runtime entry point **ompt_set_callback**, it is implementation defined
19 whether the registered callback may never or sometimes invoke this callback for the associated
20 events (see Table 4.2 on page 424)
- 21 ● **ompt_callback_cancel tool callback**: if a tool attempts to register a callback with this
22 string name using the runtime entry point **ompt_set_callback**, it is implementation defined
23 whether the registered callback may never or sometimes invoke this callback for the associated
24 events (see Table 4.2 on page 424)
- 25 ● **ompt_callback_dispatch tool callback**: if a tool attempts to register a callback with this
26 string name using the runtime entry point **ompt_set_callback**, it is implementation defined
27 whether the registered callback may never or sometimes invoke this callback for the associated
28 events (see Table 4.2 on page 424)
- 29 ● **OMP_DEBUG environment variable**: if the value is neither **disabled** nor **enabled** the
30 behavior is implementation defined (see Section 5.20 on page 612).
- 31 ● **Device tracing**: Whether a target device supports tracing or not is implementation defined; if a
32 target device does not support tracing, a **NULL** may be supplied for the *lookup* function to a
33 tool's device initializer (see Section 4.2.1.4 on page 425).
- 34 ● **ompt_set_trace_ompt runtime entry point**: it is implementation defined whether a
35 device-specific tracing interface will define this runtime entry point, indicating that it can collect
36 traces in OMPT format (see Section 4.2.1.4 on page 425).
- 37 ● **ompt_buffer_get_record_ompt runtime entry point**: it is implementation defined
38 whether a device-specific tracing interface will define this runtime entry point, indicating that it

- 1 can collect traces in OMPT format (see Section [4.2.1.4](#) on page [425](#)).
- 2
- 3
- 4 • **Memory allocators:** The storage resource that will be used by each memory allocator defined in
5 Table [2.9](#) on page [153](#) is implementation defined.
 - 6 • **allocate directive:** The effect of not being able to fulfill an allocation request specified in
7 **allocate** directive is implementation defined.
 - **allocate clause:** The effect of not being able to fulfill an allocation request specified in
allocate clause is implementation defined.

1 APPENDIX B

2 Task Frame Management for the
3 Tool Interface

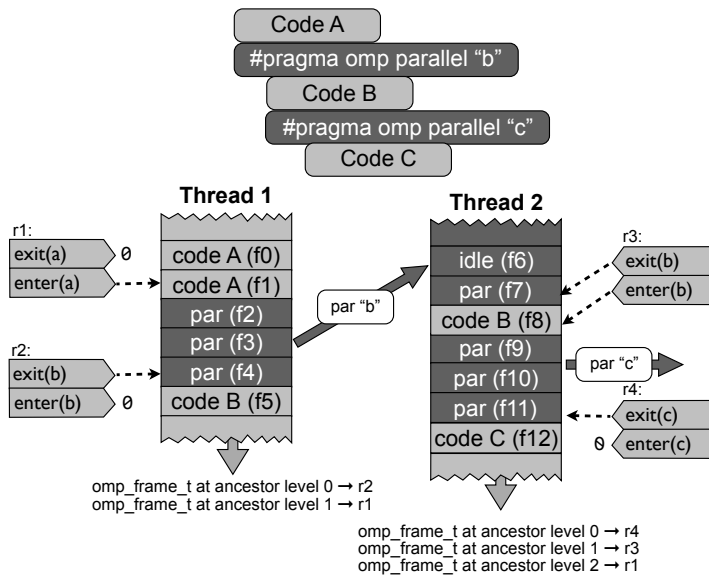


FIGURE B.1: Thread call stacks implementing nested parallelism annotated with frame information for the OMPT tool interface.

4 The top half of Figure B.1 illustrates a conceptualization of a program executing a nested parallel
5 region, where code A, B, and C represent, respectively, one or more procedure frames of code
6 associated with an initial task, an outer parallel region, and an inner parallel region. The bottom
7 half of Figure B.1 illustrates the stacks of two threads executing the nested parallel region. In the
8 illustration, stacks grow downward—a call to a function adds a new frame to the stack below the

1 frame of its caller. When thread 1 encounters the outer-parallel region “b”, it calls a routine in the
2 OpenMP runtime to create a new parallel region. The OpenMP runtime sets the *enter_frame* field
3 in the **omp_frame_t** for the initial task executing code A to the canonical frame address of frame
4 f1—the user frame in the initial task that calls the runtime. The **omp_frame_t** for the initial task
5 is labeled *r1* in Figure B.1. In this figure, three consecutive runtime system frames, labeled “par”
6 with frame identifiers f2–f4, are on the stack. Before starting the implicit task for parallel region
7 “b” in thread 1, the runtime sets the *exit_frame* in the implicit task’s **omp_frame_t** (labeled *r2*) to
8 the canonical frame address of frame f4. Execution of application code for parallel region “b”
9 begins on thread 1 when the runtime system invokes application code B (frame f5) from frame f4.

10 Let us focus now on thread 2, an OpenMP thread. Figure B.1 shows this worker executing work for
11 the outer-parallel region “b.” On the OpenMP thread’s stack is a runtime frame labeled “idle,”
12 where the OpenMP thread waits for work. When work becomes available, the runtime system
13 invokes a function to dispatch it. While dispatching parallel work might involve a chain of several
14 calls, here we assume that the length of this chain is 1 (frame f7). Before thread 2 exits the runtime
15 to execute an implicit task for parallel region “b,” the runtime sets the *exit_frame* field of the
16 implicit task’s **omp_frame_t** (labeled *r3*) to the canonical frame address of frame f7. When
17 thread 2 later encounters the inner-parallel region “c,” as execution returns to the runtime, the
18 runtime fills in the *enter_frame* field of the current task’s **omp_frame_t** (labeled *r3*) to the
19 canonical frame address of frame f8—the frame that invoked the runtime. Before the task for
20 parallel region “c” is invoked on thread 2, the runtime system sets the *exit_frame* field of the
21 **omp_frame_t** (labeled *r4*) for the implicit task for “c” to the canonical frame address of frame
22 f11. Execution of application code for parallel region “c” begins on thread 2 when the runtime
23 system invokes application code C (frame f12) from frame f11.

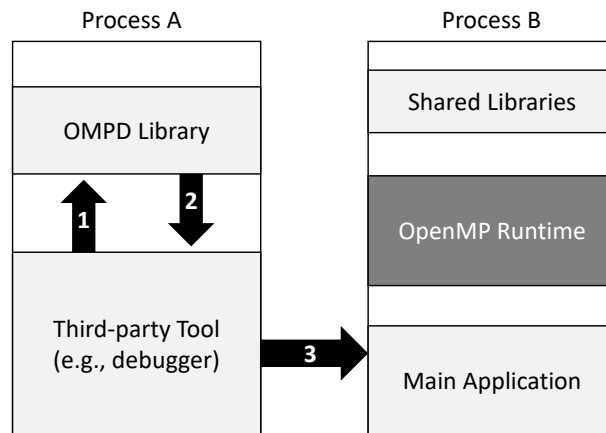
24 Below the stack for each thread in Figure B.1, the figure shows the **omp_frame_t** information
25 obtained by calls to **ompt_get_task_info** made on each thread for the stack state shown. We
26 show the ID of the **omp_frame_t** object returned at each ancestor level. Note that thread 2 has
27 task frame information for three levels of tasks, whereas thread 1 has only two.

28 Cross References

- 29 • **omp_frame_t**, see Section 4.4.1.2 on page 589.
- 30 • **ompt_get_task_info_t**, see Section 4.2.5.1.14 on page 496.

1 APPENDIX C

2 Interaction Diagram of OMPD Components



- 1 Tool makes requests to the OMPD library via an API
- 2 The OMPD library makes requests back to get information of the OpenMP program and runtime via callbacks to the tool
- 3 The tool has control over the OpenMP program so it replies to the callbacks (e.g., can lookup symbols, read/write data)

FIGURE C.1: Interaction Diagram of OMPD Components

3 The figure shows how the different components of OMPD fit together. The third-party tool loads
4 the OMPD library that matches the OpenMP runtime being used by the OpenMP program. The
5 library exports the API defined in Section 4.3 on page 518, which the tool uses to get OpenMP
6 information about the OpenMP program. The OMPD library will need to look up the symbols, or

1 read data out of the OpenMP program. It does not do this directly, but instead asks the tool to
2 perform these operations for it using a callback interface exported by the tool.

3 This architectural layout insulates the tool from the details of the internal structure of the OpenMP
4 runtime. Similarly, the OMPD library does not need to be concerned about how to access the
5 OpenMP program. Decoupling the library and tool in this way allows for flexibility in how the
6 OpenMP program and tool are deployed, so that, for example, there is no requirement that tool and
7 OpenMP program execute on the same machine.

8 **Cross References**

- 9 • See Section [4.3](#) on page [518](#).

1 APPENDIX D

2 Features History

3 This appendix summarizes the major changes between recent versions of the OpenMP API since
4 version 2.5.

5 D.1 Deprecated Features

6 The following features have been deprecated:

- 7 • the *nest-var* ICV;
- 8 • the **OMP_NESTED** environment variable;
- 9 • the **omp_set_nested** and **omp_get_nested** routines;
- 10 • the C/C++ type **omp_lock_hint_t** and the corresponding Fortran kind
11 **omp_hint_hint_kind**;
- 12 • and the lock hint constants **omp_lock_hint_none**, **omp_lock_hint_uncontended**,
13 **omp_lock_hint_contended**, **omp_lock_hint_nonspeculative**, and
14 **omp_lock_hint_speculative**.

1 D.2 Version 4.5 to 5.0 Differences

- 2 • Stubs for Runtime Library Routines (previously Appendix A) were moved to a separate document.
- 3 • Interface Declarations (previously Appendix B) were moved to a separate document.
- 4 • The memory model was extended to distinguish different types of flush operations according to
5 specified flush properties (see Section 1.4.4 on page 24) and to define a happens before order
6 based on synchronizing flush operations (see Section 1.4.5 on page 26).
- 7 • Various changes throughout the specification were made to provide initial support of C11,
8 C++11, C++14, C++17 and Fortran 2008 (see Section 1.7 on page 30).
- 9 • Support for several features of Fortran 2003 was added (see Section 1.7 on page 30 for features
10 that are still not supported).
- 11 • The list items allowable in a **depend** clause on a task generating construct was extended,
12 including for C/C++ allowing any *lvalue* expression (see Section 2.1 on page 36 and
13 Section 2.20.11 on page 248).
- 14 • The **requires** directive (see Section 2.3 on page 43) was added to support applications that
15 require implementation-specific features.
- 16 • The *target-offload-var* internal control variable (see Section 2.4 on page 47) and the
17 **OMP_TARGET_OFFLOAD** environment variable (see Section 5.17 on page 610) were added to
18 support runtime control of the execution of device constructs.
- 19 • The default value of the *nest-var* ICV was changed from *false* to implementation defined (see
20 Section 2.4.2 on page 49). The *nest-var* ICV (see Section 2.4.1 on page 47), the **OMP_NESTED**
21 environment variable (see Section 5.6 on page 601), and the **omp_set_nested** and
22 **omp_get_nested** routines were deprecated (see Section 3.2.10 on page 341 and
23 Section 3.2.11 on page 342).
- 24 • Support for array shaping (see Section 2.5 on page 58) and for array sections with non-unit strides
25 in C and C++ (see Section 2.6 on page 59) were added to facilitate specification of discontinuous
26 storage and the **target update** construct (see Section 2.15.6 on page 174) and the **depend**
27 clause (see Section 2.20.11 on page 248) were extended to allow the use of shape-operators (see
28 Section 2.5 on page 58); further, the **target update** construct (see Section 2.15.6 on
29 page 174) was modified to allow array sections that specify discontinuous storage.
- 30 • Iterators (see Section 2.7 on page 61) were added to express that an expression in a list may
31 expand to multiple expressions.
- 32 • The **teams** construct (see Section 2.10 on page 81) was extended to support execution on the
33 host device without surrounding **target** construct (see Section 2.15.5 on page 168).
- 34 • The **declare variant** directive (see Section 2.8.4 on page 67) and the metadirective
35 meta-directive (see Section 2.8.5 on page 69) were added to support selection of declared

- 1 function variants at a callsite or directive variants, respectively, based on compile-time traits for
2 the enclosing context.
- 3 • The canonical loop form was defined for Fortran and, for all base languages, extended to permit
4 non-rectangular loop nests (see Section 2.12.1 on page 95).
 - 5 • The *relational-op* in the canonical loop form for C/C++ was extended to include != (see
6 Section 2.12.1 on page 95).
 - 7 • The collapse of associated loops that are imperfectly nested loops was defined for the
8 worksharing-loop (see Section 2.12.2 on page 100), **simd** (see Section 2.12.3.1 on page 111),
9 **taskloop** (see Section 2.13.2 on page 138) and **distribute** (see Section 2.12.4.2 on
10 page 121) constructs.
 - 11 • SIMD constructs (see Section 2.12.3 on page 111) were extended to allow the use of **atomic**
12 constructs within them.
 - 13 • The **if** and **nontemporal** clauses were added to the **simd** construct (see Section 2.12.3.1 on
14 page 111).
 - 15 • The **scan** directive (see Section 2.12.6 on page 129) and the **inscan** modifier for the
16 **reduction** clause (see Section 2.22.5.4 on page 297) were added to support inclusive and
17 exclusive scan computations.
 - 18 • The **loop** construct was added to support compiler optimization and parallelization of loops for
19 which iterations may execute in any order, including concurrently (see Section 2.12.5 on
20 page 126).
 - 21 • To support task reductions, the **task** (see Section 2.13.1 on page 133) and **target** (see
22 Section 2.15.5 on page 168) constructs were extended to accept the the **in_reduction** clause
23 (see Section 2.22.5.6 on page 300), the **taskgroup** construct (see Section 2.20.6 on page 225)
24 was extended to accept the **task_reduction** clause Section 2.22.5.5 on page 299), and the
25 **task** modifier was added to the **reduction** clause (see Section 2.22.5.4 on page 297).
 - 26 • The **affinity** clause was added to the **task** construct (see Section 2.13.1 on page 133) to
27 support hints that indicate data affinity of explicit tasks.
 - 28 • To support taskloop reductions, the **taskloop** (see Section 2.13.2 on page 138) and
29 **taskloop simd** (see Section 2.13.3 on page 144) constructs were extended to accept the
30 **reduction** (see Section 2.22.5.4 on page 297) and **in_reduction** (see Section 2.22.5.6 on
31 page 300) clauses.
 - 32 • To support mutually exclusive inout sets, a **mutexinoutset** *dependence-type* was added to
33 the **depend** clause (see Section 2.13.6 on page 147 and Section 2.20.11 on page 248).
 - 34 • Predefined memory spaces (see Section 2.14.1 on page 150), predefined memory allocators and
35 allocator traits (see Section 2.14.2 on page 151) and directives, clauses (see Section 2.14 on
36 page 150 and API routines (see Section 3.7 on page 403) to use them were added to support
37 different kinds of memories.

- 1 • To support reverse offload, the **ancestor** modifier was added to the **device** clause for
2 **target** constructs (see Section 2.15.5 on page 168).
- 3 • To reduce programmer effort implicit declare target directives for some functions (C, C++,
4 Fortran) and subroutines (Fortran) were added (see Section 2.15.5 on page 168 and
5 Section 2.15.7 on page 178).
- 6 • Support for nested **declare target** directives was added (see Section 2.15.7 on page 178).
- 7 • The **implements** clause was added to the **declare target** directive to support the use of
8 device-specific function implementations (see Section 2.15.7 on page 178).
- 9 • The **declare mapper** directive was added to support mapping of complicated data types (see
10 Section 2.15.8 on page 183).
- 11 • The **depend** clause was added to the **taskwait** construct (see Section 2.20.5 on page 223).
- 12 • To support acquire and release semantics with weak memory ordering, the **acq_rel**,
13 **acquire**, and **release** clauses were added to the **atomic** construct (see Section 2.20.7 on
14 page 227) and **flush** construct (see Section 2.20.8 on page 235).
- 15 • The **atomic** construct was extended with the **hint** clause (see Section 2.20.7 on page 227).
- 16 • The **depend** clause (see Section 2.20.11 on page 248) was extended to support iterators and to
17 support depend objects that can be created with the new **depobj** construct.
- 18 • Lock hints were renamed to synchronization hints, and the old names were deprecated (see
19 Section 2.20.12 on page 253).
- 20 • To support conditional assignment to lastprivate variables, the **conditional** modifier was
21 added to the **lastprivate** clause (see Section 2.22.4.5 on page 283).
- 22 • The description of the **map** clause was modified to clarify how structure members are mapped.
23 (see Section 2.22.7.1 on page 307).
- 24 • The capability to map pointer variables (C/C++) and assign the address of device memory that is
25 mapped by an array section to them was added (see Section 2.22.7.1 on page 307).
- 26 • The **defaultmap** clause (see Section 2.22.7.2 on page 315) was extended to allow selecting
27 the data-mapping or data-sharing attributes for any of the scalar, aggregate, pointer or allocatable
28 classes on a per-region basis. Additionally it accepts the **none** parameter to support the
29 requirement that all variables referenced in the construct must be explicitly mapped or privatized.
- 30 • Runtime routines (see Section 3.2.29 on page 361, Section 3.2.30 on page 363, Section 3.2.31 on
31 page 364, and Section 3.2.32 on page 365) and environment variables (see Section 5.13 on
32 page 606 and Section 5.14 on page 607) were added to provide OpenMP thread affinity
33 information.
- 34 • The **omp_get_device_num** runtime routine (see Section 3.2.36 on page 369) was added to
35 support determination of the device on which a thread is executing.

- 1 • Support for a first-party tool interface (see Section 4.2 on page 418) was added.
- 2 • Support for a third-party tool interface (see Section 4.3 on page 518) was added.
- 3 • Support for controlling offloading behavior with the **OMP_TARGET_OFFLOAD** environment
- 4 variable was added (see Section 5.17 on page 610).

5 D.3 Version 4.0 to 4.5 Differences

- 6 • Support for several features of Fortran 2003 was added (see Section 1.7 on page 30 for features
- 7 that are still not supported).
- 8 • A parameter was added to the **ordered** clause of the loop construct (see Section 2.12.2 on
- 9 page 100) and clauses were added to the **ordered** construct (see Section 2.20.9 on page 243) to
- 10 support doacross loop nests and use of the **simd** construct on loops with loop-carried backward
- 11 dependences.
- 12 • The **linear** clause was added to the loop construct (see Section 2.12.2 on page 100).
- 13 • The **simdlen** clause was added to the **simd** construct (see Section 2.12.3.1 on page 111) to
- 14 support specification of the exact number of iterations desired per SIMD chunk.
- 15 • The **priority** clause was added to the **task** construct (see Section 2.13.1 on page 133) to
- 16 support hints that specify the relative execution priority of explicit tasks. The
- 17 **omp_get_max_task_priority** routine was added to return the maximum supported
- 18 priority value (see Section 3.2.41 on page 374) and the **OMP_MAX_TASK_PRIORITY**
- 19 environment variable was added to control the maximum priority value allowed (see
- 20 Section 5.16 on page 609).
- 21 • Taskloop constructs (see Section 2.13.2 on page 138 and Section 2.13.3 on page 144) were added
- 22 to support nestable parallel loops that create OpenMP tasks.
- 23 • To support interaction with native device implementations, the **use_device_ptr** clause was
- 24 added to the **target data** construct (see Section 2.15.2 on page 159) and the
- 25 **is_device_ptr** clause was added to the **target** construct (see Section 2.15.5 on page 168).
- 26 • The **nowait** and **depend** clauses were added to the **target** construct (see Section 2.15.5 on
- 27 page 168) to improve support for asynchronous execution of **target** regions.
- 28 • The **private**, **firstprivate** and **defaultmap** clauses were added to the **target**
- 29 construct (see Section 2.15.5 on page 168).
- 30 • The **declare target** directive was extended to allow mapping of global variables to be
- 31 deferred to specific device executions and to allow an *extended-list* to be specified in C/C++ (see
- 32 Section 2.15.7 on page 178).

- 1 • To support unstructured data mapping for devices, the **target enter data** (see
2 Section 2.15.3 on page 162) and **target exit data** (see Section 2.15.4 on page 165)
3 constructs were added and the **map** clause (see Section 2.22.7.1 on page 307) was updated.
- 4 • To support a more complete set of device construct shortcuts, the **target parallel** (see
5 Section 2.16.6 on page 192), target parallel loop (see Section 2.16.7 on page 193), target parallel
6 worksharing-loop SIMD (see Section 2.16.8 on page 195), and **target simd** (see
7 Section 2.16.9 on page 196), combined constructs were added.
- 8 • The **if** clause was extended to take a *directive-name-modifier* that allows it to apply to combined
9 constructs (see Section 2.18 on page 213).
- 10 • The **hint** clause was added to the **critical** construct (see Section 2.20.1 on page 216).
- 11 • The **source** and **sink** dependence types were added to the **depend** clause (see
12 Section 2.20.11 on page 248) to support doacross loop nests.
- 13 • The implicit data-sharing attribute for scalar variables in **target** regions was changed to
14 **firstprivate** (see Section 2.22.1.1 on page 263).
- 15 • Use of some C++ reference types was allowed in some data sharing attribute clauses (see
16 Section 2.22.4 on page 276).
- 17 • Semantics for reductions on C/C++ array sections were added and restrictions on the use of
18 arrays and pointers in reductions were removed (see Section 2.22.5.4 on page 297).
- 19 • The **ref**, **val**, and **uval** modifiers were added to the **linear** clause (see Section 2.22.4.6 on
20 page 286).
- 21 • Support was added to the map clauses to handle structure elements (see Section 2.22.7.1 on
22 page 307).
- 23 • Query functions for OpenMP thread affinity were added (see Section 3.2.23 on page 356 to
24 Section 3.2.28 on page 360).
- 25 • The lock API was extended with lock routines that support storing a hint with a lock to select a
26 desired lock implementation for a lock’s intended usage by the application code (see
27 Section 3.3.2 on page 382).
- 28 • Device memory routines were added to allow explicit allocation, deallocation, memory transfers
29 and memory associations (see Section 3.6 on page 393).
- 30 • C/C++ Grammar (previously Appendix B) was moved to a separate document.

1 D.4 Version 3.1 to 4.0 Differences

- 2 • Various changes throughout the specification were made to provide initial support of Fortran
3 2003 (see Section 1.7 on page 30).
- 4 • C/C++ array syntax was extended to support array sections (see Section 2.6 on page 59).
- 5 • The **proc_bind** clause (see Section 2.9.2 on page 79), the **OMP_PLACES** environment
6 variable (see Section 5.5 on page 599), and the **omp_get_proc_bind** runtime routine (see
7 Section 3.2.22 on page 354) were added to support thread affinity policies.
- 8 • SIMD constructs were added to support SIMD parallelism (see Section 2.12.3 on page 111).
- 9 • Implementation defined task scheduling points for untied tasks were removed (see Section 2.13.6
10 on page 147).
- 11 • Device constructs (see Section 2.15 on page 158), the **OMP_DEFAULT_DEVICE** environment
12 variable (see Section 5.15 on page 609), the **omp_set_default_device**,
13 **omp_get_default_device**, **omp_get_num_devices**, **omp_get_num_teams**,
14 **omp_get_team_num**, and **omp_is_initial_device** routines were added to support
15 execution on devices.
- 16 • The **depend** clause (see Section 2.20.11 on page 248) was added to support task dependences.
- 17 • The **taskgroup** construct (see Section 2.20.6 on page 225) was added to support more flexible
18 deep task synchronization.
- 19 • The **atomic** construct (see Section 2.20.7 on page 227) was extended to support atomic swap
20 with the **capture** clause, to allow new atomic update and capture forms, and to support
21 sequentially consistent atomic operations with a new **seq_cst** clause.
- 22 • The **cancel** construct (see Section 2.21.1 on page 256), the **cancellation point**
23 construct (see Section 2.21.2 on page 260), the **omp_get_cancellation** runtime routine
24 (see Section 3.2.9 on page 340) and the **OMP_CANCELLATION** environment variable (see
25 Section 5.11 on page 604) were added to support the concept of cancellation.
- 26 • The **reduction** clause (see Section 2.22.5.4 on page 297) was extended and the
27 **declare reduction** construct (see Section 2.23.2 on page 321) was added to support user
28 defined reductions.
- 29 • The **OMP_DISPLAY_ENV** environment variable (see Section 5.12 on page 605) was added to
30 display the value of ICVs associated with the OpenMP environment variables.
- 31 • Examples (previously Appendix A) were moved to a separate document.

1 D.5 Version 3.0 to 3.1 Differences

- 2 • The *bind-var* ICV has been added, which controls whether or not threads are bound to processors
3 (see Section 2.4.1 on page 47). The value of this ICV can be set with the **OMP_PROC_BIND**
4 environment variable (see Section 5.4 on page 598).
- 5 • The **final** and **mergeable** clauses (see Section 2.13.1 on page 133) were added to the **task**
6 construct to support optimization of task data environments.
- 7 • The **taskyield** construct (see Section 2.13.4 on page 145) was added to allow user-defined
8 task scheduling points.
- 9 • The **atomic** construct (see Section 2.20.7 on page 227) was extended to include **read**, **write**,
10 and **capture** forms, and an **update** clause was added to apply the already existing form of the
11 **atomic** construct.
- 12 • Data environment restrictions were changed to allow **intent (in)** and **const**-qualified types
13 for the **firstprivate** clause (see Section 2.22.4.4 on page 281).
- 14 • Data environment restrictions were changed to allow Fortran pointers in **firstprivate** (see
15 Section 2.22.4.4 on page 281) and **lastprivate** (see Section 2.22.4.5 on page 283).
- 16 • New reduction operators **min** and **max** were added for C and C++ (see Section 2.22.5 on
17 page 290).
- 18 • The nesting restrictions in Section 2.24 on page 327 were clarified to disallow closely-nested
19 OpenMP regions within an **atomic** region. This allows an **atomic** region to be consistently
20 defined with other OpenMP regions so that they include all code in the atomic construct.
- 21 • The **omp_in_final** runtime library routine (see Section 3.2.21 on page 353) was added to
22 support specialization of final task regions.
- 23 • The *nthreads-var* ICV has been modified to be a list of the number of threads to use at each
24 nested parallel region level. The value of this ICV is still set with the **OMP_NUM_THREADS**
25 environment variable (see Section 5.2 on page 597), but the algorithm for determining the
26 number of threads used in a parallel region has been modified to handle a list (see Section 2.9.1
27 on page 77).
- 28 • Descriptions of examples (previously Appendix A) were expanded and clarified.
- 29 • Replaced incorrect use of **omp_integer_kind** in Fortran interfaces with
30 **selected_int_kind(8)**.

1 D.6 Version 2.5 to 3.0 Differences

- 2 • The definition of active **parallel** region has been changed: in Version 3.0 a **parallel**
3 region is active if it is executed by a team consisting of more than one thread (see Section 1.2.2
4 on page 3).
- 5 • The concept of tasks has been added to the OpenMP execution model (see Section 1.2.5 on
6 page 11 and Section 1.3 on page 19).
- 7 • The OpenMP memory model now covers atomicity of memory accesses (see Section 1.4.1 on
8 page 22). The description of the behavior of **volatile** in terms of **flush** was removed.
- 9 • In Version 2.5, there was a single copy of the *nest-var*, *dyn-var*, *nthreads-var* and *run-sched-var*
10 internal control variables (ICVs) for the whole program. In Version 3.0, there is one copy of
11 these ICVs per task (see Section 2.4 on page 47). As a result, the **omp_set_num_threads**,
12 **omp_set_nested** and **omp_set_dynamic** runtime library routines now have specified
13 effects when called from inside a **parallel** region (see Section 3.2.1 on page 332,
14 Section 3.2.7 on page 338 and Section 3.2.10 on page 341).
- 15 • The *thread-limit-var* ICV has been added, which controls the maximum number of threads
16 participating in the OpenMP program. The value of this ICV can be set with the
17 **OMP_THREAD_LIMIT** environment variable and retrieved with the
18 **omp_get_thread_limit** runtime library routine (see Section 2.4.1 on page 47,
19 Section 3.2.14 on page 346 and Section 5.10 on page 604).
- 20 • The *max-active-levels-var* ICV has been added, which controls the number of nested active
21 **parallel** regions. The value of this ICV can be set with the **OMP_MAX_ACTIVE_LEVELS**
22 environment variable and the **omp_set_max_active_levels** runtime library routine, and
23 it can be retrieved with the **omp_get_max_active_levels** runtime library routine (see
24 Section 2.4.1 on page 47, Section 3.2.15 on page 347, Section 3.2.16 on page 348 and
25 Section 5.9 on page 603).
- 26 • The *stacksize-var* ICV has been added, which controls the stack size for threads that the OpenMP
27 implementation creates. The value of this ICV can be set with the **OMP_STACKSIZE**
28 environment variable (see Section 2.4.1 on page 47 and Section 5.7 on page 602).
- 29 • The *wait-policy-var* ICV has been added, which controls the desired behavior of waiting threads.
30 The value of this ICV can be set with the **OMP_WAIT_POLICY** environment variable (see
31 Section 2.4.1 on page 47 and Section 5.8 on page 603).
- 32 • The rules for determining the number of threads used in a **parallel** region have been modified
33 (see Section 2.9.1 on page 77).
- 34 • In Version 3.0, the assignment of iterations to threads in a loop construct with a **static**
35 schedule kind is deterministic (see Section 2.12.2 on page 100).
- 36 • In Version 3.0, a loop construct may be associated with more than one perfectly nested loop. The
37 number of associated loops may be controlled by the **collapse** clause (see Section 2.12.2 on

1 page 100).

- 2 • Random access iterators, and variables of unsigned integer type, may now be used as loop
3 iterators in loops associated with a loop construct (see Section 2.12.2 on page 100).
- 4 • The schedule kind **auto** has been added, which gives the implementation the freedom to choose
5 any possible mapping of iterations in a loop construct to threads in the team (see Section 2.12.2
6 on page 100).
- 7 • The **task** construct (see Section 2.13 on page 133) has been added, which provides a
8 mechanism for creating tasks explicitly.
- 9 • The **taskwait** construct (see Section 2.20.5 on page 223) has been added, which causes a task
10 to wait for all its child tasks to complete.
- 11 • Fortran assumed-size arrays now have predetermined data-sharing attributes (see
12 Section 2.22.1.1 on page 263).
- 13 • In Version 3.0, static class members variables may appear in a **threadprivate** directive (see
14 Section 2.22.2 on page 268).
- 15 • Version 3.0 makes clear where, and with which arguments, constructors and destructors of
16 private and threadprivate class type variables are called (see Section 2.22.2 on page 268,
17 Section 2.22.4.3 on page 280, Section 2.22.4.4 on page 281, Section 2.22.6.1 on page 301 and
18 Section 2.22.6.2 on page 303).
- 19 • In Version 3.0, Fortran allocatable arrays may appear in **private**, **firstprivate**,
20 **lastprivate**, **reduction**, **copyin** and **copyprivate** clauses (see Section 2.22.2 on
21 page 268, Section 2.22.4.3 on page 280, Section 2.22.4.4 on page 281, Section 2.22.4.5 on
22 page 283, Section 2.22.5.4 on page 297, Section 2.22.6.1 on page 301 and Section 2.22.6.2 on
23 page 303).
- 24 • In Fortran, **firstprivate** is now permitted as an argument to the **default** clause (see
25 Section 2.22.4.1 on page 277).
- 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 2.22.4.3 on page 280).
- 30 • The runtime library routines **omp_set_schedule** and **omp_get_schedule** have been
31 added; these routines respectively set and retrieve the value of the *run-sched-var* ICV (see
32 Section 3.2.12 on page 343 and Section 3.2.13 on page 345).
- 33 • The **omp_get_level** runtime library routine has been added, which returns the number of
34 nested **parallel** regions enclosing the task that contains the call (see Section 3.2.17 on
35 page 349).
- 36 • The **omp_get_ancestor_thread_num** runtime library routine has been added, which
37 returns, for a given nested level of the current thread, the thread number of the ancestor (see

- 1 Section [3.2.18](#) on page [350](#)).
- 2 • The `omp_get_team_size` runtime library routine has been added, which returns, for a given
- 3 nested level of the current thread, the size of the thread team to which the ancestor belongs (see
- 4 Section [3.2.19](#) on page [351](#)).
- 5 • The `omp_get_active_level` runtime library routine has been added, which returns the
- 6 number of nested, active `parallel` regions enclosing the task that contains the call (see
- 7 Section [3.2.20](#) on page [353](#)).
- 8 • In Version 3.0, locks are owned by tasks, not by threads (see Section [3.3](#) on page [378](#)).

This page intentionally left blank

Index

Symbols

`_OPENMP` macro, 41, 605–607

A

acquire flush, 26

affinity, 79

allocate, 154, 157

array sections, 59

array shaping, 58

atomic, 227

atomic construct, 615

attribute clauses, 276

attributes, data-mapping, 305

attributes, data-sharing, 263

auto, 105

B

barrier, 219

barrier, implicit, 221

C

cancel, 256

cancellation constructs, 256

cancel, 256

cancellation point, 260

cancellation point, 260

canonical loop form, 95

capture, **atomic**, 227

clauses

allocate, 157

 attribute data-sharing, 276

collapse, 100, 102

copyin, 301

copyprivate, 303

 data copying, 301

 data-sharing, 276

default, 277

defaultmap, 315

depend, 248

firstprivate, 281

hint, 253

if Clause, 213

in_reduction, 300

lastprivate, 283

linear, 286

map, 307

private, 280

reduction, 297

schedule, 102

shared, 278

task_reduction, 299

combined constructs, 185

 parallel loop construct, 187

parallel sections, 188

parallel workshare, 190

 parallel worksharing-loop construct,
 185

 parallel worksharing-loop SIMD
 construct, 191

target parallel, 192

 target parallel worksharing-loop, 193

 target parallel worksharing-loop SIMD,
 195

target simd, 196

target teams, 198

target teams distribute, 203

 target teams distribute parallel
 worksharing-loop construct, 206

 target teams distribute parallel
 worksharing-loop SIMD construct,
 209

target teams distribute simd,
 204

- teams distribute**, 199
- teams distribute parallel
 - worksharing-loop construct, 205
- teams distribute parallel
 - worksharing-loop SIMD construct, 208
- teams distribute simd**, 201
- teams loop**, 200
- compilation sentinels, 41, 42
- compliance, 30
- conditional compilation, 41
- constructs
 - atomic**, 227
 - barrier**, 219
 - cancel**, 256
 - cancellation constructs, 256
 - cancellation point**, 260
 - combined constructs, 185
 - critical**, 216
 - declare mapper**, 183
 - declare target**, 178
 - depobj**, 247
 - device constructs, 158
 - distribute**, 117
 - distribute parallel do**, 122
 - distribute parallel do simd**, 124
 - distribute parallel for**, 122
 - distribute parallel for simd**, 124
 - distribute parallel worksharing-loop, 122
 - distribute parallel worksharing-loop SIMD, 124
 - distribute simd**, 121
 - do Fortran**, 100
 - flush**, 235
 - for, C/C++**, 100
 - loop**, 126
 - master**, 214
 - ordered**, 243
 - parallel**, 72
 - parallel do Fortran**, 185
 - parallel for C/C++**, 185
 - parallel loop construct, 187
 - parallel sections**, 188
 - parallel workshare**, 190
 - parallel worksharing-loop construct, 185
 - parallel worksharing-loop SIMD construct, 191
 - sections**, 86
 - simd**, 111
 - single**, 89
 - target**, 168
 - target data**, 159
 - target enter data**, 162
 - target exit data**, 165
 - target parallel**, 192
 - target parallel do**, 193
 - target parallel do simd**, 195
 - target parallel for**, 193
 - target parallel for simd**, 195
 - target parallel worksharing-loop, 193
 - target parallel worksharing-loop SIMD, 195
 - target simd**, 196
 - target teams**, 198
 - target teams distribute**, 203
 - target teams distribute parallel
 - worksharing-loop construct, 206
 - target teams distribute parallel
 - worksharing-loop SIMD construct, 209
 - target teams distribute simd**, 204
 - target update**, 174
 - task**, 133
 - taskgroup**, 225
 - tasking constructs, 133
 - taskloop**, 138
 - taskloop simd**, 144
 - taskwait**, 223
 - taskyield**, 145
 - teams**, 81
 - teams distribute**, 199

- teams distribute parallel
 - worksharing-loop construct, 205
- teams distribute parallel
 - worksharing-loop SIMD construct, 208
- teams distribute simd**, 201
- teams loop**, 200
- workshare**, 91
 - worksharing, 85
 - worksharing-loop*, 100
 - Worksharing-Loop SIMD, 116
- controlling OpenMP thread affinity, 79
- copyin**, 301
- copyprivate**, 303
- critical**, 216

D

- data copying clauses, 301
- data environment, 263
- data terminology, 13
- data-mapping rules and clauses, 305
- data-sharing attribute clauses, 276
- data-sharing attribute rules, 263
- declare**, 317
- declare mapper**, 183
- declare reduction**, 321
- declare simd**, 317
- declare target**, 178
- declare variant**, 67
- default**, 277
- defaultmap**, 315
- depend**, 248
- Depend Object, 246
- depobj**, 247
- deprecated features, 625
- device constructs, 158
 - declare mapper**, 183
 - declare target**, 178
 - device constructs, 158
 - distribute**, 117
 - distribute parallel worksharing-loop, 122
 - distribute parallel worksharing-loop SIMD, 124

- distribute simd**, 121
 - target**, 168
 - target update**, 174
 - teams**, 81
- device data environments, 23, 162, 165
- device memory routines, 393
- directive format, 36
- directives, 35
 - allocate**, 154
 - declare, 317
 - declare mapper**, 183
 - declare reduction**, 321
 - declare simd**, 317
 - declare target**, 178
 - declare variant**, 67
 - memory management directives, 150
 - metadirective**, 69
 - requires**, 43
 - scan** Directive, 129
 - threadprivate**, 268
 - variant directives, 63
- distribute**, 117
 - distribute parallel worksharing-loop
 - construct, 122
 - distribute parallel worksharing-loop SIMD
 - construct, 124
- distribute simd**, 121
- do*, Fortran, 100
- do simd**, 116
- dynamic**, 104
- dynamic thread adjustment, 613

E

- environment variables, 595
 - OMP_AFFINITY_FORMAT**, 607
 - OMP_ALLOCATOR**, 612
 - OMP_CANCELLATION**, 604
 - OMP_DEBUG**, 612
 - OMP_DEFAULT_DEVICE**, 609
 - OMP_DISPLAY_AFFINITY**, 606
 - OMP_DISPLAY_ENV**, 605
 - OMP_DYNAMIC**, 598
 - OMP_MAX_ACTIVE_LEVELS**, 603
 - OMP_MAX_TASK_PRIORITY**, 609

- `OMP_NESTED`, 601
- `OMP_NUM_THREADS`, 597
- `OMP_PLACES`, 599
- `OMP_PROC_BIND`, 598
- `OMP_SCHEDULE`, 596
- `OMP_STACKSIZE`, 602
- `OMP_TARGET_OFFLOAD`, 610
- `OMP_THREAD_LIMIT`, 604
- `OMP_TOOL`, 611
- `OMP_TOOL_LIBRARIES`, 611
- `OMP_WAIT_POLICY`, 603

event, 392

event callback registration, 422

event callback signatures, 446

event routines, 392

execution environment routines, 332

execution model, 19

F

features history, 625

`firstprivate`, 281

fixed source form conditional compilation
sentinels, 41

fixed source form directives, 38

`flush`, 235

flush operation, 24

flush synchronization, 26

flush-set, 24

`for`, C/C++, 100

`for simd`, 116

frames, 589

free source form conditional compilation
sentinel, 42

free source form directives, 39

G

glossary, 2

`guided`, 104

H

happens before, 26

header files, 330

history of features, 625

I

ICVs (internal control variables), 47

`if` Clause, 213

implementation, 613

implementation terminology, 15

implicit barrier, 221

implicit flushes, 240

`in_reduction`, 300

include files, 330

internal control variables, 613

internal control variables (ICVs), 47

introduction, 1

iterators, 61

L

`lastprivate`, 283

`linear`, 286

list item privatization, 273

lock routines, 378

`loop`, 126

loop terminology, 9

M

`map`, 307

`master`, 214

memory allocators, 151

memory management, 150

memory management directives
memory management directives, 150

memory management routines, 403

memory model, 22

memory spaces, 150

meta-directive, 69

modifying and retrieving ICV values, 51

modifying ICV's, 49

N

nesting of regions, 327

normative references, 30

O

`OMP_AFFINITY_FORMAT`, 607

`omp_alloc`, 410

`OMP_ALLOCATOR`, 612

OMP_CANCELLATION, 604
omp_capture_affinity, 365
OMP_DEBUG, 612
OMP_DEFAULT_DEVICE, 609
omp_destroy_allocator, 407
omp_destroy_lock, 383
omp_destroy_nest_lock, 383
OMP_DISPLAY_AFFINITY, 606
omp_display_affinity, 364
OMP_DISPLAY_ENV, 605
OMP_DYNAMIC, 598
omp_free, 411
omp_fulfill_event, 392
omp_get_active_level, 353
omp_get_affinity_format, 363
omp_get_ancestor_thread_num,
350
omp_get_cancellation, 340
omp_get_default_allocator, 409
omp_get_default_device, 368
omp_get_device_num, 369
omp_get_dynamic, 339
omp_get_initial_device, 373
omp_get_level, 349
omp_get_max_active_levels, 348
omp_get_max_task_priority, 374
omp_get_max_threads, 334
omp_get_nested, 342
omp_get_num_devices, 369
omp_get_num_places, 356
omp_get_num_procs, 337
omp_get_num_teams, 370
omp_get_num_threads, 333
omp_get_partition_num_places,
360
omp_get_partition_place_nums,
360
omp_get_place_num, 359
omp_get_place_num_procs, 357
omp_get_place_proc_ids, 358
omp_get_proc_bind, 354
omp_get_schedule, 345
omp_get_team_num, 371
omp_get_team_size, 351
omp_get_thread_limit, 346
omp_get_thread_num, 336
omp_get_wtick, 391
omp_get_wtime, 390
omp_in_final, 353
omp_in_parallel, 337
omp_init_allocator, 406
omp_init_lock, 380, 382
omp_init_nest_lock, 380, 382
omp_is_initial_device, 372
OMP_MAX_ACTIVE_LEVELS, 603
OMP_MAX_TASK_PRIORITY, 609
OMP_NESTED, 601
OMP_NUM_THREADS, 597
omp_pause_resource, 375
omp_pause_resource_all, 377
OMP_PLACES, 599
OMP_PROC_BIND, 598
OMP_SCHEDULE, 596
omp_set_affinity_format, 361
omp_set_default_allocator, 408
omp_set_default_device, 367
omp_set_dynamic, 338
omp_set_lock, 384
omp_set_max_active_levels, 347
omp_set_nest_lock, 384
omp_set_nested, 341
omp_set_num_threads, 332
omp_set_schedule, 343
OMP_STACKSIZE, 602
omp_target_alloc, 393
omp_target_associate_ptr, 399
omp_target_disassociate_ptr,
401
omp_target_free, 395
omp_target_is_present, 396
omp_target_memcpy, 396
omp_target_memcpy_rect, 398
OMP_TARGET_OFFLOAD, 610
omp_test_lock, 388
omp_test_nest_lock, 388
OMP_THREAD_LIMIT, 604

OMP_TOOL, 611
OMP_TOOL_LIBRARIES, 611
omp_unset_lock, 386
omp_unset_nest_lock, 386
OMP_WAIT_POLICY, 603
ompd diagram, 623
ompd_bp_device_begin, 583
ompd_bp_device_end, 584
ompd_bp_parallel_begin, 577
ompd_bp_parallel_end, 578
ompd_bp_task_begin, 580
ompd_bp_task_end, 581
ompd_bp_thread_begin, 582
ompd_bp_thread_end, 582
ompd_callback_device_host_fn_t, 536
ompd_callback_get_thread_context_for_discretionary_thread_id_fn_t, 530
ompd_callback_memory_alloc_fn_t, 528
ompd_callback_memory_free_fn_t, 529
ompd_callback_memory_read_fn_t, 533
ompd_callback_memory_write_fn_t, 535
ompd_callback_print_string_fn_t, 537
ompd_callback_sizeof_fn_t, 531
ompd_callback_symbol_addr_fn_t, 532
ompd_callbacks_t, 538
ompd_dll_locations_valid, 594
ompd_dll_locations, 593
ompt_callback_buffer_complete_t, 475
ompt_callback_buffer_request_t, 474
ompt_callback_cancel_t, 477
ompt_callback_control_tool_t, 476
ompt_callback_device_finalize_t, 479
ompt_callback_device_initialize_t, 478
ompt_callback_dispatch_t, 464
ompt_callback_flush_t, 463
ompt_callback_implicit_task_t, 455
ompt_callback_master_t, 450
ompt_callback_mutex_acquire_t, 458
ompt_callback_mutex_t, 459
ompt_callback_nest_lock_t, 460
ompt_callback_parallel_begin_t, 447
ompt_callback_parallel_end_t, 449
ompt_callback_sync_region_t, 457
ompt_callback_device_load_t, 466
ompt_callback_device_unload_t, 467
ompt_callback_target_data_op_t, 468
ompt_callback_target_map_t, 471
ompt_callback_target_submit_t, 472
ompt_callback_target_t, 465
ompt_callback_task_create_t, 451
ompt_callback_task_dependence_t, 454
ompt_callback_task_dependences_t, 453
ompt_callback_task_schedule_t, 454
ompt_callback_thread_begin_t, 446
ompt_callback_thread_end_t, 447
ompt_callback_work_t, 461
OpenMP compliance, 30
ordered, 243

P
parallel, 72

parallel loop construct, 187
parallel sections, 188
parallel workshare, 190
parallel worksharing-loop construct, 185
parallel worksharing-loop SIMD construct,
191
private, 280

R

read, atomic, 227
reduction, 297
reduction clauses, 290
release flush, 26
requires, 43
runtime library definitions, 330
runtime library routines, 329

S

scan Directive, 129
scheduling, 147
sec:Declare Directives, 317
sections, 86
shared, 278
simd, 111
SIMD Constructs, 111
Simple Lock Routines, 379
single, 89
stand-alone directives, 40
strong flush, 24
synchronization constructs, 216
synchronization constructs and clauses, 216
synchronization hints, 253
synchronization terminology, 9

T

target, 168
target data, 159
target memory routines, 393
target parallel, 192
target parallel worksharing-loop construct,
193
target parallel worksharing-loop SIMD
construct, 195
target simd, 196

target teams, 198
target teams distribute, 203
target teams distribute parallel
worksharing-loop construct, 206
target teams distribute parallel
worksharing-loop SIMD construct,
209

target teams distribute simd, 204

target update, 174

task, 133

task frame management, 621

task scheduling, 147

task_reduction, 299

taskgroup, 225

tasking constructs, 133

tasking terminology, 11

taskloop, 138

taskloop simd, 144

taskwait, 223

taskyield, 145

teams, 81

teams distribute, 199

teams distribute parallel worksharing-loop
construct, 205

teams distribute parallel worksharing-loop
SIMD construct, 208

teams distribute simd, 201

teams loop, 200

thread affinity, 79

threadprivate, 268

timer, 390

timing routines, 390

tool control, 413

tool initialization, 420

tool interfaces definitions, 418

tool support, 417

tools header files, 418

tracing device activity, 425

U

update, atomic, 227

V

variables, environment, 595

variant directives, [63](#)

W

wait identifier, [591](#)

wall clock timer, [390](#)

workshare, [91](#)

worksharing

 constructs, [85](#)

 parallel, [185](#), [187](#)

 scheduling, [110](#)

worksharing constructs, [85](#)

worksharing-loop, [100](#)

worksharing-loop SIMD Construct, [116](#)

write, **atomic**, [227](#)