



# OpenMP Technical Report 8: Version 5.1 Preview

This Technical Report is a preview of the OpenMP Application Programming Specification Version 5.1 that augments the OpenMP API Version 5.0 with full support for C11 and C++11/14/17, syntax extensions to support C++ attribute specifiers, the tile, error and assume directives, iterator support for data motion clauses, the interop directive, extensions to task dependences, extensions to SIMD constructs and the declare variant directive as well as clarifications and improvements to existing features. See appendix B.2 for the full list.

## EDITORS

*Bronis R. de Supinski*

*Michael Klemm*

November 7, 2019

Expires November 11, 2020

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

OpenMP Architecture Review Board – [www.openmp.org](http://www.openmp.org) – [info@openmp.org](mailto: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 5.1 Preview, November 2019**

Copyright ©1997-2019 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 the OpenMP Architecture Review Board.

This is a preview of the OpenMP API Specification Version 5.1 and includes the following internal GitHub issues (corresponding Trac ticket numbers in parentheses when they exist) applied to the 5.0 LaTeX sources: 1239 (307) 1327 (395), 1438 (506), 1476 (544), 1497 (565), 1725 (793), 1727 (795), 1780 (848) 1783-1785 (851-853), 1793 (861), 1796 (864), 1802 (870), 1804 (872), 1814-1817 (882-885), 1819-1820 (887-888) 1822-1832 (890-900), 1834-1837 (902-905), 1842 (910), 1845 (913), 1847 (915), 1849-1850 (917-918), 1854 (922), 1856 (924), 1859-1867 (927-935), 1871-1876 (939-944), 1878 (946), 1880, 1883, 1887, 1889, 1892-1893, 1896-1898, 1901-1902, 1907, 1909-1911, 1914-1920, 1922-1923, 1925, 1929-1930, 1937-1940, 1947, 1954, 1960-1961, 1970, 2013, 2024

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

# Contents

<b>1 Introduction</b>	<b>1</b>
1.1 Scope	1
1.2 Glossary	2
1.2.1 Threading Concepts	2
1.2.2 OpenMP Language Terminology	2
1.2.3 Loop Terminology	9
1.2.4 Synchronization Terminology	9
1.2.5 Tasking Terminology	11
1.2.6 Data Terminology	13
1.2.7 Implementation Terminology	17
1.2.8 Tool Terminology	18
1.3 Execution Model	21
1.4 Memory Model	24
1.4.1 Structure of the OpenMP Memory Model	24
1.4.2 Device Data Environments	25
1.4.3 Memory Management	25
1.4.4 The Flush Operation	26
1.4.5 Flush Synchronization and <i>Happens Before</i>	27
1.4.6 OpenMP Memory Consistency	29
1.5 Tool Interfaces	30
1.5.1 OMPT	30
1.5.2 OMPD	31
1.6 OpenMP Compliance	31
1.7 Normative References	32
1.8 Organization of this Document	33

<b>2 Directives</b>	<b>35</b>
2.1 Directive Format	36
2.1.1 Fixed Source Form Directives	40
2.1.2 Free Source Form Directives	40
2.1.3 Stand-Alone Directives	41
2.1.4 Array Shaping	42
2.1.5 Array Sections	42
2.1.6 Iterators	46
2.2 Conditional Compilation	48
2.2.1 Fixed Source Form Conditional Compilation Sentinels	49
2.2.2 Free Source Form Conditional Compilation Sentinel	49
2.3 Variant Directives	50
2.3.1 OpenMP Context	50
2.3.2 Context Selectors	52
2.3.3 Matching and Scoring Context Selectors	53
2.3.4 Metadirectives	54
2.3.5 Declare Variant Directive	56
2.4 <b>requires</b> Directive	60
2.5 <b>assume</b> Directive	63
2.6 Internal Control Variables	66
2.6.1 ICV Descriptions	66
2.6.2 ICV Initialization	68
2.6.3 Modifying and Retrieving ICV Values	71
2.6.4 How ICVs are Scoped	74
2.6.4.1 How the Per-Data Environment ICVs Work	75
2.6.5 ICV Override Relationships	76
2.7 Utility Directives	77
2.7.1 <b>nothing</b> Directive	77
2.7.2 <b>error</b> Directive	78
2.8 <b>parallel</b> Construct	79
2.8.1 Determining the Number of Threads for a <b>parallel</b> Region	84
2.8.2 Controlling OpenMP Thread Affinity	85
2.9 <b>teams</b> Construct	87

2.10	<b>scope</b> Construct . . . . .	91
2.11	Worksharing Constructs . . . . .	93
2.11.1	<b>sections</b> Construct . . . . .	93
2.11.2	<b>single</b> Construct . . . . .	96
2.11.3	<b>workshare</b> Construct . . . . .	98
2.12	Loop-Related Directives . . . . .	101
2.12.1	Canonical Loop Form . . . . .	101
2.12.2	Worksharing-Loop Construct . . . . .	106
2.12.2.1	Determining the Schedule of a Worksharing-Loop . . . . .	114
2.12.3	SIMD Directives . . . . .	115
2.12.3.1	<b>simd</b> Construct . . . . .	115
2.12.3.2	Worksharing-Loop SIMD Construct . . . . .	120
2.12.3.3	Declare SIMD Directive . . . . .	121
2.12.4	<b>distribute</b> Loop Constructs . . . . .	125
2.12.4.1	<b>distribute</b> Construct . . . . .	125
2.12.4.2	<b>distribute simd</b> Construct . . . . .	129
2.12.4.3	Distribute Parallel Worksharing-Loop Construct . . . . .	130
2.12.4.4	Distribute Parallel Worksharing-Loop SIMD Construct . . . . .	131
2.12.5	<b>loop</b> Construct . . . . .	133
2.12.6	<b>scan</b> Directive . . . . .	137
2.12.7	Loop Transformation Constructs . . . . .	140
2.12.7.1	<b>tile</b> Construct . . . . .	140
2.13	Tasking Constructs . . . . .	142
2.13.1	<b>task</b> Construct . . . . .	142
2.13.2	<b>taskloop</b> Construct . . . . .	147
2.13.3	<b>taskloop simd</b> Construct . . . . .	152
2.13.4	<b>taskyield</b> Construct . . . . .	154
2.13.5	Initial Task . . . . .	155
2.13.6	Task Scheduling . . . . .	156
2.14	Memory Management Directives . . . . .	158
2.14.1	Memory Spaces . . . . .	158
2.14.2	Memory Allocators . . . . .	159
2.14.3	<b>allocate</b> Directive . . . . .	162

2.14.4	<b>allocate</b> Clause	164
2.15	Device Directives	165
2.15.1	Device Initialization	165
2.15.2	<b>target data</b> Construct	166
2.15.3	<b>target enter data</b> Construct	169
2.15.4	<b>target exit data</b> Construct	172
2.15.5	<b>target</b> Construct	174
2.15.6	<b>target update</b> Construct	181
2.15.7	Declare Target Directive	185
2.15.8	<b>interop</b> Construct	191
2.16	Combined Constructs	193
2.16.1	Parallel Worksharing-Loop Construct	193
2.16.2	<b>parallel loop</b> Construct	194
2.16.3	<b>parallel sections</b> Construct	195
2.16.4	<b>parallel workshare</b> Construct	196
2.16.5	Parallel Worksharing-Loop SIMD Construct	197
2.16.6	<b>parallel master</b> Construct	198
2.16.7	<b>master taskloop</b> Construct	199
2.16.8	<b>master taskloop simd</b> Construct	200
2.16.9	<b>parallel master taskloop</b> Construct	201
2.16.10	<b>parallel master taskloop simd</b> Construct	202
2.16.11	<b>teams distribute</b> Construct	203
2.16.12	<b>teams distribute simd</b> Construct	204
2.16.13	Teams Distribute Parallel Worksharing-Loop Construct	205
2.16.14	Teams Distribute Parallel Worksharing-Loop SIMD Construct	206
2.16.15	<b>teams loop</b> Construct	207
2.16.16	<b>target parallel</b> Construct	208
2.16.17	Target Parallel Worksharing-Loop Construct	210
2.16.18	Target Parallel Worksharing-Loop SIMD Construct	211
2.16.19	<b>target parallel loop</b> Construct	212
2.16.20	<b>target simd</b> Construct	213
2.16.21	<b>target teams</b> Construct	215
2.16.22	<b>target teams distribute</b> Construct	215



2.16.23	<b>target teams distribute simd</b> Construct	216
2.16.24	<b>target teams loop</b> Construct	218
2.16.25	Target Teams Distribute Parallel Worksharing-Loop Construct	219
2.16.26	Target Teams Distribute Parallel Worksharing-Loop SIMD Construct	220
2.17	Clauses on Combined and Composite Constructs	222
2.18	<b>if</b> Clause	224
2.19	<b>master</b> Construct	225
2.20	Synchronization Constructs and Clauses	226
2.20.1	<b>critical</b> Construct	226
2.20.2	<b>barrier</b> Construct	229
2.20.3	Implicit Barriers	231
2.20.4	Implementation-Specific Barriers	233
2.20.5	<b>taskwait</b> Construct	233
2.20.6	<b>taskgroup</b> Construct	235
2.20.7	<b>atomic</b> Construct	237
2.20.8	<b>flush</b> Construct	245
2.20.8.1	Implicit Flushes	249
2.20.9	<b>ordered</b> Construct	252
2.20.10	Depend Objects	256
2.20.10.1	<b>depobj</b> Construct	256
2.20.11	<b>depend</b> Clause	257
2.20.12	Synchronization Hints	262
2.21	Cancellation Constructs	265
2.21.1	<b>cancel</b> Construct	265
2.21.2	<b>cancellation point</b> Construct	269
2.22	Data Environment	271
2.22.1	Data-Sharing Attribute Rules	271
2.22.1.1	Variables Referenced in a Construct	271
2.22.1.2	Variables Referenced in a Region but not in a Construct	275
2.22.2	<b>threadprivate</b> Directive	276
2.22.3	List Item Privatization	281
2.22.4	Data-Sharing Attribute Clauses	284
2.22.4.1	<b>default</b> Clause	284

2.22.4.2	<b>shared</b> Clause . . . . .	285
2.22.4.3	<b>private</b> Clause . . . . .	287
2.22.4.4	<b>firstprivate</b> Clause . . . . .	288
2.22.4.5	<b>lastprivate</b> Clause . . . . .	290
2.22.4.6	<b>linear</b> Clause . . . . .	292
2.22.5	Reduction Clauses and Directives . . . . .	295
2.22.5.1	Properties Common To All Reduction Clauses . . . . .	295
2.22.5.2	Reduction Scoping Clauses . . . . .	302
2.22.5.3	Reduction Participating Clauses . . . . .	302
2.22.5.4	<b>reduction</b> Clause . . . . .	302
2.22.5.5	<b>task_reduction</b> Clause . . . . .	305
2.22.5.6	<b>in_reduction</b> Clause . . . . .	306
2.22.5.7	<b>declare_reduction</b> Directive . . . . .	307
2.22.6	Data Copying Clauses . . . . .	313
2.22.6.1	<b>copyin</b> Clause . . . . .	313
2.22.6.2	<b>copyprivate</b> Clause . . . . .	315
2.22.7	Data-Mapping Attribute Rules, Clauses, and Directives . . . . .	317
2.22.7.1	<b>map</b> Clause . . . . .	318
2.22.7.2	Pointer Initialization for Device Data Environments . . . . .	327
2.22.7.3	<b>defaultmap</b> Clause . . . . .	328
2.22.7.4	<b>declare_mapper</b> Directive . . . . .	330
2.23	Nesting of Regions . . . . .	332

### **3 Runtime Library Routines 335**

3.1	Runtime Library Definitions . . . . .	335
3.2	Execution Environment Routines . . . . .	338
3.2.1	<b>omp_set_num_threads</b> . . . . .	338
3.2.2	<b>omp_get_num_threads</b> . . . . .	339
3.2.3	<b>omp_get_max_threads</b> . . . . .	340
3.2.4	<b>omp_get_thread_num</b> . . . . .	341
3.2.5	<b>omp_get_num_procs</b> . . . . .	342
3.2.6	<b>omp_in_parallel</b> . . . . .	342
3.2.7	<b>omp_set_dynamic</b> . . . . .	343
3.2.8	<b>omp_get_dynamic</b> . . . . .	344

3.2.9	<code>omp_get_cancellation</code>	345
3.2.10	<code>omp_set_nested</code>	345
3.2.11	<code>omp_get_nested</code>	346
3.2.12	<code>omp_set_schedule</code>	347
3.2.13	<code>omp_get_schedule</code>	349
3.2.14	<code>omp_get_thread_limit</code>	350
3.2.15	<code>omp_get_supported_active_levels</code>	351
3.2.16	<code>omp_set_max_active_levels</code>	352
3.2.17	<code>omp_get_max_active_levels</code>	353
3.2.18	<code>omp_get_level</code>	353
3.2.19	<code>omp_get_ancestor_thread_num</code>	354
3.2.20	<code>omp_get_team_size</code>	355
3.2.21	<code>omp_get_active_level</code>	356
3.2.22	<code>omp_in_final</code>	357
3.2.23	<code>omp_get_proc_bind</code>	357
3.2.24	<code>omp_get_num_places</code>	359
3.2.25	<code>omp_get_place_num_procs</code>	360
3.2.26	<code>omp_get_place_proc_ids</code>	360
3.2.27	<code>omp_get_place_num</code>	361
3.2.28	<code>omp_get_partition_num_places</code>	362
3.2.29	<code>omp_get_partition_place_nums</code>	363
3.2.30	<code>omp_set_affinity_format</code>	364
3.2.31	<code>omp_get_affinity_format</code>	365
3.2.32	<code>omp_display_affinity</code>	366
3.2.33	<code>omp_capture_affinity</code>	367
3.2.34	<code>omp_set_default_device</code>	369
3.2.35	<code>omp_get_default_device</code>	370
3.2.36	<code>omp_get_num_devices</code>	370
3.2.37	<code>omp_get_device_num</code>	371
3.2.38	<code>omp_get_num_teams</code>	372
3.2.39	<code>omp_get_team_num</code>	372
3.2.40	<code>omp_is_initial_device</code>	373
3.2.41	<code>omp_get_initial_device</code>	374

3.2.42	<code>omp_get_max_task_priority</code>	374
3.2.43	<code>omp_pause_resource</code>	375
3.2.44	<code>omp_pause_resource_all</code>	377
3.2.45	<code>omp_set_num_teams</code>	378
3.2.46	<code>omp_get_max_teams</code>	379
3.2.47	<code>omp_set_teams_thread_limit</code>	380
3.2.48	<code>omp_get_teams_thread_limit</code>	381
3.3	Lock Routines	382
3.3.1	<code>omp_init_lock</code> and <code>omp_init_nest_lock</code>	384
3.3.2	<code>omp_init_lock_with_hint</code> and <code>omp_init_nest_lock_with_hint</code>	385
3.3.3	<code>omp_destroy_lock</code> and <code>omp_destroy_nest_lock</code>	387
3.3.4	<code>omp_set_lock</code> and <code>omp_set_nest_lock</code>	388
3.3.5	<code>omp_unset_lock</code> and <code>omp_unset_nest_lock</code>	389
3.3.6	<code>omp_test_lock</code> and <code>omp_test_nest_lock</code>	391
3.4	Timing Routines	393
3.4.1	<code>omp_get_wtime</code>	393
3.4.2	<code>omp_get_wtick</code>	394
3.5	Event Routine	395
3.5.1	<code>omp_fulfill_event</code>	395
3.6	Device Memory Routines	396
3.6.1	<code>omp_target_alloc</code>	396
3.6.2	<code>omp_target_free</code>	397
3.6.3	<code>omp_target_is_present</code>	398
3.6.4	<code>omp_target_memcpy</code>	399
3.6.5	<code>omp_target_memcpy_rect</code>	400
3.6.6	<code>omp_target_memcpy_async</code>	402
3.6.7	<code>omp_target_memcpy_rect_async</code>	404
3.6.8	<code>omp_target_associate_ptr</code>	406
3.6.9	<code>omp_get_mapped_ptr</code>	408
3.6.10	<code>omp_target_disassociate_ptr</code>	409
3.7	Interoperability Routines	410
3.7.1	<code>omp_get_num_interop_properties</code>	410

3.7.2	<code>omp_get_interop_property</code>	410
3.7.3	<code>omp_get_interop_property_name</code>	412
3.8	Memory Management Routines	413
3.8.1	Memory Management Types	413
3.8.2	<code>omp_init_allocator</code>	415
3.8.3	<code>omp_destroy_allocator</code>	417
3.8.4	<code>omp_set_default_allocator</code>	417
3.8.5	<code>omp_get_default_allocator</code>	418
3.8.6	<code>omp_alloc</code>	419
3.8.7	<code>omp_free</code>	420
3.9	Tool Control Routine	421
<b>4</b>	<b>OMPT Interface</b>	<b>425</b>
4.1	OMPT Interfaces Definitions	425
4.2	Activating a First-Party Tool	425
4.2.1	<code>ompt_start_tool</code>	425
4.2.2	Determining Whether a First-Party Tool Should be Initialized	427
4.2.3	Initializing a First-Party Tool	428
4.2.3.1	Binding Entry Points in the OMPT Callback Interface	429
4.2.4	Monitoring Activity on the Host with OMPT	430
4.2.5	Tracing Activity on Target Devices with OMPT	432
4.3	Finalizing a First-Party Tool	437
4.4	OMPT Data Types	437
4.4.1	Tool Initialization and Finalization	438
4.4.2	Callbacks	438
4.4.3	Tracing	439
4.4.3.1	Record Type	440
4.4.3.2	Native Record Kind	440
4.4.3.3	Native Record Abstract Type	440
4.4.3.4	Record Type	441
4.4.4	Miscellaneous Type Definitions	442
4.4.4.1	<code>ompt_callback_t</code>	442
4.4.4.2	<code>ompt_set_result_t</code>	443
4.4.4.3	<code>ompt_id_t</code>	444

4.4.4.4	<code>ompt_data_t</code>	444
4.4.4.5	<code>ompt_device_t</code>	445
4.4.4.6	<code>ompt_device_time_t</code>	445
4.4.4.7	<code>ompt_buffer_t</code>	446
4.4.4.8	<code>ompt_buffer_cursor_t</code>	446
4.4.4.9	<code>ompt_dependence_t</code>	446
4.4.4.10	<code>ompt_thread_t</code>	447
4.4.4.11	<code>ompt_scope_endpoint_t</code>	447
4.4.4.12	<code>ompt_dispatch_t</code>	447
4.4.4.13	<code>ompt_sync_region_t</code>	448
4.4.4.14	<code>ompt_target_data_op_t</code>	448
4.4.4.15	<code>ompt_work_t</code>	449
4.4.4.16	<code>ompt_mutex_t</code>	449
4.4.4.17	<code>ompt_native_mon_flag_t</code>	449
4.4.4.18	<code>ompt_task_flag_t</code>	450
4.4.4.19	<code>ompt_task_status_t</code>	451
4.4.4.20	<code>ompt_target_t</code>	451
4.4.4.21	<code>ompt_parallel_flag_t</code>	452
4.4.4.22	<code>ompt_target_map_flag_t</code>	453
4.4.4.23	<code>ompt_dependence_type_t</code>	453
4.4.4.24	<code>ompt_cancel_flag_t</code>	453
4.4.4.25	<code>ompt_hwid_t</code>	454
4.4.4.26	<code>ompt_state_t</code>	454
4.4.4.27	<code>ompt_frame_t</code>	457
4.4.4.28	<code>ompt_frame_flag_t</code>	458
4.4.4.29	<code>ompt_wait_id_t</code>	459
4.5	OMPT Tool Callback Signatures and Trace Records	459
4.5.1	Initialization and Finalization Callback Signature	460
4.5.1.1	<code>ompt_initialize_t</code>	460
4.5.1.2	<code>ompt_finalize_t</code>	461
4.5.2	Event Callback Signatures and Trace Records	461
4.5.2.1	<code>ompt_callback_thread_begin_t</code>	462
4.5.2.2	<code>ompt_callback_thread_end_t</code>	462

4.5.2.3	<code>ompt_callback_parallel_begin_t</code> . . . . .	463
4.5.2.4	<code>ompt_callback_parallel_end_t</code> . . . . .	465
4.5.2.5	<code>ompt_callback_work_t</code> . . . . .	466
4.5.2.6	<code>ompt_callback_dispatch_t</code> . . . . .	467
4.5.2.7	<code>ompt_callback_task_create_t</code> . . . . .	469
4.5.2.8	<code>ompt_callback_dependences_t</code> . . . . .	470
4.5.2.9	<code>ompt_callback_task_dependence_t</code> . . . . .	471
4.5.2.10	<code>ompt_callback_task_schedule_t</code> . . . . .	472
4.5.2.11	<code>ompt_callback_implicit_task_t</code> . . . . .	473
4.5.2.12	<code>ompt_callback_master_t</code> . . . . .	474
4.5.2.13	<code>ompt_callback_sync_region_t</code> . . . . .	475
4.5.2.14	<code>ompt_callback_mutex_acquire_t</code> . . . . .	477
4.5.2.15	<code>ompt_callback_mutex_t</code> . . . . .	478
4.5.2.16	<code>ompt_callback_nest_lock_t</code> . . . . .	479
4.5.2.17	<code>ompt_callback_flush_t</code> . . . . .	481
4.5.2.18	<code>ompt_callback_cancel_t</code> . . . . .	481
4.5.2.19	<code>ompt_callback_device_initialize_t</code> . . . . .	482
4.5.2.20	<code>ompt_callback_device_finalize_t</code> . . . . .	484
4.5.2.21	<code>ompt_callback_device_load_t</code> . . . . .	485
4.5.2.22	<code>ompt_callback_device_unload_t</code> . . . . .	486
4.5.2.23	<code>ompt_callback_buffer_request_t</code> . . . . .	486
4.5.2.24	<code>ompt_callback_buffer_complete_t</code> . . . . .	487
4.5.2.25	<code>ompt_callback_target_data_op_t</code> . . . . .	488
4.5.2.26	<code>ompt_callback_target_t</code> . . . . .	490
4.5.2.27	<code>ompt_callback_target_map_t</code> . . . . .	491
4.5.2.28	<code>ompt_callback_target_submit_t</code> . . . . .	493
4.5.2.29	<code>ompt_callback_control_tool_t</code> . . . . .	495
4.6	OMPT Runtime Entry Points for Tools . . . . .	496
4.6.1	Entry Points in the OMPT Callback Interface . . . . .	497
4.6.1.1	<code>ompt_enumerate_states_t</code> . . . . .	497
4.6.1.2	<code>ompt_enumerate_mutex_impls_t</code> . . . . .	498
4.6.1.3	<code>ompt_set_callback_t</code> . . . . .	499
4.6.1.4	<code>ompt_get_callback_t</code> . . . . .	501

4.6.1.5	<code>ompt_get_thread_data_t</code>	501
4.6.1.6	<code>ompt_get_num_procs_t</code>	502
4.6.1.7	<code>ompt_get_num_places_t</code>	503
4.6.1.8	<code>ompt_get_place_proc_ids_t</code>	503
4.6.1.9	<code>ompt_get_place_num_t</code>	504
4.6.1.10	<code>ompt_get_partition_place_nums_t</code>	505
4.6.1.11	<code>ompt_get_proc_id_t</code>	506
4.6.1.12	<code>ompt_get_state_t</code>	506
4.6.1.13	<code>ompt_get_parallel_info_t</code>	508
4.6.1.14	<code>ompt_get_task_info_t</code>	509
4.6.1.15	<code>ompt_get_task_memory_t</code>	511
4.6.1.16	<code>ompt_get_target_info_t</code>	512
4.6.1.17	<code>ompt_get_num_devices_t</code>	513
4.6.1.18	<code>ompt_get_unique_id_t</code>	514
4.6.1.19	<code>ompt_finalize_tool_t</code>	514
4.6.2	Entry Points in the OMPT Device Tracing Interface	515
4.6.2.1	<code>ompt_get_device_num_procs_t</code>	515
4.6.2.2	<code>ompt_get_device_time_t</code>	515
4.6.2.3	<code>ompt_translate_time_t</code>	516
4.6.2.4	<code>ompt_set_trace_ompt_t</code>	517
4.6.2.5	<code>ompt_set_trace_native_t</code>	518
4.6.2.6	<code>ompt_start_trace_t</code>	519
4.6.2.7	<code>ompt_pause_trace_t</code>	520
4.6.2.8	<code>ompt_flush_trace_t</code>	521
4.6.2.9	<code>ompt_stop_trace_t</code>	521
4.6.2.10	<code>ompt_advance_buffer_cursor_t</code>	522
4.6.2.11	<code>ompt_get_record_type_t</code>	523
4.6.2.12	<code>ompt_get_record_ompt_t</code>	524
4.6.2.13	<code>ompt_get_record_native_t</code>	525
4.6.2.14	<code>ompt_get_record_abstract_t</code>	526
4.6.3	Lookup Entry Points: <code>ompt_function_lookup_t</code>	526

## 5 OMPD Interface 529

5.1	OMPD Interfaces Definitions	530
-----	-----------------------------	-----



5.2	Activating an OMPD Tool . . . . .	530
5.2.1	Enabling the Runtime for OMPD . . . . .	530
5.2.2	<code>ompd_dll_locations</code> . . . . .	530
5.2.3	<code>ompd_dll_locations_valid</code> . . . . .	531
5.3	OMP Data Types . . . . .	532
5.3.1	Size Type . . . . .	532
5.3.2	Wait ID Type . . . . .	532
5.3.3	Basic Value Types . . . . .	532
5.3.4	Address Type . . . . .	533
5.3.5	Frame Information Type . . . . .	533
5.3.6	System Device Identifiers . . . . .	534
5.3.7	Native Thread Identifiers . . . . .	534
5.3.8	OMP Handle Types . . . . .	535
5.3.9	OMP Scope Types . . . . .	535
5.3.10	ICV ID Type . . . . .	536
5.3.11	Tool Context Types . . . . .	536
5.3.12	Return Code Types . . . . .	537
5.3.13	Primitive Type Sizes . . . . .	538
5.4	OMP Tool Callback Interface . . . . .	538
5.4.1	Memory Management of OMPD Library . . . . .	539
5.4.1.1	<code>ompd_callback_memory_alloc_fn_t</code> . . . . .	539
5.4.1.2	<code>ompd_callback_memory_free_fn_t</code> . . . . .	540
5.4.2	Context Management and Navigation . . . . .	540
5.4.2.1	<code>ompd_callback_get_thread_context_for_thread_id_fn_t</code> . . . . .	541
5.4.2.2	<code>ompd_callback_sizeof_fn_t</code> . . . . .	542
5.4.3	Accessing Memory in the OpenMP Program or Runtime . . . . .	543
5.4.3.1	<code>ompd_callback_symbol_addr_fn_t</code> . . . . .	543
5.4.3.2	<code>ompd_callback_memory_read_fn_t</code> . . . . .	544
5.4.3.3	<code>ompd_callback_memory_write_fn_t</code> . . . . .	546
5.4.4	Data Format Conversion: <code>ompd_callback_device_host_fn_t</code> . . . . .	547
5.4.5	Output: <code>ompd_callback_print_string_fn_t</code> . . . . .	548
5.4.6	The Callback Interface . . . . .	548

5.5	OMPD Tool Interface Routines . . . . .	550
5.5.1	Per OMPD Library Initialization and Finalization . . . . .	550
5.5.1.1	<code>ompd_initialize</code> . . . . .	550
5.5.1.2	<code>ompd_get_api_version</code> . . . . .	551
5.5.1.3	<code>ompd_get_version_string</code> . . . . .	552
5.5.1.4	<code>ompd_finalize</code> . . . . .	552
5.5.2	Per OpenMP Process Initialization and Finalization . . . . .	553
5.5.2.1	<code>ompd_process_initialize</code> . . . . .	553
5.5.2.2	<code>ompd_device_initialize</code> . . . . .	554
5.5.2.3	<code>ompd_rel_address_space_handle</code> . . . . .	555
5.5.3	Thread and Signal Safety . . . . .	556
5.5.4	Address Space Information . . . . .	556
5.5.4.1	<code>ompd_get_omp_version</code> . . . . .	556
5.5.4.2	<code>ompd_get_omp_version_string</code> . . . . .	556
5.5.5	Thread Handles . . . . .	557
5.5.5.1	<code>ompd_get_thread_in_parallel</code> . . . . .	557
5.5.5.2	<code>ompd_get_thread_handle</code> . . . . .	558
5.5.5.3	<code>ompd_rel_thread_handle</code> . . . . .	559
5.5.5.4	<code>ompd_thread_handle_compare</code> . . . . .	560
5.5.5.5	<code>ompd_get_thread_id</code> . . . . .	560
5.5.6	Parallel Region Handles . . . . .	561
5.5.6.1	<code>ompd_get_curr_parallel_handle</code> . . . . .	561
5.5.6.2	<code>ompd_get_enclosing_parallel_handle</code> . . . . .	562
5.5.6.3	<code>ompd_get_task_parallel_handle</code> . . . . .	563
5.5.6.4	<code>ompd_rel_parallel_handle</code> . . . . .	564
5.5.6.5	<code>ompd_parallel_handle_compare</code> . . . . .	564
5.5.7	Task Handles . . . . .	565
5.5.7.1	<code>ompd_get_curr_task_handle</code> . . . . .	565
5.5.7.2	<code>ompd_get_generating_task_handle</code> . . . . .	566
5.5.7.3	<code>ompd_get_scheduling_task_handle</code> . . . . .	567
5.5.7.4	<code>ompd_get_task_in_parallel</code> . . . . .	567
5.5.7.5	<code>ompd_rel_task_handle</code> . . . . .	568
5.5.7.6	<code>ompd_task_handle_compare</code> . . . . .	569

5.5.7.7	<code>ompd_get_task_function</code>	570
5.5.7.8	<code>ompd_get_task_frame</code>	570
5.5.7.9	<code>ompd_enumerate_states</code>	571
5.5.7.10	<code>ompd_get_state</code>	573
5.5.8	Display Control Variables	573
5.5.8.1	<code>ompd_get_display_control_vars</code>	573
5.5.8.2	<code>ompd_rel_display_control_vars</code>	574
5.5.9	Accessing Scope-Specific Information	575
5.5.9.1	<code>ompd_enumerate_icvs</code>	575
5.5.9.2	<code>ompd_get_icv_from_scope</code>	576
5.5.9.3	<code>ompd_get_icv_string_from_scope</code>	577
5.5.9.4	<code>ompd_get_tool_data</code>	579
5.6	Runtime Entry Points for OMPD	580
5.6.1	Beginning Parallel Regions	580
5.6.2	Ending Parallel Regions	580
5.6.3	Beginning Task Regions	581
5.6.4	Ending Task Regions	582
5.6.5	Beginning OpenMP Threads	582
5.6.6	Ending OpenMP Threads	583
5.6.7	Initializing OpenMP Devices	583
5.6.8	Finalizing OpenMP Devices	584

## **6 Environment Variables 585**

6.1	<code>OMP_SCHEDULE</code>	585
6.2	<code>OMP_NUM_THREADS</code>	586
6.3	<code>OMP_DYNAMIC</code>	587
6.4	<code>OMP_PROC_BIND</code>	588
6.5	<code>OMP_PLACES</code>	588
6.6	<code>OMP_STACKSIZE</code>	591
6.7	<code>OMP_WAIT_POLICY</code>	592
6.8	<code>OMP_MAX_ACTIVE_LEVELS</code>	592
6.9	<code>OMP_NESTED</code>	593
6.10	<code>OMP_THREAD_LIMIT</code>	593
6.11	<code>OMP_CANCELLATION</code>	594

6.12	<code>OMP_DISPLAY_ENV</code>	594
6.13	<code>OMP_DISPLAY_AFFINITY</code>	595
6.14	<code>OMP_AFFINITY_FORMAT</code>	596
6.15	<code>OMP_DEFAULT_DEVICE</code>	598
6.16	<code>OMP_MAX_TASK_PRIORITY</code>	598
6.17	<code>OMP_TARGET_OFFLOAD</code>	599
6.18	<code>OMP_TOOL</code>	599
6.19	<code>OMP_TOOL_LIBRARIES</code>	600
6.20	<code>OMP_DEBUG</code>	600
6.21	<code>OMP_ALLOCATOR</code>	601
6.22	<code>OMP_NUM_TEAMS</code>	601
6.23	<code>OMP_TEAMS_THREAD_LIMIT</code>	601
<b>A OpenMP Implementation-Defined Behaviors</b>		<b>603</b>
<b>B Features History</b>		<b>611</b>
B.1	Deprecated Features	611
B.2	Version 5.0 to 5.1 Differences	611
B.3	Version 4.5 to 5.0 Differences	613
B.4	Version 4.0 to 4.5 Differences	617
B.5	Version 3.1 to 4.0 Differences	618
B.6	Version 3.0 to 3.1 Differences	619
B.7	Version 2.5 to 3.0 Differences	620
<b>Index</b>		<b>623</b>

# List of Figures

2.1	Determining the <b>schedule</b> for a Worksharing-Loop . . . . .	115
4.1	First-Party Tool Activation Flow Chart . . . . .	427

# List of Tables

1.1	Map-Type Decay of Map Type Combinations . . . . .	16
2.1	ICV Initial Values . . . . .	68
2.2	Ways to Modify and to Retrieve ICV Values . . . . .	71
2.3	Scopes of ICVs . . . . .	74
2.4	ICV Override Relationships . . . . .	76
2.5	<b>schedule</b> Clause <i>kind</i> Values . . . . .	110
2.6	<b>schedule</b> Clause <i>modifier</i> Values . . . . .	111
2.7	<b>ompt_callback_task_create</b> callback flags evaluation . . . . .	146
2.8	Predefined Memory Spaces . . . . .	158
2.9	Allocator Traits . . . . .	159
2.10	Predefined Allocators . . . . .	161
2.11	Implicitly Declared C/C++ <i>reduction-identifiers</i> . . . . .	297
2.12	Implicitly Declared Fortran <i>reduction-identifiers</i> . . . . .	297
3.1	Required values of the <b>omp_interop_property</b> enum . . . . .	411
3.2	Standard Tool Control Commands . . . . .	423
4.1	OMPT Callback Interface Runtime Entry Point Names and Their Type Signatures . . . . .	431
4.2	Valid Return Codes of <b>ompt_set_callback</b> for Each Callback . . . . .	433
4.3	OMPT Tracing Interface Runtime Entry Point Names and Their Type Signatures . . . . .	435
5.1	Mapping of Scope Type and OMPD Handles . . . . .	536
6.1	Defined Abstract Names for <b>OMP_PLACES</b> . . . . .	589
6.2	Available Field Types for Formatting OpenMP Thread Affinity Information . . . . .	597

# 1 Introduction

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

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

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

The directives, library routines, environment variables, and tool support defined in this document allow users to create, to manage, to debug and to analyze parallel programs while permitting portability. The directives extend the C, C++ and Fortran base languages with single program multiple data (SPMD) constructs, tasking constructs, device constructs, worksharing constructs, and synchronization constructs, and they provide support for sharing, mapping and privatizing data. The functionality to control the runtime environment is provided by library routines and environment variables. Compilers that support the OpenMP API often include command line options to enable or to disable interpretation of OpenMP directives.

## 1.1 Scope

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

## 1.2 Glossary

### 1.2.1 Threading Concepts

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

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

**thread number** A number that the OpenMP implementation assigns to an OpenMP thread. For threads within the same team, zero identifies the master thread and consecutive numbers identify the other threads of this team.

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

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

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

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

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

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

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

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

### 1.2.2 OpenMP Language Terminology

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

COMMENT: See Section 1.7 on page 32 for a listing of current *base languages* for the OpenMP API.

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

**program order** An ordering of operations performed by the same thread as determined by the execution sequence of operations specified by the *base language*.

COMMENT: For C11 and C++11, *program order* corresponds to the *sequenced before* relation between operations performed by the same thread.



1	<b>structured block</b>	For C/C++, an executable statement, possibly compound, with a single entry at the top and a single exit at the bottom, or an OpenMP <i>construct</i> .
2		
3		For Fortran, a block of executable statements with a single entry at the top and a
4		single exit at the bottom, or an OpenMP <i>construct</i> .
5		COMMENT: See Section 2.1 on page 36 for restrictions on <i>structured</i>
6		<i>blocks</i> .
7	<b>structured block sequence</b>	A <i>structured block</i> , or, for C/C++, a sequence of two or more executable statements that together have a single entry at the top and a single exit at the bottom.
8		
9	<b>compilation unit</b>	For C/C++, a translation unit.
10		For Fortran, a program unit.
11	<b>enclosing context</b>	For C/C++, the innermost scope enclosing an OpenMP <i>directive</i> .
12		For Fortran, the innermost scoping unit enclosing an OpenMP <i>directive</i> .
13	<b>directive</b>	For C/C++, a <b>#pragma</b> , for C++, an attribute specifier and, for Fortran, a comment, that specifies <i>OpenMP program</i> behavior.
14		
15		COMMENT: See Section 2.1 on page 36 for a description of OpenMP
16		<i>directive</i> syntax.
17	<b>metadirective</b>	A <i>directive</i> that conditionally resolves to another <i>directive</i> at compile time.
18	<b>white space</b>	A non-empty sequence of space and/or horizontal tab characters.
19	<b>OpenMP program</b>	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
20		
21	<b>conforming program</b>	An <i>OpenMP program</i> that follows all rules and restrictions of the OpenMP specification.
22		
23	<b>declarative directive</b>	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. For C++, if a declarative directive applies to a function declaration or definition and it is specified with one or more C++ attribute specifiers, the specified attributes must be applied to a function as permitted by the base language. For Fortran, a declarative directive must appear after any <b>USE</b> statement, any <b>IMPORT</b> statement and any <b>IMPLICIT</b> statement in a declarative context.
24		
25		
26		
27		
28		
29		
30		
31	<b>executable directive</b>	An OpenMP <i>directive</i> that is not declarative. That is, it may be placed in an executable context.
32		
33	<b>stand-alone directive</b>	An OpenMP <i>executable directive</i> that has no associated user code except for that which appears in clauses in the directive.
34		

1	<b>construct</b>	An OpenMP <i>executable directive</i> (and for Fortran, the paired <b>end directive</b> , 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> .
2		
3		
4	<b>combined construct</b>	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.
5		
6		
7		
8	<b>composite construct</b>	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.
9		
10		
11		
12	<b>combined target construct</b>	A <i>combined construct</i> that is composed of a <b>target</b> construct along with another construct.
13		
14	<b>region</b>	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> . However, an <i>explicit task region</i> corresponding to a <i>task generating construct</i> is not part of the <i>region</i> of the <i>encountering thread</i> unless it is an <i>included task region</i> . The point where a <b>target</b> or <b>teams</b> 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 <b>target</b> or <b>teams</b> directive is not.
15		
16		
17		
18		
19		
20		
21		
22		
23		COMMENTS:
24		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.
25		
26		During the execution of an <i>OpenMP program</i> , a <i>construct</i> may give rise to many <i>regions</i> .
27		
28	<b>active parallel region</b>	A <b>parallel</b> <i>region</i> that is executed by a <i>team</i> consisting of more than one <i>thread</i> .
29	<b>inactive parallel region</b>	A <b>parallel</b> <i>region</i> that is executed by a <i>team</i> of only one <i>thread</i> .
30	<b>active target region</b>	A <b>target</b> <i>region</i> that is executed on a <i>device</i> other than the <i>device</i> that encountered the <b>target</b> <i>construct</i> .
31		
32	<b>inactive target region</b>	A <b>target</b> <i>region</i> that is executed on the same <i>device</i> that encountered the <b>target</b> <i>construct</i> .
33		

1	<b>sequential part</b>	All code encountered during the execution of an <i>initial task region</i> that is not part of
2		a <b>parallel</b> <i>region</i> corresponding to a <b>parallel</b> <i>construct</i> or a <b>task</b> <i>region</i>
3		corresponding to a <b>task</b> <i>construct</i> .
4		COMMENTS:
5		A <i>sequential part</i> is enclosed by an <i>implicit parallel region</i> .
6		Executable statements in called routines may be in both a <i>sequential part</i>
7		and any number of explicit <b>parallel</b> <i>regions</i> at different points in the
8		program execution.
9	<b>master thread</b>	An <i>OpenMP thread</i> that has <i>thread</i> number 0. A <i>master thread</i> may be an <i>initial</i>
10		<i>thread</i> or the <i>thread</i> that encounters a <b>parallel</b> <i>construct</i> , creates a <i>team</i> ,
11		generates a set of <i>implicit tasks</i> , and then executes one of those <i>tasks</i> as <i>thread</i>
12		number 0.
13	<b>parent thread</b>	The <i>thread</i> that encountered the <b>parallel</b> <i>construct</i> and generated a <b>parallel</b>
14		<i>region</i> is the <i>parent thread</i> of each of the <i>threads</i> in the <i>team</i> of that <b>parallel</b>
15		<i>region</i> . The <i>master thread</i> of a <b>parallel</b> <i>region</i> is the same <i>thread</i> as its <i>parent</i>
16		<i>thread</i> with respect to any resources associated with an <i>OpenMP thread</i> .
17	<b>child thread</b>	When a <i>thread</i> encounters a <b>parallel</b> <i>construct</i> , each of the <i>threads</i> in the
18		generated <b>parallel</b> <i>region</i> 's <i>team</i> are <i>child threads</i> of the encountering <i>thread</i> .
19		The <b>target</b> or <b>teams</b> <i>region</i> 's <i>initial thread</i> is not a <i>child thread</i> of the <i>thread</i> that
20		encountered the <b>target</b> or <b>teams</b> <i>construct</i> .
21	<b>ancestor thread</b>	For a given <i>thread</i> , its <i>parent thread</i> or one of its <i>parent thread</i> 's <i>ancestor threads</i> .
22	<b>descendent thread</b>	For a given <i>thread</i> , one of its <i>child threads</i> or one of its <i>child threads</i> ' <i>descendent</i>
23		<i>threads</i> .
24	<b>team</b>	A set of one or more <i>threads</i> participating in the execution of a <b>parallel</b> <i>region</i> .
25		COMMENTS:
26		For an <i>active parallel region</i> , the <i>team</i> comprises the <i>master thread</i> and at
27		least one additional <i>thread</i> .
28		For an <i>inactive parallel region</i> , the <i>team</i> comprises only the <i>master thread</i> .
29	<b>league</b>	The set of <i>teams</i> created by a <b>teams</b> <i>construct</i> .
30	<b>contention group</b>	An <i>initial thread</i> and its <i>descendent threads</i> .
31	<b>implicit parallel region</b>	An <i>inactive parallel region</i> that is not generated from a <b>parallel</b> <i>construct</i> .
32		<i>Implicit parallel regions</i> surround the whole <i>OpenMP program</i> , all <b>target</b> <i>regions</i> ,
33		and all <b>teams</b> <i>regions</i> .
34	<b>initial thread</b>	The <i>thread</i> that executes an <i>implicit parallel region</i> .

1	<b>initial team</b>	The <i>team</i> that comprises an <i>initial thread</i> executing an <i>implicit parallel region</i> .
2	<b>nested construct</b>	A <i>construct</i> (lexically) enclosed by another <i>construct</i> .
3	<b>closely nested construct</b>	A <i>construct</i> nested inside another <i>construct</i> with no other <i>construct</i> nested between
4		them.
5	<b>explicit region</b>	A <i>region</i> that corresponds to either a <i>construct</i> of the same name or a library routine
6		call that explicitly appears in the program.
7	<b>nested region</b>	A <i>region</i> (dynamically) enclosed by another <i>region</i> . That is, a <i>region</i> generated from
8		the execution of another <i>region</i> or one of its <i>nested regions</i> .
9		COMMENT: Some nestings are <i>conforming</i> and some are not. See
10		Section 2.23 on page 332 for the restrictions on nesting.
11	<b>closely nested region</b>	A <i>region nested</i> inside another <i>region</i> with no <b>parallel</b> <i>region nested</i> between
12		them.
13	<b>strictly nested region</b>	A <i>region nested</i> inside another <i>region</i> with no other <i>explicit region nested</i> between
14		them.
15	<b>all threads</b>	All OpenMP <i>threads</i> participating in the <i>OpenMP program</i> .
16	<b>current team</b>	All <i>threads</i> in the <i>team</i> executing the innermost enclosing <b>parallel</b> <i>region</i> .
17	<b>encountering thread</b>	For a given <i>region</i> , the <i>thread</i> that encounters the corresponding <i>construct</i> .
18	<b>all tasks</b>	All <i>tasks</i> participating in the <i>OpenMP program</i> .
19	<b>current team tasks</b>	All <i>tasks</i> encountered by the corresponding <i>team</i> . The <i>implicit tasks</i> constituting the
20		<b>parallel</b> <i>region</i> and any <i>descendent tasks</i> encountered during the execution of
21		these <i>implicit tasks</i> are included in this set of tasks.
22	<b>generating task</b>	For a given <i>region</i> , the task for which execution by a <i>thread</i> generated the <i>region</i> .
23	<b>binding thread set</b>	The set of <i>threads</i> that are affected by, or provide the context for, the execution of a
24		<i>region</i> .
25		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>
26		in a <i>contention group</i> , <i>all master threads</i> executing an enclosing <b>teams</b> <i>region</i> , the
27		<i>current team</i> , or the <i>encountering thread</i> .
28		COMMENT: The <i>binding thread set</i> for a particular <i>region</i> is described in
29		its corresponding subsection of this specification.
30	<b>binding task set</b>	The set of <i>tasks</i> that are affected by, or provide the context for, the execution of a
31		<i>region</i> .

1                                   The *binding task* set for a given *region* can be *all tasks*, the *current team tasks*, *all*  
2                                   *tasks of the current team that are generated in the region*, the *binding implicit task*, or  
3                                   the *generating task*.

4                                   COMMENT: The *binding task* set for a particular *region* (if applicable) is  
5                                   described in its corresponding subsection of this specification.

6       **binding region**   The enclosing *region* that determines the execution context and limits the scope of  
7                                   the effects of the bound *region* is called the *binding region*.

8                                   *Binding region* is not defined for *regions* for which the *binding thread* set is *all*  
9                                   *threads* or the *encountering thread*, nor is it defined for *regions* for which the *binding*  
10                                  *task set* is *all tasks*.

11                                  COMMENTS:

12                                  The *binding region* for an **ordered** *region* is the innermost enclosing  
13                                  *loop region*.

14                                  The *binding region* for a **taskwait** *region* is the innermost enclosing  
15                                  *task region*.

16                                  The *binding region* for a **cancel** *region* is the innermost enclosing  
17                                  *region* corresponding to the *construct-type-clause* of the **cancel**  
18                                  construct.

19                                  The *binding region* for a **cancellation point** *region* is the  
20                                  innermost enclosing *region* corresponding to the *construct-type-clause* of  
21                                  the **cancellation point** construct.

22                                  For all other *regions* for which the *binding thread set* is the *current team*  
23                                  or the *binding task set* is the *current team tasks*, the *binding region* is the  
24                                  innermost enclosing **parallel** *region*.

25                                  For *regions* for which the *binding task set* is the *generating task*, the  
26                                  *binding region* is the *region* of the *generating task*.

27                                  A **parallel** *region* need not be *active* nor explicit to be a *binding*  
28                                  *region*.

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

30                                  A *region* never binds to any *region* outside of the innermost enclosing  
31                                  **parallel** *region*.

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

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

1		For C/C++, <i>worksharing constructs</i> are <b>for</b> , <b>sections</b> , and <b>single</b> .
2		For Fortran, <i>worksharing constructs</i> are <b>do</b> , <b>sections</b> , <b>single</b> and
3		<b>workshare</b> .
4	<b>device construct</b>	An OpenMP <i>construct</i> that accepts the <b>device</b> clause.
5	<b>device routine</b>	A function (for C/C++ and Fortran) or subroutine (for Fortran) that can be executed
6		on a <i>target device</i> , as part of a <b>target</b> region.
7	<b>foreign runtime environment</b>	A runtime environment that exists outside the OpenMP runtime with which the
8		OpenMP implementation may interoperate.
9	<b>foreign execution context</b>	A context that is instantiated from a <i>foreign runtime environment</i> in order to facilitate
10		execution on a given device.
11	<b>foreign task</b>	A unit of work executed in a <i>foreign execution context</i> .
12	<b>indirect device invocation</b>	An indirect call to the device version of a procedure on a device other than the host
13		device, through a function pointer (C/C++), a pointer to a member function (C++) or
14		a procedure pointer (Fortran) that refers to the host version of the procedure.
15	<b>place</b>	An unordered set of <i>processors</i> on a device.
16	<b>place list</b>	The ordered list that describes all OpenMP <i>places</i> available to the execution
17		environment.
18	<b>place partition</b>	An ordered list that corresponds to a contiguous interval in the OpenMP <i>place list</i> . It
19		describes the <i>places</i> currently available to the execution environment for a given
20		parallel <i>region</i> .
21	<b>place number</b>	A number that uniquely identifies a <i>place</i> in the <i>place list</i> , with zero identifying the
22		first <i>place</i> in the <i>place list</i> , and each consecutive whole number identifying the next
23		<i>place</i> in the <i>place list</i> .
24	<b>thread affinity</b>	A binding of <i>threads</i> to <i>places</i> within the current <i>place partition</i> .
25	<b>SIMD instruction</b>	A single machine instruction that can operate on multiple data elements.
26	<b>SIMD lane</b>	A software or hardware mechanism capable of processing one data element from a
27		<i>SIMD instruction</i> .
28	<b>SIMD chunk</b>	A set of iterations executed concurrently, each by a <i>SIMD lane</i> , by a single <i>thread</i> by
29		means of <i>SIMD instructions</i> .
30	<b>memory</b>	A storage resource to store and to retrieve variables accessible by OpenMP threads.
31	<b>memory space</b>	A representation of storage resources from which <i>memory</i> can be allocated or
32		deallocated. More than one memory space may exist.
33	<b>memory allocator</b>	An OpenMP object that fulfills requests to allocate and to deallocate <i>memory</i> for
34		program variables from the storage resources of its associated <i>memory space</i> .

1           **handle**   An opaque reference that uniquely identifies an abstraction.

## 2           **1.2.3 Loop Terminology**

3           **loop-associated**   An OpenMP *executable* directive for which the associated user code must be a loop  
4           **directive**            nest that is a *structured block*.

5           **associated loop(s)**   The loop(s) controlled by a *loop-associated directive* of which some may be the  
6                                    generated loop(s) of enclosed loop transformation constructs.

7                                    COMMENT: Loop transformation constructs and other *loop-associated*  
8                                    *directives* that contain a **collapse** or an **ordered** (*n*) clause may have  
9                                    more than one *associated loop*.

10          **loop transformation**   A construct that is replaced by the loop(s) that result from applying the  
11           **construct**               transformation as defined by its directive to its associated loop(s).

12          **generated loop(s)**   The loop(s) that result from a *loop transformation construct* and replace the  
13                                    construct.

14           **SIMD loop**            A loop that includes at least one *SIMD chunk*.

15          **non-rectangular loop**   A loop nest for which the iteration count of a loop inside the loop nest is the not same  
16           **nest**                   for all occurrences of the loop in the loop nest.

17          **doacross loop nest**   A loop nest that has cross-iteration dependence. An iteration is dependent on one or  
18                                    more lexicographically earlier iterations.

19                                    COMMENT: The **ordered** clause parameter on a worksharing-loop  
20                                    directive identifies the loop(s) associated with the *doacross loop nest*.

## 21          **1.2.4 Synchronization Terminology**

22           **barrier**            A point in the execution of a program encountered by a *team of threads*, beyond  
23                                    which no *thread* in the team may execute until all *threads* in the *team* have reached  
24                                    the barrier and all *explicit tasks* generated by the *team* have executed to completion.  
25                                    If *cancellation* has been requested, threads may proceed to the end of the canceled  
26                                    *region* even if some threads in the team have not reached the *barrier*.

27           **cancellation**   An action that cancels (that is, aborts) an OpenMP *region* and causes executing  
28                                    *implicit* or *explicit* tasks to proceed to the end of the canceled *region*.

1	<b>cancellation point</b>	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 265.
5	<b>flush</b>	An operation that a <i>thread</i> performs to enforce consistency between its view and
6		other <i>threads</i> ' view of memory.
7	<b>flush property</b>	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	<b>strong flush</b>	A <i>flush</i> operation that has the <i>strong flush property</i> .
17	<b>release flush</b>	A <i>flush</i> operation that has the <i>release flush property</i> .
18	<b>acquire flush</b>	A <i>flush</i> operation that has the <i>acquire flush property</i> .
19	<b>atomic operation</b>	An operation that is specified by an <b>atomic</b> construct or is implicitly performed by
20		the implementation and that atomically accesses and/or modifies a specific storage
21		location.
22	<b>atomic read</b>	An <i>atomic operation</i> that is specified by an <b>atomic</b> construct on which the <b>read</b>
23		clause is present.
24	<b>atomic write</b>	An <i>atomic operation</i> that is specified by an <b>atomic</b> construct on which the <b>write</b>
25		clause is present.
26	<b>atomic update</b>	An <i>atomic operation</i> that is specified by an <b>atomic</b> construct on which the
27		<b>update</b> clause is present.
28	<b>atomic captured</b>	An <i>atomic operation</i> that is specified by an <b>atomic</b> construct on which the
29	<b>update</b>	<b>capture</b> clause is present.
30	<b>read-modify-write</b>	An <i>atomic operation</i> that reads and writes to a given storage location.
31		COMMENT: All <i>atomic update</i> and <i>atomic captured update</i> operations
32		are <i>read-modify-write</i> operations.



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

## 4 **1.2.5 Tasking Terminology**

- 5 **task** A specific instance of executable code and its data environment that the OpenMP  
6 implementation can schedule for execution by threads.
- 7 **task region** A *region* consisting of all code encountered during the execution of a *task*.  
8 COMMENT: A **parallel** *region* consists of one or more implicit *task*  
9 *regions*.
- 10 **implicit task** A *task* generated by an *implicit parallel region* or generated when a **parallel**  
11 *construct* is encountered during execution.
- 12 **binding implicit task** The *implicit task* of the current thread team assigned to the encountering thread.
- 13 **explicit task** A *task* that is not an *implicit task*.
- 14 **initial task** An *implicit task* associated with an *implicit parallel region*.
- 15 **current task** For a given *thread*, the *task* corresponding to the *task region* in which it is executing.
- 16 **child task** A *task* is a *child task* of its generating *task region*. A *child task region* is not part of  
17 its generating *task region*.
- 18 **sibling tasks** *Tasks* that are *child tasks* of the same *task region*.
- 19 **descendent task** A *task* that is the *child task* of a *task region* or of one of its *descendent task regions*.
- 20 **task completion** A condition that is satisfied when a thread reaches the end of the executable code that  
21 is associated with the *task* and any *allow-completion-event* that is created for the *task*  
22 is fulfilled.  
23 COMMENT: Completion of the *initial task* that is generated when the  
24 program begins occurs at program exit.

1	<b>task scheduling point</b>	A point during the execution of the current <i>task region</i> at which it can be suspended to be resumed later; or the point of <i>task completion</i> , after which the executing thread may switch to a different <i>task region</i> .
2		
3		
4		COMMENT: For a list of <i>task scheduling points</i> , see Section 2.13.6 on
5		page 156.
6	<b>task switching</b>	The act of a <i>thread</i> switching from the execution of one <i>task</i> to another <i>task</i> .
7	<b>tied task</b>	A <i>task</i> that, when its <i>task region</i> is suspended, can be resumed only by the same <i>thread</i> that suspended it. That is, the <i>task</i> is tied to that <i>thread</i> .
8		
9	<b>untied task</b>	A <i>task</i> that, when its <i>task region</i> is suspended, can be resumed by any <i>thread</i> in the team. That is, the <i>task</i> is not tied to any <i>thread</i> .
10		
11	<b>undelayed task</b>	A <i>task</i> for which execution is not deferred with respect to its generating <i>task region</i> . That is, its generating <i>task region</i> is suspended until execution of the structured block associated with the <i>undelayed task</i> is completed.
12		
13		
14	<b>included task</b>	A <i>task</i> for which execution is sequentially included in the generating <i>task region</i> . That is, an <i>included task</i> is <i>undelayed</i> and executed by the <i>encountering thread</i> .
15		
16	<b>merged task</b>	A <i>task</i> for which the <i>data environment</i> , inclusive of ICVs, is the same as that of its generating <i>task region</i> .
17		
18	<b>mergeable task</b>	A <i>task</i> that may be a <i>merged task</i> if it is an <i>undelayed task</i> or an <i>included task</i> .
19	<b>final task</b>	A <i>task</i> that forces all of its <i>child tasks</i> to become <i>final</i> and <i>included tasks</i> .
20	<b>task dependence</b>	An ordering relation between two <i>sibling tasks</i> : the <i>dependent task</i> and a previously generated <i>predecessor task</i> . The <i>task dependence</i> is fulfilled when the <i>predecessor task</i> has completed.
21		
22		
23	<b>dependent task</b>	A <i>task</i> that because of a <i>task dependence</i> cannot be executed until its <i>predecessor tasks</i> have completed.
24		
25	<b>mutually exclusive tasks</b>	<i>Tasks</i> that may be executed in any order, but not at the same time.
26	<b>predecessor task</b>	A <i>task</i> that must complete before its <i>dependent tasks</i> can be executed.
27	<b>task synchronization construct</b>	A <b>taskwait</b> , <b>taskgroup</b> , or <b>barrier</b> <i>construct</i> .
28	<b>task generating construct</b>	A <i>construct</i> that generates one or more <i>explicit tasks</i> .
29	<b>target task</b>	A <i>mergeable</i> and <i>untied task</i> that is generated by a <b>target</b> , <b>target enter data</b> , <b>target exit data</b> , or <b>target update</b> <i>construct</i> .
30		
31	<b>taskgroup set</b>	A set of tasks that are logically grouped by a <b>taskgroup</b> <i>region</i> .

## 1.2.6 Data Terminology

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

COMMENT: An array element or structure element is a variable that is part of another variable.

**scalar variable** For C/C++, a scalar variable, as defined by the base language.

For Fortran, a scalar variable with intrinsic type, as defined by the base language, excluding character type.

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

**array section** A designated subset of the elements of an array that is specified using a subscript notation that can select more than one element.

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

**shape-operator** For C/C++, an array shaping operator that reinterprets a pointer expression as an array with one or more specified dimensions.

**implicit array** For C/C++, the set of array elements of non-array type  $T$  that may be accessed by applying a sequence of [] operators to a given pointer that is either a pointer to type  $T$  or a pointer to a multidimensional array of elements of type  $T$ .

For Fortran, the set of array elements for a given array pointer.

COMMENT: For C/C++, the implicit array for pointer  $p$  with type  $T$   $(*)[10]$  consists of all accessible elements  $p[i][j]$ , for all  $i$  and  $j=0..9$ .

**base pointer** For C/C++, an lvalue pointer expression that is used by a given lvalue expression or array section to refer indirectly to its storage, where the lvalue expression or array section is part of the implicit array for that lvalue pointer expression.

For Fortran, a data pointer that appears last in the designator for a given variable or array section, where the variable or array section is part of the pointer target for that data pointer.

COMMENT: For the array section  $(*p0).x0[k1].p1->p2[k2].x1[k3].x2[4][0:n]$ , where identifiers  $p_i$  have a pointer type declaration and identifiers  $x_i$  have an array type declaration, the *base pointer* is:  $(*p0).x0[k1].p1->p2$ .

**named pointer** For C/C++, the *base pointer* of a given lvalue expression or array section, or the *base pointer* of one of its *named pointers*.

For Fortran, the *base pointer* of a given variable or array section, or the *base pointer* of one of its *named pointers*.

1 COMMENT: For the array section  
2  $(*p0).x0[k1].p1 \rightarrow p2[k2].x1[k3].x2[4][0:n]$ , where identifiers  $pi$  have a  
3 pointer type declaration and identifiers  $xi$  have an array type declaration,  
4 the *named pointers* are:  $p0$ ,  $(*p0).x0[k1].p1$ , and  $(*p0).x0[k1].p1 \rightarrow p2$ .

5 **containing array** For C/C++, a non-subscripted array (a *containing array*) that appears in a given  
6 lvalue expression or array section, where the lvalue expression or array section is part  
7 of that *containing array*.

8 For Fortran, an array (a *containing array*) without the **POINTER** attribute and  
9 without a subscript list that appears in the designator of a given variable or array  
10 section, where the variable or array section is part of that *containing array*.

11 COMMENT: For the array section  
12  $(*p0).x0[k1].p1 \rightarrow p2[k2].x1[k3].x2[4][0:n]$ , where identifiers  $pi$  have a  
13 pointer type declaration and identifiers  $xi$  have an array type declaration,  
14 the *containing arrays* are:  $(*p0).x0[k1].p1 \rightarrow p2[k2].x1$  and  
15  $(*p0).x0[k1].p1 \rightarrow p2[k2].x1[k3].x2$ .

16 **base array** For C/C++, a *containing array* of a given lvalue expression or array section that does  
17 not appear in the expression of any of its other *containing arrays*.

18 For Fortran, a *containing array* of a given variable or array section that does not  
19 appear in the designator of any of its other *containing arrays*.

20 COMMENT: For the array section  
21  $(*p0).x0[k1].p1 \rightarrow p2[k2].x1[k3].x2[4][0:n]$ , where identifiers  $pi$  have a  
22 pointer type declaration and identifiers  $xi$  have an array type declaration,  
23 the *base array* is:  $(*p0).x0[k1].p1 \rightarrow p2[k2].x1[k3].x2$ .

24 **named array** For C/C++, a *containing array* of a given lvalue expression or array section, or a  
25 *containing array* of one of its *named pointers*.

26 For Fortran, a *containing array* of a given variable or array section, or a *containing*  
27 *array* of one of its *named pointers*.

28 COMMENT: For the array section  
29  $(*p0).x0[k1].p1 \rightarrow p2[k2].x1[k3].x2[4][0:n]$ , where identifiers  $pi$  have a  
30 pointer type declaration and identifiers  $xi$  have an array type declaration,  
31 the *named arrays* are:  $(*p0).x0$ ,  $(*p0).x0[k1].p1 \rightarrow p2[k2].x1$ , and  
32  $(*p0).x0[k1].p1 \rightarrow p2[k2].x1[k3].x2$ .

33 **base expression** The *base array* of a given array section or array element, if it exists; otherwise, the  
34 *base pointer* of the array section or array element.

35 COMMENT: For the array section  
36  $(*p0).x0[k1].p1 \rightarrow p2[k2].x1[k3].x2[4][0:n]$ , where identifiers  $pi$  have a  
37 pointer type declaration and identifiers  $xi$  have an array type declaration,  
38 the *base expression* is:  $(*p0).x0[k1].p1 \rightarrow p2[k2].x1[k3].x2$ .

1		More examples for C/C++:
2		• The <i>base expression</i> for $x[i]$ and for $x[i:n]$ is $x$ , if $x$ is an array or pointer.
3		• The <i>base expression</i> for $x[5][i]$ and for $x[5][i:n]$ is $x$ , if $x$ is a pointer to
4		an array or $x$ is 2-dimensional array.
5		• The <i>base expression</i> for $y[5][i]$ and for $y[5][i:n]$ is $y[5]$ , if $y$ is an array
6		of pointers or $y$ is a pointer to a pointer.
7		Examples for Fortran:
8		• The <i>base expression</i> for $x(i)$ and for $x(i:j)$ is $x$ .
9	<b>attached pointer</b>	A pointer variable in a device data environment to which the effect of a <b>map</b> clause
10		assigns the address of an object, minus some offset, that is created in the device data
11		environment. The pointer is an attached pointer for the remainder of its lifetime in
12		the device data environment.
13	<b>simply contiguous</b>	An array section that statically can be determined to have contiguous storage or that,
14	<b>array section</b>	in Fortran, has the <b>CONTIGUOUS</b> attribute.
15	<b>structure</b>	A structure is a variable that contains one or more variables.
16		For C/C++: Implemented using struct types.
17		For C++: Implemented using class types.
18		For Fortran: Implemented using derived types.
19	<b>private variable</b>	With respect to a given set of <i>task regions</i> or <i>SIMD lanes</i> that bind to the same
20		<b>parallel region</b> , a <i>variable</i> for which the name provides access to a different
21		block of storage for each <i>task region</i> or <i>SIMD lane</i> .
22		A <i>variable</i> that is part of another variable (as an array or structure element) cannot be
23		made private independently of other components.
24	<b>shared variable</b>	With respect to a given set of <i>task regions</i> that bind to the same <b>parallel region</b> , a
25		<i>variable</i> for which the name provides access to the same block of storage for each
26		<i>task region</i> .
27		A <i>variable</i> that is part of another variable (as an array or structure element) cannot be
28		<i>shared</i> independently of the other components, except for static data members of
29		C++ classes.
30	<b>threadprivate variable</b>	A <i>variable</i> that is replicated, one instance per <i>thread</i> , by the OpenMP
31		implementation. Its name then provides access to a different block of storage for each
32		<i>thread</i> .

1 A *variable* that is part of another variable (as an array or structure element) cannot be  
2 made *threadprivate* independently of the other components, except for static data  
3 members of C++ classes.

4 **threadprivate memory** The set of *threadprivate variables* associated with each *thread*.

5 **data environment** The *variables* associated with the execution of a given *region*.

6 **device data environment** The initial *data environment* associated with a device.

7 **device address** An *implementation-defined* reference to an address in a *device data environment*.

8 **device pointer** A *variable* that contains a *device address*.

9 **mapped variable** An original *variable* in a *data environment* with a corresponding *variable* in a device  
10 *data environment*.

11 COMMENT: The original and corresponding *variables* may share storage.

12 **map-type decay** The process used to determine the final map type when mapping a variable with a  
13 user defined mapper. Table 1.1 shows the final map type that the combination of the  
14 two map types determines.

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

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

20 For C, the type must be a complete type.

21 For C++, the type must be a complete type.

22 In addition, for class types:

23 • All member functions accessed in any **target** region must appear in a declare  
24 target directive.

25 For Fortran, no restrictions on the type except that for derived types:

**TABLE 1.1:** Map-Type Decay of Map Type Combinations

	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

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

## 14                   **1.2.7 Implementation Terminology**

- 15           **supporting  $n$  active** Implies allowing an *active parallel region* to be enclosed by  $n-1$  *active parallel*  
16           **levels of parallelism** *regions*.
- 17                   **supporting the** Supporting at least one active level of parallelism.  
18                   **OpenMP API**
- 18           **supporting nested** Supporting more than one active level of parallelism.  
19                   **parallelism**
- 19                   **internal control** A conceptual variable that specifies runtime behavior of a set of *threads* or *tasks* in  
20                   **variable** an *OpenMP program*.
- 21                                COMMENT: The acronym ICV is used interchangeably with the term  
22                                *internal control variable* in the remainder of this specification.
- 23                   **compliant** An implementation of the OpenMP specification that compiles and executes any  
24                   **implementation** *conforming program* as defined by the specification.
- 25                                COMMENT: A *compliant implementation* may exhibit *unspecified*  
26                                *behavior* when compiling or executing a *non-conforming program*.
- 27           **unspecified behavior** A behavior or result that is not specified by the OpenMP specification or not known  
28                                prior to the compilation or execution of an *OpenMP program*.
- 29                                Such *unspecified behavior* may result from:

- Issues documented by the OpenMP specification as having *unspecified behavior*.
- A *non-conforming program*.
- A *conforming program* exhibiting an *implementation-defined* behavior.

**implementation defined** Behavior that must be documented by the implementation, and is allowed to vary among different *compliant implementations*. An implementation is allowed to define this behavior as *unspecified*.

COMMENT: All features that have *implementation-defined* behavior are documented in Appendix A.

**deprecated** For a construct, clause, or other feature, the property that it is normative in the current specification but is considered obsolescent and will be removed in the future.

## 1.2.8 Tool Terminology

**tool** Code that can observe and/or modify the execution of an application.

**first-party tool** A tool that executes in the address space of the program that it is monitoring.

**third-party tool** A tool that executes as a separate process from the process that it is monitoring and potentially controlling.

**activated tool** A *first-party tool* that successfully completed its initialization.

**event** A point of interest in the execution of a thread.

**native thread** A thread defined by an underlying thread implementation.

**tool callback** A function that a tool provides to an OpenMP implementation to invoke when an associated event occurs.

**registering a callback** Providing a *tool callback* to an OpenMP implementation.

**dispatching a callback at an event** Processing a callback when an associated *event* occurs in a manner consistent with the return code provided when a *first-party tool* registered the callback.

**thread state** An enumeration type that describes the current OpenMP activity of a *thread*. A *thread* can be in only one state at any time.

**wait identifier** A unique opaque handle associated with each data object (for example, a lock) used by the OpenMP runtime to enforce mutual exclusion that may cause a thread to wait actively or passively.

**frame** A storage area on a thread's stack associated with a procedure invocation. A frame includes space for one or more saved registers and often also includes space for saved arguments, local variables, and padding for alignment.



1	<b>canonical frame</b>	An address associated with a procedure <i>frame</i> on a call stack that was the value of the
2	<b>address</b>	stack pointer immediately prior to calling the procedure for which the invocation is
3		represented by the frame.
4	<b>runtime entry point</b>	A function interface provided by an OpenMP runtime for use by a tool. A runtime
5		entry point is typically not associated with a global function symbol.
6	<b>trace record</b>	A data structure in which to store information associated with an occurrence of an
7		<i>event</i> .
8	<b>native trace record</b>	A <i>trace record</i> for an OpenMP device that is in a device-specific format.
9	<b>signal</b>	A software interrupt delivered to a <i>thread</i> .
10	<b>signal handler</b>	A function called asynchronously when a <i>signal</i> is delivered to a <i>thread</i> .
11	<b>async signal safe</b>	The guarantee that interruption by <i>signal</i> delivery will not interfere with a set of
12		operations. An async signal safe <i>runtime entry point</i> is safe to call from a <i>signal</i>
13		<i>handler</i> .
14	<b>code block</b>	A contiguous region of memory that contains code of an OpenMP program to be
15		executed on a device.
16	<b>OMPT</b>	An interface that helps a <i>first-party tool</i> monitor the execution of an OpenMP
17		program.
18	<b>OMPT interface state</b>	A state that indicates the permitted interactions between a first-party tool and the
19		OpenMP implementation.
20	<b>OMPT active</b>	An <i>OMPT interface state</i> in which the OpenMP implementation is prepared to accept
21		runtime calls from a <i>first party tool</i> and it dispatches any registered callbacks and in
22		which a first-party tool can invoke <i>runtime entry points</i> if not otherwise restricted.
23	<b>OMPT pending</b>	An <i>OMPT interface state</i> in which the OpenMP implementation can only call
24		functions to initialize a <i>first party tool</i> and in which a <i>first-party tool</i> cannot invoke
25		<i>runtime entry points</i> .
26	<b>OMPT inactive</b>	An <i>OMPT interface state</i> in which the OpenMP implementation will not make any
27		callbacks and in which a <i>first-party tool</i> cannot invoke <i>runtime entry points</i> .
28	<b>OMPD</b>	An interface that helps a <i>third-party tool</i> inspect the OpenMP state of a program that
29		has begun execution.
30	<b>OMPD library</b>	A dynamically loadable library that implements the <i>OMPD</i> interface.
31	<b>image file</b>	An executable or shared library.
32	<b>address space</b>	A collection of logical, virtual, or physical memory address ranges that contain code,
33		stack, and/or data. Address ranges within an address space need not be contiguous.
34		An address space consists of one or more <i>segments</i> .

1	<b>segment</b>	A portion of an address space associated with a set of address ranges.
2	<b>OpenMP architecture</b>	The architecture on which an OpenMP <i>region</i> executes.
3	<b>tool architecture</b>	The architecture on which an <i>OMPD</i> tool executes.
4	<b>OpenMP process</b>	A collection of one or more <i>threads</i> and <i>address spaces</i> . A process may contain
5		<i>threads</i> and <i>address spaces</i> for multiple <i>OpenMP architectures</i> . At least one thread
6		in an OpenMP process is an OpenMP <i>thread</i> . A process may be live or a core file.
7	<b>address space handle</b>	A <i>handle</i> that refers to an <i>address space</i> within an OpenMP process.
8	<b>thread handle</b>	A <i>handle</i> that refers to an OpenMP <i>thread</i> .
9	<b>parallel handle</b>	A <i>handle</i> that refers to an OpenMP parallel <i>region</i> .
10	<b>task handle</b>	A <i>handle</i> that refers to an OpenMP task <i>region</i> .
11	<b>descendent handle</b>	An output <i>handle</i> that is returned from the <i>OMPD</i> library in a function that accepts
12		an input <i>handle</i> : the output <i>handle</i> is a descendent of the input <i>handle</i> .
13	<b>ancestor handle</b>	An input <i>handle</i> that is passed to the <i>OMPD</i> library in a function that returns an
14		output <i>handle</i> : the input <i>handle</i> is an ancestor of the output <i>handle</i> . For a given
15		<i>handle</i> , the ancestors of the <i>handle</i> are also the ancestors of the handle's descendent.
16		COMMENT: A <i>handle</i> cannot be used by the tool in an <i>OMPD</i> call if any
17		ancestor of the <i>handle</i> has been released, except for <i>OMPD</i> calls that
18		release the <i>handle</i> .
19	<b>tool context</b>	An opaque reference provided by a tool to an <i>OMPD</i> library. A <i>tool context</i> uniquely
20		identifies an abstraction.
21	<b>address space context</b>	A <i>tool context</i> that refers to an <i>address space</i> within a process.
22	<b>thread context</b>	A <i>tool context</i> that refers to a <i>native thread</i> .
23	<b>native thread identifier</b>	An identifier for a native thread defined by a thread implementation.
24	<b>preprocessed code</b>	For C/C++, a sequence of preprocessing tokens that result from the first six phases of
25		translation, as defined by the base language.

## 1.3 Execution Model

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

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

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

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

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

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

The **teams** construct creates a *league of teams*, where each team is an initial team that comprises an initial thread that executes the **teams** region. Each initial thread executes sequentially, as if the code encountered is part of an initial task region that is generated by an implicit parallel region associated with each team. Whether the initial threads concurrently execute the **teams** region is unspecified, and a program that relies on their concurrent execution for the purposes of synchronization may deadlock.

If a construct creates a data environment, the data environment is created at the time the construct is

1 encountered. The description of a construct defines whether it creates a data environment.

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 the **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 processing with the assistance of OpenMP synchronization constructs or library routines.  
31 For the case where each thread accesses a different file, no synchronization by the programmer is  
32 necessary.

33 All concurrency semantics defined by the base language with respect to threads of execution apply  
34 to OpenMP threads, unless specified otherwise.

# 1.4 Memory Model

## 1.4.1 Structure of the OpenMP Memory Model

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

A directive that accepts data-sharing attribute clauses determines two kinds of access to variables used in the directive's associated structured block: shared and private. Each variable referenced in the structured block has an original variable, which is the variable by the same name that exists in the program immediately outside the construct. Each reference to a shared variable in the structured block becomes a reference to the original variable. For each private variable referenced in the structured block, a new version of the original variable (of the same type and size) is created in memory for each task or SIMD lane that contains code associated with the directive. Creation of the new version does not alter the value of the original variable. However, the impact of attempts to access the original variable during the region corresponding to the directive is unspecified; see Section 2.22.4.3 on page 287 for additional details. References to a private variable in the structured block refer to the private version of the original variable for the current task or SIMD lane. The relationship between the value of the original variable and the initial or final value of the private version depends on the exact clause that specifies it. Details of this issue, as well as other issues with privatization, are provided in Section 2.22 on page 271.

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

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

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

A private variable in a task region that subsequently generates an inner nested **parallel** region is permitted to be made shared by implicit tasks in the inner **parallel** region. A private variable in

1 a task region can also be shared by an explicit task region generated during its execution. However,  
2 it is the programmer's responsibility to ensure through synchronization that the lifetime of the  
3 variable does not end before completion of the explicit task region sharing it. Any other access by  
4 one task to the private variables of another task results in unspecified behavior.

## 5 1.4.2 Device Data Environments

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

11 When an original variable is mapped to a device data environment and a corresponding variable is  
12 not present in the device data environment, a new corresponding variable (of the same type and size  
13 as the original variable) is created in the device data environment. Conversely, the original variable  
14 becomes the new variable's corresponding variable in the device data environment of the device  
15 that performs the mapping operation.

16 The corresponding variable in the device data environment may share storage with the original  
17 variable. Writes to the corresponding variable may alter the value of the original variable. The  
18 impact of this possibility on memory consistency is discussed in Section 1.4.6 on page 29. When a  
19 task executes in the context of a device data environment, references to the original variable refer to  
20 the corresponding variable in the device data environment. If an original variable is not currently  
21 mapped and a corresponding variable does not exist in the device data environment then accesses to  
22 the original variable result in unspecified behavior unless the **unified\_shared\_memory**  
23 clause is specified on a **requires** directive for the compilation unit.

24 The relationship between the value of the original variable and the initial or final value of the  
25 corresponding variable depends on the *map-type*. Details of this issue, as well as other issues with  
26 mapping a variable, are provided in Section 2.22.7.1 on page 318.

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

29 If a variable has a corresponding variable with which it does not share storage, a write to a storage  
30 location designated by the variable causes the value at the corresponding storage location to  
31 become undefined.

## 32 1.4.3 Memory Management

33 The host device, and target devices that an implementation may support, have attached storage  
34 resources where program variables are stored. These resources can have different traits. A memory  
35 space in an OpenMP program represents a set of these storage resources. Memory spaces are  
36 defined according to a set of traits, and a single resource may be exposed as multiple memory

1 spaces with different traits or may be part of multiple memory spaces. In any device, at least one  
2 memory space is guaranteed to exist.

3 An OpenMP program can use a *memory allocator* to allocate *memory* in which to store variables.  
4 This *memory* will be allocated from the storage resources of the *memory space* associated with the  
5 memory allocator. Memory allocators are also used to deallocate previously allocated *memory*.  
6 When an OpenMP memory allocator is not used to allocate memory, OpenMP does not prescribe  
7 the storage resource for the allocation; the memory for the variables may be allocated in any storage  
8 resource.

## 9 1.4.4 The Flush Operation

10 The memory model has relaxed-consistency because a thread's temporary view of memory is not  
11 required to be consistent with memory at all times. A value written to a variable can remain in the  
12 thread's temporary view until it is forced to memory at a later time. Likewise, a read from a  
13 variable may retrieve the value from the thread's temporary view, unless it is forced to read from  
14 memory. OpenMP flush operations are used to enforce consistency between a thread's temporary  
15 view of memory and memory, or between multiple threads' view of memory.

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

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

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



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

5 If a flush operation is a release flush or acquire flush, it can enforce consistency between the views  
6 of memory of two synchronizing threads. A release flush guarantees that any prior operation that  
7 writes or reads a shared variable will appear to be completed before any operation that writes or  
8 reads the same shared variable and follows an acquire flush with which the release flush  
9 synchronizes (see Section 1.4.5 on page 27 for more details on flush synchronization). A release  
10 flush will propagate the values of all shared variables in its temporary view to memory prior to the  
11 thread performing any subsequent atomic operation that may establish a synchronization. An  
12 acquire flush will discard any value of a shared variable in its temporary view to which the thread  
13 has not written since last performing a release flush, so that it may subsequently read a value  
14 propagated by a release flush that synchronizes with it. Therefore, release and acquire flushes may  
15 also be used to guarantee that a value written to a variable by one thread may be read by a second  
16 thread. To accomplish this, the programmer must ensure that the second thread has not written to  
17 the variable since its last acquire flush, and that the following sequence of events happen in this  
18 specific order:

- 19 1. The value is written to the variable by the first thread;
- 20 2. The first thread performs a release flush;
- 21 3. The second thread performs an acquire flush; and
- 22 4. The value is read from the variable by the second thread.

23  
24 **Note** – OpenMP synchronization operations, described in Section 2.20 on page 226 and in  
25 Section 3.3 on page 382, are recommended for enforcing this order. Synchronization through  
26 variables is possible but is not recommended because the proper timing of flushes is difficult.  
27

28 The flush properties that define whether a flush operation is a strong flush, a release flush, or an  
29 acquire flush are not mutually disjoint. A flush operation may be a strong flush and a release flush;  
30 it may be a strong flush and an acquire flush; it may be a release flush and an acquire flush; or it  
31 may be all three.

## 32 **1.4.5 Flush Synchronization and *Happens Before***

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

1 A release flush has one or more associated *release sequences* that define the set of modifications  
2 that may be used to establish a synchronization. A release sequence starts with an atomic operation  
3 that follows the release flush and modifies a shared variable and additionally includes any  
4 read-modify-write atomic operations that read a value taken from some modification in the release  
5 sequence. The following rules determine the atomic operation that starts an associated release  
6 sequence.

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

15 An acquire flush is associated with one or more prior atomic operations that read a shared variable  
16 and that may be used to establish a synchronization. The following rules determine the associated  
17 atomic operation that may establish a synchronization.

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

26 A release flush synchronizes with an acquire flush if an atomic operation associated with the  
27 acquire flush reads a value written by a modification from a release sequence associated with the  
28 release flush.

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

- 31 1. *X* and *Y* are performed by the same thread, and *X* precedes *Y* in the thread's program order;
- 32 2. *X* synchronizes with *Y* according to the flush synchronization conditions explained above or  
33 according to the base language's definition of *synchronizes with*, if such a definition exists; or
- 34 3. There exists another operation *Z*, such that *X* simply happens before *Z* and *Z* simply happens  
35 before *Y*.

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

- 1 1.  $X$  happens before  $Y$  according to the base language's definition of *happens before*, if such a  
2 definition exists; or
- 3 2.  $X$  simply happens before  $Y$ .

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

## 6 1.4.6 OpenMP Memory Consistency

7 The following rules guarantee the observable completion order of memory operations, as seen by  
8 all threads.

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

22 The flush operation can be specified using the **flush** directive, and is also implied at various  
23 locations in an OpenMP program: see Section 2.20.8 on page 245 for details.

24 ▼  
25 **Note** – Since flush operations by themselves cannot prevent data races, explicit flush operations are  
26 only useful in combination with non-sequentially consistent atomic directives.  
27 ▲

28 OpenMP programs that:

- 29 • Do not use non-sequentially consistent atomic directives;
- 30 • Do not rely on the accuracy of a *false* result from **omp\_test\_lock** and  
31 **omp\_test\_nest\_lock**; and
- 32 • Correctly avoid data races as required in Section 1.4.1 on page 24,

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

## 4 1.5 Tool Interfaces

5 The OpenMP API includes two tool interfaces, OMPT and OMPD, to enable development of  
6 high-quality, portable, tools that support monitoring, performance, or correctness analysis and  
7 debugging of OpenMP programs developed using any implementation of the OpenMP API,

8 An implementation of the OpenMP API may differ from the abstract execution model described by  
9 its specification. The ability of tools that use the OMPT or OMPD interfaces to observe such  
10 differences does not constrain implementations of the OpenMP API in any way.

### 11 1.5.1 OMPT

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

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

21 OpenMP implementations may differ with respect to the *thread states* that they support, the mutual  
22 exclusion implementations that they employ, and the OpenMP events for which tool callbacks are  
23 invoked. For some OpenMP events, OpenMP implementations must guarantee that a registered  
24 callback will be invoked for each occurrence of the event. For other OpenMP events, OpenMP  
25 implementations are permitted to invoke a registered callback for some or no occurrences of the  
26 event; for such OpenMP events, however, OpenMP implementations are encouraged to invoke tool  
27 callbacks on as many occurrences of the event as is practical. Section 4.2.4 specifies the subset of  
28 OMPT callbacks that an OpenMP implementation must support for a minimal implementation of  
29 the OMPT interface.

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

## 1.5.2 OMPD

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

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

Section 5 describes OMPD in detail.

## 1.6 OpenMP Compliance

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

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

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

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

## 1.7 Normative References

- ISO/IEC 9899:1990, *Information Technology - Programming Languages - C*.  
This OpenMP API specification refers to ISO/IEC 9899:1990 as C90.
- ISO/IEC 9899:1999, *Information Technology - Programming Languages - C*.  
This OpenMP API specification refers to ISO/IEC 9899:1999 as C99.
- ISO/IEC 9899:2011, *Information Technology - Programming Languages - C*.  
This OpenMP API specification refers to ISO/IEC 9899:2011 as C11.
- ISO/IEC 9899:2018, *Information Technology - Programming Languages - C*.  
This OpenMP API specification refers to ISO/IEC 9899:2018 as C18.
- ISO/IEC 14882:1998, *Information Technology - Programming Languages - C++*.  
This OpenMP API specification refers to ISO/IEC 14882:1998 as C++98.
- ISO/IEC 14882:2011, *Information Technology - Programming Languages - C++*.  
This OpenMP API specification refers to ISO/IEC 14882:2011 as C++11.
- ISO/IEC 14882:2014, *Information Technology - Programming Languages - C++*.  
This OpenMP API specification refers to ISO/IEC 14882:2014 as C++14.
- ISO/IEC 14882:2017, *Information Technology - Programming Languages - C++*.  
This OpenMP API specification refers to ISO/IEC 14882:2017 as C++17.
- ISO/IEC 1539:1980, *Information Technology - Programming Languages - Fortran*.  
This OpenMP API specification refers to ISO/IEC 1539:1980 as Fortran 77.
- ISO/IEC 1539:1991, *Information Technology - Programming Languages - Fortran*.  
This OpenMP API specification refers to ISO/IEC 1539:1991 as Fortran 90.
- ISO/IEC 1539-1:1997, *Information Technology - Programming Languages - Fortran*.  
This OpenMP API specification refers to ISO/IEC 1539-1:1997 as Fortran 95.
- ISO/IEC 1539-1:2004, *Information Technology - Programming Languages - Fortran*.  
This OpenMP API specification refers to ISO/IEC 1539-1:2004 as Fortran 2003.
- ISO/IEC 1539-1:2010, *Information Technology - Programming Languages - Fortran*.  
This OpenMP API specification refers to ISO/IEC 1539-1:2010 as Fortran 2008. While future versions of the OpenMP specification are expected to address the following features, currently their use may result in unspecified behavior.

- 1           – DO CONCURRENT
- 2           – Allocatable components of recursive type
- 3           – Simply contiguous arrays rank remapping to rank>1 target
- 4           – Polymorphic assignment
- 5           – Accessing real and imaginary parts
- 6           – Pointer function reference is a variable
- 7           – The BLOCK construct

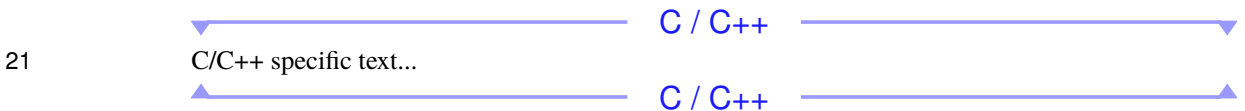
8           Where this OpenMP API specification refers to C, C++ or Fortran, reference is made to the base  
 9           language supported by the implementation.

## 10           **1.8 Organization of this Document**

11           The remainder of this document is structured as follows:

- 12           • Chapter 2 “Directives”
- 13           • Chapter 3 “Runtime Library Routines”
- 14           • Chapter 4 “OMPT Interface”
- 15           • Chapter 5 “OMPD Interface”
- 16           • Chapter 6 “Environment Variables”
- 17           • Appendix A “OpenMP Implementation-Defined Behaviors”
- 18           • Appendix B “Features History”

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



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



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

1 C90 specific text... C90

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

3 C99 specific text... C99

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

5 C++ specific text... C++

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

7 Fortran specific text..... Fortran

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

Fortran (cont.)

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

C/C++ (cont.)

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

13 Note – Non-normative text...  
14  
15



## 2 Directives

This chapter describes the syntax and behavior of OpenMP directives.

C / C++

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

C / C++

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.

Fortran

Compilers can therefore ignore OpenMP directives and conditionally compiled code if support of the OpenMP API is not provided or enabled. A compliant implementation must provide an option or interface that ensures that underlying support of all OpenMP directives and OpenMP conditional 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.

### Restrictions

C / C++

The following restriction applies to all OpenMP declarative directives:

- A declarative directive may not be used in place of the statement following an **if**, **else**, **while**, **do**, **for**, **switch**, or a label.

C / C++

Fortran

The following restrictions apply to all OpenMP directives:

- OpenMP directives, except **simd** and any declarative directive, may not appear in pure procedures.
- OpenMP directives may not appear in the WHERE and FORALL constructs.

Fortran

## 2.1 Directive Format

C / C++

OpenMP directives for C/C++ may be specified with **#pragma** directives as follows:

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

Each **#pragma** 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 **#pragma omp** are subject to macro replacement.

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

C / C++

C++

Some OpenMP directives for C++ may be specified with C++ attribute specifiers as follows:

```
[[ omp :: directive-name( [clause[, clause]... ] ) ] ]
```

or

```
[[ using omp : directive-name( [clause[, clause]... ] ) ] ]
```

The above two forms are interchangeable for any OpenMP directives that can be specified with C++ attribute specifiers. Some OpenMP directives may be composed of consecutive attribute specifiers if specified in their syntax. Any two consecutive attribute specifiers may be reordered or expressed as a single attribute specifier, as permitted by the base language, without changing the behavior of the OpenMP directive.

C++

C / C++

Directives are case-sensitive.

Each of the expressions used in the OpenMP syntax inside of the clauses must be a valid *assignment-expression* of the base language unless otherwise specified.

C / C++

C++

Directives may not appear in **constexpr** functions or in constant expressions.

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 form and free form source files, as described in Section 2.1.1 on page 40 and Section 2.1.2 on page 40.

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

Each of the expressions used in the OpenMP syntax inside of the clauses must be a valid *expression* of the base language unless otherwise specified.

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 193). 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 clauses accept a *list*, an *extended-list*, or a *locator-list*. A *list* consists of a comma-separated collection of one or more *list items*. An *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 an array section. An *extended list item* is a *list item* or a function name. A *locator list item* is any lvalue expression, including variables, an array section, or a reserved locator.

## C / C++

## Fortran

A *list item* is a variable that is not coindexed, an array section that is not coindexed, or a common block name (enclosed in slashes). An *extended list item* is a *list item* or a procedure name. A *locator list item* is a *list item*, or a reserved locator.

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

Although variables in common blocks can be accessed by use association or host association, common block names cannot. As a result, a common block name specified in a data-sharing

1 attribute, a data copying or a data-mapping attribute clause must be declared to be a common block  
2 in the same scoping unit in which the clause appears.

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

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

### Fortran

9 For all base languages, a *list item*, an *extended list item*, or a *locator list item* is subject to the  
10 restrictions specified in Section 2.1.5 on page 42 and in each of the sections describing clauses and  
11 directives for which the *list*, the *extended-list*, or the *locator-list* appears.

12 Some clauses and directives accept the use of reserved locators as special identifiers representing  
13 system storage not necessarily bound to any base language storage item. Reserved locators may  
14 only appear in clauses and directives where they are explicitly allowed and may not otherwise be  
15 referenced in the program. The list of reserved locators is:

16 `omp_all_memory`

17 The reserved locator `omp_all_memory` is a reserved identifier that denotes a list item treated as  
18 having storage that corresponds to the storage of all other objects in memory.

19 Some executable directives include a structured block or a structured block sequence.

### C / C++

20 A structured block sequence that consists of more than one statement may appear only for  
21 executable directives that explicitly allow it. The corresponding compound statement obtained by  
22 enclosing the sequence in { and } must be a structured block and the structured block sequence  
23 then should be considered to be a structured block with all of its restrictions.

### C / C++

24 A structured block:

- 25 • may contain infinite loops where the point of exit is never reached;
- 26 • may halt due to an IEEE exception;

### C / C++

- 27 • may contain calls to `exit()`, `_Exit()`, `quick_exit()`, `abort()` or functions with a  
28 `_Noreturn` specifier (in C) or a `noreturn` attribute (in C/C++);
- 29 • may be an expression statement, iteration statement, selection statement, or try block, provided  
30 that the corresponding compound statement obtained by enclosing it in { and } would be a  
31 structured block; and

### C / C++

▼ Fortran ▼

- may contain **STOP** or **ERROR STOP** statements.

▲ Fortran ▲

## Restrictions

Restrictions to structured blocks are as follows:

- Entry to a structured block must not be the result of a branch.
- The point of exit cannot be a branch out of the structured block.

▼ C / C++ ▼

- The point of entry to a structured block must not be a call to **set jmp ()**.
- **longjmp ()** and **throw ()** must not violate the entry/exit criteria.

▲ C / C++ ▲

▼ Fortran ▼

Restrictions to explicit OpenMP regions (that arise from executable directives) are as follows:

- If more than one image is executing the program, any image control statement, **ERROR STOP** statement, or access to a coindexed object that appears in an explicit OpenMP region will result in unspecified behavior.

## 2.1.1 Fixed Source Form Directives

The following sentinels are recognized in fixed form source files:

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

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

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

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):

```
c23456789
!$omp parallel do shared(a,b,c)

c$omp parallel do
c$omp+shared(a,b,c)

c$omp paralleldoshared(a,b,c)
```

## 2.1.2 Free Source Form Directives

The following sentinel is recognized in free form source files:

```
!$omp
```

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

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

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

3  
4 **Note** – In the following example the three formats for specifying the directive are equivalent (the  
5 first line represents the position of the first 9 columns):

```
6 !23456789  
7     !$omp parallel do &  
8         !$omp shared(a,b,c)  
9  
10    !$omp parallel &  
11    !$omp&do shared(a,b,c)  
12  
13    !$omp paralleldo shared(a,b,c)
```

14  
15  
▲────────────────── Fortran ───────────────────▶

## 16 2.1.3 Stand-Alone Directives

### 17 Summary

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

### 19 Description

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

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

### 25 Restrictions

- 26 • A stand-alone directive may not be used in place of the statement following an **if**, **else**,  
27 **while**, **do**, **for**, **switch**, or a label.

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

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

## 2.1.4 Array Shaping

If an expression has a type of pointer to  $T$ , 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 := ([ $s_1$ ] [ $s_2$ ] . . . [ $s_n$ ]) cast-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 to the shape-operator are as follows:

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

▼ C++ ▼

- If the type  $T$  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++ ▲

## 2.1.5 Array Sections

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

▼ C / C++ ▼

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



```

1      [ lower-bound : length : stride ] or
2      [ lower-bound : length : ] or
3      [ lower-bound : length ] or
4      [ lower-bound : : stride ] or
5      [ lower-bound : : ] or
6      [ lower-bound : ] or
7      [ : length : stride ] or
8      [ : length : ] or
9      [ : length ] or
10     [ : : stride ]
11     [ : : ]
12     [ : ]

```

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

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

Each of the *lower-bound*, *length*, and *stride* expressions if specified must be an integral type expression of the base language. When evaluated they represent a set of integer values as follows:

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

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

The *stride* must evaluate to a positive integer.

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

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

When the *length* is absent it defaults to  $\lceil (size - lower-bound) / stride \rceil$ , where *size* is the size of the array dimension.

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

The precedence of a subscript operator that uses the array section syntax is the same as the precedence of a subscript operator that does not use the array section syntax.

Note – The following are examples of array sections:

```

a[0:6]
a[0:6:1]

```

```

1  a[1:10]
2  a[1:]
3  a[:10:2]
4  b[10][:][:]
5  b[10][:][:0]
6  c[42][0:6][:]
7  c[42][0:6:2][:]
8  c[1:10][42][0:6]
9  S.c[:100]
10 p->y[:10]
11 this->a[:N]
12 (p+10)[:N]

```

Assume **a** is declared to be a 1-dimensional array with dimension size 11. The first two examples are equivalent, and the third and fourth examples are equivalent. The fifth example specifies a stride of 2 and therefore is not contiguous.

Assume **b** is declared to be a pointer to a 2-dimensional array with dimension sizes 10 and 10. The sixth example refers to all elements of the 2-dimensional array given by **b[10]**. The seventh example is a zero-length array section.

Assume **c** is declared to be a 3-dimensional array with dimension sizes 50, 50, and 50. The eighth example is contiguous, while the ninth and tenth examples are not contiguous.

The final four examples show array sections that are formed from more general base expressions.

The following are examples that are non-conforming array sections:

```

s[:10].x
p[:10]->y
*(xp[:10])

```

For all three examples, a base language operator is applied in an undefined manner to an array section. The only operator that may be applied to an array section is a subscript operator for which the array section appears as the postfix expression.

C / C++

## Fortran

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

## Fortran

### Restrictions

3 Restrictions to array sections are as follows:

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

## C / C++

- 6 • An element of an array section with a non-zero size must have a complete type.
- 7 • The base expression of an array section must have an array or pointer type.
- 8 • If a consecutive sequence of array subscript expressions appears in an array section, and the first  
9 subscript expression in the sequence uses the extended array section syntax defined in this  
10 section, then only the last subscript expression in the sequence may select array elements that  
11 have a pointer type.

## C / C++

### C++

- 12 • If the type of the base expression of an array section is a reference to a type  $T$ , then the type will  
13 be considered to be  $T$  for all purposes of the array section.
- 14 • An array section cannot be used in an overloaded `[]` operator.

### C++

## Fortran

- 15 • If a stride expression is specified, it must be positive.
- 16 • The upper bound for the last dimension of an assumed-size dummy array must be specified.
- 17 • If a list item is an array section with vector subscripts, the first array element must be the lowest  
18 in the array element order of the array section.
- 19 • If a list item is an array section, the last *part-ref* of the list item must have a section subscript list.

## Fortran

## 2.1.6 Iterators

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

The syntax of the **iterator** modifier is as follows:

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

where *iterators-definition* is one of the following:

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

where *iterator-specifier* is one of the following:

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

where:

- *identifier* is a base language identifier.

▼ C / C++ ▼

- *iterator-type* is a type name.

▲ C / C++ ▲

▼ Fortran ▼

- *iterator-type* is a type specifier.

▲ Fortran ▲

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

▼ C / C++ ▼

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

▲ C / C++ ▲

▼ Fortran ▼

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

▲ Fortran ▲

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

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

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

## C / C++

- 1 The values of the iterator are the set of values  $i_0, \dots, i_{N-1}$  where:
- 2 •  $i_0 = (\text{iterator-type}) \text{ begin}$ ,
- 3 •  $i_j = (\text{iterator-type}) (i_{j-1} + \text{step})$ , and
- 4 • if  $\text{step} > 0$ ,
- 5     –  $i_0 < (\text{iterator-type}) \text{ end}$ ,
- 6     –  $i_{N-1} < (\text{iterator-type}) \text{ end}$ , and
- 7     –  $(\text{iterator-type}) (i_{N-1} + \text{step}) \geq (\text{iterator-type}) \text{ end}$ ;
- 8 • if  $\text{step} < 0$ ,
- 9     –  $i_0 > (\text{iterator-type}) \text{ end}$ ,
- 10     –  $i_{N-1} > (\text{iterator-type}) \text{ end}$ , and
- 11     –  $(\text{iterator-type}) (i_{N-1} + \text{step}) \leq (\text{iterator-type}) \text{ end}$ .

## C / C++









## Fortran

- 12 The values of the iterator are the set of values  $i_1, \dots, i_N$  where:
- 13 •  $i_1 = \text{begin}$ ,
- 14 •  $i_j = i_{j-1} + \text{step}$ , and
- 15 • if  $\text{step} > 0$ ,
- 16     –  $i_1 \leq \text{end}$ ,
- 17     –  $i_N \leq \text{end}$ , and
- 18     –  $i_N + \text{step} > \text{end}$ ;
- 19 • if  $\text{step} < 0$ ,
- 20     –  $i_1 \geq \text{end}$ ,
- 21     –  $i_N \geq \text{end}$ , and
- 22     –  $i_N + \text{step} < \text{end}$ .

## Fortran

1 The set of values will be empty if no possible value complies with the conditions above.  
 2 For those clauses that contain expressions that contain iterator identifiers, the effect is as if the list  
 3 item is instantiated within the clause for each value of the iterator in the set defined above,  
 4 substituting each occurrence of the iterator identifier in the expression with the iterator value. If the  
 5 set of values of the iterator is empty then the effect is as if the clause was not specified.  
 6 The behavior is unspecified if  $i_j + step$  cannot be represented in *iterator-type* in any of the  
 7  $i_j + step$  computations for any  $0 \leq j < N$  in C/C++ or  $0 < j \leq N$  in Fortran.

## 8 Restrictions

- 9 • An expression that contains an iterator identifier can only appear in clauses that explicitly allow  
 10 expressions that contain iterators.
  - 11 • The *iterator-type* must not declare a new type.
- 12  C / C++ 
- 13 • The *iterator-type* must be an integral or pointer type.
  - 13 • The *iterator-type* must not be **const** qualified.
- 14  C / C++   
 Fortran 
- 14 • The *iterator-type* must be an integer type.
- 15  Fortran 
- 15 • If the *step* expression of a *range-specification* equals zero, the behavior is unspecified.
  - 16 • Each iterator identifier can only be defined once in an *iterators-definition*.
  - 17 • Iterators cannot appear in the *range-specification*.

## 18 2.2 Conditional Compilation

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

22 If a **#define** or a **#undef** preprocessing directive in user code defines or undefines the  
 23 **\_OPENMP** macro name, the behavior is unspecified.

 Fortran 

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

## 2.2.1 Fixed Source Form Conditional Compilation Sentinels

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

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

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

- 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.

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

```
c23456789
!$ 10 iam = omp_get_thread_num() +
!$   &           index

#ifdef _OPENMP
    10 iam = omp_get_thread_num() +
    &           index
#endif
```

## 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.

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

```
c23456789
  !$ iam = omp_get_thread_num() +      &
  !$&   index

#ifdef _OPENMP
    iam = omp_get_thread_num() +      &
    index
#endif
```

Fortran

## 2.3 Variant Directives

### 2.3.1 OpenMP Context

At any point in a program, an OpenMP context exists that defines traits that describe the active OpenMP constructs, the execution devices, and functionality supported by the implementation. The traits are grouped into 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 constructs at that point in the program up to a **target** construct. Combined and composite constructs are added to the set as distinct constructs in the same nesting order specified by the original construct. The set is ordered by nesting level in ascending order. Specifically, the ordering of the set of constructs is  $c_1, \dots, c_N$ , where  $c_1$  is the construct at the outermost nesting level and  $c_N$  is the construct at the innermost nesting level. In addition, if the point in the program is not enclosed by a **target** construct, the following rules are applied in order:

- For functions with a declare SIMD directive, the *simd* trait is added to the beginning of the set as  $c_1$  for any generated SIMD versions so the total size of the set is increased by 1.



- 1           2. For functions that are determined to be function variants by a declare variant directive, the  
 2           selectors  $c_1, \dots, c_M$  of the **construct** selector set are added in the same order to the  
 3           beginning of the set as  $c_1, \dots, c_M$  so the total size of the set is increased by  $M$ .
- 4           3. For functions within a declare target block, the *target* trait is added to the beginning of the set as  
 5            $c_1$  for any versions of the function that are generated for **target** regions so the total size of the  
 6           set is increased by 1.

7           The *simd* trait can be further defined with properties that match the clauses accepted by the declare  
 8           SIMD directive with the same name and semantics. The *simd* trait must define at least the *simdlen*  
 9           property and one of the *inbranch* or *notinbranch* properties.

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

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

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

- 24          • The *vendor(vendor-name-list)* trait, which specifies the vendor identifiers of the implementation.  
 25          OpenMP defined values for *vendor-name* are defined in the “OpenMP Context Definitions”  
 26          document, which is available at <http://www.openmp.org/>.
- 27          • The *extension(extension-name-list)* trait, which specifies vendor specific extensions to the  
 28          OpenMP specification. The accepted *extension-name* values are implementation defined.
- 29          • A trait with a name that is identical to the name of any clause that can be supplied to the  
 30          **requires** directive.

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

```
33          identifier [ (context-element [, context-element [, ...]]) ]
34
35          context-element :
36            identifier [ (context-element [, context-element [, ...]]) ]
```

```

1   or
2   context-value
3
4   context-value :
5   constant string
6   or
7   constant integer expression

```

8 where *identifier* is a base language identifier.

## 9 2.3.2 Context Selectors

10 Context selectors are used to define the properties of an OpenMP context that a directive or clause  
11 can match. OpenMP defines different sets of selectors, each containing different selectors.

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

```


13 trait-set-selector[, trait-set-selector[, ...]]
14
15 trait-set-selector :
16   trait-set-selector-name={trait-selector[, trait-selector[, ...]]}
17
18 trait-selector :
19   trait-selector-name[ ([trait-score: ] trait-property[, trait-property[, ...]]) ]
20
21 trait-score :
22   score (score-expression)

```

23 The **construct** selector set defines the *construct* traits that should be active in the OpenMP  
24 context. The following selectors can be defined in the **construct** set: **target**; **teams**;  
25 **parallel**; **for** (in C/C++); **do** (in Fortran); and **simd**. The properties of each selector are the  
26 same properties that are defined for the corresponding trait. The **construct** selector is an  
27 ordered list  $c_1, \dots, c_N$ .

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

32 The **user** selector set defines the **condition** selector that provides additional user-defined  
33 conditions.

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



C++

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

C++

Fortran

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

Fortran

5 A *score-expression* must be an constant integer expression.

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

## 8 **Restrictions**

- 9 • Each *trait-set-selector-name* can only be specified once.
- 10 • Each *trait-selector-name* can only be specified once.
- 11 • A *trait-score* cannot be specified in traits from the **construct** or **device** *trait-selector-sets*.

## 12 **2.3.3 Matching and Scoring Context Selectors**

13 A given context selector is compatible with a given OpenMP context if the following conditions are  
14 satisfied:

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

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

- 23 • The **simdlen** (*N*) property of the selector matches the *simdlen*(*M*) trait of the OpenMP context  
24 if  $M\%N$  equals zero; and
- 25 • The **aligned** (*list:N*) property of the selector matches the *aligned*(*list:M*) trait of the OpenMP  
26 context if  $N\%M$  equals zero.

27 Among compatible context selectors, a score is computed using the following algorithm:

- 1 1. Each trait that appears in the *construct* trait set in the OpenMP context is given the value  $2^{p-1}$   
2 where  $p$  is the position of the construct trait,  $c_p$ , in the set;
- 3 2. The **kind**, **arch**, and **isa** selectors are given the values  $2^l$ ,  $2^{l+1}$  and  $2^{l+2}$ , respectively, where  
4  $l$  is the number of traits in the *construct* set;
- 5 3. Traits for which a *trait-score* is specified are given the value specified by the *trait-score*  
6 *score-expression*;
- 7 4. The values given to any additional selectors allowed by the implementation are implemented  
8 defined;
- 9 5. Other selectors are given a value of zero; and
- 10 6. A context selector that is a strict subset of another context selector has a score of zero. For other  
11 context selectors, the final score is the sum of the values of all specified selectors plus 1. If the  
12 traits that correspond to the **construct** selectors appear multiple times in the OpenMP  
13 context, the highest valued subset of traits that contains all selectors in the same order are used.

## 2.3.4 Metadirectives

### Summary

A metadirective is a directive that can specify multiple directive variants of which one may be conditionally selected to replace the metadirective based on the enclosing OpenMP context.

### Syntax

C / C++

The syntax of a metadirective is as follows:

```
#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++

The syntax of a metadirective is as follows:

```
!$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])
```

In the **when** clause, *context-selector-specification* specifies a context selector (see Section 2.3.2).

In the **when** and **default** clauses, *directive-variant* has the following form and specifies a directive variant that specifies an OpenMP directive with clauses that apply to it.

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

## Description

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

The OpenMP context for a given metadirective is defined according to Section 2.3.1. For each **when** clause that appears on a metadirective, the specified directive variant, if present, is a candidate to replace the metadirective if the corresponding context selector is compatible with the OpenMP context according to the matching rules defined in Section 2.3.3. If only one compatible context selector specified by a **when** clause has the highest score and it specifies a directive variant, the directive variant will replace the metadirective. If more than one **when** clause specifies a compatible context selector that has the highest computed score and at least one specifies a directive variant, the first directive variant specified in the lexical order of those **when** clauses will replace the metadirective.

If no context selector from any **when** clause is compatible with the OpenMP context and a **default** clause that specifies a directive variant is present, that directive variant will replace the metadirective.

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

The **begin metadirective** directive behaves identically to the **metadirective** directive, except that the directive syntax for the specified directive variants must accept a paired

1 **end directive**. For any directive variant that is selected to replace the **begin metadirective**  
2 directive, the **end metadirective** directive will be implicitly replaced by its paired  
3 **end directive** to demarcate the statements that are affected by or are associated with the directive  
4 variant. If no directive variant is selected to replace the **begin metadirective** directive, its  
5 paired **end metadirective** directive is ignored.

## 6 **Restrictions**

7 Restrictions to metadirectives are as follows:

- 8 • The directive variant appearing in a **when** or **default** clause must not specify a  
9 **metadirective**, **begin metadirective**, or **end metadirective** directive.

10  C / C++

- 11 • The directive variant that appears in a **when** or **default** clause must not specify a  
12 **begin declare variant** or **end declare variant**.

13  C / C++

- 14 • The context selector that appears in a **when** clause must not specify any properties for the **simd**  
15 selector.
- 16 • Any replacement that occurs for a metadirective must not result in a non-conforming OpenMP  
17 program.
- 18 • Any directive variant that is specified by a **when** or **default** clause on a  
19 **begin metadirective** directive must be an OpenMP directive that has a paired  
20 **end directive**, and the **begin metadirective** directive must have a paired  
21 **end metadirective** directive.
- 22 • The **default** clause may appear at most once on a metadirective.

## 21 **2.3.5 Declare Variant Directive**

### 22 **Summary**

23 The declare variant directive declares a specialized variant of a base function and specifies the  
24 context in which that specialized variant is used. The declare variant directive is a declarative  
25 directive.

### 26 **Syntax**

27  C / C++

28 The syntax of the declare variant directive is as follows:

```
29 #pragma omp declare variant (variant-func-id) clause new-line  
30 [#pragma omp declare variant (variant-func-id) clause new-line]  
31 [ ... ]  
32 function definition or declaration
```

33 or







## C / C++

- 1 • If the base function associated with a **declare variant** directive has any declarations then  
2 the **declare variant** directives for any declarations that have one must be equivalent. If the  
3 function definition has a **declare variant** directive, it must also be equivalent. Otherwise,  
4 the result is unspecified.
- 5 • The type of the function variant must be compatible with the type of the base function after the  
6 implementation-defined transformation for its OpenMP context.
- 7 • The **match** clause of a **begin declare variant** directive may not contain a **simd**  
8 *trait-selector-name*.
- 9 • Matching pairs of **begin declare variant** and **end declare variant** directives shall  
10 encompass disjoint source ranges or they shall be nested perfectly.

## C / C++

### C++

- 11 • The **declare variant** directive cannot be specified for a virtual function.

### C++

## Fortran

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

## Fortran

### Cross References

- 24 • OpenMP Context Specification, see Section [2.3.1](#) on page [50](#).
- 25 • Context Selectors, see Section [2.3.2](#) on page [52](#).
- 26

## 2.4 requires Directive

### Summary

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

### Syntax

C / C++

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

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

C / C++

Fortran

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

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

Fortran

Where *clause* is either one of the requirement clauses listed below or a clause of the form **ext\_implementation-defined-requirement** for an implementation defined requirement clause.

```
reverse_offload  
unified_address  
unified_shared_memory  
atomic_default_mem_order(seq_cst | acq_rel | relaxed)  
dynamic_allocators
```

### Description

The **requires** directive specifies features that an implementation must support for correct execution. The behavior that a requirement clause specifies may override the normal behavior specified elsewhere in this document. Whether an implementation supports the feature that a given requirement clause specifies is implementation defined.

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

1  
2 **Note** – Use of this directive makes your code less portable. Users should be aware that not all  
3 devices or implementations support all requirements.  
4

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

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

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

27 The **unified\_shared\_memory** clause makes the **map** clause optional on **target** constructs  
28 and the **declare target** directive optional for static lifetime variables accessed inside functions to  
29 which a **declare target** directive is applied. Scalar variables are still firstprivate by default when  
30 referenced inside **target** constructs. Values stored into memory by one device may not be visible  
31 to another device until those two devices synchronize with each other or both devices synchronize  
32 with the host.

33 The **atomic\_default\_mem\_order** clause specifies the default memory ordering behavior for  
34 **atomic** constructs that must be provided by an implementation. If the default memory ordering is  
35 specified as **seq\_cst**, all **atomic** constructs on which *memory-order-clause* is not specified  
36 behave as if the **seq\_cst** clause appears. If the default memory ordering is specified as  
37 **relaxed**, all **atomic** constructs on which *memory-order-clause* is not specified behave as if the  
38 **relaxed** clause appears.

39 If the default memory ordering is specified as **acq\_rel**, **atomic** constructs on which  
40 *memory-order-clause* is not specified behave as if the **release** clause appears if the atomic write

1 or atomic update operation is specified, as if the **acquire** clause appears if the atomic read  
2 operation is specified, and as if the **acq\_rel** clause appears if the atomic captured update  
3 operation is specified.

4 The **dynamic\_allocators** clause removes certain restrictions on the use of memory allocators  
5 in **target** regions. It makes the **uses\_allocators** clause optional on **target** constructs for  
6 the purpose of using allocators in the corresponding **target** regions. It allows calls to the  
7 **omp\_init\_allocator** and **omp\_destroy\_allocator** API routines in **target** regions.  
8 Finally, it allows default allocators to be used by **allocate** directives, **allocate** clauses, and  
9 **omp\_alloc** API routines in **target** regions.

10 Implementers are allowed to include additional implementation defined requirement clauses. All  
11 implementation defined requirements should begin with **ext\_**. Requirement names that do not  
12 start with **ext\_** are reserved.

## 13 Restrictions

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

- 15 • Each of the clauses can appear at most once on the directive.
- 16 • At most one **requires** directive with **atomic\_default\_mem\_order** clause can appear in  
17 a single compilation unit.
- 18 • A **requires** directive with a **unified\_address**, **unified\_shared\_memory**, or  
19 **reverse\_offload** clause must appear lexically before any device constructs or device  
20 routines.
- 21 • A **requires** directive with any of the following clauses must appear in all *compilation units* of  
22 a program that contain device constructs or device routines or in none of them:
  - 23 – **reverse\_offload**
  - 24 – **unified\_address**
  - 25 – **unified\_shared\_memory**
- 26 • The **requires** directive with **atomic\_default\_mem\_order** clause may not appear  
27 lexically after any **atomic** construct on which *memory-order-clause* is not specified.

- 28 • The **requires** directive may only appear at file scope.  
▲ C ▼
- 29 • The **requires** directive may only appear at file or namespace scope.  
▲ C++ ▼

## 2.5 assume Directive

### Summary

The assume directive provides invariants to the implementation that may be used for optimization purposes. If the invariants do not hold at runtime, the behavior is unspecified.

### Syntax

The syntax of the assume directive is as follows:

C / C++

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

or

```
#pragma omp begin assumes clause[ [ [, ] clause] ... ] new-line
    declaration-definition-seq
#pragma omp end assumes new-line
```

or

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

Where *clause* is either *assumption-clause* or a clause of the form **ext\_implementation-defined-assumption** for an implementation-defined assumption clause.

Where *assumption-clause* is one of the following:

```
absent (directive-name [ [, ] directive-name]...)
contains (directive-name [ [, ] directive-name]...)
holds (scalar-expression)
no_openmp
no_openmp_routines
no_parallelism
```

C / C++

Fortran

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

or

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

1 Where *clause* is either *assumption-clause* or a clause of the form  
2 **ext\_implementation-defined-assumption** for an implementation-defined assumption clause.

3 Where *assumption-clause* is one of the following:

```
4     absent (directive-name [[, directive-name]...])  
5     contains (directive-name [[, directive-name]...])  
6     holds (scalar-logical-expression)  
7     no_openmp  
8     no_openmp_routines  
9     no_parallelism
```

Fortran

## 10 Description

11 The assume directive gives the implementation additional information about the expected  
12 properties of the program that can optionally be used to optimize the implementation. An  
13 implementation may ignore this information without altering the behavior of the program.

14 The **assume** form is an executable directive while the **assumes** and **begin assumes** forms are  
15 declarative directives.

16 The scope of the **assumes** directive is the code executed and reached from the current compilation  
17 unit. The scope of the **assume** construct is the code executed in the corresponding region or in any  
18 region that is nested in the corresponding region.

C / C++

19 The scope of the **begin assumes** directive is the code that is executed and reached from any of  
20 the declared functions in *declaration-definition-seq*.

C / C++

21 The **absent** and **contains** clauses both accept a list of one or more directive names that may  
22 match a construct that is encountered within the scope of the directive. An encountered construct  
23 matches the directive name if it has the same name as one of the specified directive names or if it is  
24 a combined or composite construct for which a constituent construct has the same name as one of  
25 the specified directive names.

26 When the **absent** clause appears on an assume directive, the application guarantees that no  
27 constructs that match a listed directive name are encountered in the scope of the assume directive.

28 When the **contained** clause appears on an assume directive, the application provides a hint that  
29 constructs that match the listed directive names are likely to be encountered in the scope of the  
30 assume directive.

31 When the **holds** clause appears on an assume directive, the application guarantees that the listed  
32 expression evaluates to *true* in the scope of the directive. The effect of the clause does not include  
33 an evaluation of the expression that is observable.

1 The `no_openmp` clause guarantees that no OpenMP related code is executed in the scope of the  
2 directive.

3 The `no_openmp_routines` clause guarantees that no explicit OpenMP runtime library calls are  
4 executed in the scope of the directive.



5 The `no_parallelism` clause guarantees that no OpenMP tasks (explicit or implicit) will be  
6 generated and that no SIMD constructs will be executed in the scope of the directive.

7 Implementers are allowed to include additional implementation-defined assumption clauses. All  
8 implementation-defined assumptions should begin with `ext_`. Assumption names that do not start  
9 with `ext_` are reserved.



## 10 Restrictions

11 The restrictions for the assume directive are as follows:

- 12 • Each clause except the `absent`, `contains` and `holds` clause can appear at most once on the  
13 directive.
- 14 • Each *directive-name* listed in the clauses can appear at most once on the directive.
- 15 • A *directive-name* that appears in an `absent` or `contains` clause may not be a combined or  
16 composite directive name.
- 17 • A *directive-name* that appears in an `absent` or `contains` clause may not be a directive that is  
18 not associated with the execution of user or implementation code, i.e., a `nothing` directive, a  
19 declarative directive, a metadirective, or a loop transformation directive.

20 ▼  C  ▲  
• The declarative versions of the assume directive may only appear at file scope.

▲  C  ▲

21 ▼  C++  ▲  
• The declarative versions of the assume directive may only appear at file or namespace scope.

▲  C++  ▲

## 2.6 Internal Control Variables

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

For purposes of exposition, this document refers to the ICVs by certain names, but an implementation is not required to use these names or to offer any way to access the variables other than through the ways shown in Section 2.6.2 on page 68.

### 2.6.1 ICV Descriptions

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

- *dyn-var* - controls whether dynamic adjustment of the number of threads is enabled for encountered **parallel** regions. There is one copy of this ICV per data environment.
- *nthreads-var* - controls the number of threads requested for encountered **parallel** regions. There is one copy of this ICV per data environment.
- *thread-limit-var* - controls the maximum number of threads participating in the contention group. There is one copy of this ICV per data environment.
- *max-active-levels-var* - controls the maximum number of nested active **parallel** regions. There is one copy of this ICV per device.
- *place-partition-var* - controls the place partition available to the execution environment for encountered **parallel** regions. There is one copy of this ICV per implicit task.
- *active-levels-var* - the number of nested active **parallel** regions that enclose the current task such that all of the **parallel** regions are enclosed by the outermost initial task region on the current device. There is one copy of this ICV per data environment.
- *levels-var* - the number of nested parallel regions that enclose the current task such that all of the **parallel** regions are enclosed by the outermost initial task region on the current device. There is one copy of this ICV per data environment.
- *bind-var* - controls the binding of OpenMP threads to places. When binding is requested, the variable indicates that the execution environment is advised not to move threads between places. The variable can also provide default thread affinity policies. There is one copy of this ICV per data environment.

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



- 1           • *run-sched-var* - controls the schedule that is used for worksharing-loop regions when the  
2           **runtime** schedule kind is specified. There is one copy of this ICV per data environment.
- 3           • *def-sched-var* - controls the implementation defined default scheduling of worksharing-loop  
4           regions. There is one copy of this ICV per device.

5           The following ICVs store values that affect program execution.

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

22          The following ICVs store values that affect the operation of the OMPT tool interface.

- 23          • *tool-var* - controls whether an OpenMP implementation will try to register a tool. There is one  
24          copy of this ICV for the whole program.
- 25          • *tool-libraries-var* - specifies a list of absolute paths to tool libraries for OpenMP devices. There  
26          is one copy of this ICV for the whole program.

27          The following ICVs store values that affect the operation of the OMPD tool interface.

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

31          The following ICVs store values that may be queried by interface routines.

- 32          • *num-procs-var* - the number of processors that are available to the device. There is one copy of  
33          this ICV per device.

- *thread-num-var* - the thread number of an implicit task within its binding team. There is one copy of this ICV per data environment.
- *final-task-var* - whether the current task is a final task. There is one copy of this ICV per data environment.
- *implicit-task-var* - whether the current task is an implicit task. There is one copy of this ICV per data environment.
- *team-size-var* - the size of the team. There is one copy of this ICV per data environment.

The following ICVs store values that affect default memory allocation.

- *def-allocator-var* - controls the memory allocator to be used by memory allocation routines, directives and clauses when a memory allocator is not specified by the user. There is one copy of this ICV per implicit task.

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

- *nteams-var* - controls the number of teams requested for encountered **teams** regions. There is one copy of this ICV per device.
- *teams-thread-limit-var* - controls the maximum number of threads participating in each contention group created by a **teams** construct. There is one copy of this ICV per device.

## 2.6.2 ICV Initialization

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>	<b>OMP_DYNAMIC</b>	See description below
<i>nthreads-var</i>	<b>OMP_NUM_THREADS</b>	Implementation defined
<i>run-sched-var</i>	<b>OMP_SCHEDULE</b>	Implementation defined
<i>def-sched-var</i>	(none)	Implementation defined
<i>bind-var</i>	<b>OMP_PROC_BIND</b>	Implementation defined
<i>stacksize-var</i>	<b>OMP_STACKSIZE</b>	Implementation defined
<i>wait-policy-var</i>	<b>OMP_WAIT_POLICY</b>	Implementation defined

*table continued on next page*

*table continued from previous page*

<b>ICV</b>	<b>Environment Variable</b>	<b>Initial value</b>
<i>thread-limit-var</i>	<b>OMP_THREAD_LIMIT</b>	Implementation defined
<i>max-active-levels-var</i>	<b>OMP_MAX_ACTIVE_LEVELS,</b> <b>OMP_NESTED,</b> <b>OMP_NUM_THREADS,</b> <b>OMP_PROC_BIND</b>	Implementation defined
<i>active-levels-var</i>	(none)	<i>zero</i>
<i>levels-var</i>	(none)	<i>zero</i>
<i>place-partition-var</i>	<b>OMP_PLACES</b>	Implementation defined
<i>cancel-var</i>	<b>OMP_CANCELLATION</b>	<i>false</i>
<i>display-affinity-var</i>	<b>OMP_DISPLAY_AFFINITY</b>	<i>false</i>
<i>affinity-format-var</i>	<b>OMP_AFFINITY_FORMAT</b>	Implementation defined
<i>default-device-var</i>	<b>OMP_DEFAULT_DEVICE</b>	Implementation defined
<i>target-offload-var</i>	<b>OMP_TARGET_OFFLOAD</b>	<b>DEFAULT</b>
<i>max-task-priority-var</i>	<b>OMP_MAX_TASK_PRIORITY</b>	<i>zero</i>
<i>tool-var</i>	<b>OMP_TOOL</b>	<i>enabled</i>
<i>tool-libraries-var</i>	<b>OMP_TOOL_LIBRARIES</b>	<i>empty string</i>
<i>debug-var</i>	<b>OMP_DEBUG</b>	<i>disabled</i>
<i>num-procs-var</i>	(none)	Implementation defined
<i>thread-num-var</i>	(none)	<i>zero</i>
<i>final-task-var</i>	(none)	<i>false</i>
<i>implicit-task-var</i>	(none)	<i>true</i>
<i>team-size-var</i>	(none)	<i>one</i>
<i>def-allocator-var</i>	<b>OMP_ALLOCATOR</b>	Implementation defined
<i>nteams-var</i>	<b>OMP_NUM_TEAMS</b>	<i>zero</i>
<i>teams-thread-limit-var</i>	<b>OMP_TEAMS_THREAD_LIMIT</b>	<i>zero</i>

1 Each ICV that does not have global scope (see Table 2.3 on page 74) has a set of device-specific  
2 environment variables that extend the variables defined in Table 2.1 on page 68 with the following  
3 syntax.

4 `<ENVIRONMENT VARIABLE>_DEV[_<device>]`

5 Where `<ENVIRONMENT VARIABLE>` is one of the variables from Table 2.1 on page 68 and  
6 `<device>` is the device number as specified in the `device` clause (see Section 2.15 on page 165).

## 7 Description

- 8 • Each device has its own ICVs.
- 9 • The initial value of `dyn-var` is implementation defined if the implementation supports dynamic  
10 adjustment of the number of threads; otherwise, the initial value is *false*.
- 11 • The value of the `nthreads-var` ICV is a list.
- 12 • The value of the `bind-var` ICV is a list.

13 The host and target device ICVs are initialized before any OpenMP API construct or OpenMP API  
14 routine executes. After the initial values are assigned, the values of any OpenMP environment  
15 variables that were set by the user are read and the associated ICVs are modified accordingly. If no  
16 `<device>` number is specified on the device-specific environment variable then the value is applied  
17 to all non-host devices.

## 18 Cross References

- 19 • `OMP_SCHEDULE` environment variable, see Section 6.1 on page 585.
- 20 • `OMP_NUM_THREADS` environment variable, see Section 6.2 on page 586.
- 21 • `OMP_DYNAMIC` environment variable, see Section 6.3 on page 587.
- 22 • `OMP_PROC_BIND` environment variable, see Section 6.4 on page 588.
- 23 • `OMP_PLACES` environment variable, see Section 6.5 on page 588.
- 24 • `OMP_STACKSIZE` environment variable, see Section 6.6 on page 591.
- 25 • `OMP_WAIT_POLICY` environment variable, see Section 6.7 on page 592.
- 26 • `OMP_MAX_ACTIVE_LEVELS` environment variable, see Section 6.8 on page 592.
- 27 • `OMP_NESTED` environment variable, see Section 6.9 on page 593.
- 28 • `OMP_THREAD_LIMIT` environment variable, see Section 6.10 on page 593.
- 29 • `OMP_CANCELLATION` environment variable, see Section 6.11 on page 594.
- 30 • `OMP_DISPLAY_AFFINITY` environment variable, see Section 6.13 on page 595.
- 31 • `OMP_AFFINITY_FORMAT` environment variable, see Section 6.14 on page 596.
- 32 • `OMP_DEFAULT_DEVICE` environment variable, see Section 6.15 on page 598.

- 1       • **OMP\_MAX\_TASK\_PRIORITY** environment variable, see Section 6.16 on page 598.
- 2       • **OMP\_TARGET\_OFFLOAD** environment variable, see Section 6.17 on page 599.
- 3       • **OMP\_TOOL** environment variable, see Section 6.18 on page 599.
- 4       • **OMP\_TOOL\_LIBRARIES** environment variable, see Section 6.19 on page 600.
- 5       • **OMP\_DEBUG** environment variable, see Section 6.20 on page 600.
- 6       • **OMP\_ALLOCATOR** environment variable, see Section 6.21 on page 601.
- 7       • **OMP\_NUM\_TEAMS** environment variable, see Section 6.22 on page 601.
- 8       • **OMP\_TEAMS\_THREAD\_LIMIT** environment variable, see Section 6.23 on page 601.

### 9       2.6.3 Modifying and Retrieving ICV Values

10       Table 2.2 shows the method for modifying and retrieving the values of ICVs through OpenMP API  
 11       routines. If an ICV is not listed in this table, no OpenMP API routine modifies or retrieves this ICV.

**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>nthreads-var</i>	<code>omp_set_num_threads()</code>	<code>omp_get_max_threads()</code>
<i>run-sched-var</i>	<code>omp_set_schedule()</code>	<code>omp_get_schedule()</code>
<i>bind-var</i>	(none)	<code>omp_get_proc_bind()</code>
<i>thread-limit-var</i>	<code>teams</code> construct	<code>omp_get_thread_limit()</code>
<i>max-active-levels-var</i>	<code>omp_set_max_active_levels()</code> , <code>omp_set_nested()</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>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>

*table continued on next page*

table continued from previous page

ICV	Ways to Modify Value	Ways to Retrieve Value
<i>max-task-priority-var</i>	(none)	<code>omp_get_max_task_priority()</code>
<i>num-procs-var</i>	(none)	<code>omp_get_num_procs()</code>
<i>thread-num-var</i>	(none)	<code>omp_get_thread_num()</code>
<i>final-task-var</i>	(none)	<code>omp_in_final()</code>
<i>team-size-var</i>	(none)	<code>omp_get_num_threads()</code>
<i>def-allocator-var</i>	<code>omp_set_default_allocator()</code>	<code>omp_get_default_allocator()</code>
<i>ntteams-var</i>	<code>omp_set_num_teams()</code>	<code>omp_get_max_teams()</code>
<i>teams-thread-limit-var</i>	<code>omp_set_teams_thread_limit()</code>	<code>omp_get_teams_thread_limit()</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 2.9 on page 87.
- `omp_set_num_threads` routine, see Section 3.2.1 on page 338.
- `omp_get_num_threads` routine, see Section 3.2.2 on page 339.
- `omp_get_max_threads` routine, see Section 3.2.3 on page 340.
- `omp_get_thread_num` routine, see Section 3.2.4 on page 341.
- `omp_get_num_procs` routine, see Section 3.2.5 on page 342.
- `omp_set_dynamic` routine, see Section 3.2.7 on page 343.
- `omp_get_dynamic` routine, see Section 3.2.8 on page 344.
- `omp_get_cancellation` routine, see Section 3.2.9 on page 345.
- `omp_set_nested` routine, see Section 3.2.10 on page 345.
- `omp_set_schedule` routine, see Section 3.2.12 on page 347.

- 1       • `omp_get_schedule` routine, see Section 3.2.13 on page 349.
- 2       • `omp_get_thread_limit` routine, see Section 3.2.14 on page 350.
- 3       • `omp_get_supported_active_levels`, see Section 3.2.15 on page 351.
- 4       • `omp_set_max_active_levels` routine, see Section 3.2.16 on page 352.
- 5       • `omp_get_max_active_levels` routine, see Section 3.2.17 on page 353.
- 6       • `omp_get_level` routine, see Section 3.2.18 on page 353.
- 7       • `omp_get_active_level` routine, see Section 3.2.21 on page 356.
- 8       • `omp_in_final` routine, see Section 3.2.22 on page 357.
- 9       • `omp_get_proc_bind` routine, see Section 3.2.23 on page 357.
- 10      • `omp_get_place_num_procs` routine, see Section 3.2.25 on page 360.
- 11      • `omp_get_place_proc_ids` routine, see Section 3.2.26 on page 360.
- 12      • `omp_get_partition_num_places` routine, see Section 3.2.28 on page 362.
- 13      • `omp_get_partition_place_nums` routine, see Section 3.2.29 on page 363.
- 14      • `omp_set_affinity_format` routine, see Section 3.2.30 on page 364.
- 15      • `omp_get_affinity_format` routine, see Section 3.2.31 on page 365.
- 16      • `omp_set_default_device` routine, see Section 3.2.34 on page 369.
- 17      • `omp_get_default_device` routine, see Section 3.2.35 on page 370.
- 18      • `omp_get_max_task_priority` routine, see Section 3.2.42 on page 374.
- 19      • `omp_set_num_teams` routine, see Section 3.2.45 on page 378.
- 20      • `omp_get_max_teams` routine, see Section 3.2.46 on page 379.
- 21      • `omp_set_teams_thread_limit` routine, see Section 3.2.47 on page 380.
- 22      • `omp_get_teams_thread_limit` routine, see Section 3.2.48 on page 381.
- 23      • `omp_set_default_allocator` routine, see Section 3.8.4 on page 417.
- 24      • `omp_get_default_allocator` routine, see Section 3.8.5 on page 418.

## 1 2.6.4 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>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
<i>debug-var</i>	global

*table continued on next page*



table continued from previous page

ICV	Scope
<i>num-procs-var</i>	device
<i>thread-num-var</i>	implicit task
<i>final-task-var</i>	data environment
<i>implicit-task-var</i>	data environment
<i>team-size-var</i>	team
<i>def-allocator-var</i>	implicit task
<i>nteams-var</i>	device
<i>teams-thread-limit-var</i>	device

## Description

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

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

### 2.6.4.1 How the Per-Data Environment ICVs Work

When a **task** construct, a **parallel** construct or a **teams** construct is encountered, each generated task inherits the values of the data environment scoped ICVs from each generating task's ICV values.

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

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

1 When a *target task* executes an active **target** region, the generated initial task uses the values of  
2 the data environment scoped ICVs from the device data environment ICV values of the device that  
3 will execute the region.

4 When a *target task* executes an inactive **target** region, the generated initial task uses the values  
5 of the data environment scoped ICVs from the data environment of the task that encountered the  
6 **target** construct.

7 If a **teams** construct with a **thread\_limit** clause is encountered, the *thread-limit-var* ICV  
8 from the data environment of the initial task for each team is instead set to an implementation  
9 defined value between one and the value specified in the clause.

10 If a **teams** construct with no **thread\_limit** clause is encountered, the *thread-limit-var* ICV  
11 from the data environment of the initial task of each team is set to an implementation defined value  
12 that is greater than zero and does not exceed *teams-thread-limit-var*, if *teams-thread-limit-var* is  
13 greater than zero.

14 When encountering a worksharing-loop region for which the **runtime** schedule kind is specified,  
15 all implicit task regions that constitute the binding parallel region must have the same value for  
16 *run-sched-var* in their data environments. Otherwise, the behavior is unspecified.

## 17 2.6.5 ICV Override Relationships

18 Table 2.4 shows the override relationships among construct clauses and ICVs. The table only lists  
19 ICV that can be overwritten by a clause.

**TABLE 2.4:** ICV Override Relationships

ICV	construct clause, if used
<i>nthreads-var</i>	<b>num_threads</b>
<i>run-sched-var</i>	<b>schedule</b>
<i>def-sched-var</i>	<b>schedule</b>
<i>bind-var</i>	<b>proc_bind</b>
<i>def-allocator-var</i>	<b>allocator</b>
<i>nteams-var</i>	<b>num_teams</b>
<i>teams-thread-limit-var</i>	<b>thread_limit</b>

### 20 Description

- 21 • The **num\_threads** clause overrides the value of the first element of the *nthreads-var* ICV.
- 22 • If a **schedule** clause specifies a modifier then that modifier overrides any modifier that is  
23 specified in the *run-sched-var* ICV.

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

## Cross References

- **parallel** construct, see Section 2.8 on page 79.
- **proc\_bind** clause, Section 2.8 on page 79.
- **num\_threads** clause, see Section 2.8.1 on page 84.
- **num\_teams** clause, see Section 2.9 on page 87.
- **thread\_limit** clause, see Section 2.9 on page 87.
- Worksharing-Loop construct, see Section 2.12.2 on page 106.
- **schedule** clause, see Section 2.12.2.1 on page 114.

## 2.7 Utility Directives

Utility directives facilitate interactions with the compiler and support code readability.

### 2.7.1 **nothing** Directive

#### Summary

The **nothing** directive has no effect. It can be used in a metadirective and other contexts to indicate explicitly that the intent is no effect on the execution of the OpenMP program.

#### Syntax

 C / C++

The syntax of the **nothing** directive is as follows:

```
#pragma omp nothing new-line
```

 C / C++

 Fortran

The syntax of the **nothing** directive is as follows:

```
!$omp nothing
```

 Fortran

#### Description

The **nothing** directive has no effect on the execution of the OpenMP program.

#### Cross References

- Metadirectives, see Section 2.3.4 on page 54.

## 2.7.2 error Directive

### Summary

The **error** directive instructs the compiler or runtime to display a message and to perform an error action. If the **at** clause is specified with the **runtime** value then the **error** directive is a stand-alone directive.

### Syntax

C / C++

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

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

where *clause* is one of the following:

```
at(compile | runtime)  
severity(fatal | warning)  
message (msg-string)
```

where *msg-string* is a string of const char \* type.

C / C++

Fortran

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

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

where *clause* is one of the following:

```
at(compile | runtime)  
severity(fatal | warning)  
message (msg-string)
```

where *msg-string* is a character string of character(len=\*) type

Fortran

## Description

The **error** directive performs an error action. The error action includes the display of an implementation defined message. The **severity** clause determines if the error action includes anything other than the message display.

The **at** clause determines when the implementation performs the error action. When the **at** clause specifies **compile**, the error action is performed during compilation if the **error** directive appears in a declarative context or in an executable context that is reachable at runtime. When the **at** clause specifies **compile** and the **error** directive appears in an executable context that is not reachable at runtime, the error action may or may not be performed. When the **at** clause specifies **runtime**, the error action is performed during program execution when a thread encounters the directive. If the **at** clause is not specified then the **error** directive behaves as if the **at** clause specifies **compile**.

The **severity** clause determines the action that the implementation performs. When the **severity** clause specifies **warning**, the implementation takes no action besides displaying the message. When the **severity** clause specifies **fatal** and the **at** clause specifies **compile** then the message is displayed and compilation of the current compilation unit is aborted. When the **severity** clause specifies **fatal** and the **at** clause specifies **runtime** then the message is displayed and program execution is aborted. If no **severity** clause is specified then the **error** directive behaves as if the **severity** clause specifies **fatal**.

If the **message** clause is specified then *msg-string* is included in the implementation defined message.

## Restrictions

Restrictions to the **error** directive are as follows:

- At most one **at** clause can appear on the directive.
- At most one **severity** clause can appear on the directive.
- At most one **message** clause can appear on the directive.

## Cross References

- Stand-Alone Directives, see Section [2.1.3](#) on page [41](#).

## 2.8 parallel Construct

### Summary

The parallel construct creates a team of OpenMP threads that execute the region.

## Syntax

C / C++

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

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

where *clause* is one of the following:

```
if ([ parallel : ] scalar-expression)  
    num_threads (integer-expression)  
    default (data-sharing-attribute)  
    private (list)  
    firstprivate (list)  
    shared (list)  
    copyin (list)  
    reduction ([reduction-modifier , ] reduction-identifier : list)  
    proc_bind(master | close | spread)  
    allocate ([allocator : ] list)
```

C / C++

Fortran

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

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

where *clause* is one of the following:

```
if ([ parallel : ] scalar-logical-expression)  
    num_threads (scalar-integer-expression)  
    default (data-sharing-attribute)  
    private (list)  
    firstprivate (list)  
    shared (list)  
    copyin (list)  
    reduction ([reduction-modifier , ] reduction-identifier : list)  
    proc_bind(master | close | spread)  
    allocate ([allocator : ] list)
```

## 1 Binding

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

## 4 Description

5 When a thread encounters a **parallel** construct, a team of threads is created to execute the  
6 **parallel** region (see Section 2.8.1 on page 84 for more information about how the number of  
7 threads in the team is determined, including the evaluation of the **if** and **num\_threads** clauses).  
8 The thread that encountered the **parallel** construct becomes the master thread of the new team,  
9 with a thread number of zero for the duration of the new **parallel** region. All threads in the new  
10 team, including the master thread, execute the region. Once the team is created, the number of  
11 threads in the team remains constant for the duration of that **parallel** region.

12 The optional **proc\_bind** clause, described in Section 2.8.2 on page 85, specifies the mapping of  
13 OpenMP threads to places within the current place partition, that is, within the places listed in the  
14 *place-partition-var* ICV for the implicit task of the encountering thread.

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

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

25 The implementation may cause any thread to suspend execution of its implicit task at a task  
26 scheduling point, and to switch to execution of any explicit task generated by any of the threads in  
27 the team, before eventually resuming execution of the implicit task (for more details see  
28 Section 2.13 on page 142).

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

31 If a thread in a team executing a **parallel** region encounters another **parallel** directive, it  
32 creates a new team, according to the rules in Section 2.8.1 on page 84, and it becomes the master of  
33 that new team.

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

## Execution Model Events

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

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

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

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

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

The *parallel-end* event occurs in the thread that encounters the **parallel** construct after the thread executes its *implicit-task-end* event but before the thread resumes execution of the encountering task.

If a native thread is destroyed at the end of a **parallel** region, a *native thread-end* event occurs in the thread as the last event prior to destruction of the thread.

## Tool Callbacks

A thread dispatches a registered **ompt\_callback\_parallel\_begin** callback for each occurrence of a *parallel-begin* event in that thread. The callback occurs in the task that encounters the **parallel** construct. This callback has the type signature **ompt\_callback\_parallel\_begin\_t**. In the dispatched callback, (*flags & ompt\_parallel\_team*) evaluates to *true*.

A thread dispatches a registered **ompt\_callback\_implicit\_task** callback with **ompt\_scope\_begin** as its *endpoint* argument for each occurrence of an *implicit-task-begin* event in that thread. Similarly, a thread dispatches a registered **ompt\_callback\_implicit\_task** callback with **ompt\_scope\_end** as its *endpoint* argument for each occurrence of an *implicit-task-end* event in that thread. The callbacks occur in the context of the implicit task and have type signature **ompt\_callback\_implicit\_task\_t**. In the dispatched callback, (*flags & ompt\_task\_implicit*) evaluates to *true*.

A thread dispatches a registered **ompt\_callback\_parallel\_end** callback for each occurrence of a *parallel-end* event in that thread. The callback occurs in the task that encounters the **parallel** construct. This callback has the type signature **ompt\_callback\_parallel\_end\_t**.

A thread dispatches a registered **ompt\_callback\_thread\_begin** callback for the *native-thread-begin* event in that thread. The callback occurs in the context of the thread. The callback has type signature **ompt\_callback\_thread\_begin\_t**.



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

## 4 **Restrictions**

5 Restrictions to the `parallel` construct are as follows:

- 6 • A program that branches into or out of a `parallel` region is non-conforming.
- 7 • A program must not depend on any ordering of the evaluations of the clauses of the `parallel`  
8 directive, or on any side effects of the evaluations of the clauses.
- 9 • At most one `if` clause can appear on the directive.
- 10 • At most one `proc_bind` clause can appear on the directive.
- 11 • At most one `num_threads` clause can appear on the directive. The `num_threads`  
12 expression must evaluate to a positive integer value.

▼ C++ ▲

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

▲ C++ ▼

## 15 **Cross References**

- 16 • OpenMP execution model, see Section 1.3 on page 21.
- 17 • `num_threads` clause, see Section 2.8 on page 79.
- 18 • `proc_bind` clause, see Section 2.8.2 on page 85.
- 19 • `allocate` clause, see Section 2.14.4 on page 164.
- 20 • `if` clause, see Section 2.18 on page 224.
- 21 • `default`, `shared`, `private`, `firstprivate`, and `reduction` clauses, see  
22 Section 2.22.4 on page 284.
- 23 • `copyin` clause, see Section 2.22.6 on page 313.
- 24 • `omp_get_thread_num` routine, see Section 3.2.4 on page 341.
- 25 • `ompt_scope_begin` and `ompt_scope_end`, see Section 4.4.4.11 on page 447.
- 26 • `ompt_callback_thread_begin_t`, see Section 4.5.2.1 on page 462.
- 27 • `ompt_callback_thread_end_t`, see Section 4.5.2.2 on page 462.
- 28 • `ompt_callback_parallel_begin_t`, see Section 4.5.2.3 on page 463.
- 29 • `ompt_callback_parallel_end_t`, see Section 4.5.2.4 on page 465.
- 30 • `ompt_callback_implicit_task_t`, see Section 4.5.2.11 on page 473.

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

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

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

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

---

### Algorithm 2.1

---

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

**let** *ActiveParRegions* be the number of enclosing active parallel regions;

**if** an `if` clause exists

**then let** *IfClauseValue* be the value of the `if` clause expression;

**else let** *IfClauseValue* = `true`;

**if** a `num_threads` clause exists

**then let** *ThreadsRequested* be the value of the `num_threads` clause expression;

**else let** *ThreadsRequested* = value of the first element of *nthreads-var*;

**let** *ThreadsAvailable* = (*thread-limit-var* - *ThreadsBusy* + 1);

**if** (*IfClauseValue* = `false`)

**then** number of threads = 1;

**else if** (*ActiveParRegions* = *max-active-levels-var*)

**then** number of threads = 1;

**else if** (*dyn-var* = `true`) **and** (*ThreadsRequested* ≤ *ThreadsAvailable*)

**then** 1 ≤ number of threads ≤ *ThreadsRequested*;

**else if** (*dyn-var* = `true`) **and** (*ThreadsRequested* > *ThreadsAvailable*)

**then** 1 ≤ number of threads ≤ *ThreadsAvailable*;

```
1     else if (dyn-var = false) and (ThreadsRequested ≤ ThreadsAvailable)
2     then number of threads = ThreadsRequested;
3     else if (dyn-var = false) and (ThreadsRequested > ThreadsAvailable)
4     then behavior is implementation defined;
```

---

```
7
8     ▼
9     Note – Since the initial value of the dyn-var ICV is implementation defined, programs that depend
10    on a specific number of threads for correct execution should explicitly disable dynamic adjustment
11    of the number of threads.
12    ▲
```

## 12 **Cross References**

- 13 • *nthreads-var*, *dyn-var*, *thread-limit-var*, and *max-active-levels-var* ICVs, see Section 2.6 on  
14 page 66.
- 15 • **parallel** construct, see Section 2.8 on page 79.
- 16 • **num\_threads** clause, see Section 2.8 on page 79.
- 17 • **if** clause, see Section 2.18 on page 224.

## 18 **2.8.2 Controlling OpenMP Thread Affinity**

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

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

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

- 1 •  $T \leq P$ : The master thread executes on the place of the parent thread. The thread with the next  
2 smallest thread number executes on the next place in the place partition, and so on, with wrap  
3 around with respect to the place partition of the master thread.
- 4 •  $T > P$ : Each place  $p$  will contain  $S_p$  threads with consecutive thread numbers where  
5  $\lfloor T/P \rfloor \leq S_p \leq \lceil T/P \rceil$ . The first  $S_0$  threads (including the master thread) are assigned to the  
6 place of the parent thread. The next  $S_1$  threads are assigned to the next place in the place  
7 partition, and so on, with wrap around with respect to the place partition of the master thread.  
8 When  $P$  does not divide  $T$  evenly, the exact number of threads in a particular place is  
9 implementation defined.

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

- 17 •  $T \leq P$ : The parent thread's place partition is split into  $T$  subpartitions, where each subpartition  
18 contains  $\lfloor P/T \rfloor$  or  $\lceil P/T \rceil$  consecutive places. A single thread is assigned to each subpartition.  
19 The master thread executes on the place of the parent thread and is assigned to the subpartition  
20 that includes that place. The thread with the next smallest thread number is assigned to the first  
21 place in the next subpartition, and so on, with wrap around with respect to the original place  
22 partition of the master thread.
- 23 •  $T > P$ : The parent thread's place partition is split into  $P$  subpartitions, each consisting of a  
24 single place. Each subpartition is assigned  $S_p$  threads with consecutive thread numbers, where  
25  $\lfloor T/P \rfloor \leq S_p \leq \lceil T/P \rceil$ . The first  $S_0$  threads (including the master thread) are assigned to the  
26 subpartition containing the place of the parent thread. The next  $S_1$  threads are assigned to the  
27 next subpartition, and so on, with wrap around with respect to the original place partition of the  
28 master thread. When  $P$  does not divide  $T$  evenly, the exact number of threads in a particular  
29 subpartition is implementation defined.

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

32

---

33 **Note** – Wrap around is needed if the end of a place partition is reached before all thread  
34 assignments are done. For example, wrap around may be needed in the case of **close** and  $T \leq P$ ,  
35 if the master thread is assigned to a place other than the first place in the place partition. In this  
36 case, thread 1 is assigned to the place after the place of the master place, thread 2 is assigned to the  
37 place after that, and so on. The end of the place partition may be reached before all threads are  
38 assigned. In this case, assignment of threads is resumed with the first place in the place partition.  
39

---

## 2.9 teams Construct

### Summary

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

### Syntax

C / C++

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

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

where *clause* is one of the following:

```
num_teams ([lower-bound : ] upper-bound)  
thread_limit (integer-expression)  
default (data-sharing-attribute)  
private (list)  
firstprivate (list)  
shared (list)  
reduction ([default , ] reduction-identifier : list)  
allocate ([allocator : ] list)
```

and where *lower-bound* and *upper-bound* are scalar integer expressions.

C / C++

Fortran

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

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

where *clause* is one of the following:

```
num_teams ([lower-bound : ] upper-bound)  
thread_limit (scalar-integer-expression)  
default (data-sharing-attribute)  
private (list)  
firstprivate (list)  
shared (list)
```

```
1  reduction ([default ,] reduction-identifier : list)
2  allocate ([allocator :] list)
```

3 and where *lower-bound* and *upper-bound* are scalar integer expressions.

## Fortran

### 4 Binding

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

### 6 Description

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

9 If the **num\_teams** clause is present, *lower-bound* is the specified lower bound and *upper-bound* is  
10 the specified upper bound on the number of teams requested. If a lower bound is not specified, the  
11 lower bound is set to the specified upper bound. The number of teams created is implementation  
12 defined, but it will be greater than or equal to the lower bound and less than or equal to the upper  
13 bound.

14 If the **num\_teams** clause is not specified and the value of the *nteam*-var ICV is greater than zero,  
15 the number of teams created is less or equal to the value of the *nteam*-var ICV. Otherwise, the  
16 number of teams created is implementation defined, but it will be greater than or equal to 1.

17 A thread may obtain the number of teams created by the construct with a call to the  
18 **omp\_get\_num\_teams** routine.

19 As described in Section 2.6.4.1, the **teams** construct limits the number of threads that may  
20 participate in a contention group initiated by each team by setting the value of the *thread-limit-var*  
21 ICV for the initial task to an implementation defined value greater than zero. If the  
22 **thread\_limit** clause is specified, the number of threads will be less than or equal to the value  
23 specified in the clause. Otherwise, if the *teams-thread-limit-var* ICV is greater than zero, the  
24 number of threads will be less than or equal to that value.

25 On a combined or composite construct that includes **target** and **teams** constructs, the  
26 expressions in **num\_teams** and **thread\_limit** clauses are evaluated on the host device on  
27 entry to the **target** construct.

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

30 Within a **teams** region, initial team numbers uniquely identify each initial team. Initial team  
31 numbers are consecutive whole numbers ranging from zero to one less than the number of initial  
32 teams. A thread may obtain its own initial team number by a call to the **omp\_get\_team\_num**  
33 library routine. The policy for assigning the initial threads to places is implementation defined. The  
34 **teams** construct sets the *place-partition-var* and *default-device-var* ICVs for each initial thread to  
35 an implementation-defined value.

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

### 3 **Execution Model Events**

4 The *teams-begin* event occurs in a thread that encounters a **teams** construct before any initial task  
5 is created for the corresponding **teams** region.

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

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

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

12 The *teams-end* event occurs in the thread that encounters the **teams** construct after the thread  
13 executes its *initial-task-end* event but before it resumes execution of the encountering task.

14 If a native thread is destroyed at the end of a **teams** region, a *native-thread-end* event occurs in the  
15 thread as the last event prior to destruction of the thread.

### 16 **Tool Callbacks**

17 A thread dispatches a registered **ompt\_callback\_parallel\_begin** callback for each  
18 occurrence of a *teams-begin* event in that thread. The callback occurs in the task that encounters the  
19 **teams** construct. This callback has the type signature  
20 **ompt\_callback\_parallel\_begin\_t**. In the dispatched callback,  
21 *(flags & ompt\_parallel\_league)* evaluates to *true*.

22 A thread dispatches a registered **ompt\_callback\_implicit\_task** callback with  
23 **ompt\_scope\_begin** as its *endpoint* argument for each occurrence of an *initial-task-begin* in  
24 that thread. Similarly, a thread dispatches a registered **ompt\_callback\_implicit\_task**  
25 callback with **ompt\_scope\_end** as its *endpoint* argument for each occurrence of an  
26 *initial-task-end* event in that thread. The callbacks occur in the context of the initial task and have  
27 type signature **ompt\_callback\_implicit\_task\_t**. In the dispatched callback,  
28 *(flags & ompt\_task\_initial)* evaluates to *true*.

29 A thread dispatches a registered **ompt\_callback\_parallel\_end** callback for each  
30 occurrence of a *teams-end* event in that thread. The callback occurs in the task that encounters the  
31 **teams** construct. This callback has the type signature **ompt\_callback\_parallel\_end\_t**.

32 A thread dispatches a registered **ompt\_callback\_thread\_begin** callback for the  
33 *native-thread-begin* event in that thread. The callback occurs in the context of the thread. The  
34 callback has type signature **ompt\_callback\_thread\_begin\_t**.

35 A thread dispatches a registered **ompt\_callback\_thread\_end** callback for the  
36 *native-thread-end* event in that thread. The callback occurs in the context of the thread. The  
37 callback has type signature **ompt\_callback\_thread\_end\_t**.

## Restrictions

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

- A program that branches into or out of a **teams** region is non-conforming.
- A program must not depend on any ordering of the evaluations of the clauses of the **teams** directive, or on any side effects of the evaluation of the clauses.
- At most one **thread\_limit** clause can appear on the directive. The **thread\_limit** expression must evaluate to a positive integer value.
- At most one **num\_teams** clause can appear on the directive. The *lower-bound* and *upper-bound* specified in the **num\_teams** clause must evaluate to positive integer values.
- A **teams** region must be strictly nested within the implicit parallel region that surrounds the whole OpenMP program or a **target** region. If a **teams** region is nested inside a **target** region, the corresponding **target** construct must contain no statements, declarations or directives outside of the corresponding **teams** construct.
- **distribute**, **distribute simd**, distribute parallel worksharing-loop, distribute parallel worksharing-loop SIMD, **parallel** regions, including any **parallel** regions arising from combined constructs, **omp\_get\_num\_teams()** regions, and **omp\_get\_team\_num()** regions are the only OpenMP regions that may be strictly nested inside the **teams** region.

## Cross References

- **parallel** construct, see Section 2.8 on page 79.
- **distribute** construct, see Section 2.12.4.1 on page 125.
- **distribute simd** construct, see Section 2.12.4.2 on page 129.
- **allocate** clause, see Section 2.14.4 on page 164.
- **target** construct, see Section 2.15.5 on page 174.
- **default**, **shared**, **private**, **firstprivate**, and **reduction** clauses, see Section 2.22.4 on page 284.
- **omp\_get\_num\_teams** routine, see Section 3.2.38 on page 372.
- **omp\_get\_team\_num** routine, see Section 3.2.39 on page 372.
- **ompt\_callback\_thread\_begin\_t**, see Section 4.5.2.1 on page 462.
- **ompt\_callback\_thread\_end\_t**, see Section 4.5.2.2 on page 462.
- **ompt\_callback\_parallel\_begin\_t**, see Section 4.5.2.3 on page 463.
- **ompt\_callback\_parallel\_end\_t**, see Section 4.5.2.4 on page 465.
- **ompt\_callback\_implicit\_task\_t**, see Section 4.5.2.11 on page 473.



## 2.10 scope Construct

### Summary

The **scope** construct defines a structured block that is executed by all threads in a team but where additional OpenMP operations can be specified.

### Syntax

C / C++

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

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

where *clause* is one of the following:

```
private (list)  
reduction ([reduction-modifier ,] reduction-identifier : list)  
nowait
```

C / C++

Fortran

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

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

where *clause* is one of the following:

```
private (list)  
reduction ([reduction-modifier ,] reduction-identifier : list)
```

Fortran

### Binding

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

### Description

All encountering threads will execute the structured block associated with the **scope** construct. There is an implicit barrier at the end of a **scope** construct unless a **nowait** clause is specified.

## Execution Model Events

The *scope-begin* event occurs after an **implicit task** encounters a **scope** construct but before the task starts to execute the structured block of the **scope** region.

The *scope-end* event occurs after an implicit task finishes execution of a **scope** region but before it resumes execution of the enclosing region.

## Tool Callbacks

A thread dispatches a registered **ompt\_callback\_work** callback with **ompt\_scope\_begin** as its *endpoint* argument and **ompt\_work\_scope** as its *wstype* argument for each occurrence of a *scope-begin* event in that thread. Similarly, a thread dispatches a registered **ompt\_callback\_work** callback with **ompt\_scope\_end** as its *endpoint* argument and **ompt\_work\_scope** as its *wstype* argument for each occurrence of a *scope-end* event in that thread. The callbacks occur in the context of the implicit task. The callbacks have type signature **ompt\_callback\_work\_t**.

## Restrictions

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

- Each **scope** region must be encountered by all threads in a team or by none at all, unless cancellation has been requested for the innermost enclosing parallel region.
- The sequence of worksharing regions, **scope** regions and **barrier** regions encountered must be the same for every thread in a team.
- At most one **nowait** clause can appear on a **scope** construct.

▼ C++ ▼

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

▲ C++ ▲

## Cross References

- **reduction** clause, see Section [2.22.4](#) on page [284](#).
- **ompt\_scope\_begin** and **ompt\_scope\_end**, see Section [4.4.4.11](#) on page [447](#).
- **ompt\_work\_scope**, see Section [4.4.4.15](#) on page [449](#).
- **ompt\_callback\_work\_t**, see Section [4.5.2.5](#) on page [466](#).

## 2.11 Worksharing Constructs

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

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

The OpenMP API defines the worksharing constructs that are described in this section as well as the worksharing-loop construct, which is described in Section 2.12.2 on page 106.

### Restrictions

The following restrictions apply to worksharing constructs:

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

### 2.11.1 sections Construct

#### Summary

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

#### Syntax

C / C++

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

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

1 where *clause* is one of the following:

```
2 private (list)
3 firstprivate (list)
4 lastprivate ([ lastprivate-modifier : ] list)
5 reduction ([reduction-modifier , ] reduction-identifier : list)
6 allocate ([allocator : ] list)
7 nowait
```

C / C++

Fortran

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

```
9 !$omp sections [clause[ [, ] clause] ... ]
10     [!$omp section]
11     structured-block-sequence
12     [!$omp section
13     structured-block-sequence]
14     ...
15 !$omp end sections [nowait]
```

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 allocate ([allocator : ] list)
```

Fortran

## 22 Binding

23 The binding thread set for a **sections** region is the current team. A **sections** region binds to  
24 the innermost enclosing **parallel** region. Only the threads of the team that executes the binding  
25 **parallel** region participate in the execution of the structured block sequences and the implied  
26 barrier of the **sections** region if the barrier is not eliminated by a **nowait** clause.

## Description

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

The method of scheduling the structured block sequences 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.

## Execution Model Events

The *section-begin* event occurs after an implicit task encounters a **sections** construct but before the task executes any structured block sequences of the **sections** region.

The *sections-end* event occurs after an implicit task finishes execution of a **sections** region but before it resumes execution of the enclosing context.

The *section-begin* event occurs before an implicit task starts to execute a structured block sequence in the **sections** construct for each of those structured block sequences that the task executes.

## Tool Callbacks

A thread dispatches a registered **ompt\_callback\_work** callback with **ompt\_scope\_begin** as its *endpoint* argument and **ompt\_work\_sections** as its *wstype* argument for each occurrence of a *section-begin* event in that thread. Similarly, a thread dispatches a registered **ompt\_callback\_work** callback with **ompt\_scope\_end** as its *endpoint* argument and **ompt\_work\_sections** as its *wstype* argument for each occurrence of a *sections-end* event in that thread. The callbacks occur in the context of the implicit task. The callbacks have type signature **ompt\_callback\_work\_t**.

A thread dispatches a registered **ompt\_callback\_dispatch** callback for each occurrence of a *section-begin* event in that thread. The callback occurs in the context of the implicit task. The callback has type signature **ompt\_callback\_dispatch\_t**.

## Restrictions

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

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

▼ C++ ▲

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

▲ C++ ▼

## Cross References

- **allocate** clause, see Section 2.14.4 on page 164.
- **private**, **firstprivate**, **lastprivate**, and **reduction** clauses, see Section 2.22.4 on page 284.
- **ompt\_scope\_begin** and **ompt\_scope\_end**, see Section 4.4.4.11 on page 447.
- **ompt\_work\_sections**, see Section 4.4.4.15 on page 449.
- **ompt\_callback\_work\_t**, see Section 4.5.2.5 on page 466.
- **ompt\_callback\_dispatch\_t**, see Section 4.5.2.6 on page 467.

## 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)
allocate ([allocator : ] list)
nowait
```

C / C++

Fortran

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

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

where *clause* is one of the following:

```
1   private (list)
2   firstprivate (list)
3   allocate ([allocator :] list)
```

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

```
5   copyprivate (list)
6   nowait
```

## Fortran

### Binding

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

### Description

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

### Execution Model Events

17 The *single-begin* event occurs after an **implicit task** encounters a **single** construct but before the task starts to execute the structured block of the **single** region.

20 The *single-end* event occurs after an implicit task finishes execution of a **single** region but before it resumes execution of the enclosing region.

### Tool Callbacks

22 A thread dispatches a registered **ompt\_callback\_work** callback with **ompt\_scope\_begin** as its *endpoint* argument for each occurrence of a *single-begin* event in that thread. Similarly, a thread dispatches a registered **ompt\_callback\_work** callback with **ompt\_scope\_end** as its *endpoint* argument for each occurrence of a *single-end* event in that thread. For each of these callbacks, the *wstype* argument is **ompt\_work\_single\_executor** if the thread executes the structured block associated with the **single** region; otherwise, the *wstype* argument is **ompt\_work\_single\_other**. The callback has type signature **ompt\_callback\_work\_t**.

## Restrictions

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

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

C++

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

C++

## Cross References

- **allocate** clause, see Section 2.14.4 on page 164.
- **private** and **firstprivate** clauses, see Section 2.22.4 on page 284.
- **copyprivate** clause, see Section 2.22.6.2 on page 315.
- **ompt\_scope\_begin** and **ompt\_scope\_end**, see Section 4.4.4.11 on page 447.
- **ompt\_work\_single\_executor** and **ompt\_work\_single\_other**, see Section 4.4.4.15 on page 449.
- **ompt\_callback\_work\_t**, Section 4.5.2.5 on page 466.

Fortran

## 2.11.3 workshare Construct

### Summary

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

### Syntax

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

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

### Binding

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



## Description

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

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

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

- For array expressions within each statement, including transformational array intrinsic functions that compute scalar values from arrays:
  - Evaluation of each element of the array expression, including any references to elemental functions, is a unit of work.
  - Evaluation of transformational array intrinsic functions may be freely subdivided into any number of units of work.
- For an array assignment statement, the assignment of each element is a unit of work.
- For a scalar assignment statement, the assignment operation is a unit of work.
- For a **WHERE** statement or construct, the evaluation of the mask expression and the masked assignments are each a unit of work.
- For a **FORALL** statement or construct, the evaluation of the mask expression, expressions occurring in the specification of the iteration space, and the masked assignments are each a unit of work.
- For an **atomic** construct, the atomic operation on the storage location designated as  $x$  is a unit of work.
- For a **critical** construct, the construct is a single unit of work.
- For a **parallel** construct, the construct is a unit of work with respect to the **workshare** construct. The statements contained in the **parallel** construct are executed by a new thread team.
- If none of the rules above apply to a portion of a statement in the structured block, then that portion is a unit of work.

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

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

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

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

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

## 8 Execution Model Events

9 The *workshare-begin* event occurs after an implicit task encounters a **workshare** construct but  
 10 before the task starts to execute the structured block of the **workshare** region.

11 The *workshare-end* event occurs after an implicit task finishes execution of a **workshare** region  
 12 but before it resumes execution of the enclosing context.

## 13 Tool Callbacks

14 A thread dispatches a registered **ompt\_callback\_work** callback with **ompt\_scope\_begin**  
 15 as its *endpoint* argument and **ompt\_work\_workshare** as its *wstype* argument for each  
 16 occurrence of a *workshare-begin* event in that thread. Similarly, a thread dispatches a registered  
 17 **ompt\_callback\_work** callback with **ompt\_scope\_end** as its *endpoint* argument and  
 18 **ompt\_work\_workshare** as its *wstype* argument for each occurrence of a *workshare-end* event  
 19 in that thread. The callbacks occur in the context of the implicit task. The callbacks have type  
 20 signature **ompt\_callback\_work\_t**.

## 21 Restrictions

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

- 23 • The only OpenMP constructs that may be closely nested inside a **workshare** construct are the  
 24 **atomic**, **critical**, and **parallel** constructs.
- 25 • Base language statements that are encountered inside a **workshare** construct but that are not  
 26 enclosed within a **parallel** construct that is nested inside the **workshare** construct must  
 27 consist of only the following:
  - 28 – array assignments
  - 29 – scalar assignments
  - 30 – **FORALL** statements
  - 31 – **FORALL** constructs
  - 32 – **WHERE** statements
  - 33 – **WHERE** constructs

- 1 • All array assignments, scalar assignments, and masked array assignments that are encountered  
2 inside a **workshare** construct but are not nested inside a **parallel** construct that is nested  
3 inside the **workshare** construct must be intrinsic assignments.
- 4 • The construct must not contain any user defined function calls unless the function is  
5 **ELEMENTAL** or the function call is contained inside a **parallel** construct that is nested inside  
6 the **workshare** construct.

## 7 Cross References

- 8 • **parallel** construct, see Section 2.8 on page 79.
- 9 • **critical** construct, see Section 2.20.1 on page 226.
- 10 • **atomic** construct, see Section 2.20.7 on page 237.
- 11 • **ompt\_scope\_begin** and **ompt\_scope\_end**, see Section 4.4.4.11 on page 447.
- 12 • **ompt\_work\_workshare**, see Section 4.4.4.15 on page 449.
- 13 • **ompt\_callback\_work\_t**, see Section 4.5.2.5 on page 466.

▲ Fortran ▲

## 14 2.12 Loop-Related Directives

### 15 2.12.1 Canonical Loop Form

▼ C / C++ ▼

16 The loops associated with a loop-associated directive have *canonical loop form* if they conform to  
17 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:

---

*continued on next page*

*continued from previous page*


---

	<pre> ++var var++ -- var var -- var += incr var -= incr var = var + incr var = incr + var var = var - incr </pre>
<i>var</i>	<p>One of the following:</p> <ul style="list-style-type: none"> <li>A variable of a signed or unsigned integer type.</li> <li>For C++, a variable of a random access iterator type.</li> <li>For C, a variable of a pointer type.</li> </ul> <p>This variable must not be modified during the execution of the <i>for-loop</i> other than in <i>incr-expr</i>.</p>
<i>relational-op</i>	<p>One of the following:</p> <pre> &lt; &lt;= &gt; &gt;= != </pre>
<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> <pre> var-outer var-outer + a2 a2 + var-outer var-outer - a2 a2 - var-outer a1 * var-outer a1 * var-outer + a2 a2 + a1 * var-outer a1 * var-outer - a2 a2 - a1 * var-outer </pre>

---

*continued on next page*

*continued from previous page*

---

*var-outer* \* *a1*  
*var-outer* \* *a1* + *a2*  
*a2* + *var-outer* \* *a1*  
*var-outer* \* *a1* - *a2*  
*a2* - *var-outer* \* *a1*

*incr* An integer expression that is loop invariant with respect to the outermost associated loop.

---



- 1 The loops associated with a loop-associated directive have *canonical loop form* if each of them is a  
2 *do-loop* that is a *do-construct* or an *inner-shared-do-construct* as defined by the Fortran standard. If  
3 an **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* ]

---

*var* A variable of integer type.

*lb* and *b* Expressions of a type compatible with the type of *var* that are loop invariant with respect to the outermost associated loop or are one of the following (where *var-outer*, *a1*, and *a2* have a type compatible with the type of *var*, *var-outer* is *var* from an outer associated loop, and *a1* and *a2* are loop invariant integer expressions with respect to the outermost loop):

*var-outer*  
*var-outer* + *a2*  
*a2* + *var-outer*  
*var-outer* - *a2*  
*a2* - *var-outer*  
*a1* \* *var-outer*  
*a1* \* *var-outer* + *a2*  
*a2* + *a1* \* *var-outer*  
*a1* \* *var-outer* - *a2*  
*a2* - *a1* \* *var-outer*  
*var-outer* \* *a1*  
*var-outer* \* *a1* + *a2*

---

*continued on next page*

continued from previous page

$a2 + \text{var-outer} * a1$

$\text{var-outer} * a1 - a2$

$a2 - \text{var-outer} * a1$

*incr* 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

1 The canonical form allows the iteration count of all associated loops to be computed before  
2 executing the outermost loop. The *incr* and *range-expr* are evaluated before executing the  
3 loop-associated construct. If *b* or *lb* is loop invariant with respect to the outermost associated loop,  
4 it is evaluated before executing the loop-associated construct. If *b* or *lb* is not loop invariant with  
5 respect to the outermost associated loop, *a1* and/or *a2* are evaluated before executing the  
6 loop-associated construct. The computation is performed for each loop in an integer type. This type  
7 is derived from the type of *var* as follows:

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

C++

9 • If *var* is of a random access iterator type, then the type is the type that would be used by  
10 *std::distance* applied to variables of the type of *var*.

C++

C

11 • If *var* is of a pointer type, then the type is `ptrdiff_t`.

C

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

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

17 **Note** – Random access iterators are required to support random access to elements in constant  
18 time. Other iterators are precluded by the restrictions since they can take linear time or offer limited  
19 functionality. The use of tasks to parallelize those cases is therefore advisable.  
20  
21

## C++

1 A range-based for loop that is valid in the base language and has a begin value that satisfies the  
2 random access iterator requirement has *canonical loop form*. Range-based for loops are of the  
3 following form:

4 **for** (*range-decl*: *range-expr*) *structured-block*

5 The *begin-expr* and *end-expr* expressions are derived from *range-expr* by the base language and  
6 assigned to variables to which this specification refers as `__begin` and `__end` respectively. Both  
7 `__begin` and `__end` are privatized. For the purpose of the rest of the standard `__begin` is the  
8 iteration variable of the range-for loop.

## C++

### 9 Restrictions

10 The following restrictions also apply:

## C / C++

- 11 • If *test-expr* is of the form *var relational-op b* and *relational-op* is `<` or `<=` then *incr-expr* must  
12 cause *var* to increase on each iteration of the loop. If *test-expr* is of the form *var relational-op b*  
13 and *relational-op* is `>` or `>=` then *incr-expr* must cause *var* to decrease on each iteration of the  
14 loop.
- 15 • If *test-expr* is of the form *b relational-op var* and *relational-op* is `<` or `<=` then *incr-expr* must  
16 cause *var* to decrease on each iteration of the loop. If *test-expr* is of the form *b relational-op var*  
17 and *relational-op* is `>` or `>=` then *incr-expr* must cause *var* to increase on each iteration of the  
18 loop.
- 19 • If *test-expr* is of the form *b != var* or *var != b* then *incr-expr* must cause *var* either to increase on  
20 each iteration of the loop or to decrease on each iteration of the loop.
- 21 • If *relational-op* is `!=` and *incr-expr* is of the form that has *incr* then *incr* must be a constant  
22 expression and evaluate to -1 or 1.

## C / C++

## C++

- 23 • In the `simd` construct the only random access iterator types that are allowed for *var* are pointer  
24 types.
- 25 • The *range-expr* of a range-for loop must be loop invariant with respect to the outermost  
26 associated loop, and must not reference iteration variables of any associated loops.
- 27 • The loops associated with an ordered clause with a parameter may not include range-for loops.

## C++

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

## Cross References

- **simd** construct, see Section 2.12.3.1 on page 115.
- **lastprivate** clause, see Section 2.22.4.5 on page 290.
- **linear** clause, see Section 2.22.4.6 on page 292.

## 2.12.2 Worksharing-Loop Construct

### Summary

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

### Syntax

C / C++

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

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

where clause is one of the following:

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





## Description

The worksharing-loop construct is associated with a loop nest that consists 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 is zero and that loop does not enclose the intervening code, 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.

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

1 If any of the associated loop(s) is a generated loop of a **tile** directive, the effect of the  
2 **schedule** clause on the loop schedule is implementation defined.

3 See Section 2.12.2.1 on page 114 for details of how the schedule for a worksharing-loop region is  
4 determined.

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

6 The schedule *modifier* can be one of those specified in Table 2.6. If the **static** schedule kind is  
7 specified or if the **ordered** clause is specified, and if the **nonmonotonic** modifier is not  
8 specified, the effect is as if the **monotonic** modifier is specified. Otherwise, unless the  
9 **monotonic** modifier is specified, the effect is as if the **nonmonotonic** modifier is specified. If  
10 a **schedule** clause specifies a modifier then that modifier overrides any modifier that is specified  
11 in the *run-sched-var* ICV.

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

---

<b>static</b>	<p>When <i>kind</i> is <b>static</b>, iterations are divided into chunks of size <i>chunk_size</i>, and the chunks are assigned to the threads in the team in a round-robin fashion in the order of the thread number. Each chunk contains <i>chunk_size</i> iterations, except for the chunk that contains the sequentially last iteration, which may have fewer iterations.</p> <p>When no <i>chunk_size</i> is specified, the iteration space is divided into chunks that are approximately equal in size, and at most one chunk is distributed to each thread. The size of the chunks is unspecified in this case.</p> <p>A compliant implementation of the <b>static</b> schedule must ensure that the same assignment of logical iteration numbers to threads will be used in two 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 <i>chunk_size</i> specified, or both worksharing-loop regions have no <i>chunk_size</i> 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 <b>nowait</b> clause.</p>
<b>dynamic</b>	<p>When <i>kind</i> is <b>dynamic</b>, 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 <i>chunk_size</i> iterations, except for the chunk that contains the sequentially last iteration, which may have fewer iterations.</p> <p>When no <i>chunk_size</i> is specified, it defaults to 1.</p>
<b>guided</b>	<p>When <i>kind</i> is <b>guided</b>, 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> <p>For a <i>chunk_size</i> of 1, the size of each chunk is proportional to the number of unassigned iterations divided by the number of threads in the team, decreasing to 1. For a <i>chunk_size</i> with value <i>k</i> (greater than 1), the size of each chunk is determined in the same way, with the restriction that the chunks do not contain fewer than <i>k</i> iterations (except for the chunk that contains the sequentially last iteration, which may have fewer than <i>k</i> iterations).</p>

---

*table continued on next page*

---

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

---

1

2 **Note** – For a team of  $p$  threads and a loop of  $n$  iterations, let  $\lceil n/p \rceil$  be the integer  $q$  that satisfies

3  $n = p * q - r$ , with  $0 \leq r < p$ . One compliant implementation of the **static** schedule (with no

4 specified *chunk\_size*) would behave as though *chunk\_size* had been specified with value  $q$ . Another

5 compliant implementation would assign  $q$  iterations to the first  $p - r$  threads, and  $q - 1$  iterations to

6 the remaining  $r$  threads. This illustrates why a conforming program must not rely on the details of a

7 particular implementation.

8 A compliant implementation of the **guided** schedule with a *chunk\_size* value of  $k$  would assign

9  $q = \lceil n/p \rceil$  iterations to the first available thread and set  $n$  to the larger of  $n - q$  and  $p * k$ . It would

10 then repeat this process until  $q$  is greater than or equal to the number of remaining iterations, at

11 which time the remaining iterations form the final chunk. Another compliant implementation could

12 use the same method, except with  $q = \lceil n/(2p) \rceil$ , and set  $n$  to the larger of  $n - q$  and  $2 * p * k$ .

13

**TABLE 2.6:** `schedule` Clause *modifier* Values

---

	<b>monotonic</b>	When the <b>monotonic</b> modifier is specified then each thread executes the chunks that it is assigned in increasing logical iteration order.
	<b>nonmonotonic</b>	When the <b>nonmonotonic</b> 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.
14	<b>simd</b>	When the <b>simd</b> 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 $new\_chunk\_size$ iterations except if it is also the last chunk. The last chunk may have fewer iterations than $new\_chunk\_size$ . If the <b>simd</b> modifier is specified and the loop is not associated with a SIMD construct, the modifier is ignored.

---

## 1 Execution Model Events

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

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

6 The *ws-loop-iteration-begin* event occurs once for each iteration of a worksharing-loop before the  
7 iteration is executed by an implicit task.

## 8 Tool Callbacks

9 A thread dispatches a registered **ompt\_callback\_work** callback with **ompt\_scope\_begin**  
10 as its *endpoint* argument and **work\_loop** as its *wstype* argument for each occurrence of a  
11 *ws-loop-begin* event in that thread. Similarly, a thread dispatches a registered  
12 **ompt\_callback\_work** callback with **ompt\_scope\_end** as its *endpoint* argument and  
13 **work\_loop** as its *wstype* argument for each occurrence of a *ws-loop-end* event in that thread. The  
14 callbacks occur in the context of the implicit task. The callbacks have type signature  
15 **ompt\_callback\_work\_t**.

16 A thread dispatches a registered **ompt\_callback\_dispatch** callback for each occurrence of a  
17 *ws-loop-iteration-begin* event in that thread. The callback occurs in the context of the implicit task.  
18 The callback has type signature **ompt\_callback\_dispatch\_t**.

## 19 Restrictions

20 Restrictions to the worksharing-loop construct are as follows:

- 21 • No OpenMP directive may appear in the region between any associated loops.
- 22 • If a **collapse** clause is specified, exactly one loop must occur in the region at each nesting  
23 level up to the number of loops specified by the parameter of the **collapse** clause.
- 24 • If the **ordered** clause is present, all loops associated with the construct must be perfectly  
25 nested; that is there must be no intervening code between any two loops.
- 26 • If a **reduction** clause with the **inscan** modifier is specified, neither the **ordered** nor  
27 **schedule** clause may appear on the worksharing-loop directive.
- 28 • The values of the loop control expressions of the loops associated with the worksharing-loop  
29 construct must be the same for all threads in the team.
- 30 • At most one **schedule** clause can appear on a worksharing-loop directive.
- 31 • The **schedule** clause must not appear on the worksharing-loop directive if the associated  
32 loop(s) form a non-rectangular loop nest.
- 33 • The **ordered** clause must not appear on the worksharing-loop directive if the associated  
34 loop(s) form a non-rectangular loop nest.

- 1       • The **ordered** clause must not appear on the worksharing-loop directive if the associated loops
- 2       include the generated loops of a **tile** directive.
- 3       • At most one **collapse** clause can appear on a worksharing-loop directive.
- 4       • *chunk\_size* must be a loop invariant integer expression with a positive value.
- 5       • The value of the *chunk\_size* expression must be the same for all threads in the team.
- 6       • The value of the *run-sched-var* ICV must be the same for all threads in the team.
- 7       • When **schedule(runtime)** or **schedule(auto)** is specified, *chunk\_size* must not be
- 8       specified.
- 9       • A *modifier* may not be specified on a **linear** clause.
- 10      • At most one **ordered** clause can appear on a worksharing-loop directive.
- 11      • The **ordered** clause must be present on the worksharing-loop construct if any **ordered**
- 12      region ever binds to a worksharing-loop region arising from the worksharing-loop construct.
- 13      • The **nonmonotonic** modifier cannot be specified if an **ordered** clause is specified.
- 14      • Each *modifier* may be specified at most once on the same **schedule** clause.
- 15      • Either the **monotonic** modifier or the **nonmonotonic** modifier can be specified but not both.
- 16      • The loop iteration variable may not appear in a **threadprivate** directive.
- 17      • If both the **collapse** and **ordered** clause with a parameter are specified, the parameter of the
- 18      **ordered** clause must be greater than or equal to the parameter of the **collapse** clause.
- 19      • A **linear** clause or an **ordered** clause with a parameter can be specified on a
- 20      worksharing-loop directive but not both.
- 21      • At most one **order(concurrent)** clause can appear on a worksharing-loop directive.
- 22      • If an **order(concurrent)** clause is present, all restrictions from the **loop** construct with an
- 23      **order(concurrent)** clause also apply.
- 24      • If an **order(concurrent)** clause is present, an **ordered** clause may not appear on the
- 25      same directive.

C / C++

- 26      • The associated *for-loops* must be structured blocks.
- 27      • Only an iteration of the innermost associated loop may be curtailed by a **continue** statement.
- 28      • No statement can branch to any associated **for** statement.
- 29      • At most one **nowait** clause can appear on a **for** directive.

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

C / C++

Fortran

- The associated *do-loops* must be structured blocks.
- Only an iteration of the innermost associated loop may be curtailed by a **CYCLE** statement.
- No statement in the associated loops other than the **DO** statements can cause a branch out of the loops.
- 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.

Fortran

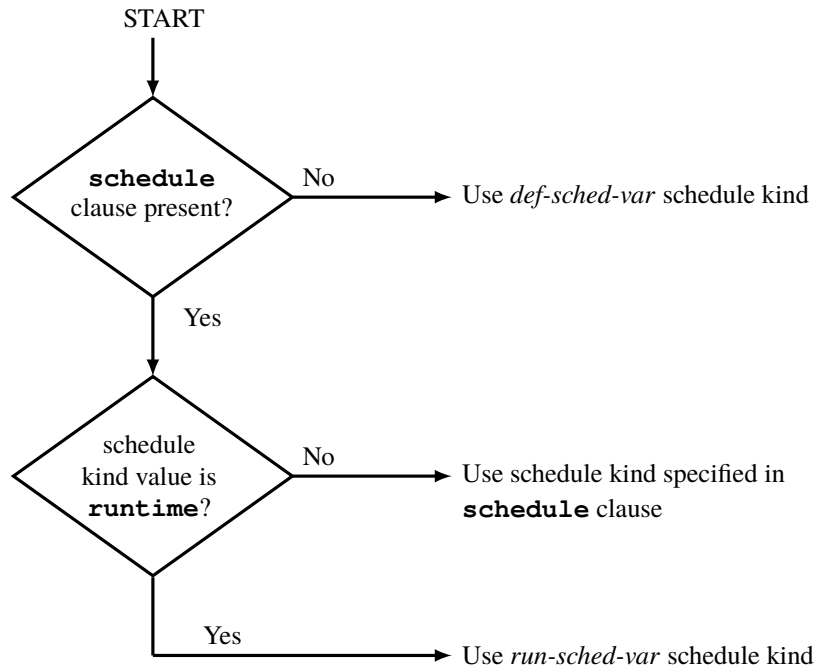
## Cross References

- **order (concurrent)** clause, see Section 2.12.5 on page 133.
- **tile** construct, see Section 2.12.7.1 on page 140.
- **ordered** construct, see Section 2.20.9 on page 252.
- **depend** clause, see Section 2.20.11 on page 257.
- Data-sharing attribute clauses, see Section 2.22.4 on page 284.
- **ompt\_scope\_begin** and **ompt\_scope\_end**, see Section 4.4.4.11 on page 447.
- **ompt\_work\_loop**, see Section 4.4.4.15 on page 449.
- **ompt\_callback\_work\_t**, see Section 4.5.2.5 on page 466.
- **OMP\_SCHEDULE** environment variable, see Section 6.1 on page 585.

### 2.12.2.1 Determining the Schedule of a Worksharing-Loop

When execution encounters a worksharing-loop directive, the **schedule** clause (if any) on the directive, and the *run-sched-var* and *def-sched-var* ICVs are used to determine how loop iterations are assigned to threads. See Section 2.6 on page 66 for details of how the values of the ICVs are determined. If the worksharing-loop directive does not have a **schedule** clause then the current value of the *def-sched-var* ICV determines the schedule. If the worksharing-loop directive has a **schedule** clause that specifies the **runtime** schedule kind then the current value of the *run-sched-var* ICV determines the schedule. Otherwise, the value of the **schedule** clause determines the schedule. Figure 2.1 describes how the schedule for a worksharing-loop is determined.





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

## Cross References

- ICVs, see Section 2.6 on page 66.

## 2.12.3 SIMD Directives

### 2.12.3.1 `simd` Construct

#### Summary

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

#### Syntax

C / C++

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

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

C / C++

C++

The syntax of the **simd** construct is also as follows:

```
[[omp :: simd [(clause[ [, ] clause] ... )] ] ]  
for-loops
```

C++

C / C++

where *clause* is one of the following:

```
if([ simd :] scalar-expression)  
safelen (length)  
simdlen (length)  
linear (list[ : linear-step])  
aligned (list[ : alignment])  
nontemporal (list)  
private (list)  
lastprivate ([ lastprivate-modifier : ] list)  
reduction ([ reduction-modifier , ]reduction-identifier : list)  
collapse (n)  
order (concurrent)
```

The **simd** directive places restrictions on the structure of the associated *for-loops*. Specifically, all associated *for-loops* must have *canonical loop form* (Section 2.12.1 on page 101).

C / C++

Fortran

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

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

where *clause* is one of the following:

```
if([ simd :] scalar-logical-expression)  
safelen (length)  
simdlen (length)  
linear (list[ : linear-step])  
aligned (list[ : alignment])  
nontemporal (list)
```

```
1      private (list)
2      lastprivate ([ lastprivate-modifier:] list)
3      reduction ([ reduction-modifier, ]reduction-identifier : list)
4      collapse (n)
5      order (concurrent)
```

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

8 The **simd** directive places restrictions on the structure of all associated *do-loops*. Specifically, all  
9 associated *do-loops* must have *canonical loop form* (see Section 2.12.1 on page 101).

---

## Fortran

### 10 Binding

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

### 13 Description

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

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

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

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

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

32 A SIMD loop has logical iterations numbered 0,1,...,N-1 where N is the number of loop iterations,  
33 and the logical numbering denotes the sequence in which the iterations would be executed if the  
34 associated loop(s) were executed with no SIMD instructions. At the beginning of each logical  
35 iteration, the loop iteration variable of each associated loop has the value that it would have if the

1 set of the associated loop(s) were executed sequentially. The number of iterations that are executed  
2 concurrently at any given time is implementation defined. Each concurrent iteration will be  
3 executed by a different SIMD lane. Each set of concurrent iterations is a SIMD chunk. Lexical  
4 forward dependencies in the iterations of the original loop must be preserved within each SIMD  
5 chunk.

6 The **safelen** clause specifies that no two concurrent iterations within a SIMD chunk can have a  
7 distance in the logical iteration space that is greater than or equal to the value given in the clause.  
8 The parameter of the **safelen** clause must be a constant positive integer expression. The  
9 **simdlen** clause specifies the preferred number of iterations to be executed concurrently unless an  
10 **if** clause is present and evaluates to *false*, in which case the preferred number of iterations to be  
11 executed concurrently is one. The parameter of the **simdlen** clause must be a constant positive  
12 integer expression.

▼ C / C++ ▼

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

▲ C / C++ ▲

▼ Fortran ▼

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

▲ Fortran ▲

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

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

## 22 Restrictions

- 23 • No OpenMP directive may appear in the region between any associated loops.
- 24 • If a **collapse** clause is specified, exactly one loop must occur in the region at each nesting  
25 level up to the number of loops specified by the parameter of the **collapse** clause.
- 26 • The associated loops must be structured blocks.
- 27 • A program that branches into or out of a **simd** region is non-conforming.
- 28 • At most one **collapse** clause can appear on a **simd** directive.
- 29 • A *list-item* cannot appear in more than one **aligned** clause.
- 30 • A *list-item* cannot appear in more than one **nontemporal** clause.
- 31 • At most one **safelen** clause can appear on a **simd** directive.

- 1       • At most one **simdlen** clause can appear on a **simd** directive.
- 2       • At most one **if** clause can appear on the directive.
- 3       • If both **simdlen** and **safelen** clauses are specified, the value of the **simdlen** parameter
- 4       must be less than or equal to the value of the **safelen** parameter.
- 5       • A *modifier* may not be specified on a **linear** clause.
- 6       • The only OpenMP constructs that can be encountered during execution of a **simd** region are the
- 7       **atomic** construct, the **loop** construct, the **simd** construct and the **ordered** construct with
- 8       the **simd** clause.
- 9       • At most one **order (concurrent)** clause can appear on a **simd** directive.
- 10      • If an **order (concurrent)** clause is present, all restrictions from the **loop** construct with an
- 11      **order (concurrent)** clause also apply.

C / C++

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

C / C++

C

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

C

C++

- 14      • The type of list items appearing in the **aligned** clause must be array, pointer, reference to
- 15      array, or reference to pointer.

- 16      • No exception can be raised in the **simd** region.

C++

Fortran

- 17      • The *do-loop* iteration variable must be of type **integer**.
- 18      • The *do-loop* cannot be a **DO WHILE** or a **DO** loop without loop control.
- 19      • If a list item on the **aligned** clause has the **ALLOCATABLE** attribute, the allocation status must
- 20      be allocated.
- 21      • If a list item on the **aligned** clause has the **POINTER** attribute, the association status must be
- 22      associated.
- 23      • If the type of a list item on the **aligned** clause is either **C\_PTR** or Cray pointer, the list item
- 24      must be defined.

Fortran

## Cross References

- **order (concurrent)** clause, see Section 2.12.5 on page 133.
- **if** Clause, see Section 2.18 on page 224.
- Data-sharing attribute clauses, see Section 2.22.4 on page 284.

## 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++

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

```
#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

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

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

### Description

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

## Execution Model Events

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

## Tool Callbacks

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

## Restrictions

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

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

## Cross References

- Worksharing-loop construct, see Section 2.12.2 on page 106.
- **simd** construct, see Section 2.12.3.1 on page 115.
- Data-sharing attribute clauses, see Section 2.22.4 on page 284.

### 2.12.3.3 Declare SIMD Directive

#### Summary

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

#### Syntax

C / C++

The syntax of the declare SIMD directive is as follows:

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

C / C++

C++

The syntax of the declare SIMD directive is also as follows:

```
[[omp :: declare_simd [(clause[ [, ] clause] ... ) ] ] ]
[[omp :: declare_simd [(clause[ [, ] clause] ... ) ] ] ] ]
[ ... ]
```

C++

C / C++

where *clause* is one of the following:

```

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

```

C / C++

Fortran

The syntax of the **declare simd** directive is as follows:

```

!$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 on 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



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

▼ C++ ▲

6 The special *this* pointer can be used as if it was one of the arguments to the function in any of the  
7 **linear**, **aligned**, or **uniform** clauses.

▲ C++ ▼

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

▼ C / C++ ▲

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

▲ C / C++ ▼

▼ Fortran ▲

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

▲ Fortran ▼

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

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

## Restrictions

- Each argument can appear in at most one **uniform** or **linear** clause.
- At most one **simdlen** clause can appear in a declare SIMD directive.
- Either **inbranch** or **notinbranch** may be specified, but not both.
- When a *linear-step* expression is specified in a **linear** clause it must be either a constant integer expression or an integer-typed parameter that is specified in a **uniform** clause on the directive.
- The function or subroutine body must be a structured block.
- The execution of the function or subroutine, when called from a SIMD loop, cannot result in the execution of an OpenMP construct except for an **ordered** construct with the **simd** clause or an **atomic** construct.
- 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 directive for any declaration that has one must be equivalent to the one specified for the definition. Otherwise, the result is unspecified.
- The function cannot contain calls to the **longjmp** or **setjmp** functions.

▲ C / C++ ▲

▼ C ▼

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

▲ C ▲

▼ C++ ▼

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

▲ C++ ▲

## Fortran

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

## Fortran

### Cross References

- 17 • **linear** clause, see Section [2.22.4.6](#) on page [292](#).
- 18 • **reduction** clause, see Section [2.22.5.4](#) on page [302](#).

## 2.12.4 distribute Loop Constructs

### 2.12.4.1 distribute Construct

#### Summary

22 The **distribute** construct specifies that the iterations of one or more loops will be executed by

23 the initial teams in the context of their implicit tasks. The iterations are distributed across the initial

24 threads of all initial teams that execute the **teams** region to which the **distribute** region binds.

25

## 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)
```

The **distribute** 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 101).

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 101).

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 any of the associated loop(s) is a generated loop of a **tile** directive, the effect of the  
25       **dist\_schedule** clause on the loop schedule is implementation defined.

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

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

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

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

## 1 Execution Model Events

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

4 The *distribute-end* event occurs after an implicit task finishes execution of a **distribute** region  
5 but before it resumes execution of the enclosing context.

## 6 Tool Callbacks

7 A thread dispatches a registered **ompt\_callback\_work** callback with **ompt\_scope\_begin**  
8 as its *endpoint* argument and **ompt\_work\_distribute** as its *wstype* argument for each  
9 occurrence of a *distribute-begin* event in that thread. Similarly, a thread dispatches a registered  
10 **ompt\_callback\_work** callback with **ompt\_scope\_end** as its *endpoint* argument and  
11 **ompt\_work\_distribute** as its *wstype* argument for each occurrence of a *distribute-end* event  
12 in that thread. The callbacks occur in the context of the implicit task. The callbacks have type  
13 signature **ompt\_callback\_work\_t**.

## 14 Restrictions

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

- 16 • The **distribute** construct inherits the restrictions of the worksharing-loop construct.
- 17 • Each **distribute** region must be encountered by the initial threads of all initial teams in a  
18 league or by none at all.
- 19 • The sequence of the **distribute** regions encountered must be the same for every initial thread  
20 of every initial team in a league.
- 21 • The region corresponding to the **distribute** construct must be strictly nested inside a **teams**  
22 region.
- 23 • A list item may appear in a **firstprivate** or **lastprivate** clause but not both.
- 24 • At most one **dist\_schedule** clause can appear on the directive.
- 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 2.9 on page 87
- 29 • Worksharing-loop construct, see Section 2.12.2 on page 106.
- 30 • **tile** construct, see Section 2.12.7.1 on page 140.
- 31 • **ompt\_work\_distribute**, see Section 4.4.4.15 on page 449.
- 32 • **ompt\_callback\_work\_t**, see Section 4.5.2.5 on page 466.

## 2.12.4.2 `distribute simd` Construct

### Summary

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

### Syntax

C / C++

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

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

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

C / C++

Fortran

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

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

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

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

Fortran

### Description

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

### Execution Model Events

This composite construct generates the same events as the `distribute` construct.

### Tool Callbacks

This composite construct dispatches the same callbacks as the `distribute` construct.

## Restrictions

- The restrictions for the **distribute** and **simd** constructs apply.
- A list item may not appear in a **linear** clause unless it is the loop iteration variable of a loop that is associated with the construct.
- The **conditional** modifier may not appear in a **lastprivate** clause.

## Cross References

- **simd** construct, see Section 2.12.3.1 on page 115.
- **distribute** construct, see Section 2.12.4.1 on page 125.
- Data-sharing attribute clauses, see Section 2.22.4 on page 284.

## 2.12.4.3 Distribute Parallel Worksharing-Loop Construct

### Summary

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

### Syntax

C / C++

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

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

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

C / C++

Fortran

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

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

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

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

Fortran



## Description

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

## Execution Model Events

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

## Tool Callbacks

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

## Restrictions

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

## Cross References

- `distribute` construct, see Section [2.12.4.1](#) on page [125](#).
- `Parallel worksharing-loop` construct, see Section [2.16.1](#) on page [193](#).
- Data-sharing attribute clauses, see Section [2.22.4](#) on page [284](#).

## 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 worksharing-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 worksharing-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

## Description

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

## Execution Model Events

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

## Tool Callbacks

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

## Restrictions

- The restrictions for the **distribute** and parallel worksharing-loop SIMD constructs apply.
- No **ordered** clause can be specified.
- A list item may not appear in a **linear** clause unless it is the loop iteration variable of a loop that is associated with the construct.
- The **conditional** modifier may not appear in a **lastprivate** clause.
- 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 **simd** *directive-name-modifier* can appear on the directive.

## Cross References

- **distribute** construct, see Section 2.12.4.1 on page 125.
- Parallel worksharing-loop SIMD construct, see Section 2.16.5 on page 197.
- Data-sharing attribute clauses, see Section 2.22.4 on page 284.

## 2.12.5 loop Construct

### Summary

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

### Syntax

C / C++

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

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

where *clause* is one of the following:

```
bind (binding)  
collapse (n)  
order (concurrent)  
private (list)  
lastprivate (list)  
reduction ([default , ]reduction-identifier : list)
```

1 where *binding* is one of the following:

```
2 teams
3 parallel
4 thread
```

5 The **loop** directive places restrictions on the structure of all associated *for-loops*. Specifically, all  
6 associated *for-loops* must have *canonical loop form* (see Section 2.12.1 on page 101).

▲ C / C++ ▲

▼ Fortran ▼

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

```
8 !$omp loop [clause[ [, ] clause] ... ]
9 do-loops
10 [!$omp end loop]
```

11 where *clause* is one of the following:

```
12 bind(binding)
13 collapse(n)
14 order(concurrent)
15 private(list)
16 lastprivate(list)
17 reduction([default , ]reduction-identifier : list)
```

18 where *binding* is one of the following:

```
19 teams
20 parallel
21 thread
```

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

24 The **loop** directive places restrictions on the structure of all associated *do-loops*. Specifically, all  
25 associated *do-loops* must have *canonical loop form* (see Section 2.12.1 on page 101).

▲ Fortran ▲

## 1 Binding

2 If the **bind** clause is present on the construct, the binding region is determined by *binding*.  
3 Specifically, if *binding* is **teams** and there exists an innermost enclosing **teams** region then the  
4 binding region is that **teams** region; if *binding* is **parallel** then the binding region is the  
5 innermost enclosing parallel region, which may be an implicit parallel region; and if *binding* is  
6 **thread** then the binding region is not defined. If the **bind** clause is not present on the construct  
7 and the **loop** construct is closely nested inside a **teams** or **parallel** construct, the binding  
8 region is the corresponding **teams** or **parallel** region. If none of those conditions hold, the  
9 binding region is not defined.

10 If the binding region is a **teams** region, then the binding thread set is the set of master threads that  
11 are executing that region. If the binding region is a parallel region, then the binding thread set is the  
12 team of threads that are executing that region. If the binding region is not defined, then the binding  
13 thread set is the encountering thread.

## 14 Description

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

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

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

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

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

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

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

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

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

## 8 **Restrictions**

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

- 10 • If the **collapse** clause is specified then there may be no intervening OpenMP directives  
11 between the associated loops.
- 12 • At most one **collapse** clause can appear on a **loop** directive.
- 13 • A list item may not appear in a **lastprivate** clause unless it is the loop iteration variable of a  
14 loop that is associated with the construct.
- 15 • If a **loop** construct is not nested inside another OpenMP construct and it appears in a procedure,  
16 the **bind** clause must be present.
- 17 • If a **loop** region binds to a **teams** or parallel region, it must be encountered by all threads in  
18 the binding thread set or by none of them.
- 19 • At most one **bind** clause can appear on a **loop** directive.
- 20 • If the **bind** clause is present and *binding* is **teams**, the **loop** region corresponding to the  
21 **loop** construct must be strictly nested inside a **teams** region.
- 22 • If the **bind** clause is present and *binding* is **parallel**, the behavior is unspecified if the **loop**  
23 region corresponding to a **loop** construct is closely nested inside a **simd** region.
- 24 • The only constructs that may be nested inside a **loop** region are the **loop** construct, the  
25 **parallel** construct, the **simd** construct, and combined constructs for which the first construct  
26 is a **parallel** construct.
- 27 • A **loop** region corresponding to a **loop** construct may not contain calls to procedures that  
28 contain OpenMP directives.
- 29 • A **loop** region corresponding to a **loop** construct may not contain calls to the OpenMP  
30 Runtime API.
- 31 • If a threadprivate variable is referenced inside a **loop** region, the behavior is unspecified.
- 32 • At most one **order (concurrent)** clause can appear on a **loop** directive.

## C / C++

- The associated *for-loops* must be structured blocks.
- No statement can branch to any associated **for** statement.

## C / C++

## Fortran

- The associated *do-loops* must be structured blocks.
- No statement in the associated loops other than the DO statements can cause a branch out of the loops.

## Fortran

### Cross References

- The **single** construct, see Section 2.11.2 on page 96.
- The Worksharing-Loop construct, see Section 2.12.2 on page 106.
- SIMD directives, see Section 2.12.3 on page 115.
- **distribute** construct, see Section 2.12.4.1 on page 125.

## 2.12.6 scan Directive

### Summary

The **scan** directive specifies that scan computations update the list items on each iteration.

### Syntax

## C / C++

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

```
loop-associated-directive
for-loop-headers
{
    structured-block-sequence
    #pragma omp scan clause new-line
    structured-block-sequence
}
```

where *clause* is one of the following:

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

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

## C / C++

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

```

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

where *clause* is one of the following:

```

inclusive (list)
exclusive (list)
    
```

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

## 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 a scan computation updates each list item on each loop iteration. 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 sequence that lexically precede the directive constitute the *input phase* and all statements in the structured block sequence that lexically follow the directive constitute the *scan phase*. If the list item appears in an **exclusive** clause, all statements in the structured block sequence that lexically precede the directive constitute the *scan phase* and all statements in the structured block sequence that lexically follow the directive constitute the *input phase*. 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 uses the result of the scan computation for that iteration.

The list items that appear in an **inclusive** or **exclusive** clause may include array sections.

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:



- 1           • if  $N = 1, a_1$
- 2           • 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
- 3            $1 \leq K + 1 = L \leq N$ .

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

8           Let *orig-val* be the value of the original list item on entry to the enclosing worksharing-loop,  
9           worksharing-loop SIMD, or **simd** construct. Let *combiner* be the combiner for the  
10          *reduction-identifier* specified by the **reduction** clause on the construct. And let  $u_I$  be the update  
11          value of a list item for iteration  $I$ . For list items appearing in an **inclusive** clause on the **scan**  
12          directive, at the beginning of the *scan phase* for iteration  $I$  the list item is assigned the result of the  
13          operation  $\text{PRESUM}_{\text{last}}(\text{combiner}, \text{orig-val}, u_0, \dots, u_I)$ . For list items appearing in an  
14          **exclusive** clause on the **scan** directive, at the beginning of the *scan phase* for iteration  $I = 0$   
15          the list item is assigned the value *orig-val*, and at the beginning of the *scan phase* for iteration  $I > 0$   
16          the list item is assigned the result of the operation  $\text{PRESUM}_{\text{last}}(\text{combiner}, \text{orig-val}, u_0, \dots, u_{I-1})$ .

17          For list items appearing in an **inclusive** clause, at the end of the enclosing worksharing-loop,  
18          worksharing-loop SIMD, or **simd** construct, the original list item is assigned the private copy from  
19          the last logical iteration of the loops associated with the enclosing construct. For list items  
20          appearing in an **exclusive** clause, let  $L$  be the last logical iteration of the loops associated with  
21          the enclosing construct. At the end of the enclosing construct, the original list item is assigned the  
22          result of the operation  $\text{PRESUM}_{\text{last}}(\text{combiner}, \text{orig-val}, u_0, \dots, u_L)$ .

## 23           **Restrictions**

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

- 25          • Exactly one **scan** directive must appear in the loop body of an enclosing worksharing-loop,  
26          worksharing-loop SIMD, or **simd** construct on which a **reduction** clause with the **inscan**  
27          modifier is present.
- 28          • A list item that appears in the **inclusive** or **exclusive** clause must appear in a  
29          **reduction** clause with the **inscan** modifier on the enclosing worksharing-loop,  
30          worksharing-loop SIMD, or **simd** construct.
- 31          • Cross-iteration dependences across different logical iterations must not exist, except for  
32          dependences for the list items specified in an **inclusive** or **exclusive** clause.
- 33          • Intra-iteration dependences from a statement in the structured block sequence preceding a **scan**  
34          directive to a statement in the structured block sequence following a **scan** directive must not  
35          exist, except for dependences for the list items specified in an **inclusive** or **exclusive**  
36          clause.
- 37          • The private copy of list items that appear in the **inclusive** or **exclusive** clause may not be  
38          modified in the *scan phase*.

## Cross References

- Worksharing-loop construct, see Section 2.12.2 on page 106.
- **simd** construct, see Section 2.12.3.1 on page 115.
- Worksharing-loop SIMD construct, see Section 2.12.3.2 on page 120.
- **reduction** clause, see Section 2.22.5.4 on page 302.

## 2.12.7 Loop Transformation Constructs

A loop transformation construct replaces an associated loop nest with another loop nest that is called the *generated loop(s)*. A loop transformation construct that is closely nested within another loop transformation construct applies before the enclosing loop transformation construct.

The associated loop(s) of a transformation construct must have *canonical loop form* (see Section 2.12.1 on page 101). Each generated loop individually has canonical form, but not when part of multiple associated loops, unless otherwise specified. Loop iteration variables of generated loops are always private in the enclosing **teams**, **parallel** or task generating construct.

## Cross References

- Canonical loop form, see Section 2.12.1 on page 101.

### 2.12.7.1 **tile** Construct

#### Summary

The **tile** construct tiles one or more loops.

#### Syntax

C / C++

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

```
#pragma omp tile sizes(size-list) new-line  
    for-loops
```

where *size-list* is a list  $s_1, \dots, s_n$  of positive integer expressions.

C / C++

Fortran

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

```
!$omp tile sizes(size-list)  
    do-loops  
[$omp end tile]
```

1 where *size-list* is a list  $s_1, \dots, s_n$  of positive integer expressions.

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

## Fortran

### Description

4 The **tile** construct is associated with a loop nest that consists of  $n$  loops that follow the directive,  
5 where  $n$  is the number of items in *size-list*. Let  $\ell_1, \dots, \ell_n$  be the associated loops, from outermost to  
6 innermost, which the construct replaces with a loop nest that consists of  $2n$  loops. Let  
7  $f_1, \dots, f_n, t_1, \dots, t_n$  be the generated loops, from outermost to innermost. The  $f_1, \dots, f_n$  are the  
8 *floor loops* and the  $t_1, \dots, t_n$  are the *tile loops*.  
9

10 The  $n$ -dimensional logical iteration space of the associated loop(s) is the set  
11  $\Omega = \{(\omega_1, \dots, \omega_n) \in \mathbb{N}^n \mid \forall i = 1 \dots n : 0 \leq \omega_i < N_i\}$  where  $N_i$  is the number of iterations of  
12 loop  $\ell_i$ . For any  $(\alpha_1, \dots, \alpha_n) \in \mathbb{N}^n$ , define a tile  $T_{\alpha_1, \dots, \alpha_n}$  as the set of iterations  
13  $\{(\omega_1, \dots, \omega_n) \in \Omega \mid \forall i = 1 \dots n : s_i \alpha_i \leq \omega_i < s_i \alpha_i + s_i\}$  and  
14  $F = \{T_{\alpha_1, \dots, \alpha_n} \mid T_{\alpha_1, \dots, \alpha_n} \neq \emptyset\}$  as the set of tiles with at least one iteration. Tiles that contain  
15  $\prod_{i=1 \dots n} s_i$  iterations are complete tiles.

16 The floor loops iterate over all tiles  $\{T_{\alpha_1, \dots, \alpha_n} \in F\}$  in lexicographic order with respect to their  
17 indices  $(\alpha_1, \dots, \alpha_n)$  and the tile loops iterate over the iterations in  $T_{\alpha_1, \dots, \alpha_n}$  in lexicographic order  
18 with respect to its elements  $(\omega_1, \dots, \omega_n)$ . A program that relies on the order of iterations in a  
19 non-complete tile relative to other iterations has unspecified behavior.

20 Multiple floor loops can be associated with a loop-associated directive to form canonical loop(s),  
21 unless the resulting associated loop nest is non-rectangular.

### Restrictions

- 22 • Exactly one loop must occur in the construct at each of the  $n$  nesting levels.
- 23
- 24 • All loops associated with the construct must be perfectly nested; that is there must be no  
25 intervening code between any two loops.
- 26 • The associated loop(s) must not form a non-rectangular loop nest.

### Cross References

- 27 • Worksharing-loop construct, see Section 2.12.2 on page 106.
- 28
- 29 • **distribute** construct, see Section 2.12.4.1 on page 125.
- 30 • **taskloop** construct, see Section 2.13.2 on page 147.

## 2.13 Tasking Constructs

### 2.13.1 task Construct

#### Summary

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

#### Syntax

C / C++

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

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

where *clause* is one of the following:

```
if ([ task : ] scalar-expression)  
final (scalar-expression)  
untied  
default (data-sharing-attribute)  
mergeable  
private (list)  
firstprivate (list)  
shared (list)  
in_reduction (reduction-identifier : list)  
depend ([depend-modifier , ] dependence-type : locator-list)  
priority (priority-value)  
allocate ([allocator : ] list)  
affinity ([aff-modifier : ] locator-list)  
detach (event-handle)
```

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

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

where *event-handle* is a variable of the **omp\_event\_handle\_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 (data-sharing-attribute)
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-handle)
```

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

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

22 where *event-handle* is an integer variable of **omp\_event\_handle\_kind** *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 *allow-completion-event* is created. The *allow-completion-event* is connected to the completion of the associated **task** region. The original *event-handle* will be updated to represent the *allow-completion-event* event before the task data environment is created. The *event-handle* will be considered as if it was specified on a **firstprivate** clause. 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 execution of the *structured block* that is associated with 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. 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 if a **final** clause is present on the same **task** construct and the **final** clause expression evaluates to *true*, or if a task is an included task.

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

4 When the **mergeable** clause is present on a **task** construct, the generated task is a *mergeable*  
5 *task*.

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

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

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

20 The list items that appear in the **affinity** clause may use shape-operators.  
C / C++

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

22  
23 Note – When storage is shared by an explicit **task** region, the programmer must ensure, by  
24 adding proper synchronization, that the storage does not reach the end of its lifetime before the  
25 explicit **task** region completes its execution.  
26

## 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.

## 30 Tool Callbacks

31 A thread dispatches a registered **ompt\_callback\_task\_create** callback for each occurrence  
32 of a *task-create* event in the context of the encountering task. This callback has the type signature  
33 **ompt\_callback\_task\_create\_t** and the *flags* argument indicates the task types shown in  
34 Table 2.7.

**TABLE 2.7: `ompt_callback_task_create` callback flags evaluation**

Operation	Evaluates to true
<code>(flags &amp; ompt_task_explicit)</code>	Always in the dispatched callback
<code>(flags &amp; ompt_task_undeferred)</code>	If the task is an undeferred task
<code>(flags &amp; ompt_task_final)</code>	If the task is a final task
<code>(flags &amp; ompt_task_untied)</code>	If the task is an untied task
<code>(flags &amp; ompt_task_mergeable)</code>	If the task is a mergeable task
<code>(flags &amp; ompt_task_merged)</code>	If the task is a merged task

## Restrictions

Restrictions to the `task` construct are as follows:

- A program that branches into or out of a `task` region is non-conforming.
- A program must not depend on any ordering of the evaluations of the clauses of the `task` directive, or on any side effects of the evaluations of the clauses.
- 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.
- At most one `detach` clause can appear on the directive.
- If a `detach` clause appears on the directive, then a `mergeable` clause cannot appear on the same directive.

C / C++

- 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++

## Cross References

- Task scheduling constraints, see Section 2.13.6 on page 156.
- `allocate` clause, see Section 2.14.4 on page 164.
- `if` clause, see Section 2.18 on page 224.
- `depend` clause, see Section 2.20.11 on page 257.
- Data-sharing attribute clauses, Section 2.22.4 on page 284.
- `in_reduction` clause, see Section 2.22.5.6 on page 306.



- `omp_fulfill_event`, see Section 3.5.1 on page 395.
- `ompt_callback_task_create_t`, see Section 4.5.2.7 on page 469.

## 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-expression)  
shared (list)  
private (list)  
firstprivate (list)  
lastprivate (list)  
reduction ([default ,] reduction-identifier : list)  
in_reduction (reduction-identifier : list)  
default (data-sharing-attribute)  
grainsize (grain-size)  
num_tasks (num-tasks)  
collapse (n)  
final (scalar-expr)  
priority (priority-value)  
untied  
mergeable  
nogroup  
allocate ([allocator :] list)
```

1 The **taskloop** directive places restrictions on the structure of all associated *for-loops*.  
2 Specifically, all associated *for-loops* must have *canonical loop form* (see Section 2.12.1 on  
3 page 101).

▲ C / C++ ▲

▼ Fortran ▼

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

```
5 !$omp taskloop [clause[[,] clause] ...]  
6 do-loops  
7 [$omp end taskloop]
```

8 where *clause* is one of the following:

```
9 if([ taskloop :] scalar-logical-expression)  
10 shared(list)  
11 private(list)  
12 firstprivate(list)  
13 lastprivate(list)  
14 reduction ([default ,] reduction-identifier : list)  
15 in_reduction(reduction-identifier : list)  
16 default (data-sharing-attribute)  
17 grainsize (grain-size)  
18 num_tasks (num-tasks)  
19 collapse (n)  
20 final (scalar-logical-expr)  
21 priority (priority-value)  
22 untied  
23 mergeable  
24 nogroup  
25 allocate ([allocator :] list)
```

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

28 The **taskloop** directive places restrictions on the structure of all associated *do-loops*.  
29 Specifically, all associated *do-loops* must have *canonical loop form* (see Section 2.12.1 on  
30 page 101).

▲ Fortran ▲

## 1           **Binding**

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

## 4           **Description**

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

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

15          If a **reduction** clause is present on the **taskloop** construct, the behavior is as if a  
16          **task\_reduction** clause with the same reduction operator and list items was applied to the  
17          implicit **taskgroup** construct enclosing the **taskloop** construct. The **taskloop** construct  
18          executes as if each generated task was defined by a **task** construct on which an **in\_reduction**  
19          clause with the same reduction operator and list items is present. Thus, the generated tasks are  
20          participants of the reduction defined by the **task\_reduction** clause that was applied to the  
21          implicit **taskgroup** construct.

22          If an **in\_reduction** clause is present on the **taskloop** construct, the behavior is as if each  
23          generated task was defined by a **task** construct on which an **in\_reduction** clause with the  
24          same reduction operator and list items is present. Thus, the generated tasks are participants of a  
25          reduction previously defined by a reduction scoping clause.

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

30          The parameter of the **grainsize** clause must be a positive integer expression. If **num\_tasks** is  
31          specified, the **taskloop** construct creates as many tasks as the minimum of the *num-tasks*  
32          expression and the number of logical loop iterations. Each task must have at least one logical loop  
33          iteration. The parameter of the **num\_tasks** clause must be a positive integer expression. If neither  
34          a **grainsize** nor **num\_tasks** clause is present, the number of loop tasks generated and the  
35          number of logical loop iterations assigned to these tasks is implementation defined.

36          If any of the associated loop(s) is a generated loop of a **tile** directive, the effect of the  
37          **grainsize** or **num\_tasks** clauses on the partitioning of iterations into tasks is implementation  
38          defined.

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

10 If more than one loop is associated with the **taskloop** construct then the number of times that  
11 any intervening code between any two associated loops will be executed is unspecified but will be  
12 at least once per iteration of the loop enclosing the intervening code and at most once per iteration  
13 of the innermost loop associated with the construct. If the iteration count of any loop that is  
14 associated with the **taskloop** construct is zero and that loop does not enclose intervening code,  
15 the behavior is unspecified.

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

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

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

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

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

33 When a **priority** clause is present on a **taskloop** construct, the generated tasks use the  
34 *priority-value* as if it was specified for each individual task. If the **priority** clause is not  
35 specified, tasks generated by the **taskloop** construct have the default task priority (zero).

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

37 When the **mergeable** clause is present on a **taskloop** construct, each generated task is a  
38 *mergeable task*.

C++

1 For **firstprivate** variables of class type, the number of invocations of copy constructors to  
2 perform the initialization is implementation-defined.

C++

3  
4 **Note** – When storage is shared by a **taskloop** region, the programmer must ensure, by adding  
5 proper synchronization, that the storage does not reach the end of its lifetime before the **taskloop**  
6 region and its descendant tasks complete their execution.  
7

## 8 Execution Model Events

9 The *taskloop-begin* event occurs after a task encounters a **taskloop** construct but before any  
10 other events that may trigger as a consequence of executing the **taskloop**. Specifically, a  
11 *taskloop-begin* event for a **taskloop** will precede the *taskgroup-begin* that occurs unless a  
12 **nogroup** clause is present. Regardless of whether an implicit taskgroup is present, a  
13 *taskloop-begin* will always precede any *task-create* events for generated tasks.

14 The *taskloop-end* event occurs after a **taskloop** region finishes execution but before resuming  
15 execution of the encountering task.

16 The *taskloop-iteration-begin* event occurs before an explicit task executes each iteration of a  
17 **taskloop**.

## 18 Tool Callbacks

19 A thread dispatches a registered **ompt\_callback\_work** callback for each occurrence of a  
20 *taskloop-begin* and *taskloop-end* event in that thread. The callback occurs in the context of the  
21 encountering task. The callback has type signature **ompt\_callback\_work\_t**. The callback  
22 receives **ompt\_scope\_begin** or **ompt\_scope\_end** as its *endpoint* argument, as appropriate,  
23 and **ompt\_work\_taskloop** as its *wstype* argument.

24 A thread dispatches a registered **ompt\_callback\_dispatch** callback for each occurrence of a  
25 *taskloop-iteration-begin* event in that thread. The callback occurs in the context of the encountering  
26 task. The callback has type signature **ompt\_callback\_dispatch\_t**.

## 27 Restrictions

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

- 29
- A program that branches into or out of a **taskloop** region is non-conforming.
  - No OpenMP directive may appear in the region between any associated loops.
  - If a **collapse** clause is specified, exactly one loop must occur in the region at each nesting  
31 level up to the number of loops specified by the parameter of the **collapse** clause.  
32

- 1 • If a **reduction** clause is present on the **taskloop** directive, the **nogroup** clause must not  
2 be specified.
- 3 • The same list item cannot appear in both a **reduction** and an **in\_reduction** clause.
- 4 • At most one **grainsize** clause can appear on a **taskloop** directive.
- 5 • At most one **num\_tasks** clause can appear on a **taskloop** directive.
- 6 • The **grainsize** clause and **num\_tasks** clause are mutually exclusive and may not appear on  
7 the same **taskloop** directive.
- 8 • At most one **collapse** clause can appear on a **taskloop** directive.
- 9 • At most one **if** clause can appear on the directive.
- 10 • At most one **final** clause can appear on the directive.
- 11 • At most one **priority** clause can appear on the directive.

## 12 Cross References

- 13 • **tile** construct, see Section [2.12.7.1](#) on page [140](#).
- 14 • **task** construct, Section [2.13.1](#) on page [142](#).
- 15 • **if** clause, see Section [2.18](#) on page [224](#).
- 16 • **taskgroup** construct, Section [2.20.6](#) on page [235](#).
- 17 • Data-sharing attribute clauses, Section [2.22.4](#) on page [284](#).
- 18 • **ompt\_scope\_begin** and **ompt\_scope\_end**, see Section [4.4.4.11](#) on page [447](#).
- 19 • **ompt\_work\_taskloop**, see Section [4.4.4.15](#) on page [449](#).
- 20 • **ompt\_callback\_work\_t**, see Section [4.5.2.5](#) on page [466](#).
- 21 • **ompt\_callback\_dispatch\_t**, see Section [4.5.2.6](#) on page [467](#).

## 22 2.13.3 taskloop simd Construct

### 23 Summary

24 The **taskloop simd** construct specifies a loop that can be executed concurrently using SIMD  
25 instructions and that those iterations will also be executed in parallel using explicit tasks. The  
26 **taskloop simd** construct is a composite construct.

## Syntax

C / C++

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

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

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

C / C++

Fortran

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

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

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

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

Fortran

## Binding

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

## Description

The **taskloop simd** construct will first distribute the iterations of the associated loop(s) across tasks in a manner consistent with any clauses that apply to the **taskloop** construct. The resulting tasks will then be converted to a SIMD loop in a manner consistent with any clauses that apply to the **simd** construct, except for the **collapse** clause. For the purposes of each task's conversion to a SIMD loop, the **collapse** clause is ignored and the effect of any **in\_reduction** clause is as if a **reduction** clause with the same reduction operator and list items is present on the **simd** construct.

## Execution Model Events

This composite construct generates the same events as the **taskloop** construct.

## Tool Callbacks

This composite construct dispatches the same callbacks as the **taskloop** construct.

## Restrictions

- The restrictions for the **taskloop** and **simd** constructs apply.
- The **conditional** modifier may not appear in a **lastprivate** clause.
- 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 **taskloop** *directive-name-modifier* can appear on the directive.
- At most one **if** clause with the **simd** *directive-name-modifier* can appear on the directive.

## Cross References

- **simd** construct, see Section 2.12.3.1 on page 115.
- **taskloop** construct, see Section 2.13.2 on page 147.
- Data-sharing attribute clauses, see Section 2.22.4 on page 284.

## 2.13.4 taskyield Construct

### Summary

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

### Syntax

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 [156](#).

## 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-begin* event occurs after an *initial-thread-begin* event but before the first OpenMP region in the initial task begins to execute.

The *initial-task-end* event occurs before an *initial-thread-end* event but after the last OpenMP region in the initial task finishes 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.

### Tool Callbacks

A thread dispatches a registered `ompt_callback_thread_begin` callback for the *initial-thread-begin* event in an initial thread. The callback occurs in the context of the initial thread. The callback has type signature `ompt_callback_thread_begin_t`. The callback receives `ompt_thread_initial` as its *thread\_type* argument.

A thread dispatches a registered `ompt_callback_implicit_task` callback with `ompt_scope_begin` as its *endpoint* argument for each occurrence of an *initial-task-begin* in that thread. Similarly, a thread dispatches a registered `ompt_callback_implicit_task` callback with `ompt_scope_end` as its *endpoint* argument for each occurrence of an *initial-task-end* event in that thread. The callbacks occur in the context of the initial task and have type signature `ompt_callback_implicit_task_t`. In the dispatched callback, `(flag & ompt_task_initial)` always evaluates to *true*.

A thread dispatches a registered `ompt_callback_thread_end` callback for the *initial-thread-end* event in that thread. The callback occurs in the context of the thread. The callback has type signature `ompt_callback_thread_end_t`. The implicit parallel region does not dispatch a `ompt_callback_parallel_end` callback; however, the implicit parallel region can be finalized within this `ompt_callback_thread_end` callback.

## Cross References

- `ompt_thread_initial`, see Section [4.4.4.10](#) on page [447](#).
- `ompt_task_initial`, see Section [4.4.4.18](#) on page [450](#).
- `ompt_callback_thread_begin_t`, see Section [4.5.2.1](#) on page [462](#).

- 1 • `ompt_callback_thread_end_t`, see Section 4.5.2.2 on page 462.
- 2 • `ompt_callback_parallel_begin_t`, see Section 4.5.2.3 on page 463.
- 3 • `ompt_callback_parallel_end_t`, see Section 4.5.2.4 on page 465.
- 4 • `ompt_callback_implicit_task_t`, see Section 4.5.2.11 on page 473.

## 5 2.13.6 Task Scheduling

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

- 9 • during the generation of an explicit task;
- 10 • the point immediately following the generation of an explicit task;
- 11 • after the point of completion of the structured block associated with a task;
- 12 • in a `taskyield` region;
- 13 • in a `taskwait` region;
- 14 • at the end of a `taskgroup` region;
- 15 • in an implicit barrier region;
- 16 • in an explicit `barrier` region;
- 17 • during the generation of a `target` region;
- 18 • the point immediately following the generation of a `target` region;
- 19 • at the beginning and end of a `target data` region;
- 20 • in a `target update` region;
- 21 • in a `target enter data` region;
- 22 • in a `target exit data` region;
- 23 • in the `omp_target_memcpy` routine;
- 24 • in the `omp_target_memcpy_async` routine;
- 25 • in the `omp_target_memcpy_rect` routine;
- 26 • in the `omp_target_memcpy_rect_async` routine.

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

- 29 • begin execution of a tied task bound to the current team;

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

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

5           *Task Scheduling Constraints* are as follows:

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

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

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

22          A program must behave correctly and consistently with all conceivable scheduling sequences that  
23          are compatible with the rules above.

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

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

32          The use of **threadprivate** variables and the use of locks or critical sections in an explicit task with an  
33          **if** clause must take into account that when the **if** clause evaluates to *false*, the task is executed  
34          immediately, without regard to *Task Scheduling Constraint 2*.  
35          ▲

## Execution Model Events

The *task-schedule* event occurs in a thread when the thread switches tasks at a task scheduling point; no event occurs when switching to or from a merged task.

## Tool Callbacks

A thread dispatches a registered `ompt_callback_task_schedule` callback for each occurrence of a *task-schedule* event in the context of the task that begins or resumes. This callback has the type signature `ompt_callback_task_schedule_t`. The argument *prior\_task\_status* is used to indicate the cause for suspending the prior task. This cause may be the completion of the prior task region, the encountering of a `taskyield` construct, or the encountering of an active cancellation point.

## Cross References

- `ompt_callback_task_schedule_t`, see Section 4.5.2.10 on page 472.

# 2.14 Memory Management Directives

## 2.14.1 Memory Spaces

OpenMP memory spaces represent storage resources where variables can be stored and retrieved. Table 2.8 shows the list of predefined memory spaces. The selection of a given memory space expresses an intent to use storage with certain traits for the allocations. The actual storage resources that each memory space represents are implementation defined.

TABLE 2.8: 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.

Note – For variables allocated in the `omp_const_mem_space` memory space OpenMP supports initializing constant memory either by means of the `firstprivate` clause or through initialization with compile time constants for static and constant variables. Implementation-defined mechanisms to provide the constant value of these variables may also be supported.

## Cross References

- `omp_init_allocator` routine, see Section 3.8.2 on page 415.

## 2.14.2 Memory Allocators

OpenMP memory allocators can be used by a program to make allocation requests. When a memory allocator receives a request to allocate storage of a certain size, an allocation of logically consecutive *memory* in the resources of its associated memory space of at least the size that was requested will be returned if possible. This allocation will not overlap with any other existing allocation from an OpenMP memory allocator.

The behavior of the allocation process can be affected by the allocator traits that the user specifies. Table 2.9 shows the allowed allocators traits, their possible values and the default value of each trait.

**TABLE 2.9:** 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 that 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>
<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>

The `sync_hint` trait describes the expected manner in which multiple threads may use the allocator. The values and their description are:

- **contended:** high contention is expected on the allocator; that is, many threads are expected to request allocations simultaneously.
- **uncontended:** low contention is expected on the allocator; that is, few threads are expected to request allocations simultaneously.

- 1 • **serialized**: only one thread at a time will request allocations with the allocator. Requesting  
2 two allocations simultaneously when specifying **serialized** results in unspecified behavior.
- 3 • **private**: the same thread will request allocations with the allocator every time. Requesting an  
4 allocation from different threads, simultaneously or not, when specifying **private** results in  
5 unspecified behavior.

6 Allocated memory will be byte aligned to at least the value specified for the **alignment** trait of  
7 the allocator. Some directives and API routines can specify additional requirements on alignment  
8 beyond those described in this section.

9 Memory allocated by allocators with the **access** trait defined to be **all** must be accessible by all  
10 threads in the device where the allocation was requested. Memory allocated by allocators with the  
11 **access** trait defined to be **cgroup** will be memory accessible by all threads in the same  
12 contention group as the thread that requested the allocation. Attempts to access the memory  
13 returned by an allocator with the **access** trait defined to be **cgroup** from a thread that is not part  
14 of the same contention group as the thread that allocated the memory result in unspecified behavior.  
15 Memory allocated by allocators with the **access** trait defined to be **pteam** will be memory  
16 accessible by all threads that bind to the same **parallel** region of the thread that requested the  
17 allocation. Attempts to access the memory returned by an allocator with the **access** trait defined  
18 to be **pteam** from a thread that does not bind to the same **parallel** region as the thread that  
19 allocated the memory result in unspecified behavior. Memory allocated by allocator with the  
20 **access** trait defined to be **thread** will be memory accessible by the *thread* that requested the  
21 allocation. Attempts to access the memory returned by an allocator with the **access** trait defined  
22 to be **thread** from a thread other than the one that allocated the memory result in unspecified  
23 behavior.

24 The total amount of storage in bytes that an allocator can use is limited by the **pool\_size** trait.  
25 For allocators with the **access** trait defined to be **all**, this limit refers to allocations from all  
26 threads that access the allocator. For allocators with the **access** trait defined to be **cgroup**, this  
27 limit refers to allocations from threads that access the allocator from the same contention group. For  
28 allocators with the **access** trait defined to be **pteam**, this limit refers to allocations from threads  
29 that access the allocator from the same parallel team. For allocators with the **access** trait defined  
30 to be **thread**, this limit refers to allocations from each thread that access the allocator. Requests  
31 that would result in using more storage than **pool\_size** will not be fulfilled by the allocator.

32 The **fallback** trait specifies how the allocator behaves when it cannot fulfill an allocation  
33 request. If the **fallback** trait is set to **null\_fb**, the allocator returns the value zero if it fails to  
34 allocate the memory. If the **fallback** trait is set to **abort\_fb**, program execution will be  
35 terminated if the allocation fails. If the **fallback** trait is set to **allocator\_fb** then when an  
36 allocation fails the request will be delegated to the allocator specified in the **fb\_data** trait. If the  
37 **fallback** trait is set to **default\_mem\_fb** then when an allocation fails another allocation will  
38 be tried in the **omp\_default\_mem\_space** memory space, which assumes all allocator traits to  
39 be set to their default values except for **fallback** trait which will be set to **null\_fb**.

40 Allocators with the **pinned** trait defined to be **true** ensure that their allocations remain in the

- 1 same storage resource at the same location for their entire lifetime.
- 2 The **partition** trait describes the partitioning of allocated memory over the storage resources  
 3 represented by the memory space associated with the allocator. The partitioning will be done in  
 4 parts with a minimum size that is implementation defined. The values are:
- 5 • **environment**: the placement of allocated memory is determined by the execution  
 6 environment.
  - 7 • **nearest**: allocated memory is placed in the storage resource that is nearest to the thread that  
 8 requests the allocation.
  - 9 • **blocked**: allocated memory is partitioned into parts of approximately the same size with at  
 10 most one part per storage resource.
  - 11 • **interleaved**: allocated memory parts are distributed in a round-robin fashion across the  
 12 storage resources.

13 Table 2.10 shows the list of predefined memory allocators and their associated memory spaces. The  
 14 predefined memory allocators have default values for their allocator traits unless otherwise  
 15 specified.

**TABLE 2.10:** 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	<b>access:cgroup</b>
<code>omp_pteam_mem_alloc</code>	Implementation defined	<b>access:pteam</b>
<code>omp_thread_mem_alloc</code>	Implementation defined	<b>access:thread</b>

▼ **Fortran** ▲

17 If any operation of the base language causes a reallocation of an array that is allocated with a  
 18 memory allocator then that memory allocator will be used to release the current memory and to  
 19 allocate the new memory.

▲ **Fortran** ▼

## Cross References

- `omp_init_allocator` routine, see Section 3.8.2 on page 415.
- `omp_destroy_allocator` routine, see Section 3.8.3 on page 417.
- `omp_set_default_allocator` routine, see Section 3.8.4 on page 417.
- `omp_get_default_allocator` routine, see Section 3.8.5 on page 418.
- `OMP_ALLOCATOR` environment variable, see Section 6.21 on page 601.

## 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] new-line
```

where *clause* is one of the following:

```
allocator (allocator)
```

where *allocator* is an expression of `omp_allocator_handle_t` type.

C / C++

Fortran

The syntax of the `allocate` directive is as follows:

```
!$omp allocate (list) [clause]
```

or

```
!$omp allocate [(list)] clause  
!$omp allocate (list) clause  
[...]]  
allocate statement
```

where *clause* is one of the following:

```
allocator (allocator)
```

where *allocator* is an integer expression of `omp_allocator_handle_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. The allocation of each *list item* will be byte aligned to at least the alignment required by the base language for the type of that *list item*.

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 **allocate** directive must appear in the same scope as the declarations of each of its list items and must follow all such declarations.
- At most one **allocator** clause can appear on the **allocate** directive.
- **allocate** directives that appear 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, the *allocator* expression in the **allocator** clause must be a constant expression that evaluates to one of the predefined memory allocator values.
- After a *list item* has been allocated, the scope that contains the **allocate** directive must not end abnormally other than through C++ exceptions, such as through a call to the **longjmp** function.

### 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 parameters 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 2.6.1 on page 66.
- Memory allocators, see Section 2.14.2 on page 159.
- **omp\_allocator\_handle\_t** and **omp\_allocator\_handle\_kind**, see Section 3.8.1 on page 413.

## 2.14.4 allocate Clause

### Summary

The **allocate** clause specifies the memory allocator to be used to obtain storage for private variables of a directive.

### Syntax

The syntax of the **allocate** clause is as follows:

```
allocate (allocator : ] list)
```

## C / C++

where *allocator* is an expression of the **omp\_allocator\_handle\_t** type.

## C / C++

## Fortran

where *allocator* is an integer expression of the **omp\_allocator\_handle\_kind** *kind*.

## Fortran

## Description

The storage for new list items that arise from list items that appear in the directive will be provided through a memory allocator. If an *allocator* is specified in the clause, that allocator will be used for allocations. For all directives except the **target** directive, if no *allocator* is specified in the clause then the memory allocator that is specified by the *def-allocator-var* ICV will be used for the list items that are specified in the **allocate** clause. The allocation of each *list item* will be byte aligned to at least the alignment required by the base language for the type of that *list item*.

For allocations that arise from this clause the **null\_fb** value of the fallback allocator trait will behave as if the **abort\_fb** had been specified.

## Restrictions

- For any list item that is specified in the **allocate** clause on a directive, a data-sharing attribute clause that may create a private copy of that list item must be specified on the same directive.
- For **task**, **taskloop** or **target** directives, allocation requests to memory allocators with the trait **access** set to **thread** result in unspecified behavior.
- **allocate** clauses that appear on a **target** construct or on constructs in a **target** region must specify an *allocator* expression unless a **requires** directive with the **dynamic\_allocators** clause is present in the same compilation unit.

## Cross References

- *def-allocator-var* ICV, see Section 2.6.1 on page 66.
- Memory allocators, see Section 2.14.2 on page 159.
- **omp\_allocator\_handle\_t** and **omp\_allocator\_handle\_kind**, see Section 3.8.1 on page 413.

# 2.15 Device Directives

## 2.15.1 Device Initialization

### Execution Model Events

The *device-initialize* event occurs in a thread that encounters the first **target**, **target data**, or **target enter data** construct or a device memory routine that is associated with a particular target device after the thread initiates initialization of OpenMP on the device and the device's OpenMP initialization, which may include device-side tool initialization, completes.

The *device-load* event for a code block for a target device occurs in some thread before any thread executes code from that code block on that target device.

The *device-unload* event for a target device occurs in some thread whenever a code block is unloaded from the device.

1 The *device-finalize* event for a target device that has been initialized occurs in some thread before  
2 an OpenMP implementation shuts down.

### 3 **Tool Callbacks**

4 A thread dispatches a registered **ompt\_callback\_device\_initialize** callback for each  
5 occurrence of a *device-initialize* event in that thread. This callback has type signature  
6 **ompt\_callback\_device\_initialize\_t**.

7 A thread dispatches a registered **ompt\_callback\_device\_load** callback for each occurrence  
8 of a *device-load* event in that thread. This callback has type signature  
9 **ompt\_callback\_device\_load\_t**.

10 A thread dispatches a registered **ompt\_callback\_device\_unload** callback for each  
11 occurrence of a *device-unload* event in that thread. This callback has type signature  
12 **ompt\_callback\_device\_unload\_t**.

13 A thread dispatches a registered **ompt\_callback\_device\_finalize** callback for each  
14 occurrence of a *device-finalize* event in that thread. This callback has type signature  
15 **ompt\_callback\_device\_finalize\_t**.

### 16 **Restrictions**

17 No thread may offload execution of an OpenMP construct to a device until a dispatched  
18 **ompt\_callback\_device\_initialize** callback completes.

19 No thread may offload execution of an OpenMP construct to a device after a dispatched  
20 **ompt\_callback\_device\_finalize** callback occurs.

### 21 **Cross References**

- 22 • **ompt\_callback\_device\_initialize\_t**, see Section [4.5.2.19](#) on page [482](#).
- 23 • **ompt\_callback\_device\_finalize\_t**, see Section [4.5.2.20](#) on page [484](#).
- 24 • **ompt\_callback\_device\_load\_t**, see Section [4.5.2.21](#) on page [485](#).
- 25 • **ompt\_callback\_device\_unload\_t**, see Section [4.5.2.22](#) on page [486](#).

## 26 **2.15.2 target data Construct**

### 27 **Summary**

28 The **target data** construct maps variables to a device data environment for the extent of the  
29 region.

## Syntax

C / C++

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

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

where *clause* is one of the following:

```
if([ target data :] scalar-expression)  
    device(integer-expression)  
    map([[map-type-modifier[ , ] [map-type-modifier[ , ] ...] map-type: ] locator-list)  
    use_device_ptr(ptr-list)  
    use_device_addr(list)
```

C / C++

Fortran

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

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

where *clause* is one of the following:

```
if([ target data :] scalar-logical-expression)  
    device(scalar-integer-expression)  
    map([[map-type-modifier[ , ] [map-type-modifier[ , ] ...] map-type: ] locator-list)  
    use_device_ptr(ptr-list)  
    use_device_addr(list)
```

Fortran

## Binding

The binding task set for a **target data** region is the generating task. The **target data** region binds to the region of the generating task.

## Description

When a **target data** construct is encountered, the encountering task executes the region. If no **device** clause is present, the default device is determined by the *default-device-var* ICV. When an **if** clause is present and the **if** clause expression evaluates to *false*, the device is the host.

Variables are mapped for the extent of the region, according to any data-mapping attribute clauses, from the data environment of the encountering task to the device data environment. Pointers that appear in a **use\_device\_ptr** clause are privatized and the device pointers to the corresponding list items in the device data environment are assigned into the private versions.

List items that appear in a **use\_device\_addr** clause have the address of the corresponding object in the device data environment inside the construct. For objects, any reference to the value of the object will be to the corresponding object on the device, while references to the address will result in a valid device address that points to that object. Array sections privatize the base of the array section and assign the private copy to the address of the corresponding array section in the device data environment.

If one or more of the **use\_device\_ptr** or **use\_device\_addr** clauses and one or more **map** clauses are present on the same construct, the address conversions of **use\_device\_addr** and **use\_device\_ptr** clauses will occur as if performed after all variables are mapped according to those **map** clauses.

## Execution Model Events

The events associated with entering a target data region are the same events as associated with a target enter data construct, described in Section 2.15.3 on page 169.

The events associated with exiting a target data region are the same events as associated with a target exit data construct, described in Section 2.15.4 on page 172.

## Tool Callbacks

The tool callbacks dispatched when entering a target data region are the same as the tool callbacks dispatched when encountering a target enter data construct, described in Section 2.15.3 on page 169.

The tool callbacks dispatched when exiting a target data region are the same as the tool callbacks dispatched when encountering a target exit data construct, described in Section 2.15.4 on page 172.

## Restrictions

- A program must not depend on any ordering of the evaluations of the clauses of the **target data** directive, except as explicitly stated for **map** clauses relative to **use\_device\_ptr** and **use\_device\_addr** clauses, or on any side effects of the evaluations of the clauses.
- 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()` or to the value of `omp_get_initial_device()`.

- 1       • At most one **if** clause can appear on the directive.
- 2       • A *map-type* in a **map** clause must be **to**, **from**, **tofrom** or **alloc**.
- 3       • At least one **map**, **use\_device\_addr** or **use\_device\_ptr** clause must appear on the
- 4       directive.
- 5       • A list item in a **use\_device\_ptr** clause must hold the address of an object that has a
- 6       corresponding list item in the device data environment.
- 7       • A list item in a **use\_device\_addr** clause must have a corresponding list item in the device
- 8       data environment.
- 9       • A list item that specifies a given variable may not appear in more than one **use\_device\_ptr**
- 10      clause.
- 11      • A reference to a list item in a **use\_device\_addr** clause must be to the address of the list item.

## 12       **Cross References**

- 13      • *default-device-var*, see Section 2.6 on page 66.
- 14      • **if** Clause, see Section 2.18 on page 224.
- 15      • **map** clause, see Section 2.22.7.1 on page 318.
- 16      • **omp\_get\_num\_devices** routine, see Section 3.2.36 on page 370.
- 17      • **ompt\_callback\_target\_t**, see Section 4.5.2.26 on page 490.

## 18       **2.15.3 target enter data Construct**

### 19       **Summary**

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

### 22       **Syntax**

C / C++

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

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

25       where *clause* is one of the following:

```
26       if( [ target enter data : ] scalar-expression )
27       device ( integer-expression )
28       map ( [ map-type-modifier[ , ] [ map-type-modifier[ , ] ... ] map-type : locator-list )
29       depend ( [ depend-modifier , ] dependence-type : locator-list )
30       nowait
```

C / C++

Fortran

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

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

3 where *clause* is one of the following:

```
4     if( [ target enter data : ] scalar-logical-expression )  
5     device ( scalar-integer-expression )  
6     map ( [ map-type-modifier [ , ] [ map-type-modifier [ , ] ... ] map-type : locator-list )  
7     depend ( [ depend-modifier , ] dependence-type : locator-list )  
8     nowait
```

Fortran

## 9 Binding

10 The binding task set for a **target enter data** region is the generating task, which is the *target*  
11 *task* generated by the **target enter data** construct. The **target enter data** region binds  
12 to the corresponding *target task* region.

## 13 Description

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

16 The **target enter data** construct is a task generating construct. The generated task is a *target*  
17 *task*. The generated task region encloses the **target enter data** region.

18 All clauses are evaluated when the **target enter data** construct is encountered. The data  
19 environment of the *target task* is created according to the data-sharing attribute clauses on the  
20 **target enter data** construct, per-data environment ICVs, and any default data-sharing  
21 attribute rules that apply to the **target enter data** construct. A variable that is mapped in the  
22 **target enter data** construct has a default data-sharing attribute of shared in the data  
23 environment of the *target task*.

24 Assignment operations associated with mapping a variable (see Section 2.22.7.1 on page 318)  
25 occur when the *target task* executes.

26 If the **nowait** clause is present, execution of the *target task* may be deferred. If the **nowait**  
27 clause is not present, the *target task* is an included task.

28 If a **depend** clause is present, it is associated with the *target task*.

29 If no **device** clause is present, the default device is determined by the *default-device-var* ICV.

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



## 1 Execution Model Events

2 Events associated with a *target task* are the same as for the **task** construct defined in  
3 Section 2.13.1 on page 142.

4 The *target-enter-data-begin* event occurs when a thread enters a **target enter data** region.

5 The *target-enter-data-end* event occurs when a thread exits a **target enter data** region.

## 6 Tool Callbacks

7 Callbacks associated with events for *target tasks* are the same as for the **task** construct defined in  
8 Section 2.13.1 on page 142; (*flags & omp\_target\_target*) always evaluates to *true* in the  
9 dispatched callback.

10 A thread dispatches a registered **omp\_callback\_target** callback with  
11 **omp\_scope\_begin** as its *endpoint* argument and **omp\_target\_enter\_data** as its *kind*  
12 argument for each occurrence of a *target-enter-data-begin* event in that thread in the context of the  
13 target task on the host. Similarly, a thread dispatches a registered **omp\_callback\_target**  
14 callback with **omp\_scope\_end** as its *endpoint* argument and **omp\_target\_enter\_data**  
15 as its *kind* argument for each occurrence of a *target-enter-data-end* event in that thread in the  
16 context of the target task on the host. These callbacks have type signature  
17 **omp\_callback\_target\_t**.

## 18 Restrictions

- 19 • A program must not depend on any ordering of the evaluations of the clauses of the  
20 **target enter data** directive, or on any side effects of the evaluations of the clauses.
- 21 • At least one **map** clause must appear on the directive.
- 22 • At most one **device** clause can appear on the directive. The **device** clause expression must  
23 evaluate to a non-negative integer value less than the value of **omp\_get\_num\_devices()** or  
24 to the value of **omp\_get\_initial\_device()**.
- 25 • At most one **if** clause can appear on the directive.
- 26 • A *map-type* must be specified in all **map** clauses and must be either **to** or **alloc**.
- 27 • At most one **nowait** clause can appear on the directive.

## 28 Cross References

- 29 • *default-device-var*, see Section 2.6.1 on page 66.
- 30 • **task**, see Section 2.13.1 on page 142.
- 31 • **task scheduling constraints**, see Section 2.13.6 on page 156.
- 32 • **target data**, see Section 2.15.2 on page 166.
- 33 • **target exit data**, see Section 2.15.4 on page 172.
- 34 • **if** Clause, see Section 2.18 on page 224.

- **map** clause, see Section 2.22.7.1 on page 318.
- **omp\_get\_num\_devices** routine, see Section 3.2.36 on page 370.
- **ompt\_callback\_target\_t**, see Section 4.5.2.26 on page 490.

## 2.15.4 target exit data Construct

### Summary

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

### Syntax

C / C++

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

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

where *clause* is one of the following:

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

C / C++

Fortran

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

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

where *clause* is one of the following:

```
if ([ target exit data : ] scalar-logical-expression)
device (scalar-integer-expression)
map ([ map-type-modifier [ , ] [ map-type-modifier [ , ] ... ] map-type : locator-list )
depend ([ depend-modifier , ] dependence-type : locator-list)
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 318)  
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 no **device** clause is present, 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 142.

26       The *target-exit-data-begin* event occurs when a thread enters a **target exit data** region.

27       The *target-exit-data-end* event occurs when a thread exits a **target exit data** region.

## 28       **Tool Callbacks**

29       Callbacks associated with events for *target tasks* are the same as for the **task** construct defined in  
30       Section 2.13.1 on page 142; (*flags & ompt\_task\_target*) always evaluates to *true* in the  
31       dispatched callback.

32       A thread dispatches a registered **ompt\_callback\_target** callback with  
33       **ompt\_scope\_begin** as its *endpoint* argument and **ompt\_target\_exit\_data** as its *kind*  
34       argument for each occurrence of a *target-exit-data-begin* event in that thread in the context of the  
35       target task on the host. Similarly, a thread dispatches a registered **ompt\_callback\_target**

1 callback with `ompt_scope_end` as its *endpoint* argument and `ompt_target_exit_data` as  
2 its *kind* argument for each occurrence of a *target-exit-data-end* event in that thread in the context of  
3 the target task on the host. These callbacks have type signature `ompt_callback_target_t`.

## 4 **Restrictions**

- 5 • A program must not depend on any ordering of the evaluations of the clauses of the  
6 **target exit data** directive, or on any side effects of the evaluations of the clauses.
- 7 • At least one **map** clause must appear on the directive.
- 8 • At most one **device** clause can appear on the directive. The **device** clause expression must  
9 evaluate to a non-negative integer value less than the value of `omp_get_num_devices()` or  
10 to the value of `omp_get_initial_device()`.
- 11 • At most one **if** clause can appear on the directive.
- 12 • A *map-type* must be specified in all **map** clauses and must be either **from**, **release**, or  
13 **delete**.
- 14 • At most one **nowait** clause can appear on the directive.

## 15 **Cross References**

- 16 • *default-device-var*, see Section [2.6.1](#) on page [66](#).
- 17 • **task**, see Section [2.13.1](#) on page [142](#).
- 18 • **task scheduling constraints**, see Section [2.13.6](#) on page [156](#).
- 19 • **target data**, see Section [2.15.2](#) on page [166](#).
- 20 • **target enter data**, see Section [2.15.3](#) on page [169](#).
- 21 • **if** Clause, see Section [2.18](#) on page [224](#).
- 22 • **map** clause, see Section [2.22.7.1](#) on page [318](#).
- 23 • `omp_get_num_devices` routine, see Section [3.2.36](#) on page [370](#).
- 24 • `ompt_callback_target_t`, see Section [4.5.2.26](#) on page [490](#).

## 25 **2.15.5 target Construct**

### 26 **Summary**

27 The **target** construct maps variables to a device data environment and executes the construct on  
28 that device.

## Syntax

C / C++

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

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

where *clause* is one of the following:

```
if ([ target : ] scalar-expression)  
device ([ device-modifier : ] integer-expression)  
private (list)  
firstprivate (list)  
in_reduction (reduction-identifier : list)  
map ([ [map-type-modifier [ , ] [map-type-modifier [ , ] ... ] map-type : ] locator-list)  
is_device_ptr (list)  
defaultmap (implicit-behavior : [variable-category])  
nowait  
depend ([ depend-modifier , ] dependence-type : locator-list)  
allocate ([ [allocator : ] list)  
uses_allocators (allocator [ (allocator-traits-array) ]  
    [ , allocator [ (allocator-traits-array) ] ... ])
```

and where *device-modifier* is one of the following:

```
ancestor  
device_num
```

and where *allocator* is an identifier of **omp\_allocator\_handle\_t** type and *allocator-traits-array* is an identifier of **const omp\_alloctrail\_t \*** type.

C / C++

Fortran

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

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

where *clause* is one of the following:

```

1  if ([ target :] scalar-logical-expression)
2  device ([ device-modifier :] scalar-integer-expression)
3  private (list)
4  firstprivate (list)
5  in_reduction (reduction-identifier : list)
6  map ([[map-type-modifier[ , ] [map-type-modifier[ , ] ...] map-type: ] locator-list)
7  is_device_ptr (list)
8  defaultmap (implicit-behavior[:variable-category])
9  nowait
10 depend ([depend-modifier, ] dependence-type : locator-list)
11 allocate ([allocator:]list)
12 uses_allocators (allocator[ (allocator-traits-array) ]
13                 [, allocator[ (allocator-traits-array) ] ...])

```

and where *device-modifier* is one of the following:

```

15 ancestor
16 device_num

```

and where *allocator* is an integer expression of `omp_allocator_handle_kind` kind and *allocator-traits-array* is an array of `type(omp_alloctrail)` type.

Fortran

## Binding

The binding task set for a **target** region is the generating task, which is the *target task* generated 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` and `use_device_addr` clauses.

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

1 appear as a list item in an **in\_reduction** clause on the construct, the variable has a default  
2 data-sharing attribute of *shared* in the data environment of the *target task*.

3 Assignment operations associated with mapping a variable (see Section 2.22.7.1 on page 318)  
4 occur when the *target task* executes.

5 If a **device** clause in which the **device\_num** *device-modifier* appears is present on the  
6 construct, the **device** clause expression specifies the device number of the target device. If  
7 *device-modifier* does not appear in the clause, the behavior of the clause is as if *device-modifier* is  
8 **device\_num**.

9 If a **device** clause in which the **ancestor** *device-modifier* appears is present on the **target**  
10 construct and the **device** clause expression evaluates to 1, execution of the **target** region occurs  
11 on the parent device of the enclosing **target** region. If the **target** construct is not encountered  
12 in a **target** region, the current device is treated as the parent device. The encountering thread  
13 waits for completion of the **target** region on the parent device before resuming. For any list item  
14 that appears in a **map** clause on the same construct, if the corresponding list item exists in the device  
15 data environment of the parent device, it is treated as if it has a reference count of positive infinity.

16 If the **nowait** clause is present, execution of the *target task* may be deferred. If the **nowait**  
17 clause is not present, the *target task* is an included task.

18 If a **depend** clause is present, it is associated with the *target task*.

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

21 The **is\_device\_ptr** clause is used to indicate that a list item is a device pointer already in the  
22 device data environment and that it should be used directly. Support for device pointers created  
23 outside of OpenMP, specifically outside of any OpenMP mechanism that returns a device pointer, is  
24 implementation defined.

25 If a function (C, C++, Fortran) or subroutine (Fortran) is referenced in a **target** construct that  
26 does not specify a **device** clause in which the **ancestor** *device-modifier* appears then that  
27 function or subroutine is treated as if its name had appeared in a **to** clause on a **declare target**  
28 directive.

29 Each memory *allocator* specified in the **uses\_allocators** clause will be made available in the  
30 **target** region. For each non-predefined allocator that is specified, a new allocator handle will be  
31 associated with an allocator that is created with the specified *traits* as if by a call to  
32 **omp\_init\_allocator** at the beginning of the **target** region. Each non-predefined allocator  
33 will be destroyed as if by a call to **omp\_destroy\_allocator** at the end of the **target** region.

▼ C / C++ ▼

34 If a list item in a **map** clause has a base pointer and it is a scalar variable with a predetermined  
35 data-sharing attribute of *firstprivate* (see Section 2.22.1.1 on page 271), then on entry to the  
36 **target** region:

- If the list item is not a zero-length array section, the corresponding private variable is initialized such that the corresponding list item in the device data environment can be accessed through the pointer in the **target** region.
- If the list item is a zero-length array section, the corresponding private variable is initialized according to Section 2.22.7.2 on page 327.

▲ C / C++ ▲

▼ Fortran ▼

When an internal procedure is called in a **target** region, any references to variables that are host associated in the procedure have unspecified behavior.

▲ Fortran ▲

## Execution Model Events

Events associated with a *target task* are the same as for the **task** construct defined in Section 2.13.1 on page 142.

Events associated with the *initial task* that executes the **target** region are defined in Section 2.13.5 on page 155.

The *target-begin* event occurs when a thread enters a **target** region.

The *target-end* event occurs when a thread exits a **target** region.

The *target-submit* event occurs prior to creating an initial task on a target device for a **target** region.

## Tool Callbacks

Callbacks associated with events for *target tasks* are the same as for the **task** construct defined in Section 2.13.1 on page 142; (*flags & ompt\_task\_target*) always evaluates to *true* in the dispatched callback.

A thread dispatches a registered **ompt\_callback\_target** callback with **ompt\_scope\_begin** as its *endpoint* argument and **ompt\_target** as its *kind* argument for each occurrence of a *target-begin* event in that thread in the context of the target task on the host. Similarly, a thread dispatches a registered **ompt\_callback\_target** callback with **ompt\_scope\_end** as its *endpoint* argument and **ompt\_target** as its *kind* argument for each occurrence of a *target-end* event in that thread in the context of the target task on the host. These callbacks have type signature **ompt\_callback\_target\_t**.

A thread dispatches a registered **ompt\_callback\_target\_submit** callback for each occurrence of a *target-submit* event in that thread. The callback has type signature **ompt\_callback\_target\_submit\_t**.



## Restrictions

- If a **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.
- At most one **if** 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.
- A list item that appears in an **is\_device\_ptr** clause must not be specified on any data-sharing attribute clause on the same target construct.
- 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()** or to the value of **omp\_get\_initial\_device()**.
- 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.

- 1 • Predefined allocators appearing in a **uses\_allocators** clause cannot have *traits* specified.
- 2 • Non-predefined allocators appearing in a **uses\_allocators** clause must have *traits* specified.
- 3 • Arrays that contain allocator traits that appear in a **uses\_allocators** clause must be
- 4 constant arrays, have constant values and be defined in the same scope as the construct in which
- 5 the clause appears.
- 6 • Any IEEE floating-point exception status flag, halting mode, or rounding mode set prior to a
- 7 **target** region is unspecified in the region.
- 8 • Any IEEE floating-point exception status flag, halting mode, or rounding mode set in a **target**
- 9 region is unspecified upon exiting the region.

C / C++

- 10 • An attached pointer must not be modified in a **target** region.

C / C++

C

- 11 • A list item that appears in an **is\_device\_ptr** clause must have a type of pointer or array.

C

C++

- 12 • A list item that appears in an **is\_device\_ptr** clause must have a type of pointer, array,
- 13 reference to pointer or reference to array.

- 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 • An attached pointer that is associated with a given pointer target must not become associated
- 17 with a different pointer target in a **target** region.
- 18 • A list item that appears in an **is\_device\_ptr** clause must be a pointer or a dummy argument
- 19 that does not have the **ALLOCATABLE** or **VALUE** attribute.
- 20 • If a list item in a **map** clause is an array section, and the array section is derived from a variable
- 21 with a **POINTER** or **ALLOCATABLE** attribute then the behavior is unspecified if the
- 22 corresponding list item's variable is modified in the region.
- 23 • A reference to a coarray that is encountered on a target device must not be coindexed or appear as
- 24 an actual argument to a procedure where the corresponding dummy argument is a coarray.

Fortran

## Cross References

- *default-device-var*, see Section 2.6 on page 66.
- **task** construct, see Section 2.13.1 on page 142.
- **task** scheduling constraints, see Section 2.13.6 on page 156
- Memory allocators, see Section 2.14.2 on page 159.
- **target data** construct, see Section 2.15.2 on page 166.
- **if** Clause, see Section 2.18 on page 224.
- **private** and **firstprivate** clauses, see Section 2.22.4 on page 284.
- Data-Mapping Attribute Rules and Clauses, see Section 2.22.7 on page 317.
- **omp\_get\_num\_devices** routine, see Section 3.2.36 on page 370.
- **omp\_alloctrail\_t** and **omp\_alloctrail** types, see Section 3.8.1 on page 413.
- **omp\_set\_default\_allocator** routine, see Section 3.8.4 on page 417.
- **omp\_get\_default\_allocator** routine, see Section 3.8.5 on page 418.
- **ompt\_callback\_target\_t**, see Section 4.5.2.26 on page 490.
- **ompt\_callback\_target\_submit\_t**, Section 4.5.2.28 on page 493.

## 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:



## 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 and there is no **present** modifier in the clause, 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. If a **present** modifier appears in the clause and the corresponding list item is not present in the device data environment then an error occurs and the program terminates.

The list items that appear in a **to** or **from** may reference iterators defined by an *iterators-definition* appearing on an **iterator** modifier.

For each list item in a **from** or a **to** clause:

- For each part of the list item that is an attached pointer:

### C / C++

- 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;

### C / C++

### Fortran

- On exit from the region that part of the original list item, if associated, will be associated with the same pointer target with which it was associated on entry to the region;
- On exit from the region that part of the corresponding list item, if associated, will be associated with the same pointer target with which it was associated on entry to the region.

### Fortran

- 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 data races:
  - 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.

1 The list items that appear in the **to** or **from** clauses may use shape-operators.

2 The list items that appear in the **to** or **from** clauses may include array sections with *stride*  
3 expressions.

4 The **target update** construct is a task generating construct. The generated task is a *target task*.  
5 The generated task region encloses the **target update** region.

6 All clauses are evaluated when the **target update** construct is encountered. The data  
7 environment of the *target task* is created according to the data-sharing attribute clauses on the  
8 **target update** construct, per-data environment ICVs, and any default data-sharing attribute  
9 rules that apply to the **target update** construct. A variable that is mapped in the  
10 **target update** construct has a default data-sharing attribute of **shared** in the data  
11 environment of the *target task*.

12 Assignment operations associated with mapping a variable (see Section 2.22.7.1 on page 318)  
13 occur when the *target task* executes.

14 If the **nowait** clause is present, execution of the *target task* may be deferred. If the **nowait**  
15 clause is not present, the *target task* is an included task.

16 If a **depend** clause is present, it is associated with the *target task*.

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

## 20 Execution Model Events

21 Events associated with a *target task* are the same as for the **task** construct defined in  
22 Section 2.13.1 on page 142.

23 The *target-update-begin* event occurs when a thread enters a **target update** region.

24 The *target-update-end* event occurs when a thread exits a **target update** region.

## 25 Tool Callbacks

26 Callbacks associated with events for *target tasks* are the same as for the **task** construct defined in  
27 Section 2.13.1 on page 142; (*flags & ompt\_task\_target*) always evaluates to *true* in the  
28 dispatched callback.

29 A thread dispatches a registered **ompt\_callback\_target** callback with  
30 **ompt\_scope\_begin** as its *endpoint* argument and **ompt\_target\_update** as its *kind*  
31 argument for each occurrence of a *target-update-begin* event in that thread in the context of the  
32 target task on the host. Similarly, a thread dispatches a registered **ompt\_callback\_target**  
33 callback with **ompt\_scope\_end** as its *endpoint* argument and **ompt\_target\_update** as its

1 *kind* argument for each occurrence of a *target-update-end* event in that thread in the context of the  
2 target task on the host. These callbacks have type signature `ompt_callback_target_t`.

### 3 **Restrictions**

- 4 • A program must not depend on any ordering of the evaluations of the clauses of the  
5 **target update** directive, or on any side effects of the evaluations of the clauses.
- 6 • Each of the *motion-modifier* modifiers can appear at most once on a motion clause.
- 7 • At least one *motion-clause* must be specified.
- 8 • A list item can only appear in a **to** or **from** clause, but not both.
- 9 • A list item in a **to** or **from** clause must have a mappable type.
- 10 • At most one **device** clause can appear on the directive. The **device** clause expression must  
11 evaluate to a non-negative integer value less than the value of `omp_get_num_devices()` or  
12 to the value of `omp_get_initial_device()`.
- 13 • At most one **if** clause can appear on the directive.
- 14 • At most one **nowait** clause can appear on the directive.

### 15 **Cross References**

- 16 • Array shaping, Section [2.1.4](#) on page [42](#)
- 17 • Array sections, Section [2.1.5](#) on page [42](#)
- 18 • Iterators, see Section [2.1.6](#) on page [46](#).
- 19 • *default-device-var*, see Section [2.6](#) on page [66](#).
- 20 • **task** construct, see Section [2.13.1](#) on page [142](#).
- 21 • **task** scheduling constraints, see Section [2.13.6](#) on page [156](#)
- 22 • **target data**, see Section [2.15.2](#) on page [166](#).
- 23 • **if** Clause, see Section [2.18](#) on page [224](#).
- 24 • `omp_get_num_devices` routine, see Section [3.2.36](#) on page [370](#).
- 25 • `ompt_callback_task_create_t`, see Section [4.5.2.7](#) on page [469](#).
- 26 • `ompt_callback_target_t`, see Section [4.5.2.26](#) on page [490](#).

## 27 **2.15.7 Declare Target Directive**

### 28 **Summary**

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

## Syntax

### C / C++

The syntax of the `declare target` directive is as follows:

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

or

```
#pragma omp declare target (extended-list) new-line
```

or

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

or

```
#pragma omp begin declare target [clause [ [ , ] clause ] ... ] new-line  
    declaration-definition-seq  
#pragma omp end declare target new-line
```

where *clause* is one of the following:

```
to (extended-list)  
link (list)  
device_type (host | nohost | any)  
indirect [ (invoked-by-fptr) ]
```

where *invoked-by-fptr* is a constant boolean expression that evaluates to true or false at compile time.

### 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)  
device_type (host | nohost | any)  
indirect [ (invoked-by-fptr) ]
```



1 where *invoked-by-fptr* is a constant **logical** expression that evaluates to true or false at compile  
2 time.

## Fortran

### Description

3 The declare target directive ensures that procedures and global variables can be executed or  
4 accessed on a device. Variables are mapped for all device executions, or for specific device  
5 executions through a **link** clause.  
6

7 If an *extended-list* is present with no clause then the **to** clause is assumed.

8 The **device\_type** clause specifies if a version of the procedure or variable should be made  
9 available on the host, device or both. If **host** is specified only a host version of the procedure or  
10 variable is made available. If **any** is specified then both device and host versions of the procedure  
11 or variable are made available. If **nohost** is specified for a procedure then only a device version of  
12 the procedure is made available. If **nohost** is specified for a variable then that variable is not  
13 available on the host.

## C / C++

14 If a function appears in a **to** clause in the same compilation unit in which the definition of the  
15 function occurs then a device-specific version of the function is created.

16 If a variable appears in a **to** clause in the same compilation unit in which the definition of the  
17 variable occurs then the original list item is allocated a corresponding list item in the device data  
18 environment of all devices.

## C / C++

## Fortran

19 If a procedure appears in a **to** clause in the same compilation unit in which the definition of the  
20 procedure occurs then a device-specific version of the procedure is created.

21 If a variable that is host associated appears in a **to** clause then the original list item is allocated a  
22 corresponding list item in the device data environment of all devices.

## Fortran

23 If a variable appears in a **to** clause then the corresponding list item in the device data environment  
24 of each device is initialized once, in the manner specified by the program, but at an unspecified  
25 point in the program prior to the first reference to that list item. The list item is never removed from  
26 those device data environments as if its reference count is initialized to positive infinity.

27 Including list items in a **link** clause supports compilation of functions called in a **target** region  
28 that refer to the list items. The list items are not mapped by the declare target directive. Instead,  
29 they are mapped according to the data mapping rules described in Section 2.22.7 on page 317.

C / C++

1 If a function is referenced in a function that appears as a list item in a **to** clause on a **declare target**  
2 directive that does not specify a **device\_type** clause with **host** and the function reference is  
3 not enclosed in a **target** construct that specifies a **device** clause in which the **ancestor**  
4 *device-modifier* appears then the name of the referenced function is treated as if it had appeared in a  
5 **to** clause on a **declare target** directive.

6 If a variable with static storage duration or a function (except *lambda* for C++) is referenced in the  
7 initializer expression list of a variable with static storage duration that appears as a list item in a **to**  
8 clause on a **declare target** directive then the name of the referenced variable or function is treated as  
9 if it had appeared in a **to** clause on a **declare target** directive.

10 The form, preceded by either the **declare target** directive that has no clauses and no  
11 *extended-list* or the **begin declare target** directive and followed by a matching  
12 **end declare target** directive, defines an implicit *extended-list*. The implicit *extended-list*  
13 consists of the variable names of any variable declarations at file or namespace scope that appear  
14 between the two directives and of the function names of any function declarations at file,  
15 namespace or class scope that appear between the two directives. The implicit *extended-list* is  
16 converted to an implicit **to** clause.

17 The *declaration-definition-seq* preceded by either **begin declare target** directive or a  
18 **declare target** directive without any clauses or an *extended-list* and followed by an  
19 **end declare target** directive may contain **declare target** directives. If a **device\_type**  
20 clause is present on the contained **declare target** directive, then its argument determines which  
21 versions are made available. If a list item appears both in an implicit and explicit list, the explicit  
22 list determines which versions are made available.

C / C++

Fortran

23 If a procedure is referenced in a procedure that appears as a list item in a **to** clause on a  
24 **declare target** directive that does not specify a **device\_type** clause with **host** and the  
25 procedure reference is not enclosed in a **target** construct that specifies a **device** clause in  
26 which the **ancestor device-modifier** appears then the name of the referenced procedure is treated  
27 as if it had appeared in a **to** clause on a **declare target** directive.

28 If a **declare target** does not have any clauses and does not have an *extended-list* then an  
29 implicit **to** clause with one item is formed from the name of the enclosing subroutine subprogram,  
30 function subprogram or interface body to which it applies.

31 If a **declare target** directive has a **device\_type** clause then any enclosed internal  
32 procedures cannot contain any **declare target** directives. The enclosing **device\_type**  
33 clause implicitly applies to internal procedures.

Fortran

1 If the **indirect** clause is present and *invoked-by-fptr* is not specified, the effect of the clause is as  
2 if *invoked-by-fptr* evaluates to true.

3 If the **indirect** clause is present and *invoked-by-fptr* evaluates to true, any procedures that appear  
4 in a **to** clause on the directive may be called with an indirect device invocation. If the **indirect**  
5 clause is present and *invoked-by-fptr* does not evaluate to true, any procedures that appear in a **to**  
6 clause on the directive may not be called with an indirect device invocation. Unless otherwise  
7 specified by an **indirect** clause, procedures may not be called with an indirect device invocation.

C / C++

8 If a function appears in the **to** clause of a **begin declare target** directive and in the **to**  
9 clause of a declare target directive that is contained in the *declaration-definition-seq* associated with  
10 the **begin declare target** directive, and if an **indirect** clause appears on both directives,  
11 then the **indirect** clause on the **begin declare target** directive has no effect for that  
12 function.

C / C++

## Restrictions

- 14 • A **threadprivate** variable cannot appear in a declare target directive.
- 15 • A variable declared in a declare target directive must have a mappable type.
- 16 • The same list item must not appear multiple times in clauses on the same directive.
- 17 • The same list item must not explicitly appear in both a **to** clause on one declare target directive  
18 and a **link** clause on another declare target directive.
- 19 • If a declare target directive has a clause, it must contain at least one **to** clause or at least one  
20 **link** clause.
- 21 • A variable for which **nohost** is specified may not appear in a **link** clause.
- 22 • At most one **indirect** clause can be specified on the declare target directive.
- 23 • If an **indirect** clause is present and *invoked-by-fptr* evaluates to true then the only permitted  
24 **device\_type** clause is **device\_type(any)**.

C++

- 25 • A **to** clause or a **link** clause cannot appear in a **begin declare target** directive.
- 26 • The function names of overloaded functions or template functions may only be specified within  
27 an implicit *extended-list*.
- 28 • If a *lambda declaration and definition* appears between a declare target directive and the paired  
29 **end declare target** directive, all variables that are captured by the *lambda* expression must  
30 also appear in a **to** clause.

C++

- 1 • If a list item is a procedure name, it must not be a generic name, procedure pointer or entry name.
- 2 • If a **declare target** directive does not have any clause, has a **device\_type** clause, or has
- 3 clauses that contain a function name, it must appear in the specification part of a subroutine
- 4 subprogram, function subprogram or interface body.
- 5 • If a **declare target** directive has a variable name in *extended-list*, it must appear in the
- 6 specification part of a subroutine subprogram, function subprogram, program or module.
- 7 • If a **declare target** directive is specified in an interface block for a procedure, it must match
- 8 a **declare target** directive in the definition of the procedure, including the **device\_type**
- 9 clause if present.
- 10 • If an external procedure is a type-bound procedure of a derived type and a **declare target**
- 11 directive is specified in the definition of the external procedure, such a directive must appear in
- 12 the interface block that is accessible to the derived type definition.
- 13 • If any procedure is declared via a procedure declaration statement that is not in the type-bound
- 14 procedure part of a derived-type definition, any **declare target** with the procedure name
- 15 must appear in the same specification part.
- 16 • A variable that is part of another variable (as an array, structure element or type parameter
- 17 inquiry) cannot appear in a **declare target** directive.
- 18 • The **declare target** directive must appear in the declaration section of a scoping unit in
- 19 which the common block or variable is declared.
- 20 • If a **declare target** directive that specifies a common block name appears in one program
- 21 unit, then such a directive must also appear in every other program unit that contains a **COMMON**
- 22 statement that specifies the same name, after the last such **COMMON** statement in the program unit.
- 23 • If a list item is declared with the **BIND** attribute, the corresponding C entities must also be
- 24 specified in a **declare target** directive in the C program.
- 25 • A blank common block cannot appear in a **declare target** directive.
- 26 • A variable can only appear in a **declare target** directive in the scope in which it is declared.
- 27 It must not be an element of a common block or appear in an **EQUIVALENCE** statement.
- 28 • A variable that appears in a **declare target** directive must be declared in the Fortran scope
- 29 of a module or have the **SAVE** attribute, either explicitly or implicitly.

### Cross References

- 30 • **target data** construct, see Section [2.15.2](#) on page [166](#).
- 31 • **target** construct, see Section [2.15.5](#) on page [174](#).
- 32

## 2.15.8 `interop` Construct

### Summary

The **`interop`** construct retrieves interoperability properties from the OpenMP implementation to enable interoperability with foreign execution contexts. The **`interop`** construct is a stand-alone directive.

### Syntax

The syntax of the **`interop`** construct is as follows:

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

where *clause* is one of the following:

```
obj (interop-var)
init
destroy
tasksync
device (integer-expression)
depend ([depend-modifier, ] dependence-type : locator-list)
```

where *interop-var* is a variable of type `omp_interop_t`.

### Binding

The binding task set for an **`interop`** region is the generating task. The **`interop`** region binds to the region of the generating task.

### Description

When an **`interop`** construct is encountered, the encountering task executes the region. If no **`device`** clause is present, the default device is determined by the *default-device-var* ICV.

If neither the **`init`** clause nor the **`destroy`** clause is specified, an implicit **`init`** clause is assumed.

If the **`init`** clause is specified, the *interop-var* specified in the **`obj`** clause is initialized to refer to a foreign execution context with the **`type`**, **`type_name`**, **`device`** and **`device_context`** properties appropriately defined for the specified device.

If the **`tasksync`** clause is specified, the **`tasksync`** property for the *interop-var* specified by the **`obj`** clause will be initialized to enable synchronization between OpenMP tasks and foreign tasks that execute in the foreign execution context to which the *interop-var* specified in the **`obj`** clause is initialized.

1 If the implementation is unable to initialize the *interop-var* specified in the **obj** clause according to  
2 the above, *interop-var* will be initialized to the value of **OMP\_INTEROP\_NONE**.

3 If the **destroy** clause is specified, the *interop-var* specified in the **obj** clause must be non-const  
4 and will be set to the value of **OMP\_INTEROP\_NONE** after releasing resources associated with  
5 *interop-var*. The object referred to by the *interop-var* specified in the **obj** clause will be unusable  
6 after destruction until it is initialized again by an **interop** construct.

7 If the *interop-var* specified in the **obj** clause is initialized with the **tasksync** property, an empty  
8 task is generated. Any **depend** clauses that are present on the construct apply to the generated  
9 task. The **interop** construct ensures an ordered execution of the generated task relative to foreign  
10 tasks executed in the foreign execution context through the foreign synchronization object to which  
11 the *interop-var* specified in the **obj** clause is initialized. When the creation of the foreign task  
12 precedes the encountering of an **interop** construct in happens before order (see Section 1.4.5),  
13 the foreign task must complete execution before the generated task begins execution. Likewise,  
14 when the creation of a foreign task follows the encountering of an **interop** construct in happens  
15 before order, the foreign task must not begin execution until the generated task completes execution.

## 16 Restrictions

- 17 • Exactly one **obj** clause must appear on the directive.
- 18 • At most one **init** clause can appear on the directive.
- 19 • At most one **destroy** clause can appear on the directive.
- 20 • The *interop-var* must be non-const if a **init** clause appears on the directive
- 21 • Neither an **init** clause nor a **tasksync** clause can be specified if a **destroy** clause is  
22 specified on the directive.
- 23 • At most one **tasksync** clause can appear on the directive.
- 24 • A **depend** clause can only appear on the directive if a **tasksync** clause is present or if a  
25 **destroy** clause is present and the *interop-var* specified in the **obj** clause was initialized by a  
26 prior **interop** construct on which the **tasksync** clause was present.
- 27 • At most one **device** clause can appear on the directive. The **device** clause expression must  
28 evaluate to a non-negative integer value less than the value of **omp\_get\_num\_devices()** or  
29 to the value of **omp\_get\_initial\_device()**.

## 30 Cross References

- 31 • Interoperability routines, see Section 3.7 on page 410.
- 32 • **depend** clause, see Section 2.20.11 on page 257(WRONG ORDERING:  
33 directives\_device.tex line 1561).

▲ C / C++ ▲

## 2.16 Combined Constructs

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

For combined constructs, tool callbacks are invoked as if the constructs were explicitly nested.

### 2.16.1 Parallel Worksharing-Loop Construct

#### Summary

The parallel worksharing-loop construct is a shortcut for specifying a **parallel** construct containing a worksharing-loop construct with one or more associated loops and no other 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 2.8 on page 79.
- Worksharing-loop construct, see Section 2.12.2 on page 106.
- Data attribute clauses, see Section 2.22.4 on page 284.

## 2.16.2 parallel loop Construct

### Summary

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

### Syntax

C / C++

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

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

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

C / C++

Fortran

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

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

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

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

Fortran

### Description

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

### Restrictions

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



## Cross References

- **parallel** construct, see Section 2.8 on page 79.
- **loop** construct, see Section 2.12.5 on page 133.
- Data attribute clauses, see Section 2.22.4 on page 284.

## 2.16.3 parallel sections Construct

### Summary

The **parallel sections** construct is a shortcut for specifying a **parallel** construct containing a **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-sequence
  #pragma omp section new-line
    structured-block-sequence]
  ...
}
```

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-sequence
  !$omp section
    structured-block-sequence]
  ...
!$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 2.8 on page 79.
- **sections** construct, see Section 2.11.1 on page 93.
- Data attribute clauses, see Section 2.22.4 on page 284.

Fortran

## 2.16.4 parallel workshare Construct

### Summary

The **parallel workshare** construct is a shortcut for specifying a **parallel** construct containing a **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 2.8 on page 79.
- **workshare** construct, see Section 2.11.3 on page 98.
- Data attribute clauses, see Section 2.22.4 on page 284.

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 a worksharing-loop SIMD construct and no other statements.

### 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 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 **simd** *directive-name-modifier* can appear on the directive.

## Cross References

- **parallel** construct, see Section 2.8 on page 79.
- Worksharing-loop SIMD construct, see Section 2.12.3.2 on page 120.
- Data attribute clauses, see Section 2.22.4 on page 284.

## 2.16.6 parallel master Construct

### Summary

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

### Syntax

C / C++

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

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

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

C / C++

Fortran

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

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

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

Fortran

## Description

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

## Restrictions

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

## Cross References

- **parallel** construct, see Section 2.8 on page 79.
- **master** construct, see Section 2.19 on page 225.
- Data attribute clauses, see Section 2.22.4 on page 284.

## 2.16.7 master taskloop Construct

### Summary

The **master taskloop** construct is a shortcut for specifying a **master** construct containing a **taskloop** construct and no other statements.

### Syntax

C / C++

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

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

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

C / C++

Fortran

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

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

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

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

Fortran

## Description

The semantics are identical to explicitly specifying a **master** directive immediately followed by a **taskloop** directive.

## Restrictions

The restrictions for the **master** and **taskloop** constructs apply.

## Cross References

- **taskloop** construct, see Section 2.13.2 on page 147.
- **master** construct, see Section 2.19 on page 225.
- Data attribute clauses, see Section 2.22.4 on page 284.

## 2.16.8 master taskloop simd Construct

### Summary

The **master taskloop simd** construct is a shortcut for specifying a **master** construct containing a **taskloop simd** construct and no other statements.

### Syntax

C / C++

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

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

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

C / C++

Fortran

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

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

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

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

Fortran

## Description

The semantics are identical to explicitly specifying a **master** directive immediately followed by a **taskloop simd** directive.

## Restrictions

The restrictions for the **master** and **taskloop simd** constructs apply.

## Cross References

- **taskloop simd** construct, see Section 2.13.3 on page 152.
- **master** construct, see Section 2.19 on page 225.
- Data attribute clauses, see Section 2.22.4 on page 284.

## 2.16.9 parallel master taskloop Construct

### Summary

The **parallel master taskloop** construct is a shortcut for specifying a **parallel** construct containing a **master taskloop** construct and no other statements.

### Syntax

C / C++

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

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

where *clause* can be any of the clauses accepted by the **parallel** or **master taskloop** directives, except the **in\_reduction** clause, with identical meanings and restrictions.

C / C++

Fortran

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

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

where *clause* can be any of the clauses accepted by the **parallel** or **master taskloop** directives, except the **in\_reduction** clause, with identical meanings and restrictions.

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

Fortran

## Description

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

## Restrictions

The restrictions for the **parallel** construct and the **master taskloop** construct 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 **taskloop** *directive-name-modifier* can appear on the directive.

## Cross References

- **parallel** construct, see Section 2.8 on page 79.
- **master taskloop** construct, see Section 2.16.7 on page 199.
- Data attribute clauses, see Section 2.22.4 on page 284.

## 2.16.10 parallel master taskloop simd Construct

### Summary

The **parallel master taskloop simd** construct is a shortcut for specifying a **parallel** construct containing a **master taskloop simd** construct and no other statements.

### Syntax

C / C++

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

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

where *clause* can be any of the clauses accepted by the **parallel** or **master taskloop simd** directives, except the **in\_reduction** clause, with identical meanings and restrictions.

C / C++



## Fortran

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

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

where *clause* can be any of the clauses accepted by the **parallel** or **master taskloop simd** directives, except the **in\_reduction** clause, with identical meanings and restrictions.

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

## Fortran

### Description

The semantics are identical to explicitly specifying a **parallel** directive immediately followed by a **master taskloop simd** directive.

### Restrictions

The restrictions for the **parallel** construct and the **master taskloop simd** construct 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 **taskloop** *directive-name-modifier* can appear on the directive.
- At most one **if** clause with the **simd** *directive-name-modifier* can appear on the directive.

### Cross References

- **parallel** construct, see Section 2.8 on page 79.
- **master taskloop simd** construct, see Section 2.16.8 on page 200.
- Data attribute clauses, see Section 2.22.4 on page 284.

## 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

## Description

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

## Restrictions

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

## Cross References

- **teams** construct, see Section 2.9 on page 87.
- **distribute** construct, see Section 2.12.4.1 on page 125.
- Data attribute clauses, see Section 2.22.4 on page 284.

## 2.16.12 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 [2.9](#) on page [87](#).
- **distribute simd** construct, see Section [2.12.4.2](#) on page [129](#).
- Data attribute clauses, see Section [2.22.4](#) on page [284](#).

## 2.16.13 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

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

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

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

If an **end teams distribute parallel do** directive is not specified, an **end teams distribute parallel do** 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 directive.

## Restrictions

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

## Cross References

- **teams** construct, see Section 2.9 on page 87.
- Distribute parallel worksharing-loop construct, see Section 2.12.4.3 on page 130.
- Data attribute clauses, see Section 2.22.4 on page 284.

## 2.16.14 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 SIMD 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 SIMD 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 2.9 on page 87.
- Distribute parallel worksharing-loop SIMD construct, see Section 2.12.4.4 on page 131.
- Data attribute clauses, see Section 2.22.4 on page 284.

## 2.16.15 teams loop Construct

### Summary

The **teams loop** construct is a shortcut for specifying a **teams** construct containing a **loop** construct and no other statements.

## Syntax

C / C++

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

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

where *clause* can be any of the clauses accepted by the **teams** or **loop** directives with identical 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 2.9 on page 87.
- **loop** construct, see Section 2.12.5 on page 133.
- Data attribute clauses, see Section 2.22.4 on page 284.

## 2.16.16 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

## Description

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

## Restrictions

The restrictions for the **target** and **parallel** 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

- **parallel** construct, see Section 2.8 on page 79.
- **target** construct, see Section 2.15.5 on page 174.
- **if** Clause, see Section 2.18 on page 224.
- Data attribute clauses, see Section 2.22.4 on page 284.

## 2.16.17 Target Parallel Worksharing-Loop Construct

### Summary

The target parallel worksharing-loop construct is a shortcut for specifying a **target** construct 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 174.
- Parallel Worksharing-Loop construct, see Section 2.16.1 on page 193.
- **if** Clause, see Section 2.18 on page 224.
- Data attribute clauses, see Section 2.22.4 on page 284.

## 2.16.18 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

### Description

The semantics are identical to explicitly specifying a **target** directive immediately followed by a parallel worksharing-loop SIMD directive.

## Restrictions

The restrictions for the **target** and parallel worksharing-loop SIMD 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.
- At most one **if** clause with the **simd** *directive-name-modifier* can appear on the directive.

## Cross References

- **target** construct, see Section 2.15.5 on page 174.
- Parallel worksharing-loop SIMD construct, see Section 2.16.5 on page 197.
- **if** Clause, see Section 2.18 on page 224.
- Data attribute clauses, see Section 2.22.4 on page 284.

## 2.16.19 target parallel loop Construct

### Summary

The **target parallel loop** construct is a shortcut for specifying a **target** construct containing a **parallel loop** construct and no other statements.

### Syntax

C / C++

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

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

where *clause* can be any of the clauses accepted by the **target** or **parallel loop** directives, except for **copyin**, with identical meanings and restrictions.

C / C++

## Fortran

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

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

where *clause* can be any of the clauses accepted by the **target** or **parallel loop** directives, except for **copyin**, with identical meanings and restrictions.

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

## Fortran

### Description

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

### Restrictions

The restrictions for the **target** and **parallel 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 174.
- **parallel loop** construct, see Section 2.16.2 on page 194.
- Data attribute clauses, see Section 2.22.4 on page 284.

## 2.16.20 target simd Construct

### Summary

The **target simd** construct is a shortcut for specifying a **target** construct containing a **simd** 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 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 **target** *directive-name-modifier* can appear on the directive.
- At most one **if** clause with the **simd** *directive-name-modifier* can appear on the directive.

## Cross References

- **simd** construct, see Section 2.12.3.1 on page 115.
- **target** construct, see Section 2.15.5 on page 174.
- Data attribute clauses, see Section 2.22.4 on page 284.

## 2.16.21 target teams Construct

### Summary

The **target teams** construct is a shortcut for specifying a **target** construct containing a **teams** construct and no other statements.

### Syntax

C / C++

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

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

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

C / C++

Fortran

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

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

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

Fortran

### Description

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

### Restrictions

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

### Cross References

- **teams** construct, see Section 2.9 on page 87.
- **target** construct, see Section 2.15.5 on page 174.
- Data attribute clauses, see Section 2.22.4 on page 284.

## 2.16.22 target teams distribute Construct

### Summary

The **target teams distribute** construct is a shortcut for specifying a **target** construct containing a **teams distribute** construct and no other statements.

## Syntax

C / C++

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

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

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

C / C++

Fortran

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

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

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

If an **end target teams distribute** directive is not specified, an **end target teams distribute** 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** directive.

## Restrictions

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

## Cross References

- **target** construct, see Section 2.15.2 on page 166.
- **teams distribute** construct, see Section 2.16.11 on page 203.
- Data attribute clauses, see Section 2.22.4 on page 284.

## 2.16.23 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 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 **target** *directive-name-modifier* can appear on the directive.
- At most one **if** clause with the **simd** *directive-name-modifier* can appear on the directive.

## Cross References

- **target** construct, see Section [2.15.2](#) on page [166](#).
- **teams distribute simd** construct, see Section [2.16.12](#) on page [204](#).
- Data attribute clauses, see Section [2.22.4](#) on page [284](#).

## 2.16.24 target teams loop Construct

### Summary

The **target teams loop** construct is a shortcut for specifying a **target** construct containing a **teams loop** construct and no other statements.

### Syntax

C / C++

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

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

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

C / C++

Fortran

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

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

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

If an **end target teams loop** directive is not specified, an **end target teams loop** 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 loop** directive.

### Restrictions

The restrictions for the **target** and **teams loop** constructs.

### Cross References

- **target** construct, see Section 2.15.5 on page 174.
- Teams loop construct, see Section 2.16.15 on page 207.
- Data attribute clauses, see Section 2.22.4 on page 284.



## 2.16.25 Target Teams Distribute Parallel Worksharing-Loop Construct

### Summary

The target teams distribute parallel worksharing-loop construct is a shortcut for specifying a **target** construct containing a teams distribute parallel worksharing-loop construct and no other 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 174.
- **Teams** distribute parallel worksharing-loop construct, see Section 2.16.13 on page 205.
- **if** Clause, see Section 2.18 on page 224.
- Data attribute clauses, see Section 2.22.4 on page 284.

## 2.16.26 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

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

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

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

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

## Fortran

### Description

The semantics are identical to explicitly specifying a **target** directive immediately followed by a teams distribute parallel worksharing-loop SIMD directive.

### Restrictions

The restrictions for the **target** and teams distribute parallel worksharing-loop SIMD 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.
- At most one **if** clause with the **simd** *directive-name-modifier* can appear on the directive.

### Cross References

- **target** construct, see Section 2.15.5 on page 174.
- Teams distribute parallel worksharing-loop SIMD construct, see Section 2.16.14 on page 206.
- **if** Clause, see Section 2.18 on page 224.
- Data attribute clauses, see Section 2.22.4 on page 284.

## 2.17 Clauses on Combined and Composite Constructs

This section specifies the handling of clauses on combined or composite constructs and the handling of implicit clauses from variables with predetermined data sharing if they are not predetermined only on a particular construct. Some clauses are permitted only on a single construct of the constructs that constitute the combined or composite construct, in which case the effect is as if the clause is applied to that specific construct. As detailed in this section, other clauses have the effect as if they are applied to one or more constituent constructs.

The **collapse** clause is applied once to the combined or composite construct.

The effect of the **private** clause is as if it is applied only to the innermost constituent construct that permits it.

The effect of the **firstprivate** clause is as if it is applied to one or more constructs as follows:

- To the **distribute** construct if it is among the constituent constructs;
- To the **teams** construct if it is among the constituent constructs and the **distribute** construct is not;
- To the worksharing-loop construct if it is among the constituent constructs;
- To the **taskloop** construct if it is among the constituent constructs;
- To the **parallel** construct if it is among the constituent constructs and the worksharing-loop construct or the **taskloop** construct is not;
- To the **target** construct if it is among the constituent constructs and the same list item does not appear in a **lastprivate** or **map** clause.

If the **parallel** construct is among the constituent constructs and the effect is not as if the **firstprivate** clause is applied to it by the above rules, then the effect is as if the **shared** clause with the same list item is applied to the **parallel** construct. If the **teams** construct is among the constituent constructs and the effect is not as if the **firstprivate** clause is applied to it by the above rules, then the effect is as if the **shared** clause with the same list item is applied to the **teams** construct.

The effect of the **lastprivate** clause is as if it is applied to one or more constructs as follows:

- To the worksharing-loop construct if it is among the constituent constructs;
- To the **taskloop** construct if it is among the constituent constructs;
- To the **distribute** construct if it is among the constituent constructs; and
- To the innermost constituent construct that permits it unless it is a worksharing-loop or **distribute** construct.

1 If the **parallel** construct is among the constituent constructs and the list item is not also specified  
2 in the **firstprivate** clause, then the effect of the **lastprivate** clause is as if the **shared**  
3 clause with the same list item is applied to the **parallel** construct. If the **teams** construct is  
4 among the constituent constructs and the list item is not also specified in the **firstprivate**  
5 clause, then the effect of the **lastprivate** clause is as if the **shared** clause with the same list  
6 item is applied to the **teams** construct. If the **target** construct is among the constituent  
7 constructs and the list item is not specified in a **map** clause, the effect of the **lastprivate** clause  
8 is as if the same list item appears in a **map** clause with a *map-type* of **tofrom**.

9 The effect of the **shared**, **default**, **order**, or **allocate** clause is as if it is applied to all  
10 constituent constructs that permit the clause.

11 The effect of the **reduction** clause is as if it is applied to all constructs that permit the clause,  
12 except for the following constructs:

- 13 • The **parallel** construct, when combined with the **sections**, worksharing-loop, **loop**, or  
14 **taskloop** construct; and
- 15 • The **teams** construct, when combined with the **loop** construct.

16 For the **parallel** and **teams** constructs above, the effect of the **reduction** clause instead is as  
17 if each list item or, for any list item that is an array item, its corresponding base array or base  
18 pointer appears in a **shared** clause for the construct. If the **task reduction-modifier** is specified,  
19 the effect is as if it only modifies the behavior of the **reduction** clause on the innermost  
20 construct that constitutes the combined construct and that accepts the modifier (see Section 2.22.5.4  
21 on page 302). If the **inscan reduction-modifier** is specified, the effect is as if it modifies the  
22 behavior of the **reduction** clause on all constructs of the combined construct to which the clause  
23 is applied and that accept the modifier. If a construct to which the **inscan reduction-modifier** is  
24 applied is combined with the **target** construct, the effect is as if the same list item also appears in  
25 a **map** clause with a *map-type* of **tofrom**.

26 The **in\_reduction** clause is permitted on a single construct among those that constitute the  
27 combined or composite construct and the effect is as if the clause is applied to that construct, but if  
28 that construct is a **target** construct, the effect is also as if the same list item appears in a **map**  
29 clause with a *map-type* of **tofrom** and a *map-type-modifier* of **always**.

30 The effect of the **if** clause is described in Section 2.18 on page 224.

31 The effect of the **linear** clause is as if it is applied to the innermost constituent construct.  
32 Additionally, if the list item is not the iteration variable of a **simd** or worksharing-loop SIMD  
33 construct, the effect on the outer constituent constructs is as if the list item was specified in  
34 **firstprivate** and **lastprivate** clauses on the combined or composite construct, with the  
35 rules specified above applied. If a list item of the **linear** clause is the iteration variable of a  
36 **simd** or worksharing-loop SIMD construct and it is not declared in the construct, the effect on the  
37 outer constituent constructs is as if the list item was specified in a **lastprivate** clause on the  
38 combined or composite construct with the rules specified above applied.

1 The effect of the **nowait** clause is as if it is applied to the outermost constituent construct that  
2 permits it.

3 If the clauses have expressions on them, such as for various clauses where the argument of the  
4 clause is an expression, or *lower-bound*, *length*, or *stride* expressions inside array sections (or  
5 *subscript* and *stride* expressions in *subscript-triplet* for Fortran), or *linear-step* or *alignment*  
6 expressions, the expressions are evaluated immediately before the construct to which the clause has  
7 been split or duplicated per the above rules (therefore inside of the outer constituent constructs).  
8 However, the expressions inside the **num\_teams** and **thread\_limit** clauses are always  
9 evaluated before the outermost constituent construct.

10 The restriction that a list item may not appear in more than one data sharing clause with the  
11 exception of specifying a variable in both **firstprivate** and **lastprivate** clauses applies  
12 after the clauses are split or duplicated per the above rules.

## 13 2.18 if Clause

### 14 Summary

15 The semantics of an **if** clause are described in the section on the construct to which it applies. The  
16 **if** clause *directive-name-modifier* names the associated construct to which an expression applies,  
17 and is particularly useful for composite and combined constructs.

### 18 Syntax

19 C / C++

The syntax of the **if** clause is as follows:

20 **if** (*[ directive-name-modifier :] scalar-expression*)

21 C / C++

22 Fortran

The syntax of the **if** clause is as follows:

23 **if** (*[ directive-name-modifier :] scalar-logical-expression*)

24 Fortran

### 25 Description

26 The effect of the **if** clause depends on the construct to which it is applied. For combined or  
27 composite constructs, the **if** clause only applies to the semantics of the construct named in the  
28 *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.

### Description

Only the master thread of the team that executes the binding **parallel** region participates in the execution of the structured block of the **master** region. Other threads in the team do not execute the associated structured block. There is no implied barrier either on entry to, or exit from, the **master** construct.

### Execution Model Events

The *master-begin* event occurs in the master thread of a team that encounters the **master** construct on entry to the master region.

The *master-end* event occurs in the master thread of a team that encounters the **master** construct on exit from the master region.

### Tool Callbacks

A thread dispatches a registered `ompt_callback_master` callback with `ompt_scope_begin` as its *endpoint* argument for each occurrence of a *master-begin* event in that thread. Similarly, a thread dispatches a registered `ompt_callback_master` callback with `ompt_scope_end` as its *endpoint* argument for each occurrence of a *master-end* event in that thread. These callbacks occur in the context of the task executed by the master thread and have the type signature `ompt_callback_master_t`.

## Restrictions

C++

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

C++

## Cross References

- **parallel** construct, see Section 2.8 on page 79.
- **ompt\_scope\_begin** and **ompt\_scope\_end**, see Section 4.4.4.11 on page 447.
- **ompt\_callback\_master\_t**, see Section 4.5.2.12 on page 474.

## 2.20 Synchronization Constructs and Clauses

A synchronization construct orders the completion of code executed by different threads. This ordering is imposed by synchronizing flush operations that are executed as part of the region that corresponds to the construct.

Synchronization through the use of synchronizing flush operations and atomic operations is described in Section 1.4.4 on page 26 and Section 1.4.6 on page 29. Section 2.20.8.1 on page 249 defines the behavior of synchronizing flush operations that are implied at various other locations in an OpenMP program.

### 2.20.1 critical Construct

#### Summary

The **critical** construct restricts execution of the associated structured block to a single thread at a time.

#### Syntax

C / C++

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

```
#pragma omp critical [(name) [[,] hint (hint-expression)]] new-line
    structured-block
```

where *hint-expression* is an integer constant expression that evaluates to a valid synchronization hint (as described in Section 2.20.12 on page 262).

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 262).

## Fortran

### Binding

The binding thread set for a **critical** region is all threads in the contention group.

### Description

The region that corresponds to a **critical** construct is executed as if only a single thread at a time among all threads in the contention group enters the region for execution, without regard to the team(s) to which the threads belong. 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 executes 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.

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.

## Execution Model Events

The *critical-acquiring* event occurs in a thread that encounters the **critical** construct on entry to the **critical** region before initiating synchronization for the region.

The *critical-acquired* event occurs in a thread that encounters the **critical** construct after it enters the region, but before it executes the structured block of the **critical** region.

The *critical-released* event occurs in a thread that encounters the **critical** construct after it completes any synchronization on exit from the **critical** region.

## Tool Callbacks

A thread dispatches a registered **ompt\_callback\_mutex\_acquire** callback for each occurrence of a *critical-acquiring* event in that thread. This callback has the type signature **ompt\_callback\_mutex\_acquire\_t**.

A thread dispatches a registered **ompt\_callback\_mutex\_acquired** callback for each occurrence of a *critical-acquired* event in that thread. This callback has the type signature **ompt\_callback\_mutex\_t**.

A thread dispatches a registered **ompt\_callback\_mutex\_released** callback for each occurrence of a *critical-released* event in that thread. This callback has the type signature **ompt\_callback\_mutex\_t**.

The callbacks occur in the task that encounters the critical construct. The callbacks should receive **ompt\_mutex\_critical** as their *kind* argument if practical, but a less specific kind is acceptable.

## Restrictions

The following restrictions apply to the critical construct:

- Unless the effect is as if **hint(omp\_sync\_hint\_none)** was specified, the **critical** construct must specify a name.
- If the **hint** clause is specified, each of the **critical** constructs with the same *name* must have a **hint** clause for which the *hint-expression* evaluates to the same value.

C++

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

C++

Fortran

- If a *name* is specified on a **critical** directive, the same *name* must also be specified on the **end critical** directive.
- If no *name* appears on the **critical** directive, no *name* can appear on the **end critical** directive.

Fortran

## Cross References

- Synchronization Hints, see Section 2.20.12 on page 262.
- `ompt_mutex_critical`, see Section 4.4.4.16 on page 449.
- `ompt_callback_mutex_acquire_t`, see Section 4.5.2.14 on page 477.
- `ompt_callback_mutex_t`, see Section 4.5.2.15 on page 478.

## 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 that is 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.

## Execution Model Events

The *explicit-barrier-begin* event occurs in each thread that encounters the **barrier** construct on entry to the **barrier** region.

The *explicit-barrier-wait-begin* event occurs when a task begins an interval of active or passive waiting in a **barrier** region.

The *explicit-barrier-wait-end* event occurs when a task ends an interval of active or passive waiting and resumes execution in a **barrier** region.

The *explicit-barrier-end* event occurs in each thread that encounters the **barrier** construct after the barrier synchronization on exit from the **barrier** region.

A *cancellation* event occurs if cancellation is activated at an implicit cancellation point in a **barrier** region.

## Tool Callbacks

A thread dispatches a registered **ompt\_callback\_sync\_region** callback with **ompt\_sync\_region\_barrier\_explicit** — or **ompt\_sync\_region\_barrier**, if the implementation cannot make a distinction — as its *kind* argument and **ompt\_scope\_begin** as its *endpoint* argument for each occurrence of an *explicit-barrier-begin* event in the task that encounters the **barrier** construct. Similarly, a thread dispatches a registered **ompt\_callback\_sync\_region** callback with **ompt\_sync\_region\_barrier\_explicit** — or **ompt\_sync\_region\_barrier**, if the implementation cannot make a distinction — as its *kind* argument and **ompt\_scope\_end** as its *endpoint* argument for each occurrence of an *explicit-barrier-end* event in the task that encounters the **barrier** construct. These callbacks occur in the task that encounters the **barrier** construct and have the type signature **ompt\_callback\_sync\_region\_t**.

A thread dispatches a registered **ompt\_callback\_sync\_region\_wait** callback with **ompt\_sync\_region\_barrier\_explicit** — or **ompt\_sync\_region\_barrier**, if the implementation cannot make a distinction — as its *kind* argument and **ompt\_scope\_begin** as its *endpoint* argument for each occurrence of an *explicit-barrier-wait-begin* event. Similarly, a thread dispatches a registered **ompt\_callback\_sync\_region\_wait** callback with **ompt\_sync\_region\_barrier\_explicit** — or **ompt\_sync\_region\_barrier**, if the implementation cannot make a distinction — as its *kind* argument and **ompt\_scope\_end** as its *endpoint* argument for each occurrence of an *explicit-barrier-wait-end* event. These callbacks occur in the context of the task that encountered the **barrier** construct and have type signature **ompt\_callback\_sync\_region\_t**.

A thread dispatches a registered **ompt\_callback\_cancel** callback with **ompt\_cancel\_detected** as its *flags* argument 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**.

## Restrictions

The following restrictions apply to the **barrier** construct:

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

## Cross References

- `ompt_scope_begin` and `ompt_scope_end`, see Section 4.4.4.11 on page 447.
- `ompt_sync_region_barrier`, see Section 4.4.4.13 on page 448.
- `ompt_callback_sync_region_t`, see Section 4.5.2.13 on page 475.
- `ompt_callback_cancel_t`, see Section 4.5.2.18 on page 481.

## 2.20.3 Implicit Barriers

This section describes the OMPT events and tool callbacks associated with implicit barriers, which occur at the end of various regions as defined in the description of the constructs to which they correspond. Implicit barriers are task scheduling points. For a description of task scheduling points, associated events, and tool callbacks, see Section 2.13.6 on page 156.

### Execution Model Events

The *implicit-barrier-begin* event occurs in each implicit task at the beginning of an implicit barrier region.

The *implicit-barrier-wait-begin* event occurs when a task begins an interval of active or passive waiting in an implicit barrier region.

The *implicit-barrier-wait-end* event occurs when a task ends an interval of active or waiting and resumes execution of an implicit barrier region.

The *implicit-barrier-end* event occurs in each implicit task after the barrier synchronization on exit from an implicit barrier region.

A *cancellation* event occurs if cancellation is activated at an implicit cancellation point in an implicit barrier region.

## 1 Tool Callbacks

2 A thread dispatches a registered `ompt_callback_sync_region` callback with  
3 `ompt_sync_region_barrier_implicit` — or `ompt_sync_region_barrier`, if the  
4 implementation cannot make a distinction — as its *kind* argument and `ompt_scope_begin` as  
5 its *endpoint* argument for each occurrence of an *implicit-barrier-begin* event in that thread.

6 Similarly, a thread dispatches a registered `ompt_callback_sync_region` callback with  
7 `ompt_sync_region_barrier_implicit` — or `ompt_sync_region_barrier`, if the  
8 implementation cannot make a distinction — as its *kind* argument and `ompt_scope_end` as its  
9 *endpoint* argument for each occurrence of an *implicit-barrier-end* event in that thread. These  
10 callbacks occur in the implicit task that executes the parallel region and have the type signature  
11 `ompt_callback_sync_region_t`.

12 A thread dispatches a registered `ompt_callback_sync_region_wait` callback with  
13 `ompt_sync_region_barrier_implicit` — or `ompt_sync_region_barrier`, if the  
14 implementation cannot make a distinction — as its *kind* argument and `ompt_scope_begin` as  
15 its *endpoint* argument for each occurrence of a *implicit-barrier-wait-begin* event in that thread.

16 Similarly, a thread dispatches a registered `ompt_callback_sync_region_wait` callback  
17 with `ompt_sync_region_barrier_explicit` — or `ompt_sync_region_barrier`,  
18 if the implementation cannot make a distinction — as its *kind* argument and `ompt_scope_end`  
19 as its *endpoint* argument for each occurrence of an *implicit-barrier-wait-end* event in that thread.  
20 These callbacks occur in the implicit task that executes the parallel region and have type signature  
21 `ompt_callback_sync_region_t`.

22 A thread dispatches a registered `ompt_callback_cancel` callback with  
23 `ompt_cancel_detected` as its *flags* argument for each occurrence of a *cancellation* event in  
24 that thread. The callback occurs in the context of the encountering task. The callback has type  
25 signature `ompt_callback_cancel_t`.

## 26 Restrictions

27 If a thread is in the state `ompt_state_wait_barrier_implicit_parallel`, a call to  
28 `ompt_get_parallel_info` may return a pointer to a copy of the data object associated with  
29 the parallel region rather than a pointer to the associated data object itself. Writing to the data  
30 object returned by `omp_get_parallel_info` when a thread is in the  
31 `ompt_state_wait_barrier_implicit_parallel` results in unspecified behavior.

## 32 Cross References

- 33 • `ompt_scope_begin` and `ompt_scope_end`, see Section 4.4.4.11 on page 447.
- 34 • `ompt_sync_region_barrier`, see Section 4.4.4.13 on page 448
- 35 • `ompt_cancel_detected`, see Section 4.4.4.24 on page 453.
- 36 • `ompt_callback_sync_region_t`, see Section 4.5.2.13 on page 475.
- 37 • `ompt_callback_cancel_t`, see Section 4.5.2.18 on page 481.

## 2.20.4 Implementation-Specific Barriers

An OpenMP implementation can execute implementation-specific barriers that are not implied by the OpenMP specification; therefore, no *execution model events* are bound to these barriers. The implementation can handle these barriers like implicit barriers and dispatch all events as for implicit barriers. These callbacks are dispatched with `omp_sync_region_barrier_implementation` — or `omp_sync_region_barrier`, if the implementation cannot make a distinction — as the *kind* argument.

## 2.20.5 taskwait Construct

### Summary

The `taskwait` construct specifies a wait on the completion of child tasks of the current task. The `taskwait` construct is a stand-alone directive.

### Syntax

C / C++

The syntax of the `taskwait` construct is as follows:

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

where *clause* is one of the following:

```
depend ([depend-modifier, ]dependence-type : locator-list)
```

C / C++

Fortran

The syntax of the `taskwait` construct is as follows:

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

where *clause* is one of the following:

```
depend ([depend-modifier, ]dependence-type : locator-list)
```

Fortran

### Binding

The `taskwait` region binds to the current task region. The binding thread set of the `taskwait` region is the current team.

## Description

If no **depend** clause is present on the **taskwait** construct, the current task region is suspended at an implicit task scheduling point associated with the construct. The current task region remains suspended until all child tasks that it generated before the **taskwait** region complete execution.

Otherwise, if one or more **depend** clauses are present on the **taskwait** construct, the behavior is as if these clauses were applied to a **task** construct with an empty associated structured block that generates a *mergeable* and *included task*. Thus, the current task region is suspended until the *predecessor tasks* of this task complete execution.

## Execution Model Events

The *taskwait-begin* event occurs in each thread that encounters the **taskwait** construct on entry to the **taskwait** region.

The *taskwait-wait-begin* event occurs when a task begins an interval of active or passive waiting in a **taskwait** region.

The *taskwait-wait-end* event occurs when a task ends an interval of active or passive waiting and resumes execution in a **taskwait** region.

The *taskwait-end* event occurs in each thread that encounters the **taskwait** construct after the **taskwait** synchronization on exit from the **taskwait** region.

## Tool Callbacks

A thread dispatches a registered **ompt\_callback\_sync\_region** callback with **ompt\_sync\_region\_taskwait** as its *kind* argument and **ompt\_scope\_begin** as its *endpoint* argument for each occurrence of a *taskwait-begin* event in the task that encounters the **taskwait** construct. Similarly, a thread dispatches a registered **ompt\_callback\_sync\_region** callback with **ompt\_sync\_region\_taskwait** as its *kind* argument and **ompt\_scope\_end** as its *endpoint* argument for each occurrence of a *taskwait-end* event in the task that encounters the **taskwait** construct. These callbacks occur in the task that encounters the **taskwait** construct and have the type signature **ompt\_callback\_sync\_region\_t**.

A thread dispatches a registered **ompt\_callback\_sync\_region\_wait** callback with **ompt\_sync\_region\_taskwait** as its *kind* argument and **ompt\_scope\_begin** as its *endpoint* argument for each occurrence of a *taskwait-wait-begin* event. Similarly, a thread dispatches a registered **ompt\_callback\_sync\_region\_wait** callback with **ompt\_sync\_region\_taskwait** as its *kind* argument and **ompt\_scope\_end** as its *endpoint* argument for each occurrence of a *taskwait-wait-end* event. These callbacks occur in the context of the task that encounters the **taskwait** construct and have type signature **ompt\_callback\_sync\_region\_t**.



## Restrictions

The following restrictions apply to the `taskwait` construct:

- The `mutexinoutset` *dependence-type* may not appear in a `depend` clause on a `taskwait` construct.
- If the *dependence-type* of a `depend` clause is `depobj` then the dependence objects cannot represent dependences of the `mutexinoutset` dependence type.

## Cross References

- `task` construct, see Section 2.13.1 on page 142.
- Task scheduling, see Section 2.13.6 on page 156.
- `depend` clause, see Section 2.20.11 on page 257.
- `ompt_scope_begin` and `ompt_scope_end`, see Section 4.4.4.11 on page 447.
- `ompt_sync_region_taskwait`, see Section 4.4.4.13 on page 448.
- `ompt_callback_sync_region_t`, see Section 4.5.2.13 on page 475.

## 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 The binding task set of a **taskgroup** region is all tasks of the current team that are generated in  
10 the region. A **taskgroup** region binds to the 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 that encounters the **taskgroup** construct on  
20 entry 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 that encounters the **taskgroup** construct after the  
26 taskgroup synchronization on exit from the **taskgroup** region.

## Tool Callbacks

A thread dispatches a registered `ompt_callback_sync_region` callback with `ompt_sync_region_taskgroup` as its *kind* argument and `ompt_scope_begin` as its *endpoint* argument for each occurrence of a *taskgroup-begin* event in the task that encounters the `taskgroup` construct. Similarly, a thread dispatches a registered `ompt_callback_sync_region` callback with `ompt_sync_region_taskgroup` as its *kind* argument and `ompt_scope_end` as its *endpoint* argument for each occurrence of a *taskgroup-end* event in the task that encounters the `taskgroup` construct. These callbacks occur in the task that encounters the `taskgroup` construct and have the type signature `ompt_callback_sync_region_t`.

A thread dispatches a registered `ompt_callback_sync_region_wait` callback with `ompt_sync_region_taskgroup` as its *kind* argument and `ompt_scope_begin` as its *endpoint* argument for each occurrence of a *taskgroup-wait-begin* event. Similarly, a thread dispatches a registered `ompt_callback_sync_region_wait` callback with `ompt_sync_region_taskgroup` as its *kind* argument and `ompt_scope_end` as its *endpoint* argument for each occurrence of a *taskgroup-wait-end* event. These callbacks occur in the context of the task that encounters the `taskgroup` construct and have type signature `ompt_callback_sync_region_t`.

## Cross References

- Task scheduling, see Section [2.13.6](#) on page [156](#).
- `task_reduction` Clause, see Section [2.22.5.5](#) on page [305](#).
- `ompt_scope_begin` and `ompt_scope_end`, see Section [4.4.4.11](#) on page [447](#).
- `ompt_sync_region_taskgroup`, see Section [4.4.4.13](#) on page [448](#).
- `ompt_callback_sync_region_t`, see Section [4.5.2.13](#) on page [475](#).

## 2.20.7 atomic Construct

### Summary

The `atomic` construct ensures that a specific storage location is accessed atomically, rather than exposing it to the possibility of multiple, simultaneous reading and writing threads that may result in indeterminate values.

### Syntax

In the following syntax, *atomic-clause* is a clause that indicates the semantics for which atomicity is enforced, *memory-order-clause* is a clause that indicates the memory ordering behavior of the construct and *clause* is a clause other than *atomic-clause*. Specifically, *atomic-clause* is one of the 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.
- $x$  must not have an atomic 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.
- *hint-expression* is a constant integer expression that evaluates to a valid synchronization hint.

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
```

```

1      or
2      !$omp atomic [clause[[[, ] clause] ... ] [, ]] capture [[[, ] clause [[[, ] clause] ... ]]
3          capture-statement
4          update-statement
5      !$omp end atomic

```

```

6      or
7      !$omp atomic [clause[[[, ] clause] ... ] [, ]] capture [[[, ] clause [[[, ] clause] ... ]]
8          capture-statement
9          write-statement
10     !$omp end atomic

```

where *write-statement* has the following form (if *atomic-clause* is **capture** or **write**):

```

12     x = expr

```

where *capture-statement* has the following form (if *atomic-clause* is **capture** or **read**):

```

14     v = x

```

and where *update-statement* has one of the following forms (if *atomic-clause* is **update**, **capture**, or not present):

```

17     x = x operator expr
18
19     x = expr operator x
20
21     x = intrinsic_procedure_name (x, expr_list)
22
23     x = intrinsic_procedure_name (expr_list, x)

```

In the preceding statements:

- *x* and *v* (as applicable) are both scalar variables of intrinsic type.
- *x* must not have the **ALLOCATABLE** attribute.
- During the execution of an atomic region, multiple syntactic occurrences of *x* must designate the same storage location.
- None of *v*, *expr*, and *expr\_list* (as applicable) may access the same storage location as *x*.
- None of *x*, *expr*, and *expr\_list* (as applicable) may access the same storage location as *v*.
- *expr* is a scalar expression.

- 1 • *expr\_list* is a comma-separated, non-empty list of scalar expressions. If
- 2 *intrinsic\_procedure\_name* refers to **IAND**, **IOR**, or **IEOR**, exactly one expression must appear in
- 3 *expr\_list*.
- 4 • *intrinsic\_procedure\_name* is one of **MAX**, **MIN**, **IAND**, **IOR**, or **IEOR**.
- 5 • *operator* is one of **+**, **\***, **-**, **/**, **.AND.**, **.OR.**, **.EQV.**, or **.NEQV.**
- 6 • The expression *x operator expr* must be numerically equivalent to *x operator (expr)*. This
- 7 requirement is satisfied if the operators in *expr* have precedence greater than *operator*, or by
- 8 using parentheses around *expr* or subexpressions of *expr*.
- 9 • The expression *expr operator x* must be numerically equivalent to *(expr) operator x*. This
- 10 requirement is satisfied if the operators in *expr* have precedence equal to or greater than
- 11 *operator*, or by using parentheses around *expr* or subexpressions of *expr*.
- 12 • *intrinsic\_procedure\_name* must refer to the intrinsic procedure name and not to other program
- 13 entities.
- 14 • *operator* must refer to the intrinsic operator and not to a user-defined operator.
- 15 • All assignments must be intrinsic assignments.
- 16 • For forms that allow multiple occurrences of *x*, the number of times that *x* is evaluated is
- 17 unspecified.
- 18 • *hint-expression* is a constant expression that evaluates to a scalar value with kind
- 19 **omp\_sync\_hint\_kind** and a value that is a valid synchronization hint.

## 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 results in an atomic read of the location designated by *x* regardless of the native machine word size.

The **atomic** construct with the **write** clause results in an atomic write of the location designated by *x* regardless of the native machine word size.

The **atomic** construct with the **update** clause results in 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



1 atomic with respect to the read or write of the location designated by  $x$ . No task scheduling points  
2 are allowed between the read and the write of the location designated by  $x$ .

3 The **atomic** construct with the **capture** clause results in an atomic captured update — an  
4 atomic update of the location designated by  $x$  using the designated operator or intrinsic while also  
5 capturing the original or final value of the location designated by  $x$  with respect to the atomic  
6 update. The original or final value of the location designated by  $x$  is written in the location  
7 designated by  $v$  based on the base language semantics of structured block or statements of the  
8 **atomic** construct. Only the read and write of the location designated by  $x$  are performed mutually  
9 atomically. Neither the evaluation of  $expr$  or  $expr\_list$ , nor the write to the location designated by  $v$ ,  
10 need be atomic with respect to the read or write of the location designated by  $x$ . No task scheduling  
11 points are allowed between the read and the write of the location designated by  $x$ .

12 The **atomic** construct may be used to enforce memory consistency between threads, based on the  
13 guarantees provided by Section 1.4.6 on page 29. A strong flush on the location designated by  $x$  is  
14 performed on entry to and exit from the atomic operation, ensuring that the set of all atomic  
15 operations in the program applied to the same location has a total completion order. If the **write**,  
16 **update**, or **capture** clause is specified and the **release**, **acq\_rel**, or **seq\_cst** clause is  
17 specified then the strong flush on entry to the atomic operation is also a release flush. If the **read**  
18 or **capture** clause is specified and the **acquire**, **acq\_rel**, or **seq\_cst** clause is specified  
19 then the strong flush on exit from the atomic operation is also an acquire flush. Therefore, if  
20 *memory-order-clause* is specified and is not **relaxed**, release and/or acquire flush operations are  
21 implied and permit synchronization between the threads without the use of explicit **flush**  
22 directives.

23 For all forms of the **atomic** construct, any combination of two or more of these **atomic**  
24 constructs enforces mutually exclusive access to the locations designated by  $x$  among threads in the  
25 binding thread set. To avoid data races, all accesses of the locations designated by  $x$  that could  
26 potentially occur in parallel must be protected with an **atomic** construct.

27 **atomic** regions do not guarantee exclusive access with respect to any accesses outside of  
28 **atomic** regions to the same storage location  $x$  even if those accesses occur during a **critical**  
29 or **ordered** region, while an OpenMP lock is owned by the executing task, or during the  
30 execution of a **reduction** clause.

31 However, other OpenMP synchronization can ensure the desired exclusive access. For example, a  
32 barrier that follows a series of atomic updates to  $x$  guarantees that subsequent accesses do not form  
33 a race with the atomic accesses.

34 A compliant implementation may enforce exclusive access between **atomic** regions that update  
35 different storage locations. The circumstances under which this occurs are implementation defined.

36 If the storage location designated by  $x$  is not size-aligned (that is, if the byte alignment of  $x$  is not a  
37 multiple of the size of  $x$ ), then the behavior of the **atomic** region is implementation defined.

38 If present, the **hint** clause gives the implementation additional information about the expected  
39 properties of the atomic operation that can optionally be used to optimize the implementation. The

1 presence of a **hint** clause does not affect the semantics of the **atomic** construct, and all hints  
2 may be ignored. If no **hint** clause is specified, the effect is as if  
3 **hint (omp\_sync\_hint\_none)** had been specified.

## 4 Execution Model Events

5 The *atomic-acquiring* event occurs in the thread that encounters the **atomic** construct on entry to  
6 the atomic region before initiating synchronization for the region.

7 The *atomic-acquired* event occurs in the thread that encounters the **atomic** construct after it  
8 enters the region, but before it executes the structured block of the **atomic** region.

9 The *atomic-released* event occurs in the thread that encounters the **atomic** construct after it  
10 completes any synchronization on exit from the **atomic** region.

## 11 Tool Callbacks

12 A thread dispatches a registered **omp\_callback\_mutex\_acquire** callback for each  
13 occurrence of an *atomic-acquiring* event in that thread. This callback has the type signature  
14 **omp\_callback\_mutex\_acquire\_t**.

15 A thread dispatches a registered **omp\_callback\_mutex\_acquired** callback for each  
16 occurrence of an *atomic-acquired* event in that thread. This callback has the type signature  
17 **omp\_callback\_mutex\_t**.

18 A thread dispatches a registered **omp\_callback\_mutex\_released** callback with  
19 **omp\_mutex\_atomic** as the *kind* argument if practical, although a less specific *kind* may be  
20 used, for each occurrence of an *atomic-released* event in that thread. This callback has the type  
21 signature **omp\_callback\_mutex\_t** and occurs in the task that encounters the atomic  
22 construct.

## 23 Restrictions

24 The following restrictions apply to the **atomic** construct:

- 25 • OpenMP constructs may not be encountered during execution of an **atomic** region.
- 26 • At most one *memory-order-clause* may appear on the construct.
- 27 • At most one **hint** clause may appear on the construct.
- 28 • If *atomic-clause* is **read** then *memory-order-clause* must not be **acq\_rel** or **release**.
- 29 • If *atomic-clause* is **write** then *memory-order-clause* must not be **acq\_rel** or **acquire**.
- 30 • If *atomic-clause* is **update** or not present then *memory-order-clause* must not be **acq\_rel** or  
31 **acquire**.

32  C / C++

- 33 • All atomic accesses to the storage locations designated by *x* throughout the program are required  
to have a compatible type.

 C / C++

- 1     • All atomic accesses to the storage locations designated by *x* throughout the program are required  
2     to have the same type and type parameters.

### Cross References

- 3     • **critical** construct, see Section 2.20.1 on page 226.  
4  
5     • **barrier** construct, see Section 2.20.2 on page 229.  
6  
6     • **flush** construct, see Section 2.20.8 on page 245.  
7  
7     • **ordered** construct, see Section 2.20.9 on page 252.  
8  
8     • Synchronization Hints, see Section 2.20.12 on page 262.  
9  
9     • **reduction** clause, see Section 2.22.5.4 on page 302.  
10  
10    • lock routines, see Section 3.3 on page 382.  
11  
11    • **ompt\_mutex\_atomic**, see Section 4.4.4.16 on page 449.  
12  
12    • **ompt\_callback\_mutex\_acquire\_t**, see Section 4.5.2.14 on page 477.  
13  
13    • **ompt\_callback\_mutex\_t**, see Section 4.5.2.15 on page 478.

## 2.20.8 flush Construct

### Summary

15     The **flush** construct executes the OpenMP flush operation. This operation makes a thread's  
16     temporary view of memory consistent with memory and enforces an order on the memory  
17     operations of the variables explicitly specified or implied. See the memory model description in  
18     Section 1.4 on page 24 for more details. The **flush** construct is a stand-alone directive.  
19

### Syntax

20  
21     The syntax of the **flush** construct is as follows:

22     **#pragma omp flush** [*memory-order-clause*] [*(list)*] *new-line*

23     where *memory-order-clause* is one of the following:

24     **acq\_rel**  
25     **release**  
26     **acquire**

## Fortran

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

```
2 !$omp flush [memory-order-clause] [(list)]
```

3 where *memory-order-clause* is one of the following:

```
4 acq_rel  
5 release  
6 acquire
```

## Fortran

### 7 Binding

8 The binding thread set for a **flush** region is the encountering thread. Execution of a **flush**  
9 region affects the memory and the temporary view of memory of only the thread that executes the  
10 region. It does not affect the temporary view of other threads. Other threads must themselves  
11 execute a flush operation in order to be guaranteed to observe the effects of the flush operation of  
12 the encountering thread.

### 13 Description

14 If *memory-order-clause* is not specified then the **flush** construct results in a strong flush operation  
15 with the following behavior. A **flush** construct without a list, executed on a given thread, operates  
16 as if the whole thread-visible data state of the program, as defined by the base language, is flushed.  
17 A **flush** construct with a list applies the flush operation to the items in the list, and the flush  
18 operation does not complete until the operation is complete for all specified list items. An  
19 implementation may implement a **flush** with a list by ignoring the list, and treating it the same as  
20 a **flush** without a list.

21 If no list items are specified, the flush operation has the release and/or acquire flush properties:

- 22 • If *memory-order-clause* is not specified or is **acq\_rel**, the flush operation is both a release  
23 flush and an acquire flush.
- 24 • If *memory-order-clause* is **release**, the flush operation is a release flush.
- 25 • If *memory-order-clause* is **acquire**, the flush operation is an acquire flush.

## C / C++

26 If a pointer is present in the list, the pointer itself is flushed, not the memory block to which the  
27 pointer refers.

## C / C++

## Fortran

1 If the list item or a subobject of the list item has the **POINTER** attribute, the allocation or  
2 association status of the **POINTER** item is flushed, but the pointer target is not. If the list item is a  
3 Cray pointer, the pointer is flushed, but the object to which it points is not. If the list item is of type  
4 **C\_PTR**, the variable is flushed, but the storage that corresponds to that address is not flushed. If the  
5 list item or the subobject of the list item has the **ALLOCATABLE** attribute and has an allocation  
6 status of allocated, the allocated variable is flushed; otherwise the allocation status is flushed.

## Fortran

7  
8 **Note** – Use of a **flush** construct with a list is extremely error prone and users are strongly  
9 discouraged from attempting it. The following examples illustrate the ordering properties of the  
10 flush operation. In the following incorrect pseudocode example, the programmer intends to prevent  
11 simultaneous execution of the protected section by the two threads, but the program does not work  
12 properly because it does not enforce the proper ordering of the operations on variables **a** and **b**.  
13 Any shared data accessed in the protected section is not guaranteed to be current or consistent  
14 during or after the protected section. The atomic notation in the pseudocode in the following two  
15 examples indicates that the accesses to **a** and **b** are atomic write and atomic read operations.  
16 Otherwise both examples would contain data races and automatically result in unspecified behavior.  
17 The *flush* operations are strong flushes that are applied to the specified flush lists

*Incorrect example:*

**a = b = 0**

*thread 1*

*thread 2*

**atomic(b = 1)**

**atomic(a = 1)**

*flush(b)*

*flush(a)*

*flush(a)*

*flush(b)*

**atomic(tmp = a)**

**atomic(tmp = b)**

**if (tmp == 0) then**

**if (tmp == 0) then**

*protected section*

*protected section*

**end if**

**end if**

18  
19 The problem with this example is that operations on variables **a** and **b** are not ordered with respect  
20 to each other. For instance, nothing prevents the compiler from moving the flush of **b** on thread 1 or  
21 the flush of **a** on thread 2 to a position completely after the protected section (assuming that the  
22 protected section on thread 1 does not reference **b** and the protected section on thread 2 does not  
23 reference **a**). If either re-ordering happens, both threads can simultaneously execute the protected  
24 section.

1 The following pseudocode example correctly ensures that the protected section is executed by not  
2 more than one of the two threads at any one time. Execution of the protected section by neither  
3 thread is considered correct in this example. This occurs if both flushes complete prior to either  
4 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**

6 The compiler is prohibited from moving the flush at all for either thread, ensuring that the  
7 respective assignment is complete and the data is flushed before the **if** statement is executed.



## 10 Execution Model Events

11 The *flush* event occurs in a thread that encounters the **flush** construct.

## 12 Tool Callbacks

13 A thread dispatches a registered **ompt\_callback\_flush** callback for each occurrence of a  
14 *flush* event in that thread. This callback has the type signature **ompt\_callback\_flush\_t**.

## 15 Restrictions

16 The following restrictions apply to the **flush** construct:

- 17 • If *memory-order-clause* is **release**, **acquire**, or **acq\_rel**, list items must not be specified  
18 on the **flush** directive.

## 19 Cross References

- 20 • **ompt\_callback\_flush\_t**, see Section [4.5.2.17](#) on page 481.

## 2.20.8.1 Implicit Flushes

Flush operations implied when executing an **atomic** region are described in Section [2.20.7](#).

A **flush** region that corresponds to a **flush** directive with the **release** clause present is implied at the following locations:

- During a barrier region;
- At entry to a **parallel** region;
- At entry to a **teams** region;
- At exit from a **critical** region;
- During an **omp\_unset\_lock** region;
- During an **omp\_unset\_nest\_lock** region;
- Immediately before every task scheduling point;
- At exit from the task region of each implicit task;
- At exit from an **ordered** region, if a **threads** clause or a **depend** clause with a **source** dependence type is present, or if no clauses are present; and
- During a **cancel** region, if the *cancel-var* ICV is *true*.

A **flush** region that corresponds to a **flush** directive with the **acquire** clause present is implied at the following locations:

- During a barrier region;
- At exit from a **teams** region;
- At entry to a **critical** region;
- If the region causes the lock to be set, during:
  - an **omp\_set\_lock** region;
  - an **omp\_test\_lock** region;
  - an **omp\_set\_nest\_lock** region; and
  - an **omp\_test\_nest\_lock** region;
- Immediately after every task scheduling point;
- At entry to the task region of each implicit task;
- At entry to an **ordered** region, if a **threads** clause or a **depend** clause with a **sink** dependence type is present, or if no clauses are present; and
- Immediately before a cancellation point, if the *cancel-var* ICV is *true* and cancellation has been activated.

1  
2 Note – A **flush** region is not implied at the following locations:

- 3 • At entry to worksharing regions; and
- 4 • At entry to or exit from **master** regions.

5  
6 The synchronization behavior of implicit flushes is as follows:

- 7 • When a thread executes an **atomic** region for which the corresponding construct has the  
8 **release**, **acq\_rel**, or **seq\_cst** clause and specifies an atomic operation that starts a given  
9 release sequence, the release flush that is performed on entry to the atomic operation  
10 synchronizes with an acquire flush that is performed by a different thread and has an associated  
11 atomic operation that reads a value written by a modification in the release sequence.
- 12 • When a thread executes an **atomic** region for which the corresponding construct has the  
13 **acquire**, **acq\_rel**, or **seq\_cst** clause and specifies an atomic operation that reads a value  
14 written by a given modification, a release flush that is performed by a different thread and has an  
15 associated release sequence that contains that modification synchronizes with the acquire flush  
16 that is performed on exit from the atomic operation.
- 17 • When a thread executes a **critical** region that has a given name, the behavior is as if the  
18 release flush performed on exit from the region synchronizes with the acquire flush performed on  
19 entry to the next **critical** region with the same name that is performed by a different thread,  
20 if it exists.
- 21 • When a thread team executes a **barrier** region, the behavior is as if the release flush  
22 performed by each thread within the region synchronizes with the acquire flush performed by all  
23 other threads within the region.
- 24 • When a thread executes a **taskwait** region that does not result in the creation of a dependent  
25 task and the task that encounters the corresponding **taskwait** construct has at least one child  
26 task, the behavior is as if each thread that executes a child task that is generated before the  
27 **taskwait** region performs a release flush upon completion of the child task that synchronizes  
28 with an acquire flush performed in the **taskwait** region.
- 29 • When a thread executes a **taskgroup** region, the behavior is as if each thread that executes a  
30 remaining descendant task performs a release flush upon completion of the descendant task that  
31 synchronizes with an acquire flush performed on exit from the **taskgroup** region.
- 32 • When a thread executes an **ordered** region that does not arise from a stand-alone **ordered**  
33 directive, the behavior is as if the release flush performed on exit from the region synchronizes  
34 with the acquire flush performed on entry to an **ordered** region encountered in the next logical  
35 iteration to be executed by a different thread, if it exists.
- 36 • When a thread executes an **ordered** region that arises from a stand-alone **ordered** directive,  
37 the behavior is as if the release flush performed in the **ordered** region from a given source



- 1 iteration synchronizes with the acquire flush performed in all **ordered** regions executed by a  
2 different thread that are waiting for dependences on that iteration to be satisfied.
- 3 ● When a thread team begins execution of a **parallel** region, the behavior is as if the release  
4 flush performed by the master thread on entry to the **parallel** region synchronizes with the  
5 acquire flush performed on entry to each implicit task that is assigned to a different thread.
  - 6 ● When an initial thread begins execution of a **target** region that is generated by a different  
7 thread from a target task, the behavior is as if the release flush performed by the generating  
8 thread in the target task synchronizes with the acquire flush performed by the initial thread on  
9 entry to its initial task region.
  - 10 ● When an initial thread completes execution of a **target** region that is generated by a different  
11 thread from a target task, the behavior is as if the release flush performed by the initial thread on  
12 exit from its initial task region synchronizes with the acquire flush performed by the generating  
13 thread in the target task.
  - 14 ● When a thread encounters a **teams** construct, the behavior is as if the release flush performed by  
15 the thread on entry to the **teams** region synchronizes with the acquire flush performed on entry  
16 to each initial task that is executed by a different initial thread that participates in the execution of  
17 the **teams** region.
  - 18 ● When a thread that encounters a **teams** construct reaches the end of the **teams** region, the  
19 behavior is as if the release flush performed by each different participating initial thread at exit  
20 from its initial task synchronizes with the acquire flush performed by the thread at exit from the  
21 **teams** region.
  - 22 ● When a task generates an explicit task that begins execution on a different thread, the behavior is  
23 as if the thread that is executing the generating task performs a release flush that synchronizes  
24 with the acquire flush performed by the thread that begins to execute the explicit task.
  - 25 ● When an undeferred task completes execution on a given thread that is different from the thread  
26 on which its generating task is suspended, the behavior is as if a release flush performed by the  
27 thread that completes execution of the undeferred task synchronizes with an acquire flush  
28 performed by the thread that resumes execution of the generating task.
  - 29 ● When a dependent task with one or more predecessor tasks begins execution on a given thread,  
30 the behavior is as if each release flush performed by a different thread on completion of a  
31 predecessor task synchronizes with the acquire flush performed by the thread that begins to  
32 execute the dependent task.
  - 33 ● When a task begins execution on a given thread and it is mutually exclusive with respect to  
34 another sibling task that is executed by a different thread, the behavior is as if each release flush  
35 performed on completion of the sibling task synchronizes with the acquire flush performed by  
36 the thread that begins to execute the task.
  - 37 ● When a thread executes a **cancel** region, the *cancel-var* ICV is *true*, and cancellation is not  
38 already activated for the specified region, the behavior is as if the release flush performed during

1 the **cancel** region synchronizes with the acquire flush performed by a different thread  
2 immediately before a cancellation point in which that thread observes cancellation was activated  
3 for the region.

- 4 • When a thread executes an **omp\_unset\_lock** region that causes the specified lock to be unset,  
5 the behavior is as if a release flush is performed during the **omp\_unset\_lock** region that  
6 synchronizes with an acquire flush that is performed during the next **omp\_set\_lock** or  
7 **omp\_test\_lock** region to be executed by a different thread that causes the specified lock to be  
8 set.
- 9 • When a thread executes an **omp\_unset\_nest\_lock** region that causes the specified nested  
10 lock to be unset, the behavior is as if a release flush is performed during the  
11 **omp\_unset\_nest\_lock** region that synchronizes with an acquire flush that is performed  
12 during the next **omp\_set\_nest\_lock** or **omp\_test\_nest\_lock** region to be executed by  
13 a different thread that causes the specified nested lock to be set.

## 14 2.20.9 ordered Construct

### 15 Summary

16 The **ordered** construct either specifies a structured block in a worksharing-loop, **simd**, or  
17 worksharing-loop SIMD region that will be executed in the order of the loop iterations, or it is a  
18 stand-alone directive that specifies cross-iteration dependences in a doacross loop nest. The  
19 **ordered** construct sequentializes and orders the execution of **ordered** regions while allowing  
20 code outside the region to run in parallel.

### 21 Syntax

C / C++

22 The syntax of the **ordered** construct is as follows:

```
23 #pragma omp ordered [clause[ [, ] clause] ] new-line  
24     structured-block
```

25 where *clause* is one of the following:

```
26     threads  
27     simd
```

28 or

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

30 where *clause* is one of the following:

```
31     depend(source)  
32     depend(sink : vec)
```

C / C++

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

```

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

where *clause* is one of the following:

```

threads
simd
    
```

or

```

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

where *clause* is one of the following:

```

depend (source)
depend (sink : vec)
    
```

If the **depend** clause is specified, the **ordered** construct is a stand-alone directive.

### Binding

The binding thread set for an **ordered** region is the current team. An **ordered** region binds to the innermost enclosing **simd** or worksharing-loop SIMD region if the **simd** clause is present, and otherwise it binds to the innermost enclosing worksharing-loop region. **ordered** regions that bind to different regions execute independently of each other.

### Description

If no clause is specified, the **ordered** construct behaves as if the **threads** clause had been specified. If the **threads** clause is specified, the threads in the team that is executing the worksharing-loop region execute **ordered** regions sequentially in the order of the loop iterations. If any **depend** clauses are specified then those clauses specify the order in which the threads in the team execute **ordered** regions. If the **simd** clause is specified, the **ordered** regions encountered by any thread will execute one at a time in the order of the loop iterations.

When the thread that is executing the first iteration of the loop encounters an **ordered** construct, it can enter the **ordered** region without waiting. When a thread that is executing any subsequent iteration encounters an **ordered** construct without a **depend** clause, it waits at the beginning of the **ordered** region until execution of all **ordered** regions belonging to all previous iterations has completed. When a thread that is executing any subsequent iteration encounters an **ordered** construct with one or more **depend (sink : vec)** clauses, it waits until its dependences on all valid iterations specified by the **depend** clauses are satisfied before it completes execution of the **ordered** region. A specific dependence is satisfied when a thread that is executing the corresponding iteration encounters an **ordered** construct with a **depend (source)** clause.

## 1 Execution Model Events

2 The *ordered-acquiring* event occurs in the task that encounters the **ordered** construct on entry to  
3 the ordered region before it initiates synchronization for the region.

4 The *ordered-acquired* event occurs in the task that encounters the **ordered** construct after it  
5 enters the region, but before it executes the structured block of the **ordered** region.

6 The *ordered-released* event occurs in the task that encounters the **ordered** construct after it  
7 completes any synchronization on exit from the **ordered** region.

8 The *doacross-sink* event occurs in the task that encounters a **ordered** construct for each  
9 **depend (sink:vec)** clause after the dependence is fulfilled.

10 The *doacross-source* event occurs in the task that encounters a **ordered** construct with a  
11 **depend (source:vec)** clause before signaling the dependence to be fulfilled.

## 12 Tool Callbacks

13 A thread dispatches a registered **ompt\_callback\_mutex\_acquire** callback for each  
14 occurrence of an *ordered-acquiring* event in that thread. This callback has the type signature  
15 **ompt\_callback\_mutex\_acquire\_t**.

16 A thread dispatches a registered **ompt\_callback\_mutex\_acquired** callback for each  
17 occurrence of an *ordered-acquired* event in that thread. This callback has the type signature  
18 **ompt\_callback\_mutex\_t**.

19 A thread dispatches a registered **ompt\_callback\_mutex\_released** callback with  
20 **ompt\_mutex\_ordered** as the *kind* argument if practical, although a less specific kind may be  
21 used, for each occurrence of an *ordered-released* event in that thread. This callback has the type  
22 signature **ompt\_callback\_mutex\_t** and occurs in the task that encounters the atomic  
23 construct.

24 A thread dispatches a registered **ompt\_callback\_dependences** callback with all vector  
25 entries listed as **ompt\_dependence\_type\_sink** in the *deps* argument for each occurrence of a  
26 *doacross-sink* event in that thread. A thread dispatches a registered  
27 **ompt\_callback\_dependences** callback with all vector entries listed as  
28 **ompt\_dependence\_type\_source** in the *deps* argument for each occurrence of a  
29 *doacross-source* event in that thread. These callbacks have the type signature  
30 **ompt\_callback\_dependences\_t**.

## 31 Restrictions

32 Restrictions to the **ordered** construct are as follows:

- 33 • At most one **threads** clause can appear on an **ordered** construct.
- 34 • At most one **simd** clause can appear on an **ordered** construct.
- 35 • At most one **depend (source)** clause can appear on an **ordered** construct.

- 1       • The construct corresponding to the binding region of an **ordered** region must not specify a  
2       **reduction** clause with the **inscan** modifier.
- 3       • Either **depend(sink:vec)** clauses or **depend(source)** clauses may appear on an  
4       **ordered** construct, but not both.
- 5       • The worksharing-loop or worksharing-loop SIMD region to which an **ordered** region  
6       corresponding to an **ordered** construct without a **depend** clause binds must have an  
7       **ordered** clause without the parameter specified on the corresponding worksharing-loop or  
8       worksharing-loop SIMD directive.
- 9       • The worksharing-loop region to which an **ordered** region corresponding to an **ordered**  
10       construct with any **depend** clauses binds must have an **ordered** clause with the parameter  
11       specified on the corresponding worksharing-loop directive.
- 12       • An **ordered** construct with the **depend** clause specified must be closely nested inside a  
13       worksharing-loop (or parallel worksharing-loop) construct.
- 14       • An **ordered** region corresponding to an **ordered** construct without the **simd** clause  
15       specified must be closely nested inside a loop region.
- 16       • An **ordered** region corresponding to an **ordered** construct with the **simd** clause specified  
17       must be closely nested inside a **simd** or worksharing-loop SIMD region.
- 18       • An **ordered** region corresponding to an **ordered** construct with both the **simd** and  
19       **threads** clauses must be closely nested inside a worksharing-loop SIMD region or must be  
20       closely nested inside a worksharing-loop and **simd** region.
- 21       • During execution of an iteration of a worksharing-loop or a loop nest within a worksharing-loop,  
22       **simd**, or worksharing-loop SIMD region, a thread must not execute more than one **ordered**  
23       region corresponding to an **ordered** construct without a **depend** clause.

C++

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

C++

## Cross References

- 26       • worksharing-loop construct, see Section [2.12.2](#) on page [106](#).
- 27       • **simd** construct, see Section [2.12.3.1](#) on page [115](#).
- 28       • parallel Worksharing-loop construct, see Section [2.16.1](#) on page [193](#).
- 29       • **depend** Clause, see Section [2.20.11](#) on page [257](#)
- 30       • **ompt\_mutex\_ordered**, see Section [4.4.4.16](#) on page [449](#).
- 31       • **ompt\_callback\_mutex\_acquire\_t**, see Section [4.5.2.14](#) on page [477](#).
- 32       • **ompt\_callback\_mutex\_t**, see Section [4.5.2.15](#) on page [478](#).
- 33

## 2.20.10 Depend Objects

This section describes constructs that support OpenMP depend objects that can be used to supply user-computed dependences to **depend** clauses. OpenMP depend objects must be accessed only through the **depobj** construct or through the **depend** clause; programs that otherwise access OpenMP depend objects are non-conforming.

An OpenMP depend object can be in one of the following states: *uninitialized* or *initialized*. Initially OpenMP depend objects are in the *uninitialized* state.

### 2.20.10.1 depobj Construct

#### Summary

The **depobj** construct initializes, updates or destroys an OpenMP depend object. The **depobj** construct is a stand-alone directive.

#### Syntax

C / C++

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

```
#pragma omp depobj(depobj) clause new-line
```

where *depobj* is an lvalue expression of type **omp\_depend\_t**.

where *clause* is one of the following:

```
depend(dependence-type : locator)
destroy
update(dependence-type)
```

C / C++

Fortran

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

```
!$omp depobj(depobj) clause
```

where *depobj* is a scalar integer variable of the **omp\_depend\_kind** kind.

where *clause* is one of the following:

```
depend(dependence-type : locator)
destroy
update(dependence-type)
```

Fortran

## Binding

The binding thread set for **depobj** regions is the encountering thread.

## Description

A **depobj** construct with a **depend** clause present sets the state of *depobj* to initialized. The *depobj* is initialized to represent the dependence that the **depend** clause specifies.

A **depobj** construct with a **destroy** clause present changes the state of the *depobj* to uninitialized.

A **depobj** construct with an **update** clause present changes the dependence type of the dependence represented by *depobj* to the one specified by the *update* clause.

## Restrictions

- A **depend** clause on a **depobj** construct must not have **source** or **sink** as *dependence-type*.
- An **update** clause on a **depobj** construct must not have **source**, **sink** or **depobj** as *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 [257](#).

## 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:

```
1  in
2  out
3  inout
4  mutexinoutset
5  inoutset
6  depobj
```

7 where *depend-modifier* is one of the following:

```
8  iterator (iterators-definition)
```

9 or

```
10 depend (dependence-type)
```

11 where *dependence-type* is:

```
12  source
```

13 or

```
14 depend (dependence-type : vec)
```

15 where *dependence-type* is:

```
16  sink
```

17 and where *vec* is the iteration vector, which has the form:

18  $x_1 [\pm d_1], x_2 [\pm d_2], \dots, x_n [\pm d_n]$

19 where  $n$  is the value specified by the **ordered** clause in the worksharing-loop directive,  $x_i$  denotes  
20 the loop iteration variable of the  $i$ -th nested loop associated with the worksharing-loop directive,  
21 and  $d_i$  is a constant non-negative integer.

## 22 Description

23 Task dependencies are derived from the *dependence-type* of a **depend** clause and its list items  
24 when *dependence-type* is **in**, **out**, **inout**, **mutexinoutset** or **inoutset**. When the  
25 *dependence-type* is **depobj**, the task dependencies are derived from the dependences represented  
26 by the depend objects specified in the **depend** clause as if the **depend** clauses of the **depobj**  
27 constructs were specified in the current construct.

28 The storage location of a list item matches the storage location of another list item if they have the  
29 same storage location, or if any of the list items is **omp\_all\_memory**.

30 For the **in** *dependence-type*, if the storage location of at least one of the list items matches the  
31 storage location of a list item appearing in a **depend** clause with an **out**, **inout**,



1 **mutexinoutset**, or **inoutset** *dependence-type* on a construct from which a sibling task was  
2 previously generated, then the generated task will be a dependent task of that sibling task.

3 For the **out** and **inout** *dependence-types*, if the storage location of at least one of the list items  
4 matches the storage location of a list item appearing in a **depend** clause with an **in**, **out**, **inout**,  
5 **mutexinoutset**, or **inoutset** *dependence-type* on a construct from which a sibling task was  
6 previously generated, then the generated task will be a dependent task of that sibling task.

7 For the **mutexinoutset** *dependence-type*, if the storage location of at least one of the list items  
8 matches the storage location of a list item appearing in a **depend** clause with an **in**, **out**, **inout**,  
9 or **inoutset** *dependence-type* on a construct from which a sibling task was previously generated,  
10 then the generated task will be a dependent task of that sibling task.

11 If a list item appearing in a **depend** clause with a **mutexinoutset** *dependence-type* on a  
12 task-generating construct matches a list item appearing in a **depend** clause with a  
13 **mutexinoutset** *dependence-type* on a different task generating construct, and both constructs  
14 generate sibling tasks, the sibling tasks will be mutually exclusive tasks.

15 For the **inoutset** *dependence-type*, if the storage location of at least one of the list items matches  
16 the storage location of a list item appearing in a **depend** clause with an **in**, **out**, **inout**, or  
17 **mutexinoutset** *dependence-type* on a construct from which a sibling task was previously  
18 generated, then the generated task will be a dependent task of that sibling task.

19 The list items that appear in the **depend** clause may reference iterators defined by an  
20 *iterators-definition* appearing on an **iterator** modifier.

21 The list items that appear in the **depend** clause may include array sections or the  
22 **omp\_all\_memory** reserved locator.

▼ Fortran ▼

23 If a list item has the **ALLOCATABLE** attribute and its allocation status is unallocated, the behavior  
24 is unspecified. If a list item has the **POINTER** attribute and its association status is disassociated or  
25 undefined, the behavior is unspecified.

▲ Fortran ▲

▼ C / C++ ▼

26 The list items that appear in a **depend** clause may use shape-operators.

▲ C / C++ ▲

27 ▼

1 Note – The enforced task dependence establishes a synchronization of memory accesses  
2 performed by a dependent task with respect to accesses performed by the predecessor tasks.  
3 However, it is the responsibility of the programmer to synchronize properly with respect to other  
4 concurrent accesses that occur outside of those tasks.

5 

---

6 The **source** *dependence-type* specifies the satisfaction of cross-iteration dependences that arise  
7 from the current iteration.

8 The **sink** *dependence-type* specifies a cross-iteration dependence, where the iteration vector *vec*  
9 indicates the iteration that satisfies the dependence.

10 If the iteration vector *vec* does not occur in the iteration space, the **depend** clause is ignored. If all  
11 **depend** clauses on an **ordered** construct are ignored then the construct is ignored.

12 

---

13 Note – An iteration vector *vec* that does not indicate a lexicographically earlier iteration may cause  
14 a deadlock.

## 15 Execution Model Events

16 The *task-dependences* event occurs in a thread that encounters a task generating construct or a  
17 **taskwait** construct with a **depend** clause immediately after the *task-create* event for the new  
18 task or the *taskwait-begin* event.

19 The *task-dependence* event indicates an unfulfilled dependence for the generated task. This event  
20 occurs in a thread that observes the unfulfilled dependence before it is satisfied.

## 21 Tool Callbacks

22 A thread dispatches the **ompt\_callback\_dependences** callback for each occurrence of the  
23 *task-dependences* event to announce its dependences with respect to the list items in the **depend**  
24 clause. This callback has type signature **ompt\_callback\_dependences\_t**.

25 A thread dispatches the **ompt\_callback\_task\_dependence** callback for a *task-dependence*  
26 event to report a dependence between a predecessor task (*src\_task\_data*) and a dependent task  
27 (*sink\_task\_data*). This callback has type signature **ompt\_callback\_task\_dependence\_t**.

## 28 Restrictions

29 Restrictions to the **depend** clause are as follows:

- 30 • List items, other than reserved locators, used in **depend** clauses of the same task or sibling tasks  
31 must indicate identical storage locations or disjoint storage locations.
  - 32 • List items used in **depend** clauses cannot be zero-length array sections.
  - 33 • The **omp\_all\_memory** reserved locator can only be used in a **depend** clause with an **out** or  
34 **inout** *dependence-type*.
- 35

- 1           • Array sections cannot be specified in **depend** clauses with the **depobj** dependence type.
- 2           • List items used in **depend** clauses with the **depobj** dependence type must be depend objects
- 3           in the initialized state.
- ▼ C / C++ ▼
- 4           • List items used in **depend** clauses with the **depobj** dependence type must be expressions of
- 5           the **omp\_depend\_t** type.
- 6           • List items that are expressions of the **omp\_depend\_t** type can only be used in **depend**
- 7           clauses with the **depobj** dependence type.
- ▲ C / C++ ▲
- ▼ Fortran ▼
- 8           • A common block name cannot appear in a **depend** clause.
- 9           • List items used in **depend** clauses with the **depobj** dependence type must be integer
- 10          expressions of the **omp\_depend\_kind** *kind*.
- ▲ Fortran ▲
- 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
- 12          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
- 13          loop iteration variable  $x_i$  that can encounter the **ordered** construct must be computable without
- 14          overflow in the type of the loop iteration variable.
- ▼ C++ ▼
- 15          • For a *vec* element of **sink dependence-type** of the form  $x_i + d_i$  or  $x_i - d_i$  if the loop iteration
- 16          variable  $x_i$  is of a random access iterator type other than pointer type, the expression
- 17           $(x_i - lb_i) + d_i$  or  $(x_i - lb_i) - d_i$  for any value of the loop iteration variable  $x_i$  that can
- 18          encounter the **ordered** construct must be computable without overflow in the type that would
- 19          be used by *std::distance* applied to variables of the type of  $x_i$ .
- ▲ C++ ▲
- ▼ C / C++ ▼
- 20          • A bit-field cannot appear in a **depend** clause.
- ▲ C / C++ ▲

## Cross References

- Array sections, see Section 2.1.5 on page 42.
- Iterators, see Section 2.1.6 on page 46.
- **task** construct, see Section 2.13.1 on page 142.
- Task scheduling constraints, see Section 2.13.6 on page 156.
- **target enter data** construct, see Section 2.15.3 on page 169.
- **target exit data** construct, see Section 2.15.4 on page 172.
- **target** construct, see Section 2.15.5 on page 174.
- **target update** construct, see Section 2.15.6 on page 181.
- **ordered** construct, see Section 2.20.9 on page 252.
- **depobj** construct, see Section 2.20.10.1 on page 256.
- **ompt\_callback\_dependences\_t**, see Section 4.5.2.8 on page 470.
- **ompt\_callback\_task\_dependence\_t**, see Section 4.5.2.9 on page 471.

## 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 does not change the semantics of the associated construct; if ignoring the hint changes the program semantics, the result is unspecified.

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++
typedef enum omp_sync_hint_t {
    omp_sync_hint_none = 0x0,
    omp_lock_hint_none = omp_sync_hint_none,
    omp_sync_hint_uncontended = 0x1,
    omp_lock_hint_uncontended = omp_sync_hint_uncontended,
    omp_sync_hint_contended = 0x2,
    omp_lock_hint_contended = omp_sync_hint_contended,
    omp_sync_hint_nonspeculative = 0x4,
    omp_lock_hint_nonspeculative = omp_sync_hint_nonspeculative,
    omp_sync_hint_speculative = 0x8
    omp_lock_hint_speculative = omp_sync_hint_speculative
}
```

```

1  } omp_sync_hint_t;
2
3  typedef omp_sync_hint_t omp_lock_hint_t;

```

```

4  integer, parameter :: omp_lock_hint_kind = omp_sync_hint_kind
5
6  integer (kind=omp_sync_hint_kind), &
7      parameter :: omp_sync_hint_none = &
8          int(Z'0', kind=omp_sync_hint_kind)
9  integer (kind=omp_lock_hint_kind), &
10     parameter :: omp_lock_hint_none = omp_sync_hint_none
11 integer (kind=omp_sync_hint_kind), &
12     parameter :: omp_sync_hint_uncontended = &
13         int(Z'1', kind=omp_sync_hint_kind)
14 integer (kind=omp_lock_hint_kind), &
15     parameter :: omp_lock_hint_uncontended = &
16         omp_sync_hint_uncontended
17 integer (kind=omp_sync_hint_kind), &
18     parameter :: omp_sync_hint_contended = &
19         int(Z'2', kind=omp_sync_hint_kind)
20 integer (kind=omp_lock_hint_kind), &
21     parameter :: omp_lock_hint_contended = &
22         omp_sync_hint_contended
23 integer (kind=omp_sync_hint_kind), &
24     parameter :: omp_sync_hint_nonspeculative = &
25         int(Z'4', kind=omp_sync_hint_kind)
26 integer (kind=omp_lock_hint_kind), &
27     parameter :: omp_lock_hint_nonspeculative = &
28         omp_sync_hint_nonspeculative
29 integer (kind=omp_sync_hint_kind), &
30     parameter :: omp_sync_hint_speculative = &
31         int(Z'8', kind=omp_sync_hint_kind)
32 integer (kind=omp_lock_hint_kind), &
33     parameter :: omp_lock_hint_speculative = &
34         omp_sync_hint_speculative

```

Fortran

35 The hints can be combined by using the + or | operators in C/C++ or the + operator in Fortran.  
36 Combining `omp_sync_hint_none` with any other hint is equivalent to specifying the other hint.  
37 The intended meaning of each hint is:

- 1 • **omp\_sync\_hint\_uncontended**: low contention is expected in this operation, that is, few  
2 threads are expected to perform the operation simultaneously in a manner that requires  
3 synchronization;
- 4 • **omp\_sync\_hint\_contended**: high contention is expected in this operation, that is, many  
5 threads are expected to perform the operation simultaneously in a manner that requires  
6 synchronization;
- 7 • **omp\_sync\_hint\_speculative**: the programmer suggests that the operation should be  
8 implemented using speculative techniques such as transactional memory; and
- 9 • **omp\_sync\_hint\_nonspeculative**: the programmer suggests that the operation should  
10 not be implemented using speculative techniques such as transactional memory.

11

---

12 Note – Future OpenMP specifications may add additional hints to the **omp\_sync\_hint\_t** type  
13 and the **omp\_sync\_hint\_kind** kind. Implementers are advised to add implementation-defined  
14 hints starting from the most significant bit of the **omp\_sync\_hint\_t** type and  
15 **omp\_sync\_hint\_kind** kind and to include the name of the implementation in the name of the  
16 added hint to avoid name conflicts with other OpenMP implementations.

---

17

18 The **omp\_sync\_hint\_t** and **omp\_lock\_hint\_t** enumeration types and the equivalent types  
19 in Fortran are synonyms for each other. The type **omp\_lock\_hint\_t** has been deprecated.

## 20 Restrictions

21 Restrictions to the synchronization hints are as follows:

- 22 • The hints **omp\_sync\_hint\_uncontended** and **omp\_sync\_hint\_contended** cannot  
23 be combined.
- 24 • The hints **omp\_sync\_hint\_nonspeculative** and **omp\_sync\_hint\_speculative**  
25 cannot be combined.

26 The restrictions for combining multiple values of **omp\_sync\_hint** apply equally to the  
27 corresponding values of **omp\_lock\_hint**, and expressions that mix the two types.

## 28 Cross References

- 29 • **critical** construct, see Section [2.20.1](#) on page [226](#).
- 30 • **atomic** construct, see Section [2.20.7](#) on page [237](#)
- 31 • **omp\_init\_lock\_with\_hint** and **omp\_init\_nest\_lock\_with\_hint**, see  
32 Section [3.3.2](#) on page [385](#).

## 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

## 1 Binding

2 The binding thread set of the **cancel** region is the current team. The binding region of the  
3 **cancel** region is the innermost enclosing region of the type corresponding to the  
4 *construct-type-clause* specified in the directive (that is, the innermost **parallel**, **sections**,  
5 worksharing-loop, or **taskgroup** region).

## 6 Description

7 The **cancel** construct activates cancellation of the binding region only if the *cancel-var* ICV is  
8 *true*, in which case the **cancel** construct causes the encountering task to continue execution at the  
9 end of the binding region if *construct-type-clause* is **parallel**, **for**, **do**, or **sections**. If the  
10 *cancel-var* ICV is *true* and *construct-type-clause* is **taskgroup**, the encountering task continues  
11 execution at the end of the current task region. If the *cancel-var* ICV is *false*, the **cancel**  
12 construct is ignored.

13 Threads check for active cancellation only at cancellation points that are implied at the following  
14 locations:

- 15 • **cancel** regions;
- 16 • **cancellation point** regions;
- 17 • **barrier** regions;
- 18 • implicit barriers regions.

19 When a thread reaches one of the above cancellation points and if the *cancel-var* ICV is *true*, then:

- 20 • If the thread is at a **cancel** or **cancellation point** region and *construct-type-clause* is  
21 **parallel**, **for**, **do**, or **sections**, the thread continues execution at the end of the canceled  
22 region if cancellation has been activated for the innermost enclosing region of the type specified.
- 23 • If the thread is at a **cancel** or **cancellation point** region and *construct-type-clause* is  
24 **taskgroup**, the encountering task checks for active cancellation of all of the *taskgroup sets* to  
25 which the encountering task belongs, and continues execution at the end of the current task  
26 region if cancellation has been activated for any of the *taskgroup sets*.
- 27 • If the encountering task is at a barrier region, the encountering task checks for active cancellation  
28 of the innermost enclosing **parallel** region. If cancellation has been activated, then the  
29 encountering task continues execution at the end of the canceled region.

30 ▼

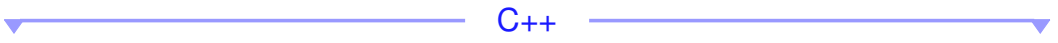
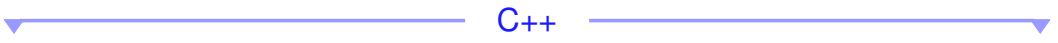
31 **Note** – If one thread activates cancellation and another thread encounters a cancellation point, the  
32 order of execution between the two threads is non-deterministic. Whether the thread that  
33 encounters a cancellation point detects the activated cancellation depends on the underlying  
34 hardware and operating system.

35 ▲

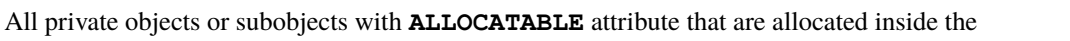
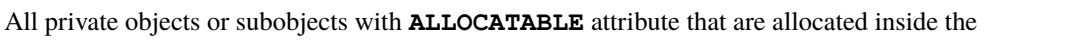


1 When cancellation of tasks is activated through a **cancel** construct with the **taskgroup**  
2 *construct-type-clause*, the tasks that belong to the *taskgroup set* of the innermost enclosing  
3 **taskgroup** region will be canceled. The task that encountered that construct continues execution  
4 at the end of its task region, which implies completion of that task. Any task that belongs to the  
5 innermost enclosing **taskgroup** and has already begun execution must run to completion or until  
6 a cancellation point is reached. Upon reaching a cancellation point and if cancellation is active, the  
7 task continues execution at the end of its task region, which implies the task's completion. Any task  
8 that belongs to the innermost enclosing **taskgroup** and that has not begun execution may be  
9 discarded, which implies its completion.



10 When cancellation is active for a **parallel**, **sections**, or worksharing-loop region, each  
11 thread of the binding thread set resumes execution at the end of the canceled region if a cancellation  
12 point is encountered. If the canceled region is a **parallel** region, any tasks that have been  
13 created by a **task** or a **taskloop** construct and their descendent tasks are canceled according to  
14 the above **taskgroup** cancellation semantics. If the canceled region is a **sections**, or  
15 worksharing-loop region, no task cancellation occurs.

16  C++   
The usual C++ rules for object destruction are followed when cancellation is performed.

17  Fortran   
All private objects or subobjects with **ALLOCATABLE** attribute that are allocated inside the  
18 canceled construct are deallocated.

19  Fortran   
If the canceled construct contains a **reduction**, **task\_reduction** or **lastprivate** clause,  
20 the final value of the list items that appeared in those clauses are undefined.

21 When an **if** clause is present on a **cancel** construct and the **if** expression evaluates to *false*, the  
22 **cancel** construct does not activate cancellation. The cancellation point associated with the  
23 **cancel** construct is always encountered regardless of the value of the **if** expression.

24   
25 **Note** – The programmer is responsible for releasing locks and other synchronization data  
26 structures that might cause a deadlock when a **cancel** construct is encountered and blocked  
27 threads cannot be canceled. The programmer is also responsible for ensuring proper  
28 synchronizations to avoid deadlocks that might arise from cancellation of OpenMP regions that  
29 contain OpenMP synchronization constructs.  
30 

## 31 Execution Model Events

32 If a task encounters a **cancel** construct that will activate cancellation then a *cancel* event occurs.

33 A *discarded-task* event occurs for any discarded tasks.

## 1 Tool Callbacks

2 A thread dispatches a registered `ompt_callback_cancel` callback for each occurrence of a  
3 `cancel` event in the context of the encountering task. This callback has type signature  
4 `ompt_callback_cancel_t`; `(flags & ompt_cancel_activated)` always evaluates to  
5 `true` in the dispatched callback; `(flags & ompt_cancel_parallel)` evaluates to `true` in the  
6 dispatched callback if `construct-type-clause` is `parallel`;  
7 `(flags & ompt_cancel_sections)` evaluates to `true` in the dispatched callback if  
8 `construct-type-clause` is `sections`; `(flags & ompt_cancel_loop)` evaluates to `true` in the  
9 dispatched callback if `construct-type-clause` is `for` or `do`; and  
10 `(flags & ompt_cancel_taskgroup)` evaluates to `true` in the dispatched callback if  
11 `construct-type-clause` is `taskgroup`.

12 A thread dispatches a registered `ompt_callback_cancel` callback with the `ompt_data_t`  
13 associated with the discarded task as its `task_data` argument and  
14 `ompt_cancel_discarded_task` as its `flags` argument for each occurrence of a  
15 `discarded-task` event. The callback occurs in the context of the task that discards the task and has  
16 type signature `ompt_callback_cancel_t`.

## 17 Restrictions

18 The restrictions to the `cancel` construct are as follows:

- 19 • The behavior for concurrent cancellation of a region and a region nested within it is unspecified.
- 20 • If `construct-type-clause` is `taskgroup`, the `cancel` construct must be closely nested inside a  
21 `task` or a `taskloop` construct and the `cancel` region must be closely nested inside a  
22 `taskgroup` region. If `construct-type-clause` is `sections`, the `cancel` construct must be  
23 closely nested inside a `sections` or `section` construct. Otherwise, the `cancel` construct  
24 must be closely nested inside an OpenMP construct that matches the type specified in  
25 `construct-type-clause` of the `cancel` construct.
- 26 • A worksharing construct that is canceled must not have a `nowait` clause.
- 27 • A worksharing-loop construct that is canceled must not have an `ordered` clause.
- 28 • A worksharing-loop construct that is canceled must not have a `reduction` clause with the  
29 `inscan` modifier.
- 30 • During execution of a construct that may be subject to cancellation, a thread must not encounter  
31 an orphaned cancellation point. That is, a cancellation point must only be encountered within  
32 that construct and must not be encountered elsewhere in its region.

## 33 Cross References

- 34 • `cancel-var` ICV, see Section 2.6.1 on page 66.
- 35 • `if` clause, see Section 2.18 on page 224.
- 36 • `cancellation point` construct, see Section 2.21.2 on page 269.

- `omp_get_cancellation` routine, see Section 3.2.9 on page 345.
- `omp_cancel_flag_t` enumeration type, see Section 4.4.4.24 on page 453.
- `ompt_callback_cancel_t`, see Section 4.5.2.18 on page 481.

## 2.21.2 cancellation point Construct

### Summary

The **cancellation point** construct introduces a user-defined cancellation point at which implicit or explicit tasks check if cancellation of the innermost enclosing region of the type specified has been activated. The **cancellation point** construct is a stand-alone directive.

### Syntax

C / C++

The syntax of the **cancellation point** construct is as follows:

```
#pragma omp cancellation point construct-type-clause new-line
```

where *construct-type-clause* is one of the following:

```
parallel  
sections  
for  
taskgroup
```

C / C++

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**, worksharing-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 *cancel* event in the context of the encountering task. This callback has type signature **ompt\_callback\_cancel\_t**; (*flags & ompt\_cancel\_detected*) always evaluates to *true* in the dispatched callback; (*flags & ompt\_cancel\_parallel*) evaluates to *true* in the dispatched callback if *construct-type-clause* of the encountered **cancellation point** construct is **parallel**; (*flags & ompt\_cancel\_sections*) evaluates to *true* in the dispatched callback if *construct-type-clause* of the encountered **cancellation point** construct is **sections**; (*flags & ompt\_cancel\_loop*) evaluates to *true* in the dispatched callback if *construct-type-clause* of the encountered **cancellation point** construct is **for** or **do**; and (*flags & ompt\_cancel\_taskgroup*) evaluates to *true* in the dispatched callback if *construct-type-clause* of the encountered **cancellation point** construct is **taskgroup**.

## Restrictions

- A **cancellation point** construct for which *construct-type-clause* is **taskgroup** must be closely nested inside a **task** or **taskloop** 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.
- A **cancellation point** construct for which *construct-type-clause* is neither **sections** nor **taskgroup** must be closely nested inside an OpenMP construct that matches the type specified in *construct-type-clause*.

## Cross References

- *cancel-var* ICV, see Section 2.6.1 on page 66.
- **cancel** construct, see Section 2.21.1 on page 265.
- **omp\_get\_cancellation** routine, see Section 3.2.9 on page 345.
- **ompt\_callback\_cancel\_t**, see Section 4.5.2.18 on page 481.

## 2.22 Data Environment

This section presents directives and clauses for controlling data environments.

### 2.22.1 Data-Sharing Attribute Rules

This section describes how the data-sharing attributes of variables referenced in data environments are determined. The following two cases are described separately:

- Section 2.22.1.1 on page 271 describes the data-sharing attribute rules for variables referenced in a construct.
- Section 2.22.1.2 on page 275 describes the data-sharing attribute rules for variables referenced in a region, but outside any construct.

#### 2.22.1.1 Variables Referenced in a Construct

The data-sharing attributes of variables that are referenced in a construct can be *predetermined*, *explicitly determined*, or *implicitly determined*, according to the rules outlined in this section.

Specifying a variable in a data-sharing attribute clause, except for the **private** clause, or **copyprivate** clause of an enclosed construct causes an implicit reference to the variable in the enclosing construct. Specifying a variable in a **map** clause of an enclosed construct may cause an implicit reference to the variable in the enclosing construct. Such implicit references are also subject to the data-sharing attribute rules outlined in this section.

Certain variables and objects have *predetermined* data-sharing attributes as follows:



- Variables that appear in **threadprivate** directives or variables with the **\_Thread\_local** (in C) or **thread\_local** (in C++) storage-class specifier are threadprivate.
- Variables with automatic storage duration that are declared in a scope inside the construct are private.
- Objects with dynamic storage duration are shared.
- Static data members are shared.

- The loop iteration variable(s) in the associated *for-loop(s)* of a **for**, **parallel for**, **taskloop**, or **distribute** construct is (are) private.
- The loop iteration variable in the associated *for-loop* of a **simd** 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** construct with multiple associated *for-loops* are lastprivate.
- The loop iteration variable(s) in the associated *for-loop(s)* of a **loop** construct is (are) lastprivate.
- Variables with static storage duration that are declared in a scope inside the construct are shared.
- If a list item in a **map** clause on the **target** construct has a base pointer, and the base pointer is a scalar variable that does not appear in a **map** clause on the construct, the base pointer is firstprivate.
- If a list item in a **reduction** or **in\_reduction** clause on a construct has a base pointer then the base pointer is private.
- The `__func__` variable and similar function-local predefined variables are shared.

▲ C / C++ ▲

▼ Fortran ▼

- Variables and common blocks that appear 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** 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** construct with multiple associated *do-loops* are lastprivate.
- The loop iteration variable(s) in the associated *do-loop(s)* of a **loop** construct is (are) lastprivate.
- Loop iteration variables inside **parallel** or task generating constructs are private in the innermost such construct that encloses the loop.
- Implied-do indices and **forall** indices are private.
- Cray pointees have the same 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 ▲

1 Variables with predetermined data-sharing attributes may not be listed in data-sharing attribute  
2 clauses, except for the cases listed below. For these exceptions only, listing a predetermined  
3 variable in a data-sharing attribute clause is allowed and overrides the variable's predetermined  
4 data-sharing attributes.

### C / C++

- 5 • The loop iteration variable(s) in the associated *for-loop(s)* of a **for**, **parallel for**,  
6 **taskloop**, **distribute**, or **loop** construct may be listed in a **private** or  
7 **lastprivate** clause.
- 8 • The loop iteration variable in the associated *for-loop* of a **simd** construct with just one  
9 associated *for-loop* may be listed in a **private**, **lastprivate**, or **linear** clause with a  
10 *linear-step* that is the increment of the associated *for-loop*.
- 11 • The loop iteration variables in the associated *for-loops* of a **simd** construct with multiple  
12 associated *for-loops* may be listed in a **private** or **lastprivate** clause.
- 13 • Variables with **const**-qualified type with no mutable members may be listed in a  
14 **firstprivate** clause, even if they are static data members.
- 15 • The **\_\_func\_\_** variable and similar function-local predefined variables may be listed in a  
16 **shared** or **firstprivate** clause.

### C / C++

### Fortran

- 17 • The loop iteration variable(s) in the associated *do-loop(s)* of a **do**, **parallel do**, **taskloop**,  
18 **distribute**, or **loop** construct may be listed in a **private** or **lastprivate** clause.
- 19 • The loop iteration variable in the associated *do-loop* of a **simd** construct with just one  
20 associated *do-loop* may be listed in a **private**, **lastprivate**, or **linear** clause with a  
21 *linear-step* that is the increment of the associated loop.
- 22 • The loop iteration variables in the associated *do-loops* of a **simd** construct with multiple  
23 associated *do-loops* may be listed in a **private** or **lastprivate** clause.
- 24 • Loop iteration variables of loops that are not associated with any OpenMP directive may be  
25 listed in data-sharing attribute clauses on the surrounding **teams**, **parallel** or task generating  
26 construct, and on enclosed constructs, subject to other restrictions.
- 27 • Assumed-size arrays may be listed in a **shared** clause.

### Fortran

1 Additional restrictions on the variables that may appear in individual clauses are described with  
2 each clause in Section 2.22.4 on page 284.

3 Variables with *explicitly determined* data-sharing attributes are those that are referenced in a given  
4 construct and are listed in a data-sharing attribute clause on the construct.

5 Variables with *implicitly determined* data-sharing attributes are those that are referenced in a given  
6 construct, do not have predetermined data-sharing attributes, and are not listed in a data-sharing  
7 attribute clause on the construct.

8 Rules for variables with *implicitly determined* data-sharing attributes are as follows:

- 9 • In a **parallel**, **teams**, or task generating construct, the data-sharing attributes of these  
10 variables are determined by the **default** clause, if present (see Section 2.22.4.1 on page 284).
- 11 • In a **parallel** construct, if no **default** clause is present, these variables are shared.
- 12 • For constructs other than task generating constructs, if no **default** clause is present, these  
13 variables reference the variables with the same names that exist in the enclosing context.
- 14 • In a **target** construct, variables that are not mapped after applying data-mapping attribute  
15 rules (see Section 2.22.7 on page 317) are firstprivate.

▼ C++ ▼

- 16 • In an orphaned task generating construct, if no **default** clause is present, formal arguments  
17 passed by reference are firstprivate.

▲ C++ ▲

▼ Fortran ▼

- 18 • In an orphaned task generating construct, if no **default** clause is present, dummy arguments  
19 are firstprivate.

▲ Fortran ▲

- 20 • In a task generating construct, if no **default** clause is present, a variable for which the  
21 data-sharing attribute is not determined by the rules above and that in the enclosing context is  
22 determined to be shared by all implicit tasks bound to the current team is shared.
- 23 • In a task generating construct, if no **default** clause is present, a variable for which the  
24 data-sharing attribute is not determined by the rules above is firstprivate.

25 Additional restrictions on the variables for which data-sharing attributes cannot be implicitly  
26 determined in a task generating construct are described in Section 2.22.4.4 on page 288.



## 2.22.1.2 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++

- Variables with static storage duration that are declared in called routines in the region are shared.
- File-scope or namespace-scope variables referenced in called routines in the region are shared unless they appear in a **threadprivate** directive.
- Objects with dynamic storage duration are shared.
- Static data members are shared unless they appear in a **threadprivate** directive.
- In C++, formal arguments of called routines in the region that are passed by reference have the same data-sharing attributes as the associated actual arguments.
- Other variables declared in called routines in the region are private.

### C / C++

### Fortran

- Local variables declared in called routines in the region and that have the **save** attribute, or that are data initialized, are shared unless they appear in a **threadprivate** directive.
- Variables belonging to common blocks, or accessed by host or use association, and referenced in called routines in the region are shared unless they appear in a **threadprivate** directive.
- Dummy arguments of called routines in the region that have the **VALUE** attribute are private.
- Dummy arguments of called routines in the region that do not have the **VALUE** attribute are private if the associated actual argument is not shared.
- Dummy arguments of called routines in the region that do not have the **VALUE** attribute are shared if the actual argument is shared and it is a scalar variable, structure, an array that is not a pointer or assumed-shape array, or a simply contiguous array section. Otherwise, the data-sharing attribute of the dummy argument is implementation-defined if the associated actual argument is shared.
- Cray pointees have the same data-sharing attribute as the storage with which their Cray pointers are associated.
- Implied-do indices, **forall** indices, and other local variables declared in called routines in the region are private.

### Fortran

## 2.22.2 threadprivate Directive

### Summary

The **threadprivate** directive specifies that variables are replicated, with each thread having its own copy. The **threadprivate** directive is a declarative directive.

### Syntax

C / C++

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

```
#pragma omp threadprivate (list) new-line
```

where *list* is a comma-separated list of file-scope, namespace-scope, or static block-scope variables that do not have incomplete types.

C / C++

Fortran

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

```
!$omp threadprivate (list)
```

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

Fortran

### Description

Unless otherwise specified, each copy of a **threadprivate** variable is initialized once, in the manner specified by the program, but at an unspecified point in the program prior to the first reference to that copy. The storage of all copies of a **threadprivate** variable is freed according to how static variables are handled in the base language, but at an unspecified point in the program.

C++

Each copy of a block-scope **threadprivate** variable that has a dynamic initializer is initialized the first time its thread encounters its definition; if its thread does not encounter its definition, its initialization is unspecified.

C++

1 A program in which a thread references another thread's copy of a threadprivate variable is  
2 non-conforming.

3 The content of a threadprivate variable can change across a task scheduling point if the executing  
4 thread switches to another task that modifies the variable. For more details on task scheduling, see  
5 Section 1.3 on page 21 and Section 2.13 on page 142.

6 In **parallel** regions, references by the master thread will be to the copy of the variable in the  
7 thread that encountered the **parallel** region.

8 During a sequential part references will be to the initial thread's copy of the variable. The values of  
9 data in the initial thread's copy of a threadprivate variable are guaranteed to persist between any  
10 two consecutive references to the variable in the program provided that no **teams** construct that is  
11 not nested inside of a **target** construct is encountered between the references and that the initial  
12 thread is not nested inside of a **teams** region. For initial threads nested inside of a **teams** region,  
13 the values of data in the copies of a threadprivate variable of those initial threads are guaranteed to  
14 persist between any two consecutive references to the variable inside of that **teams** region.

15 The values of data in the threadprivate variables of threads that are not initial threads are  
16 guaranteed to persist between two consecutive active **parallel** regions only if all of the  
17 following conditions hold:

- 18 • Neither **parallel** region is nested inside another explicit **parallel** region;
- 19 • The number of threads used to execute both **parallel** regions is the same;
- 20 • The thread affinity policies used to execute both **parallel** regions are the same;
- 21 • The value of the *dyn-var* internal control variable in the enclosing task region is *false* at entry to  
22 both **parallel** regions; and
- 23 • No **teams** construct that is not nested inside of a **target** construct is encountered between  
24 both **parallel** regions.
- 25 • Neither the **omp\_pause\_resource** nor **omp\_pause\_resource\_all** routine is called.

26 If these conditions all hold, and if a threadprivate variable is referenced in both regions, then  
27 threads with the same thread number in their respective regions will reference the same copy of that  
28 variable.

C / C++

29 If the above conditions hold, the storage duration, lifetime, and value of a thread's copy of a  
30 threadprivate variable that does not appear in any **copyin** clause on the second region will be  
31 retained. Otherwise, the storage duration, lifetime, and value of a thread's copy of the variable in  
32 the second region is unspecified.

C / C++

## Fortran

1 If the above conditions hold, the definition, association, or allocation status of a thread's copy of a  
2 threadprivate variable or a variable in a threadprivate common block that is not affected by any  
3 **copyin** clause that appears on the second region (a variable is affected by a **copyin** clause if the  
4 variable appears in the **copyin** clause or it is in a common block that appears in the **copyin**  
5 clause) will be retained. Otherwise, the definition and association status of a thread's copy of the  
6 variable in the second region are undefined, and the allocation status of an allocatable variable will  
7 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, each copy will have the same association status as the initial  
16 association status.
- 17 • If it does not have either the **POINTER** or the **ALLOCATABLE** attribute:
  - 18 – If it is initially defined, either through explicit initialization or default initialization, each copy  
19 created is so defined;
  - 20 – Otherwise, each copy created is undefined.

## Fortran

## C / C++

21 The address of a threadprivate variable is not an address constant.

## 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++

## Restrictions

The restrictions to the **threadprivate** directive are as follows:

- A **threadprivate** variable must not appear in any clause except the **copyin**, **copyprivate**, **schedule**, **num\_threads**, **thread\_limit**, and **if** clauses.
- A program in which an untied task accesses **threadprivate** storage is non-conforming.

C / C++

- If the value of a variable referenced in an explicit initializer of a **threadprivate** variable is modified prior to the first reference to any instance of the **threadprivate** variable, then the behavior is unspecified.
- A variable that is part of another variable (as an array or structure element) cannot appear in a **threadprivate** clause unless it is a static data member of a C++ class.
- A **threadprivate** directive for file-scope variables must appear outside any definition or declaration, and must lexically precede all references to any of the variables in its list.
- A **threadprivate** directive for namespace-scope variables must appear outside any definition or declaration other than the namespace definition itself, and must lexically precede all references to any of the variables in its list.
- Each variable in the list of a **threadprivate** directive at file, namespace, or class scope must refer to a variable declaration at file, namespace, or class scope that lexically precedes the directive.
- A **threadprivate** directive for static block-scope variables must appear in the scope of the variable and not in a nested scope. The directive must lexically precede all references to any of the variables in its list.
- Each variable in the list of a **threadprivate** directive in block scope must refer to a variable declaration in the same scope that lexically precedes the directive. The variable declaration must use the static storage-class specifier.
- If a variable is specified in a **threadprivate** directive in one translation unit, it must be specified in a **threadprivate** directive in every translation unit in which it is declared.

C / C++

C++

- A **threadprivate** directive for static class member variables must appear in the class definition, in the same scope in which the member variables are declared, and must lexically precede all references to any of the variables in its list.
- A **threadprivate** variable must not have an incomplete type or a reference type.
- A **threadprivate** variable with class type must have:
  - An accessible, unambiguous default constructor in the case of default initialization without a given initializer;

- An accessible, unambiguous constructor that accepts the given argument in the case of direct initialization; and
- An accessible, unambiguous copy constructor in the case of copy initialization with an explicit initializer.

C++

Fortran

- A variable that is part of another variable (as an array, structure element or type parameter inquiry) cannot appear in a **threadprivate** clause.
- A coarray cannot appear in a **threadprivate** directive.
- The **threadprivate** directive must appear in the declaration section of a scoping unit in which the common block or variable is declared.
- If a **threadprivate** directive that specifies a common block name appears in one program unit, then such a directive must also appear in every other program unit that contains a **COMMON** statement that specifies the same name. It must appear after the last such **COMMON** statement in the program unit.
- If a threadprivate variable or a threadprivate common block is declared with the **BIND** attribute, the corresponding C entities must also be specified in a **threadprivate** directive in the C program.
- A blank common block cannot appear in a **threadprivate** directive.
- A variable can only appear in a **threadprivate** directive in the scope in which it is declared. It must not be an element of a common block or appear in an **EQUIVALENCE** statement.
- A variable that appears in a **threadprivate** directive must be declared in the scope of a module or have the **SAVE** attribute, either explicitly or implicitly.

Fortran

## Cross References

- *dyn-var* ICV, see Section 2.6 on page 66.
- Number of threads used to execute a **parallel** region, see Section 2.8.1 on page 84.
- **copyin** clause, see Section 2.22.6.1 on page 313.

## 2.22.3 List Item Privatization

For any construct, a list item that appears in a data-sharing attribute clause, including a reduction clause, may be privatized. Each task that references a privatized list item in any statement in the construct receives at least one new list item if the construct has one or more associated loops, and otherwise each such task receives one new list item. Each SIMD lane used in a `simd` construct that references a privatized list item in any statement in the construct receives at least one new list item. Language-specific attributes for new list items are derived from the corresponding original list item. Inside the construct, all references to the original list item are replaced by references to a new list item received by the task or SIMD lane.

If the construct has one or more associated loops, within the same logical iteration of the loop(s) the same new list item replaces all references to the original list item. For any two logical iterations, if the references to the original list item are replaced by the same list item then the logical iterations must execute in some sequential order.

In the rest of the region, it is unspecified whether references are to a new list item or the original list item. Therefore, if an attempt is made to reference the original item, its value after the region is also unspecified. If a task or a SIMD lane does not reference a privatized list item, it is unspecified whether the task or SIMD lane receives a new list item.

The value and/or allocation status of the original list item will change only:

- If accessed and modified via pointer;
- If possibly accessed in the region but outside of the construct;
- As a side effect of directives or clauses; or

Fortran

- If accessed and modified via construct association.

Fortran

C++

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++

A new list item of the same type, with automatic storage duration, is allocated for the construct. The storage and thus lifetime of these list items last until the block in which they are created exits. The size and alignment of the new list item are determined by the type of the variable. This allocation occurs once for each task generated by the construct and once for each SIMD lane used by the construct.

The new list item is initialized, or has an undefined initial value, as if it had been locally declared without an initializer.

C / C++

C++

1 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  
2 purposes of this clause.

3 The order in which any default constructors for different private variables of class type are called is  
4 unspecified. The order in which any destructors for different private variables of class type are  
5 called is unspecified.

C++

Fortran

6 If any statement of the construct references a list item, a new list item of the same type and type  
7 parameters is allocated. This allocation occurs once for each task generated by the construct and  
8 once for each SIMD lane used by the construct. If the type of the list item has default initialization,  
9 the new list item has default initialization. Otherwise, the initial value of the new list item is  
10 undefined. The initial status of a private pointer is undefined.

11 For a list item or the subobject of a list item with the **ALLOCATABLE** attribute:

- 12 • If the allocation status is unallocated, the new list item or the subobject of the new list item will  
13 have an initial allocation status of unallocated;
- 14 • If the allocation status is allocated, the new list item or the subobject of the new list item will  
15 have an initial allocation status of allocated; and
- 16 • If the new list item or the subobject of the new list item is an array, its bounds will be the same as  
17 those of the original list item or the subobject of the original list item.

18 A privatized list item may be storage-associated with other variables when the data-sharing  
19 attribute clause is encountered. Storage association may exist because of constructs such as  
20 **EQUIVALENCE** or **COMMON**. If  $A$  is a variable that is privatized by a construct and  $B$  is a variable  
21 that is storage-associated with  $A$ , then:

- 22 • The contents, allocation, and association status of  $B$  are undefined on entry to the region;
- 23 • Any definition of  $A$ , or of its allocation or association status, causes the contents, allocation, and  
24 association status of  $B$  to become undefined; and
- 25 • Any definition of  $B$ , or of its allocation or association status, causes the contents, allocation, and  
26 association status of  $A$  to become undefined.

27 A privatized list item may be a selector of an **ASSOCIATE** or **SELECT TYPE** construct. If the  
28 construct association is established prior to a **parallel** region, the association between the  
29 associate name and the original list item will be retained in the region.

30 Finalization of a list item of a finalizable type or subobjects of a list item of a finalizable type  
31 occurs at the end of the region. The order in which any final subroutines for different variables of a  
32 finalizable type are called is unspecified.

Fortran



1 If a list item appears in both **firstprivate** and **lastprivate** clauses, the update required  
2 for the **lastprivate** clause occurs after all initializations for the **firstprivate** clause.

### 3 **Restrictions**

4 The following restrictions apply to any list item that is privatized unless otherwise stated for a given  
5 data-sharing attribute clause:

6 **C**

- A variable that is part of another variable (as an array or structure element) cannot be privatized.

7 **C**

8 **C++**

- A variable that is part of another variable (as an array or structure element) cannot be privatized  
9 except if the data-sharing attribute clause is associated with a construct within a class non-static  
10 member function and the variable is an accessible data member of the object for which the  
non-static member function is invoked.

- A variable of class type (or array thereof) that is privatized requires an accessible, unambiguous  
11 default constructor for the class type.

12 **C++**

13 **C / C++**

- A variable that is privatized must not have a **const**-qualified type unless it is of class type with  
14 a **mutable** member. This restriction does not apply to the **firstprivate** clause.
- A variable that is privatized must not have an incomplete type or be a reference to an incomplete  
15 type.

16 **C / C++**

17 **Fortran**

- A variable that is part of another variable (as an array or structure element) cannot be privatized.
- A variable that is privatized must either be definable, or an allocatable variable. This restriction  
18 does not apply to the **firstprivate** clause.
- Variables that appear in namelist statements, in variable format expressions, and in expressions  
19 for statement function definitions, may not be privatized.
- Pointers with the **INTENT (IN)** attribute may not be privatized. This restriction does not apply  
20 to the **firstprivate** clause.
- A private variable must not be coindexed or appear as an actual argument to a procedure where  
21 the corresponding dummy argument is a coarray.
- Assumed-size arrays may not be privatized in a **target**, **teams**, or **distribute** construct.

22 **Fortran**

## 2.22.4 Data-Sharing Attribute Clauses

Several constructs accept clauses that allow a user to control the data-sharing attributes of variables referenced in the construct. Not all of the clauses listed in this section are valid on all directives.

The set of clauses that is valid on a particular directive is described with the directive.

Most of the clauses accept a comma-separated list of list items (see Section 2.1 on page 36). All list items that appear in a clause must be visible, according to the scoping rules of the base language.

With the exception of the **default** clause, clauses may be repeated as needed. A list item may not appear in more than one clause on the same directive, except that it may be specified in both

**firstprivate** and **lastprivate** clauses.

The reduction data-sharing attribute clauses are explained in Section 2.22.5 on page 295.

▼ C++ ▼

If a variable referenced in a data-sharing attribute clause has a type derived from a template, and the program does not otherwise reference that variable then any behavior related to that variable is unspecified.

▲ C++ ▲

▼ Fortran ▼

When a named common block appears in a **private**, **firstprivate**, **lastprivate**, or **shared** clause of a directive, none of its members may be declared in another data-sharing attribute clause in that directive. When individual members of a common block appear in a **private**, **firstprivate**, **lastprivate**, **reduction**, or **linear** clause of a directive, the storage of the specified variables is no longer Fortran associated with the storage of the common block itself.

▲ Fortran ▲

### 2.22.4.1 default Clause

#### Summary

The **default** clause explicitly determines the data-sharing attributes of variables that are referenced in a **parallel**, **teams**, or task generating construct and would otherwise be implicitly determined (see Section 2.22.1.1 on page 271).

## Syntax

The syntax of the **default** clause is as follows:

```
default (data-sharing-attribute)
```

where *data-sharing-attribute* is one of the following:

```
shared  
firstprivate  
private  
none
```

## Description

If *data-sharing-attribute* is **shared** or, for Fortran, **firstprivate** or **private**, the data-sharing attribute of all variables referenced in the construct that have implicitly determined data-sharing attributes will be *data-sharing-attribute*.

▼ C / C++ ▼

If *data-sharing-attribute* is **firstprivate** or **private**, each variable with static storage duration that is declared in a namespace or global scope and referenced in the construct, and that does not have a predetermined data-sharing attribute, must have its data-sharing attribute explicitly determined by being listed in a data-sharing attribute clause. The data-sharing attribute of all other variables that are referenced in the construct and have implicitly determined data-sharing attributes will be *data-sharing-attribute*.

▲ C / C++ ▲

The **default (none)** clause requires that each variable that is referenced in the construct, and that does not have a predetermined data-sharing attribute, must have its data-sharing attribute explicitly determined by being listed in a data-sharing attribute clause.

## Restrictions

The restrictions to the **default** clause are as follows:

- Only a single **default** clause may be specified on a **parallel**, **task**, **taskloop** or **teams** directive.

### 2.22.4.2 shared Clause

#### Summary

The **shared** clause declares one or more list items to be shared by tasks generated by a **parallel**, **teams**, or task generating construct.

## Syntax

The syntax of the **shared** clause is as follows:

```
shared (list)
```

## Description

All references to a list item within a task refer to the storage area of the original variable at the point the directive was encountered.

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 exit from the **parallel**, **teams**, or task generating construct if it is associated with a target or a subobject of a target that appears as a privatized list item in a data-sharing attribute clause on 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 275. 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 is copied 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 data races.

### Fortran

## Restrictions

The restrictions for the **shared** clause are as follows:

- A variable that is part of another variable (as an array or structure element) cannot appear in a **shared** clause.

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++

Fortran

- A variable that is part of another variable (as an array, structure element or type parameter inquiry) cannot appear in a **shared** clause.

Fortran

### 2.22.4.3 private Clause

#### Summary

The **private** clause declares one or more list items to be private to a task or to a SIMD lane.

#### Syntax

The syntax of the **private** clause is as follows:

```
private (list)
```

#### Description

The **private** clause specifies that its list items are to be privatized according to Section 2.22.3 on page 281. Each task or SIMD lane that references a list item in the construct receives only one new list item, unless the construct has one or more associated loops and the **order (concurrent)** clause is also present.

List items that appear in a **private**, **firstprivate**, or **reduction** clause in a **parallel** construct may also appear in a **private** clause in an enclosed **parallel**, worksharing, **loop**, **task**, **taskloop**, **simd**, **target** or **scope** construct.

List items that appear in a **private** or **firstprivate** clause in a **task** or **taskloop** construct may also appear in a **private** clause in an enclosed **parallel**, **loop**, **task**, **taskloop**, **simd**, or **target** construct.

List items that appear in a **private**, **firstprivate**, **lastprivate**, or **reduction** clause in a worksharing construct may also appear in a **private** clause in an enclosed **parallel**, **loop**, **task**, **simd**, or **target** construct.

List items that appear in a **private** clause on a **loop** construct may also appear in a **private** clause in an enclosed **loop**, **parallel**, or **simd** construct.

#### Restrictions

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 281.

### 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 287, 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 data races, 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.

▼ C / C++ ▼

For variables of non-array type, the initialization occurs by copy assignment. For an array of elements of non-array type, each element is initialized as if by assignment from an element of the original array to the corresponding element of the new array.

▲ C / C++ ▲

1 For each variable of class type:

- 2 • If the **firstprivate** clause is not on a **target** construct then a copy constructor is invoked
- 3 to perform the initialization; and
- 4 • If the **firstprivate** clause is on a **target** construct then it is unspecified how many copy
- 5 constructors, if any, are invoked.

6 If copy constructors are called, the order in which copy constructors for different variables of class

7 type are called is unspecified.

8 If the original list item does not have the **POINTER** attribute, initialization of the new list items

9 occurs as if by intrinsic assignment unless the list item has a type bound procedure as a defined

10 assignment. If the original list item that does not have the **POINTER** attribute has the allocation

11 status of unallocated, the new list items will have the same status.

12 If the original list item has the **POINTER** attribute, the new list items receive the same association

13 status of the original list item as if by pointer assignment.

## 14 Restrictions

15 The restrictions to the **firstprivate** clause are as follows:

- 16 • A list item that is private within a **parallel** region must not appear in a **firstprivate**
- 17 clause on a worksharing construct if any of the worksharing regions arising from the worksharing
- 18 construct ever bind to any of the **parallel** regions arising from the **parallel** construct.
- 19 • A list item that is private within a **teams** region must not appear in a **firstprivate** clause
- 20 on a **distribute** construct if any of the **distribute** regions arising from the
- 21 **distribute** construct ever bind to any of the **teams** regions arising from the **teams**
- 22 construct.
- 23 • A list item that appears in a **reduction** clause of a **parallel** construct must not appear in a
- 24 **firstprivate** clause on a worksharing, **task**, or **taskloop** construct if any of the
- 25 worksharing or **task** regions arising from the worksharing, **task**, or **taskloop** construct
- 26 ever bind to any of the **parallel** regions arising from the **parallel** construct.
- 27 • A list item that appears in a **reduction** clause of a **teams** construct must not appear in a
- 28 **firstprivate** clause on a **distribute** construct if any of the **distribute** regions
- 29 arising from the **distribute** construct ever bind to any of the **teams** regions arising from the
- 30 **teams** construct.
- 31 • A list item that appears in a **reduction** clause of a worksharing construct must not appear in a
- 32 **firstprivate** clause in a **task** construct encountered during execution of any of the
- 33 worksharing regions arising from the worksharing construct.

## C++

- A variable of class type (or array thereof) that appears in a **firstprivate** clause requires an accessible, unambiguous copy constructor for the class type.

## C++

## C / C++

- If a list item in a **firstprivate** clause on a worksharing construct has a reference type then it must bind to the same object for all threads of the team.

## C / C++

## Fortran

- If the list item is a polymorphic variable with the **ALLOCATABLE** attribute, the behavior is unspecified.

## Fortran

### 2.22.4.5 lastprivate Clause

#### Summary

The **lastprivate** clause declares one or more list items to be private to an implicit task or to a 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 287. 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 data races, 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.

Otherwise, If the **lastprivate** clause is used on a construct to which the **nowait** or the **nogroup** clauses are applied, accesses to the original list item may create a data race. To avoid this data race, if an assignment to the original list item occurs then synchronization must be inserted to ensure that the assignment completes and the original list item is flushed to memory.

If a list item that appears in a **lastprivate** clause with the **conditional** modifier is modified in the region by an assignment outside the construct or not to the list item then the value assigned to the original list item is unspecified.

## Restrictions

The restrictions to the **lastprivate** clause are as follows:

- A list item that is private within a **parallel** region, or that appears in the **reduction** clause of a **parallel** construct, must not appear in a **lastprivate** clause on a worksharing construct if any of the corresponding worksharing regions ever binds to any of the corresponding **parallel** regions.
- A list item that appears in a **lastprivate** clause with the **conditional** modifier must be a scalar variable.

C++

- A variable of class type (or array thereof) that appears in a **lastprivate** clause requires an accessible, unambiguous default constructor for the class type, unless the list item is also specified in a **firstprivate** clause.

- A variable of class type (or array thereof) that appears in a **lastprivate** clause requires an accessible, unambiguous copy assignment operator for the class type. The order in which copy assignment operators for different variables of class type are called is unspecified.

C++

C / C++

- If a list item in a **lastprivate** clause on a worksharing construct has a reference type then it must bind to the same object for all threads of the team.

C / C++

Fortran

- A variable that appears in a **lastprivate** clause must be definable.
- If the original list item has the **ALLOCATABLE** attribute, the corresponding list item whose value is assigned to the original list item must have an allocation status of allocated upon exit from the sequentially last iteration or lexically last **section** construct.
- If the list item is a polymorphic variable with the **ALLOCATABLE** attribute, the behavior is unspecified.

Fortran

### 2.22.4.6 linear Clause

#### Summary

The **linear** clause declares one or more list items to be private and to have a linear relationship with respect to the iteration space of a loop associated with the construct on which the clause appears.

## Syntax

C

The syntax of the **linear** clause is as follows:

```
linear (linear-list [ : linear-step ])
```

where *linear-list* is one of the following

```
list
```

```
modifier (list)
```

where *modifier* is one of the following:

```
val
```

C

C++

The syntax of the **linear** clause is as follows:

```
linear (linear-list [ : linear-step ])
```

where *linear-list* is one of the following

```
list
```

```
modifier (list)
```

where *modifier* is one of the following:

```
ref
```

```
val
```

```
uval
```

C++

Fortran

The syntax of the **linear** clause is as follows:

```
linear (linear-list [ : linear-step ])
```

where *linear-list* is one of the following

```
list
```

```
modifier (list)
```

where *modifier* is one of the following:

```
ref
```

```
val
```

```
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 287 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 that corresponds to the construct. Otherwise, the execution results in unspecified behavior.
- 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++

- A *list-item* that appears in a **linear** clause without the **ref** modifier must be of integral or pointer type, or must be a reference to an integral or pointer type.
- The **ref** or **uval** modifier can only be used if the *list-item* is of a reference type.
- If a list item in a **linear** clause on a worksharing construct has a reference type then it must bind to the same object for all threads of the team.
- If the list item is of a reference type and the **ref** modifier is not specified and if any write to the list item occurs before any read of the list item then the result is unspecified.

C++

## Fortran

- 1 • A *list-item* that appears in a **linear** clause without the **ref** modifier must be of type  
2 **integer**.
- 3 • The **ref** or **uval** modifier can only be used if the *list-item* is a dummy argument without the  
4 **VALUE** attribute.
- 5 • Variables that have the **POINTER** attribute and Cray pointers may not appear in a **linear**  
6 clause.
- 7 • If the list item has the **ALLOCATABLE** attribute and the **ref** modifier is not specified, the  
8 allocation status of the list item in the sequentially last iteration must be allocated upon exit from  
9 that iteration.
- 10 • If the **ref** modifier is specified, variables with the **ALLOCATABLE** attribute, assumed-shape  
11 arrays and polymorphic variables may not appear in the **linear** clause.
- 12 • If the list item is a dummy argument without the **VALUE** attribute and the **ref** modifier is not  
13 specified and if any write to the list item occurs before any read of the list item then the result is  
14 unspecified.
- 15 • A common block name cannot appear in a **linear** clause.

## Fortran

### 2.22.5 Reduction Clauses and Directives

16 The reduction clauses are data-sharing attribute clauses that can be used to perform some forms of  
17 recurrence calculations in parallel. Reduction clauses include reduction scoping clauses and  
18 reduction participating clauses. Reduction scoping clauses define the region in which a reduction is  
19 computed. Reduction participating clauses define the participants in the reduction. Reduction  
20 clauses specify a *reduction-identifier* and one or more list items.  
21

#### 2.22.5.1 Properties Common To All Reduction Clauses

##### Syntax

22 The syntax of a *reduction-identifier* is defined as follows:

##### C

25 A *reduction-identifier* is either an *identifier* or one of the following operators: **+**, **-**, **\***, **&**, **|**, **^**, **&&**  
26 and **||**.

##### C

##### C++

27 A *reduction-identifier* is either an *id-expression* or one of the following operators: **+**, **-**, **\***, **&**, **|**, **^**,  
28 **&&** and **||**.

##### C++

▼ **Fortran** ▼

1 A *reduction-identifier* is either a base language identifier, or a user-defined operator, or one of the  
2 following operators: **+**, **-**, **\***, **.and.**, **.or.**, **.eqv.**, **.neqv.**, or one of the following intrinsic  
3 procedure names: **max**, **min**, **iand**, **ior**, **ieor**.

▲ **Fortran** ▲

▼ **C / C++** ▼

4 Table 2.11 lists each *reduction-identifier* that is implicitly declared at every scope for arithmetic  
5 types and its semantic initializer value. The actual initializer value is that value as expressed in the  
6 data type of the reduction list item.

**TABLE 2.11:** Implicitly Declared C/C++ *reduction-identifiers*

Identifier	Initializer	Combiner
+	<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</code>
&	<code>omp_priv = ~ 0</code>	<code>omp_out &amp;= omp_in</code>
	<code>omp_priv = 0</code>	<code>omp_out  = omp_in</code>
^	<code>omp_priv = 0</code>	<code>omp_out ^= omp_in</code>
&&	<code>omp_priv = 1</code>	<code>omp_out = omp_in &amp;&amp; omp_out</code>
	<code>omp_priv = 0</code>	<code>omp_out = omp_in    omp_out</code>
max	<code>omp_priv = Least representable number in the reduction list item type</code>	<code>omp_out = omp_in &gt; omp_out ? omp_in : omp_out</code>
min	<code>omp_priv = Largest representable number in the reduction list item type</code>	<code>omp_out = omp_in &lt; omp_out ? omp_in : omp_out</code>



1 Table 2.12 lists each *reduction-identifier* that is implicitly declared for numeric and logical types  
 2 and its semantic initializer value. The actual initializer value is that value as expressed in the data  
 3 type of the reduction list item.

**TABLE 2.12:** Implicitly Declared Fortran *reduction-identifiers*

Identifier	Initializer	Combiner
+	<code>omp_priv = 0</code>	<code>omp_out = omp_in + omp_out</code>
-	<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>
.and.	<code>omp_priv = .true.</code>	<code>omp_out = omp_in .and. omp_out</code>

*table continued on next page*

table continued from previous page

Identifier	Initializer	Combiner
<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. If the list item is an array or array section, the identifiers to which `omp_in`  
3 and `omp_out` correspond each refer to an array element. `omp_out` holds the final value of the  
4 combiner operation.

5 Any *reduction-identifier* that is defined with the `declare reduction` directive is also valid. In  
6 that case, the initializer and combiner of the *reduction-identifier* are specified by the  
7 *initializer-clause* and the *combiner* in the `declare reduction` directive.

## 8 Description

9 A reduction clause specifies a *reduction-identifier* and one or more list items.

10 The *reduction-identifier* specified in a reduction clause must match a previously declared  
11 *reduction-identifier* of the same name and type for each of the list items. This match is done by  
12 means of a name lookup in the base language.

13 The list items that appear in a reduction clause may include array sections.

C++

14 If the type is a derived class, then any *reduction-identifier* that matches its base classes is also a  
15 match, if there is no specific match for the type.

16 If the *reduction-identifier* is not an *id-expression*, then it is implicitly converted to one by  
17 prepending the keyword operator (for example, `+` becomes *operator+*).

18 If the *reduction-identifier* is qualified then a qualified name lookup is used to find the declaration.



1 If the *reduction-identifier* is unqualified then an *argument-dependent name lookup* must be  
2 performed using the type of each list item.

▲ C++ ▲

3 If the list item is an array or array section, it will be treated as if a reduction clause would be applied  
4 to each separate element of the array section.

5 If the list item is an array section, the elements of any copy of the array section will be allocated  
6 contiguously.

▼ Fortran ▼

7 If the original list item has the **POINTER** attribute, any copies of the list item are associated with  
8 private targets.

▲ Fortran ▲

9 Any copies associated with the reduction are initialized with the initializer value of the  
10 *reduction-identifier*.

11 Any copies are combined using the combiner associated with the *reduction-identifier*.

## 12 Execution Model Events

13 The *reduction-begin* event occurs before a task begins to perform loads and stores that belong to the  
14 implementation of a reduction and the *reduction-end* event occurs after the task has completed  
15 loads and stores associated with the reduction. If a task participates in multiple reductions, each  
16 reduction may be bracketed by its own pair of *reduction-begin/reduction-end* events or multiple  
17 reductions may be bracketed by a single pair of events. The interval defined by a pair of  
18 *reduction-begin/reduction-end* events may not contain a task scheduling point.

## 19 Tool Callbacks

20 A thread dispatches a registered **ompt\_callback\_reduction** with  
21 **ompt\_sync\_region\_reduction** in its *kind* argument and **ompt\_scope\_begin** as its  
22 *endpoint* argument for each occurrence of a *reduction-begin* event in that thread. Similarly, a thread  
23 dispatches a registered **ompt\_callback\_reduction** with  
24 **ompt\_sync\_region\_reduction** in its *kind* argument and **ompt\_scope\_end** as its  
25 *endpoint* argument for each occurrence of a *reduction-end* event in that thread. These callbacks  
26 occur in the context of the task that performs the reduction and has the type signature  
27 **ompt\_callback\_sync\_region\_t**.

## Restrictions

The restrictions common to reduction clauses are as follows:

- Any number of reduction clauses can be specified on the directive, but a list item (or any array element in an array section) can appear only once in reduction clauses for that directive.
- For a *reduction-identifier* declared with the **declare reduction** construct, the directive must appear before its use in a reduction clause.
- If a list item is an array section or an array element, its base expression must be a base language identifier.
- If a list item is an array section, it must specify contiguous storage and it cannot be a zero-length array section.
- If a list item is an array section or an array element, accesses to the elements of the array outside the specified array section or array element result in unspecified behavior.

C

- A variable that is part of another variable, with the exception of array elements, cannot appear in a reduction clause.

C

C++

- A variable that is part of another variable, with the exception of array elements, cannot appear in a reduction clause except if the reduction 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++

C / C++

- The type of a list item that appears in a reduction clause must be valid for the *reduction-identifier*. For a **max** or **min** reduction in C, the type of the list item must be an allowed arithmetic data type: **char**, **int**, **float**, **double**, or **\_Bool**, possibly modified with **long**, **short**, **signed**, or **unsigned**. For a **max** or **min** reduction in C++, the type of the list item must be an allowed arithmetic data type: **char**, **wchar\_t**, **int**, **float**, **double**, or **bool**, possibly modified with **long**, **short**, **signed**, or **unsigned**.
- A list item that appears in a reduction clause must not be **const**-qualified.
- The *reduction-identifier* for any list item must be unambiguous and accessible.

C / C++

## Fortran

- 1       • A variable that is part of another variable, with the exception of array elements, cannot appear in  
2       a reduction clause.
- 3       • A type parameter inquiry cannot appear in a reduction clause.
- 4       • The type, type parameters and rank of a list item that appears in a reduction clause must be valid  
5       for the *combiner* and *initializer*.
- 6       • A list item that appears in a reduction clause must be definable.
- 7       • A procedure pointer may not appear in a reduction clause.
- 8       • A pointer with the **INTENT (IN)** attribute may not appear in the reduction clause.
- 9       • An original list item with the **POINTER** attribute or any pointer component of an original list  
10       item that is referenced in the *combiner* must be associated at entry to the construct that contains  
11       the reduction clause. Additionally, the list item or the pointer component of the list item must not  
12       be deallocated, allocated, or pointer assigned within the region.
- 13       • An original list item with the **ALLOCATABLE** attribute or any allocatable component of an  
14       original list item that corresponds to the special variable identifier in the *combiner* or the  
15       *initializer* must be in the allocated state at entry to the construct that contains the reduction  
16       clause. Additionally, the list item or the allocatable component of the list item must be neither  
17       deallocated nor allocated, explicitly or implicitly, within the region.
- 18       • If the *reduction-identifier* is defined in a **declare reduction** directive, the  
19       **declare reduction** directive must be in the same subprogram, or accessible by host or use  
20       association.
- 21       • If the *reduction-identifier* is a user-defined operator, the same explicit interface for that operator  
22       must be accessible as at the **declare reduction** directive.
- 23       • If the *reduction-identifier* is defined in a **declare reduction** directive, any subroutine or  
24       function referenced in the initializer clause or combiner expression must be an intrinsic function,  
25       or must have an explicit interface where the same explicit interface is accessible as at the  
26       **declare reduction** directive.

## Fortran

### Cross References

- 27       • `ompt_scope_begin` and `ompt_scope_end`, see Section 4.4.4.11 on page 447.
- 28       • `ompt_sync_region_reduction`, see Section 4.4.4.13 on page 448.
- 29       • `ompt_callback_sync_region_t`, see Section 4.5.2.13 on page 475.
- 30

## 2.22.5.2 Reduction Scoping Clauses

Reduction scoping clauses define the region in which a reduction is computed by tasks or SIMD lanes. All properties common to all reduction clauses, which are defined in Section 2.22.5.1 on page 295, apply to reduction scoping clauses.

The number of copies created for each list item and the time at which those copies are initialized are determined by the particular reduction scoping clause that appears on the construct.

The time at which the original list item contains the result of the reduction is determined by the particular reduction scoping clause.

The location in the OpenMP program at which values are combined and the order in which values are combined are unspecified. Therefore, when comparing sequential and parallel runs, or when comparing one parallel run to another (even if the number of threads used is the same), there is no guarantee that bitwise-identical results will be obtained or that side effects (such as floating-point exceptions) will be identical or take place at the same location in the OpenMP program.

To avoid data races, 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 reduction computation.

## 2.22.5.3 Reduction Participating Clauses

A reduction participating clause specifies a task or a SIMD lane as a participant in a reduction defined by a reduction scoping clause. All properties common to all reduction clauses, which are defined in Section 2.22.5.1 on page 295, apply to reduction participating clauses.

Accesses to the original list item may be replaced by accesses to copies of the original list item created by a region that corresponds to a construct with a reduction scoping clause.

In any case, the final value of the reduction must be determined as if all tasks or SIMD lanes that participate in the reduction are executed sequentially in some arbitrary order.

## 2.22.5.4 `reduction` Clause

### Summary

The `reduction` clause specifies a *reduction-identifier* and one or more list items. For each list item, a private copy is created in each implicit task or SIMD lane and is initialized with the initializer value of the *reduction-identifier*. After the end of the region, the original list item is updated with the values of the private copies using the combiner associated with the *reduction-identifier*.

## Syntax

The syntax of the **reduction** clause is as follows:

```
reduction ([ reduction-modifier, ]reduction-identifier : list)
```

Where *reduction-identifier* is defined in Section 2.22.5.1 on page 295, and *reduction-modifier* is one of the following:

```
inscan  
task  
default
```

## Description

The **reduction** clause is a reduction scoping clause and a reduction participating clause, as described in Section 2.22.5.2 on page 302 and Section 2.22.5.3 on page 302.

If *reduction-modifier* is not present or the **default** *reduction-modifier* is present, the behavior is as follows. For **parallel**, **scope** and worksharing constructs, one or more private copies of each list item are created for each implicit task, as if the **private** clause had been used. For the **simd** construct, one or more private copies of each list item are created for each SIMD lane, as if the **private** clause had been used. For the **taskloop** construct, private copies are created according to the rules of the reduction scoping clauses. For the **teams** construct, one or more private copies of each list item are created for the initial task of each team in the league, as if the **private** clause had been used. For the **loop** construct, private copies are created and used in the construct according to the description and restrictions in Section 2.22.3 on page 281. At the end of a region that corresponds to a construct for which the **reduction** clause was specified, the original list item is updated by combining its original value with the final value of each of the private copies, using the combiner of the specified *reduction-identifier*.

If the **inscan** *reduction-modifier* is present, a scan computation is performed over updates to the list item performed in each logical iteration of the loop associated with the worksharing-loop, worksharing-loop SIMD, or **simd** construct (see Section 2.12.6 on page 137). The list items are privatized in the construct according to the description and restrictions in Section 2.22.3 on page 281. At the end of the region, each original list item is assigned the value described in Section 2.12.6 on page 137.

If the **task** *reduction-modifier* is present for a **parallel**, **scope**, or worksharing construct, then each list item is privatized according to the description and restrictions in Section 2.22.3 on page 281, and an unspecified number of additional private copies are created to support task reductions. Any copies associated with the reduction are initialized before they are accessed by the tasks that participate in the reduction, which include all implicit tasks in the corresponding region and all participating explicit tasks that specify an **in\_reduction** clause (see Section 2.22.5.6 on page 306). After the end of the region, the original list item contains the result of the reduction.

If **nowait** is not specified for the construct, the reduction computation will be complete at the end of the construct; however, if the **reduction** clause is used on a construct to which **nowait** is

1 also applied, accesses to the original list item will create a race and, thus, have unspecified effect  
2 unless synchronization ensures that they occur after all threads have executed all of their iterations  
3 or **section** constructs, and the reduction computation has completed and stored the computed  
4 value of that list item. This can most simply be ensured through a barrier synchronization.

## 5 **Restrictions**

6 The restrictions to the **reduction** clause are as follows:

- 7 • All restrictions common to all reduction clauses, which are listed in Section [2.22.5.1](#) on  
8 page [295](#), apply to this clause.
- 9 • A list item that appears in a **reduction** clause of a worksharing construct must be shared in  
10 the **parallel** region to which a corresponding worksharing region binds.
- 11 • A list item that appears in a **reduction** clause of a **scope** construct must be shared in the  
12 **parallel** region to which a corresponding **scope** region binds.
- 13 • If a list item that appears in a **reduction** clause of a worksharing construct, **scope** construct  
14 or **loop** construct for which the corresponding region binds to a parallel region is an array  
15 section or an array element, all threads that participate in the reduction must specify the same  
16 storage location.
- 17 • A list item that appears in a **reduction** clause with the **inscan** *reduction-modifier* must  
18 appear as a list item in an **inclusive** or **exclusive** clause on a **scan** directive enclosed by  
19 the construct.
- 20 • A **reduction** clause without the **inscan** *reduction-modifier* may not appear on a construct  
21 on which a **reduction** clause with the **inscan** *reduction-modifier* appears.
- 22 • A **reduction** clause with the **task** *reduction-modifier* may only appear on a **parallel**  
23 construct, a **scope** construct, a worksharing construct or a combined or composite construct for  
24 which any of the aforementioned constructs is a constituent construct and **simd** or **loop** are not  
25 constituent constructs.
- 26 • A **reduction** clause with the **inscan** *reduction-modifier* may only appear on a  
27 worksharing-loop construct, a **simd** construct or a combined or composite construct for which  
28 any of the aforementioned constructs is a constituent construct and **distribute** is not a  
29 constituent construct.
- 30 • A list item that appears in a **reduction** clause of the innermost enclosing worksharing,  
31 **parallel** or **scope** construct may not be accessed in an explicit task generated by a construct  
32 for which an **in\_reduction** clause over the same list item does not appear.
- 33 • The **task** *reduction-modifier* may not appear in a **reduction** clause if the **nowait** clause is  
34 specified on the same construct.

- 1       • If a list item in a **reduction** clause on a worksharing construct, **scope** construct or **loop**
- 2       construct for which the corresponding region binds to a parallel region has a reference type then
- 3       it must bind to the same object for all threads of the team.
  
- 4       • If a list item in a **reduction** clause on a worksharing construct, **scope** or **loop** construct
- 5       construct for which the corresponding region binds to a parallel region is an array section or an
- 6       array element then the base pointer must point to the same variable for all threads of the team.
  
- 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 default constructor for the
- 9       class type. The number of calls to the default constructor while performing the scan computation
- 10      is unspecified.
  
- 11      • A variable of class type (or array thereof) that appears in a **reduction** clause with the
- 12      **inscan** *reduction-modifier* requires an accessible, unambiguous copy assignment operator for
- 13      the class type. The number of calls to the copy assignment operator while performing the scan
- 14      computation is unspecified.

## Cross References

- 15      • **scan** directive, see Section [2.12.6](#) on page [137](#).
- 16
- 17      • List Item Privatization, see Section [2.22.3](#) on page [281](#).
- 18
- 18      • **private** clause, see Section [2.22.4.3](#) on page [287](#).

## 2.22.5.5 `task_reduction` Clause

### Summary

The `task_reduction` clause specifies a reduction among tasks.

### Syntax

The syntax of the `task_reduction` clause is as follows:

```
task_reduction (reduction-identifier : list)
```

Where *reduction-identifier* is defined in Section [2.22.5.1](#).

### Description

The `task_reduction` clause is a reduction scoping clause, as described in [2.22.5.2](#).

For each list item, the number of copies is unspecified. Any copies associated with the reduction are initialized before they are accessed by the tasks participating in the reduction. After the end of the region, the original list item contains the result of the reduction.

## Restrictions

The restrictions to the `task_reduction` clause are as follows:

- All restrictions common to all reduction clauses, which are listed in Section 2.22.5.1 on page 295, apply to this clause.

## 2.22.5.6 `in_reduction` Clause

### Summary

The `in_reduction` clause specifies that a task participates in a reduction.

### Syntax

The syntax of the `in_reduction` clause is as follows:

```
in_reduction (reduction-identifier : list)
```

where *reduction-identifier* is defined in Section 2.22.5.1 on page 295.

### Description

The `in_reduction` clause is a reduction participating clause, as described in Section 2.22.5.3 on page 302. For a given a list item, the `in_reduction` clause defines a task to be a participant in a task reduction that is defined by an enclosing region for a matching list item that appears in a `task_reduction` clause or a `reduction` clause with the `task` modifier, where either:

1. The matching list item has the same storage location as the list item in the `in_reduction` clause; or
2. A private copy, derived from the matching list item, that is used to perform the task reduction has the same storage location as the list item in the `in_reduction` clause.

For the `task` construct, the generated task becomes the participating task. For each list item, a private copy may be created as if the `private` clause had been used.

For the `target` construct, the target task becomes the participating task. For each list item, a private copy will be created in the data environment of the target task as if the `private` clause had been used, and this private copy will be implicitly mapped into the device data environment of the target device.

At the end of the task region, if a private copy was created its value is combined with a copy created by a reduction scoping clause or with the original list item.



## Restrictions

The restrictions to the `in_reduction` clause are as follows:

- All restrictions common to all reduction clauses, which are listed in Section 2.22.5.1 on page 295, apply to this clause.
- A list item that appears in a `task_reduction` clause or a `reduction` clause with the `task` modifier that is specified on a construct that corresponds to a region in which the region of the participating task is closely nested must match each list item. The construct that corresponds to the innermost enclosing region that meets this condition must specify the same *reduction-identifier* for the matching list item as the `in_reduction` clause.

## 2.22.5.7 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

C

The syntax of the `declare reduction` directive is as follows:

```
#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

C++

The syntax of the `declare reduction` directive is as follows:

```
#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: `+`, `-`, `*`, `&`, `|`, `^`, `&&` or `||`

- *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

The syntax of the **declare reduction** directive is as follows:

```
!$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 that uses 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

If a type with deferred or assumed length type parameter is specified in a **declare reduction** directive, the *reduction-identifier* of that directive can be used in a reduction clause with any variable of the same type and the same kind parameter, regardless of the length type Fortran parameters with which the variable is declared.

Fortran

The visibility and accessibility of this declaration are the same as those of a variable declared at the same point in the program. The enclosing context of the *combiner* and of the *initializer-expr* is that of the **declare reduction** directive. The *combiner* and the *initializer-expr* must be correct in the base language as if they were the body of a function defined at the same point in the program.

## Fortran

1 If the *reduction-identifier* is the same as the name of a user-defined operator or an extended  
2 operator, or the same as a generic name that is one of the allowed intrinsic procedures, and if the  
3 operator or procedure name appears in an accessibility statement in the same module, the  
4 accessibility of the corresponding **declare reduction** directive is determined by the  
5 accessibility attribute of the statement.

6 If the *reduction-identifier* is the same as a generic name that is one of the allowed intrinsic  
7 procedures and is accessible, and if it has the same name as a derived type in the same module, the  
8 accessibility of the corresponding **declare reduction** directive is determined by the  
9 accessibility of the generic name according to the base language.

## Fortran

### C++

10 The **declare reduction** directive can also appear at points in the program at which a static  
11 data member could be declared. In this case, the visibility and accessibility of the declaration are  
12 the same as those of a static data member declared at the same point in the program.

### C++

13 The *combiner* specifies how partial results can be combined into a single value. The *combiner* can  
14 use the special variable identifiers **omp\_in** and **omp\_out** that are of the type of the variables that  
15 this *reduction-identifier* reduces. Each of them will denote one of the values to be combined before  
16 executing the *combiner*. The special **omp\_out** identifier refers to the storage that holds the  
17 resulting combined value after executing the *combiner*.

18 The number of times that the *combiner* is executed, and the order of these executions, for any  
19 reduction clause is unspecified.

## Fortran

20 If the *combiner* is a subroutine name with an argument list, the *combiner* is evaluated by calling the  
21 subroutine with the specified argument list.

22 If the *combiner* is an assignment statement, the *combiner* is evaluated by executing the assignment  
23 statement.

24 If a generic name is used in the *combiner* expression and the list item in the corresponding  
25 **reduction** clause is an array or array section, it is resolved to the specific procedure that is  
26 elemental or has dummy arguments that are only scalars.

## Fortran

1 As the *initializer-expr* value of a user-defined reduction is not known *a priori* the *initializer-clause*  
2 can be used to specify one. Then the contents of the *initializer-clause* will be used as the initializer  
3 for private copies of reduction list items where the **omp\_priv** identifier will refer to the storage to  
4 be initialized. The special identifier **omp\_orig** can also appear in the *initializer-clause* and it will  
5 refer to the storage of the original variable to be reduced.

6 The number of times that the *initializer-expr* is evaluated, and the order of these evaluations, is  
7 unspecified.

▼ C / C++ ▼

8 If the *initializer-expr* is a function name with an argument list, the *initializer-expr* is evaluated by  
9 calling the function with the specified argument list. Otherwise, the *initializer-expr* specifies how  
10 **omp\_priv** is declared and initialized.

▲ C / C++ ▲

▼ C ▼

11 If no *initializer-clause* is specified, the private variables will be initialized following the rules for  
12 initialization of objects with static storage duration.

▲ C ▲

▼ C++ ▼

13 If no *initializer-expr* is specified, the private variables will be initialized following the rules for  
14 *default-initialization*.

▲ C++ ▲

▼ Fortran ▼

15 If the *initializer-expr* is a subroutine name with an argument list, the *initializer-expr* is evaluated by  
16 calling the subroutine with the specified argument list.

17 If the *initializer-expr* is an assignment statement, the *initializer-expr* is evaluated by executing the  
18 assignment statement.

19 If no *initializer-clause* is specified, the private variables will be initialized as follows:

- 20 • For **complex**, **real**, or **integer** types, the value 0 will be used.
- 21 • For **logical** types, the value **.false.** will be used.
- 22 • For derived types for which default initialization is specified, default initialization will be used.
- 23 • Otherwise, not specifying an *initializer-clause* results in unspecified behavior.

▲ Fortran ▲

▼ C / C++ ▼

24 If *reduction-identifier* is used in a **target** region then a declare target directive must be specified  
25 for any function that can be accessed through the *combiner* and *initializer-expr*.

▲ C / C++ ▲

## Fortran

1 If *reduction-identifier* is used in a **target** region then a **declare target** directive must be  
2 specified for any function or subroutine that can be accessed through the *combiner* and  
3 *initializer-expr*.

## Fortran

### Restrictions

- The only variables allowed in the *combiner* are **omp\_in** and **omp\_out**.
- The only variables allowed in the *initializer-clause* are **omp\_priv** and **omp\_orig**.
- If the variable **omp\_orig** is modified in the *initializer-clause*, the behavior is unspecified.
- If execution of the *combiner* or the *initializer-expr* results in the execution of an OpenMP construct or an OpenMP API call, then the behavior is unspecified.
- A *reduction-identifier* may not be re-declared in the current scope for the same type or for a type that is compatible according to the base language rules.
- At most one *initializer-clause* can be specified.
- The *typename-list* must not declare new types.

## C / C++

- 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++

### C

- 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

### C++

- 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++

## Fortran

- Any selectors in the designator of **omp\_in** and **omp\_out** can only be *component selectors*.
- If the *initializer-expr* is a subroutine name with an argument list, then one of the arguments must be **omp\_priv**.
- 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.
- 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.
- Any subroutine or function used in the **initializer** clause or *combiner* expression must be an intrinsic function, or must have an accessible interface.
- Any user-defined operator, defined assignment or extended operator used in the **initializer** clause or *combiner* expression must have an accessible interface.
- 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.
- 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.
- Any procedure used in the **initializer** clause or *combiner* expression must either be elemental or have dummy arguments that are scalar if the list item in the corresponding **reduction** clause is an array or array section.

## Fortran

### Cross References

- Properties Common To All Reduction Clauses, see Section [2.22.5.1](#) on page [295](#).

## 2.22.6 Data Copying Clauses

This section describes the **copyin** clause (allowed on the **parallel** construct and combined parallel worksharing constructs) and the **copyprivate** clause (allowed on the **single** construct).

These clauses support the copying of data values from private or threadprivate variables on one implicit task or thread to the corresponding variables on other implicit tasks or threads in the team.

The clauses accept a comma-separated list of list items (see Section 2.1 on page 36). All list items appearing in a clause must be visible, according to the scoping rules of the base language. Clauses may be repeated as needed, but a list item that specifies a given variable may not appear in more than one clause on the same directive.

### Fortran

An associate name preserves the association with the selector established at the **ASSOCIATE** statement. A list item that appears in a data copying clause may be a selector of an **ASSOCIATE** construct. If the construct association is established prior to a parallel region, the association between the associate name and the original list item will be retained in the region.

### Fortran

### 2.22.6.1 copyin Clause

#### Summary

The **copyin** clause provides a mechanism to copy the value of a threadprivate variable of the master thread to the threadprivate variable of each other member of the team that is executing the **parallel** 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 array of the master thread to the corresponding element of the array of the other thread.

### 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

1 The copy is done, as if by assignment, after the team is formed and prior to the start of execution of  
2 the associated structured block.

3 On entry to any **parallel** region, each thread's copy of a variable that is affected by a **copyin**  
4 clause for the **parallel** region will acquire the type parameters, allocation, association, and  
5 definition status of the copy of the master thread, according to the following rules:

- 6 • If the original list item has the **POINTER** attribute, each copy receives the same association  
7 status as that of the copy of the master thread as if by pointer assignment.
- 8 • If the original list item does not have the **POINTER** attribute, each copy becomes defined with  
9 the value of the copy of the master thread as if by intrinsic assignment unless the list item has a  
10 type bound procedure as a defined assignment. If the original list item that does not have the  
11 **POINTER** attribute has the allocation status of unallocated, each copy will have the same status.
- 12 • If the original list item is unallocated or unassociated, the copy of the other thread inherits the  
13 declared type parameters and the default type parameter values from the original list item.

## Fortran

### Restrictions

14 The restrictions to the **copyin** clause are as follows:

## C / C++

- 16 • A list item that appears in a **copyin** clause must be **threadprivate**.
- 17 • A variable of class type (or array thereof) that appears in a **copyin** clause requires an  
18 accessible, unambiguous copy assignment operator for the class type.

## C / C++

## Fortran

- 19 • A list item that appears in a **copyin** clause must be **threadprivate**. Named variables that appear  
20 in a **threadprivate** common block may be specified: it is not necessary to specify the whole  
21 common block.
- 22 • A common block name that appears in a **copyin** clause must be declared to be a common block  
23 in the same scoping unit in which the **copyin** clause appears.
- 24 • If the list item is a polymorphic variable with the **ALLOCATABLE** attribute, the behavior is  
25 unspecified.

## Fortran

### Cross References

- 26 • **parallel** construct, see Section 2.8 on page 79.
- 27 • **threadprivate** directive, see Section 2.22.2 on page 276.



## 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 that belong to the `parallel` region.

To avoid data races, 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 96), 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 that belong 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 that is 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++ ▲

▼ 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

1 If a list item does not have the **POINTER** attribute, then in all other implicit tasks that belong to the  
2 **parallel** region, the list item becomes defined as if by intrinsic assignment with the value of the  
3 corresponding list item in the implicit task that is associated with the thread that executed the  
4 structured block. If the list item has a type bound procedure as a defined assignment, the  
5 assignment is performed by the defined assignment.

6 If the list item has the **POINTER** attribute, then, in all other implicit tasks that belong to the  
7 **parallel** region, the list item receives, as if by pointer assignment, the same association status of  
8 the corresponding list item in the implicit task that is associated with the thread that executed the  
9 structured block.

10 The order in which any final subroutines for different variables of a finalizable type are called is  
11 unspecified.

## Fortran

12 **Note** – The **copyprivate** clause is an alternative to using a shared variable for the value when  
13 providing such a shared variable would be difficult (for example, in a recursion requiring a different  
14 variable at each level).

## Restrictions

17 The restrictions to the **copyprivate** clause are as follows:

- 18 • All list items that appear in the **copyprivate** clause must be either **threadprivate** or **private** in  
19 the enclosing context.
- 20 • A list item that appears in a **copyprivate** clause may not appear in a **private** or  
21 **firstprivate** clause on the **single** construct.

## C++

- 22 • A variable of class type (or array thereof) that appears in a **copyprivate** clause requires an  
23 accessible unambiguous copy assignment operator for the class type.

## C++

## Fortran

- 24 • A common block that appears in a **copyprivate** clause must be **threadprivate**.
- 25 • Pointers with the **INTENT (IN)** attribute may not appear in the **copyprivate** clause.
- 26 • The list item with the **ALLOCATABLE** attribute must have the allocation status of **allocated** when  
27 the intrinsic assignment is performed.
- 28 • If the list item is a polymorphic variable with the **ALLOCATABLE** attribute, the behavior is  
29 unspecified.

## Fortran

## Cross References

- **parallel** construct, see Section 2.8 on page 79.
- **threadprivate** directive, see Section 2.22.2 on page 276.
- **private** clause, see Section 2.22.4.3 on page 287.

## 2.22.7 Data-Mapping Attribute Rules, Clauses, and Directives

This section describes how the data-mapping and data-sharing attributes of any variable referenced in a **target** region are determined. When specified, explicit data-sharing attributes, **map** or **is\_device\_ptr** clauses on **target** directives determine these attributes. Otherwise, the first matching rule from the following implicit data-mapping rules applies for variables referenced in a **target** construct that are not declared in the construct and do not appear in data-sharing attribute, **map** or **is\_device\_ptr** clauses.

- If a variable appears in a **to** or **link** clause on a declare target directive that does not have a **device\_type(nohost)** clause then it is treated as if it had appeared in a **map** clause with a *map-type* of **tofrom**.
- If a list item appears in a **reduction**, **lastprivate** or **linear** clause on a combined **target** construct then it is treated as if it also appears in a **map** clause with a *map-type* of **tofrom**.
- If a list item appears in an **in\_reduction** clause on a **target** construct then it is treated as if it also appears in a **map** clause with a *map-type* of **tofrom** and a *map-type-modifier* of **always**.
- If a **defaultmap** clause is present for the category of the variable and specifies an implicit behavior other than **default**, the data-mapping attribute is determined by that clause.

C++

- If the **target** construct is within a class non-static member function, and a variable is an accessible data member of the object for which the non-static data member function is invoked, the variable is treated as if the **this[:1]** expression had appeared in a **map** clause with a *map-type* of **tofrom**. Additionally, if the variable is of a type pointer or reference to pointer, it is also treated as if it has appeared in a **map** clause as a zero-length array section.
- If the **this** keyword is referenced inside a **target** construct within a class non-static member function, it is treated as if the **this[:1]** expression had appeared in a **map** clause with a *map-type* of **tofrom**.

C++

## C / C++

- A variable that is of type pointer, but not a function pointer or (for C++) a pointer to a member function, is treated as if it is the base pointer of a zero-length array section that appeared as a list item in a **map** clause.

## C / C++

### C++

- A variable that is of type reference to pointer, but not a function pointer or a reference to a pointer to a member function is treated as if it had appeared in a **map** clause as a zero-length array section.

### C++

- If a variable is not a scalar then it is treated as if it had appeared in a **map** clause with a *map-type* of **tofrom**.

## Fortran

- If a scalar variable has the **TARGET**, **ALLOCATABLE** or **POINTER** attribute then it is treated as if it has appeared in a **map** clause with a *map-type* of **tofrom**.

## Fortran

- If none of the above rules applies then a scalar variable is not mapped, but instead has an implicit data-sharing attribute of **firstprivate** (see Section 2.22.1.1 on page 271).

### 2.22.7.1 map Clause

#### Summary

The **map** clause specifies how an original list item is mapped from the current task's data environment to a corresponding list item in the device data environment of the device identified by the construct.

#### Syntax

The syntax of the map clause is as follows:

```
map ([ [map-type-modifier[, ] [map-type-modifier[, ] ...] map-type : ] locator-list)
```

where *map-type* is one of the following:

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

1 and *map-type-modifier* is one of the following:

```
2 always  
3 close  
4 mapper (mapper-identifier)  
5 present  
6 iterator (iterators-definition)
```

## 7 Description

8 The list items that appear in a **map** clause may include array sections and structure elements.

9 The *map-type* and *map-type-modifier* specify the effect of the **map** clause, as described below.

10 For a given construct, the effect of a **map** clause with the **to**, **from**, or **tofrom** *map-type* is  
11 ordered before the effect of a **map** clause with the **alloc**, **release**, or **delete** *map-type*. If a  
12 mapper is specified for the type being mapped, or explicitly specified with the **mapper**  
13 *map-type-modifier*, then the effective **map-type** of a list item will be determined according to the  
14 rules of map-type decay. If a **map** clause with a **present** *map-type-modifier* is present in a **map**  
15 clause, then the effect of the clause is ordered before all other **map** clauses that do not have the  
16 **present** modifier. If a **mapper** *map-type-modifier* is present along with a **present**  
17 *map-type-modifier* then all elements of the object are treated as if they had a **present** modifier.

18 If a mapper is specified for the type being mapped, or explicitly specified with the **mapper**  
19 *map-type-modifier*, then all **map** clauses that appear on the **declare mapper** directive are  
20 treated as though they appeared on the construct with the **map** clause. Array sections of a mapper  
21 type are mapped as normal, then each element in the array section is mapped according to the rules  
22 of the mapper.

23 The list items that appear in the **map** clause may reference iterators defined by an  
24 *iterators-definition* appearing on an **iterator** modifier.

25  C / C++

26 If a list item in a **map** clause is a variable of structure type then it is treated as if each structure  
27 element contained in the variable is a list item in the clause.

28  C / C++

29  Fortran

30 If a list item in a **map** clause is a derived type variable then it is treated as if each component is a  
31 list item in the clause.

32 Each pointer component that is a list item that results from a mapped derived type variable is  
33 treated as if its association status is undefined, unless the pointer component appears as another list  
34 item or as the base pointer of another list item in a **map** clause on the same construct.

35  Fortran

1 If a list item in a **map** clause is a structure element then all other structure elements of the  
2 containing structure variable form a *structure sibling list*. The **map** clause and the structure sibling  
3 list are associated with the same construct. If a corresponding list item of the structure sibling list  
4 item is present in the device data environment when the construct is encountered then:

- 5 • If the structure sibling list item does not appear in a **map** clause on the construct then:
  - 6 – If the construct is a **target**, **target data**, or **target enter data** construct then the
  - 7 structure sibling list item is treated as if it is a list item in a **map** clause on the construct with a
  - 8 *map-type* of **alloc**.
  - 9 – If the construct is **target exit data** construct, then the structure sibling list item is treated
  - 10 as if it is a list item in a **map** clause on the construct with a *map-type* of **release**.

### Fortran

- 11 – If the structure sibling list item is a pointer then it is treated as if its association status is
- 12 undefined, unless it appears as the base pointer of another list item in a **map** clause on the
- 13 same construct.

### Fortran

- 14 • If the **map** clause in which the structure element appears as a list item has a *map-type* of
- 15 **delete** and the structure sibling list item does not appear as a list item in a **map** clause on the
- 16 construct with a *map-type* of **delete** then the structure sibling list item is treated as if it is a list
- 17 item in a **map** clause on the construct with a *map-type* of **delete**.

18 If *item<sub>1</sub>* is a list item in a **map** clause, and *item<sub>2</sub>* is another list item in a **map** clause on the same  
19 construct that has a base pointer that is, or is part of, *item<sub>1</sub>*, then:

- 20 • If the **map** clause(s) appear on a **target**, **target data**, or **target enter data** construct,
- 21 then on entry to the corresponding region the effect of the **map** clause on *item<sub>1</sub>* is ordered to
- 22 occur before the effect of the **map** clause on *item<sub>2</sub>*.
- 23 • If the **map** clause(s) appear on a **target**, **target data**, or **target exit data** construct
- 24 then on exit from the corresponding region the effect of the **map** clause on *item<sub>2</sub>* is ordered to
- 25 occur before the effect of the **map** clause on *item<sub>1</sub>*.

### Fortran

26 If a list item in a **map** clause is an associated pointer and the pointer is not the base pointer of  
27 another list item in a **map** clause on the same construct, then it is treated as if its pointer target is  
28 implicitly mapped in the same clause. For the purposes of the **map** clause, the mapped pointer  
29 target is treated as if its base pointer is the associated pointer.

### Fortran

30 If a list item in a **map** clause has a base pointer, and a pointer variable is present in the device data  
31 environment that corresponds to the base pointer when the effect of the **map** clause occurs, then if  
32 the corresponding pointer or the corresponding list item is created in the device data environment  
33 on entry to the construct, then:

## C / C++

1. The corresponding pointer variable is assigned an address such that the corresponding list item can be accessed through the pointer in a **target** region.

## C / C++

## Fortran

1. The corresponding pointer variable is associated with a pointer target that has the same rank and bounds as the pointer target of the original pointer, such that the corresponding list item can be accessed through the pointer in a **target** region.

## Fortran

2. The corresponding pointer variable becomes an attached pointer for the corresponding list item.
3. If the original base pointer and the corresponding attached pointer share storage, then the original list item and the corresponding list item must share storage.

## C++

9 If a *lambda* is mapped explicitly or implicitly, variables that are captured by the *lambda* behave as follows:

- 11 • the variables that are of pointer type are treated as if they had appeared in a **map** clause as zero-length array sections; and
- 13 • the variables that are of reference type are treated as if they had appeared in a **map** clause.

14 If a member variable is captured by a *lambda* in class scope, and the *lambda* is later mapped explicitly or implicitly with its full static type, the **this** pointer is treated as if it had appeared on a **map** clause.

## C++

17 The original and corresponding list items may share storage such that writes to either item by one task followed by a read or write of the other item by another task without intervening synchronization can result in data races. They are guaranteed to share storage if the **map** clause appears on a **target** construct that corresponds to an inactive **target** region, or if it appears on a **target data**, **target enter data**, or **target exit data** construct that applies to the device data environment of the host device.

23 If the **map** clause appears on a **target**, **target data**, **target enter data** or **target exit data** construct with a **present** *map-type-modifier* then on entry to the region if the corresponding list item does not appear in the device data environment an error occurs and the program terminates.

27 If the **map** clause appears on a **target**, **target data**, or **target enter data** construct then on entry to the region the following sequence of steps occurs as if performed as a single atomic operation:

- 30 1. If a corresponding list item of the original list item is not present in the device data environment, then:

- 1 a) A new list item with language-specific attributes is derived from the original list item and
- 2 created in the device data environment;
- 3 b) The new list item becomes the corresponding list item of the original list item in the device
- 4 data environment;
- 5 c) The corresponding list item has a reference count that is initialized to zero; and
- 6 d) The value of the corresponding list item is undefined;
- 7 2. If the corresponding list item's reference count was not already incremented because of the
- 8 effect of a **map** clause on the construct then:
  - 9 a) The corresponding list item's reference count is incremented by one;
- 10 3. If the corresponding list item's reference count is one or the **always map-type-modifier** is
- 11 present, and if the *map-type* is **to** or **tofrom**, then:

C / C++

- 12 a) For each part of the list item that is an attached pointer, that part of the corresponding list
- 13 item will have the value that it had immediately prior to the effect of the **map** clause; and

C / C++

Fortran

- 14 a) For each part of the list item that is an attached pointer, that part of the corresponding list
- 15 item, if associated, will be associated with the same pointer target that it was associated with
- 16 immediately prior to the effect of the **map** clause.

Fortran

- 17 b) For each part of the list item that is not an attached pointer, the value of that part of the
- 18 original list item is assigned to that part of the corresponding list item.

**Note** – If the effect of the **map** clauses on a construct would assign the value of an original list item to a corresponding list item more than once, then an implementation is allowed to ignore additional assignments of the same value to the corresponding list item.

In all cases on entry to the region, concurrent reads or updates of any part of the corresponding list item must be synchronized with any update of the corresponding list item that occurs as a result of the **map** clause to avoid data races.

If the **map** clause appears on a **target**, **target data**, or **target exit data** construct and a corresponding list item of the original list item is not present in the device data environment on exit from the region then the list item is ignored. Alternatively, if the **map** clause appears on a **target**, **target data**, or **target exit data** construct and a corresponding list item of the original list item is present in the device data environment on exit from the region, then the following sequence of steps occurs as if performed as a single atomic operation:



1. If the *map-type* is not **delete** and the corresponding list item's reference count is finite and was not already decremented because of the effect of a **map** clause on the construct then:
  - a) The corresponding list item's reference count is decremented by one;
2. If the *map-type* is **delete** and the corresponding list item's reference count is finite then:
  - a) The corresponding list item's reference count is set to zero;
3. If the *map-type* is **from** or **tofrom** and if the corresponding list item's reference count is zero or the **always** *map-type-modifier* is present then:

C / C++

- a) For each part of the list item that is an attached pointer, that part of the original list item will have the value that it had immediately prior to the effect of the **map** clause;

C / C++

Fortran

- a) For each part of the list item that is an attached pointer, that part of the corresponding list item, if associated, will be associated with the same pointer target with which it was associated immediately prior to the effect of the **map** clause; and

Fortran

- b) For each part of the list item that is not an attached pointer, the value of that part of the corresponding list item is assigned to that part of the original list item; and

4. If the corresponding list item's reference count is zero then the corresponding list item is removed from the device data environment.

**Note** – If the effect of the **map** clauses on a construct would assign the value of a corresponding list item to an original list item more than once, then an implementation is allowed to ignore additional assignments of the same value to the original list item.

In all cases on exit from the region, concurrent reads or updates of any part of the original list item must be synchronized with any update of the original list item that occurs as a result of the **map** clause to avoid data races.

If a single contiguous part of the original storage of a list item with an implicit data-mapping attribute has corresponding storage in the device data environment prior to a task encountering the construct that is associated with the **map** clause, only that part of the original storage will have corresponding storage in the device data environment as a result of the **map** clause.

If a list item with an implicit data-mapping attribute does not have any corresponding storage in the device data environment prior to a task encountering the construct associated with the **map** clause, and one or more contiguous parts of the original storage are either list items or base pointers to list items that are explicitly mapped on the construct, only those parts of the original storage will have

1 corresponding storage in the device data environment as a result of the **map** clauses on the  
2 construct.

▼ C / C++ ▼

3 If a new list item is created then a new list item of the same type, with automatic storage duration, is  
4 allocated for the construct. The size and alignment of the new list item are determined by the static  
5 type of the variable. This allocation occurs if the region references the list item in any statement.  
6 Initialization and assignment of the new list item are through bitwise copy.

▲ C / C++ ▲

▼ Fortran ▼

7 If a new list item is created then a new list item of the same type, type parameter, and rank is  
8 allocated. The new list item inherits all default values for the type parameters from the original list  
9 item. The value of the new list item becomes that of the original list item in the map initialization  
10 and assignment.

11 If the allocation status of the original list item with the **ALLOCATABLE** attribute is changed in the  
12 host device data environment and the corresponding list item is already present in the device data  
13 environment, the allocation status of the corresponding list item is unspecified until a mapping  
14 operation is performed with a **map** clause on entry to a **target**, **target data**, or  
15 **target enter data** region.

▲ Fortran ▲

16 The *map-type* determines how the new list item is initialized.

17 If a *map-type* is not specified, the *map-type* defaults to **tofrom**.

18 The **close map-type-modifier** is a hint to the runtime to allocate memory close to the target device.

## 19 Execution Model Events

20 The *target-map* event occurs when a thread maps data to or from a target device.

21 The *target-data-op* event occurs when a thread initiates a data operation on a target device.

## 22 Tool Callbacks

23 A thread dispatches a registered **ompt\_callback\_target\_map** callback for each occurrence  
24 of a *target-map* event in that thread. The callback occurs in the context of the target task and has  
25 type signature **ompt\_callback\_target\_map\_t**.

26 A thread dispatches a registered **ompt\_callback\_target\_data\_op** callback for each  
27 occurrence of a *target-data-op* event in that thread. The callback occurs in the context of the target  
28 task and has type signature **ompt\_callback\_target\_data\_op\_t**.

## Restrictions

The restrictions to the **map** clause are as follows:

- A list item cannot appear in both a **map** clause and a data-sharing attribute clause on the same construct unless the construct is a combined construct.
  - Each of the *map-type-modifier* modifiers can appear at most once on the **map** clause.
- C / C++
- List items of the **map** clauses on the same construct must not share original storage unless they are the same lvalue expression or array section.
- C / C++
- If a list item is an array section, it must specify contiguous storage.
  - If an expression that is used to form a list item in a **map** clause contains an iterator identifier, the list item instances that would result from different values of the iterator must not have the same containing array or have base pointers that share original storage.
  - If multiple list items are explicitly mapped on the same construct and have the same containing array or have base pointers that share original storage, and if any of the list items do not have corresponding list items that are present in the device data environment prior to a task encountering the construct, then the list items must refer to the same array elements of either the containing array or the implicit array of the base pointers.
  - 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 an array appears as a list item in a **map** clause, multiple parts of the array have corresponding storage in the device data environment prior to a task encountering the construct associated with the **map** clause, and the corresponding storage for those parts were created by maps from more than one earlier construct, the behavior is unspecified.
  - 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.
  - 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++

- 1 • If the type of a list item is a reference to a type *T* then the reference in the device data  
2 environment is initialized to refer to the object in the device data environment that corresponds to  
3 the object referenced by the list item. If mapping occurs, it occurs as though the object were  
4 mapped through a pointer with an array section of type *T* and length one.
- 5 • No type mapped through a reference can contain a reference to its own type, or any references to  
6 types that could produce a cycle of references.
- 7 • If the list item is a *lambda*, any pointers and references captured by the *lambda* must have the  
8 corresponding list item in the device data environment prior to the task encountering the  
9 construct.

## C++

### C / C++

- 10 • A list item cannot be a variable that is a member of a structure with a union type.
- 11 • A bit-field cannot appear in a **map** clause.
- 12 • A pointer that has a corresponding attached pointer must not be modified for the duration of the  
13 lifetime of the list item to which the corresponding pointer is attached in the device data  
14 environment.

### C / C++

### Fortran

- 15 • List items of the **map** clauses on the same construct must not share original storage unless they  
16 are the same variable or array section.
- 17 • A pointer that has a corresponding attached pointer and is associated with a given pointer target  
18 must not become associated with a different pointer target for the duration of the lifetime of the  
19 list item to which the corresponding pointer is attached in the device data environment.
- 20 • If the allocation status of a list item or any subobject of the list item with the **ALLOCATABLE**  
21 attribute is unallocated upon entry to a **target** region, the list item or any subobject of the  
22 corresponding list item must be unallocated upon exit from the region.
- 23 • If the allocation status of a list item or any subobject of the list item with the **ALLOCATABLE**  
24 attribute is allocated upon entry to a **target** region, the allocation status of the corresponding  
25 list item or any subobject of the corresponding list item must not be changed and must not be  
26 reshaped in the region.
- 27 • If an array section is mapped and the size of the section is smaller than that of the whole array,  
28 the behavior of referencing the whole array in the **target** region is unspecified.
- 29 • A list item must not be a whole array of an assumed-size array.
- 30 • If the association status of a list item with the **POINTER** attribute is associated upon entry to a  
31 **target** region, the list item must be associated with the same pointer target upon exit from the  
32 region.

- If the association status of a list item with the **POINTER** attribute is disassociated upon entry to a **target** region, the list item must be disassociated upon exit from the region.
- If the association status of a list item with the **POINTER** attribute is undefined upon entry to a **target** region, the list item must be undefined upon exit from the region.
- If the association status of a list item with the **POINTER** attribute is disassociated or undefined on entry and if the list item is associated with a pointer target inside a **target** region, then the pointer association status must become disassociated before the end of the region.

## Fortran

### Cross References

- Array sections, see Section 2.1.5 on page 42.
- Iterators, see Section 2.1.6 on page 46.
- `ompt_callback_target_data_op_t`, see Section 4.5.2.25 on page 488.
- `ompt_callback_target_map_t`, see Section 4.5.2.27 on page 491.

## C / C++

### 2.22.7.2 Pointer Initialization for Device Data Environments

This section describes how a pointer may be assigned an initial value to point to an object that exists in a device data environment and corresponds to a *matching mapped list item*.

All previously mapped list items that have corresponding storage in a given device data environment constitute the set of currently mapped list items. If a currently mapped list item has a base pointer, the *base address* of the currently mapped list item is the value of its base pointer. Otherwise, the base address is determined by the following steps:

1. Let  $X$  refer to the currently mapped list item.
2. If  $X$  refers to an array section or array element, let  $X$  refer to its base array.
3. If  $X$  refers to a structure element, let  $X$  refer to its containing structure and return to step 2.
4. The base address for the currently mapped list item is the address of  $X$ .

Additionally, each currently mapped list item has a *starting address* and an *ending address*. The starting address is the address of the first storage location associated with the list item, and the *ending address* is the address of the storage location that immediately follows the last storage location associated with the list item.

The *mapped address range* of the currently mapped list item is the range of addresses that starts from the starting address and ends with the ending address. The *extended address range* of the currently mapped list item is the range of addresses that starts from the minimum of the starting address and the base address and that ends with the maximum of the ending address and the base address.

1 If the value of a given pointer is in the mapped address range of a currently mapped list item then  
2 that currently mapped list item is a matching mapped list item. Otherwise, if the value of the  
3 pointer is in the extended address range of a currently mapped list item then that currently mapped  
4 list item is a matching mapped list item.

5 If multiple matching mapped list items are found and they all appear as part of the same containing  
6 structure, the one that has the lowest starting address is treated as the sole matching mapped list  
7 item. Otherwise, if multiple matching mapped list items are found then the behavior is unspecified.

8 If a matching mapped list item is found, the initial value that is assigned to the pointer is a device  
9 address such that the corresponding list item in the device data environment can be accessed  
10 through the pointer in a **target** region.

11 If a matching mapped list item is not found, the assigned initial value of the pointer is NULL.

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

### 12 2.22.7.3 defaultmap Clause

#### 13 Summary

14 The **defaultmap** clause explicitly determines the data-mapping attributes of variables that are  
15 referenced in a **target** construct for which the data-mapping attributes would otherwise be  
16 implicitly determined (see Section 2.22.7 on page 317).

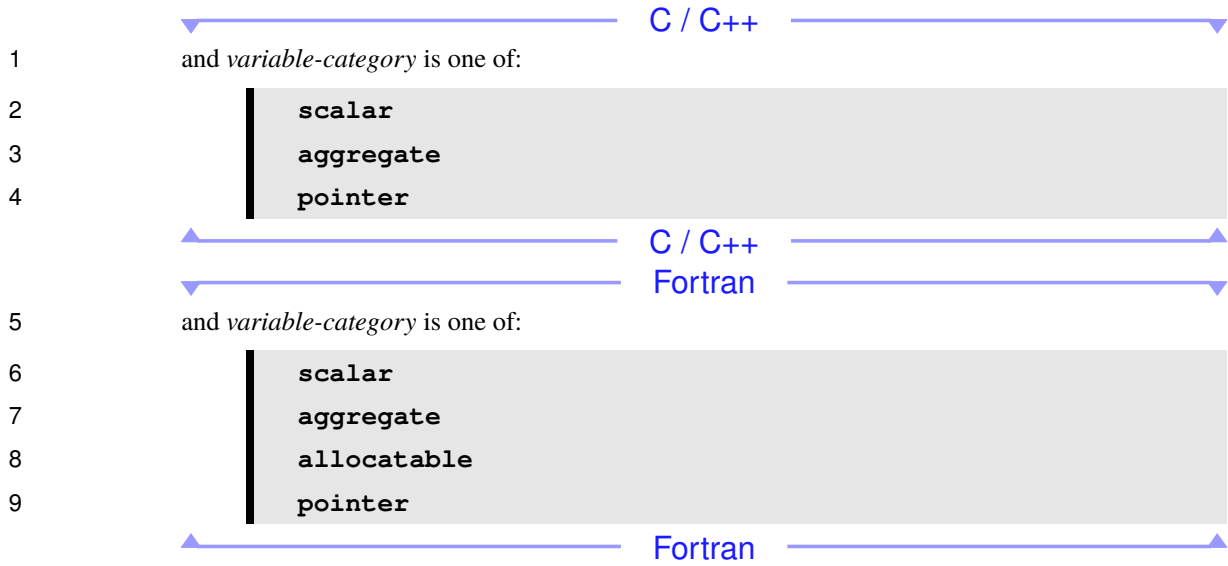
#### 17 Syntax

18 The syntax of the **defaultmap** clause is as follows:

```
19 defaultmap (implicit-behavior[:variable-category])
```

20 Where *implicit-behavior* is one of:

```
21     alloc  
22     to  
23     from  
24     tofrom  
25     firstprivate  
26     none  
27     default  
28     present
```



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 that are otherwise applied to an implicitly mapped pointer are only provided for the  
 23 **default** behavior.

24 If no *variable-category* is specified in the clause then *implicit-behavior* specifies the implicitly  
 25 determined data-mapping or data-sharing attribute for all variables referenced in the construct. If  
 26 *implicit-behavior* is **none**, each variable referenced in the construct that does not have a  
 27 predetermined data-sharing attribute and does not appear in a **to** or **link** clause on a declare  
 28 target directive must be listed in a data-mapping attribute clause, a data-sharing attribute clause  
 29 (including a data-sharing attribute clause on a combined construct where **target** is one of the  
 30 constituent constructs), or an **is\_device\_ptr** clause. If *implicit-behavior* is **default**, then  
 31 the clause has no effect for the variables in the category specified by *variable-category*. If  
 32 *implicit-behavior* is **present**, each variable referenced in the construct in the category specified

1 by *variable-category* is treated as if it had been listed in a **map** clause with the *map-type* of **alloc**  
2 and *map-type-modifier* of **present**.

### 3 2.22.7.4 declare mapper Directive

#### 4 Summary

5 The **declare mapper** directive declares a user-defined mapper for a given type, and may define  
6 a *mapper-identifier* that can be used in a **map** clause. The **declare mapper** directive is a  
7 declarative directive.

#### 8 Syntax

C / C++

9 The syntax of the **declare mapper** directive is as follows:

```
10 #pragma omp declare mapper ([mapper-identifier:] type var) \  
11 [clause[ [, ] clause] ... ] new-line
```

C / C++

Fortran

12 The syntax of the **declare mapper** directive is as follows:

```
13 !$omp declare mapper ([mapper-identifier:] type :: var) &  
14 [clause[ [, ] clause] ... ]
```

Fortran

15 where:

- 16 • *mapper-identifier* is a base-language identifier or **default**
- 17 • *type* is a valid type in scope
- 18 • *var* is a valid base-language identifier
- 19 • *clause* is **map** ([[*map-type-modifier*[ , ] [*map-type-modifier*[ , ] ...]] *map-type*: ] *list*) , where  
20 *map-type* is one of the following:

- 21 – **alloc**
- 22 – **to**
- 23 – **from**
- 24 – **tofrom**

25 and where *map-type-modifier* is one of the following:

- 26 – **always**
- 27 – **close**



## Description

User-defined mappers can be defined using the **declare mapper** directive. The type and the *mapper-identifier* uniquely identify the mapper for use in a **map** clause later in the program. If the *mapper-identifier* is not specified, then **default** is used. The visibility and accessibility of this declaration are the same as those of a variable declared at the same point in the program.

The variable declared by *var* is available for use in all **map** clauses on the directive, and no part of the variable to be mapped is mapped by default.

The default mapper for all types *T*, designated by the pre-defined *mapper-identifier* **default**, is as follows unless a user-defined mapper is specified for that type.

```
declare mapper (T v) map (tofrom: v)
```

Using the **default** *mapper-identifier* overrides the pre-defined default mapper for the given type, making it the default for all variables of *type*. All **map** clauses with this construct in scope that map a list item of *type* will use this mapper unless another is explicitly specified.

All **map** clauses on the directive are expanded into corresponding **map** clauses wherever this mapper is invoked, either by matching type or by being explicitly named in a **map** clause. A **map** clause with list item *var* maps *var* as though no mapper were specified.

C++

The **declare mapper** directive can also appear at points in the program at which a static data member could be declared. In this case, the visibility and accessibility of the declaration are the same as those of a static data member declared at the same point in the program.

C++

## Restrictions

The restrictions to the **declare mapper** directive are as follows:

- No instance of *type* can be mapped as part of the mapper, either directly or indirectly through another type, except the instance passed as the list item. If a set of **declare mapper** directives results in a cyclic definition then the behavior is unspecified.
- The *type* must be of struct, union or class type in C and C++ or a non-intrinsic type in Fortran.
- The *type* must not declare a new type.
- At least one **map** clause that maps *var* or at least one element of *var* is required.
- List-items in **map** clauses on this construct may only refer to the declared variable *var* and entities that could be referenced by a procedure defined at the same location.
- Each *map-type-modifier* can appear at most once on the **map** clause.
- A *mapper-identifier* may not be redeclared in the current scope for the same type or for a type that is compatible according to the base language rules.

- *type* must not be an abstract type.

## 2.23 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, **loop**, **task**, **taskloop**, **critical**, **ordered**, **atomic**, or **master** region.
- A **barrier** region may not be closely nested inside a worksharing, **loop**, **task**, **taskloop**, **critical**, **ordered**, **atomic**, or **master** region.
- A **master** region may not be closely nested inside a worksharing, **loop**, **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**, **loop**, **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.
- OpenMP constructs may not be encountered during execution of an **atomic** region.
- The only OpenMP constructs that can be encountered during execution of a **simd** (or worksharing-loop SIMD) region are the **atomic** construct, the **loop** construct, the **simd** construct and the **ordered** construct with the **simd** clause.
- If a **target update**, **target data**, **target enter data**, or **target exit data** construct is encountered during execution of a **target** region, the behavior is unspecified.
- If a **target** construct is encountered during execution of a **target** region and a **device** clause in which the **ancestor device-modifier** appears is not present on the construct, the behavior is unspecified.

- 1           • A **teams** region must be strictly nested within the implicit parallel region that surrounds the  
2           whole OpenMP program or a **target** region. If a **teams** construct is nested within a **target**  
3           construct, that **target** construct must contain no statements, declarations or directives outside  
4           of the **teams** construct.
- 5           • **distribute**, **distribute simd**, distribute parallel worksharing-loop, distribute parallel  
6           worksharing-loop SIMD, **loop**, **parallel** regions, including any **parallel** regions arising  
7           from combined constructs, **omp\_get\_num\_teams()** regions, and **omp\_get\_team\_num()**  
8           regions are the only OpenMP regions that may be strictly nested inside the **teams** region.
- 9           • The region corresponding to the **distribute** construct must be strictly nested inside a **teams**  
10          region.
- 11          • If *construct-type-clause* is **taskgroup**, the **cancel** construct must be closely nested inside a  
12          **task** construct and the **cancel** region must be closely nested inside a **taskgroup** region. If  
13          *construct-type-clause* is **sections**, the **cancel** construct must be closely nested inside a  
14          **sections** or **section** construct. Otherwise, the **cancel** construct must be closely nested  
15          inside an OpenMP construct that matches the type specified in *construct-type-clause* of the  
16          **cancel** construct.
- 17          • A **cancellation point** construct for which *construct-type-clause* is **taskgroup** must be  
18          closely nested inside a **task** construct, and the **cancellation point** region must be closely  
19          nested inside a **taskgroup** region. A **cancellation point** construct for which  
20          *construct-type-clause* is **sections** must be closely nested inside a **sections** or **section**  
21          construct. Otherwise, a **cancellation point** construct must be closely nested inside an  
22          OpenMP construct that matches the type specified in *construct-type-clause*.
- 23          • The only constructs that may be nested inside a **loop** region are the **loop** construct, the  
24          **parallel** construct, the **simd** construct, and combined constructs for which the first construct  
25          is a **parallel** construct.
- 26          • A **loop** region may not contain calls to procedures that contain OpenMP directives or calls to  
27          the OpenMP Runtime API.
- 28          • A **scope** region may not be closely nested inside a worksharing, **loop**, **task**, **taskloop**,  
29          **critical**, **ordered**, **atomic**, or **master** region.

*This page intentionally left blank*

# 3 Runtime Library Routines

This chapter describes the OpenMP API runtime library routines and queryable runtime states. In this chapter, *true* and *false* are used as generic terms to simplify the description of the routines.

C / C++

*true* means a nonzero integer value and *false* means an integer value of zero.

C / C++

Fortran

*true* means a logical value of **.TRUE.** and *false* means a logical value of **.FALSE.**

Fortran

Fortran

## Restrictions

The following restriction applies to all OpenMP runtime library routines:

- OpenMP runtime library routines may not be called from **PURE** or **ELEMENTAL** procedures.

Fortran

## 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 variable and a definition of each required data type listed below. In addition, each set of definitions may specify other implementation specific values.

## C / C++

1 The library routines are external functions with “C” linkage.

2 Prototypes for the C/C++ runtime library routines described in this chapter shall be provided in a  
3 header file named `omp.h`. This file also defines the following:

- 4 • The type `omp_allocator_handle_t`, which must be an implementation-defined enum type  
5 with at least the `omp_null_allocator` enumerator with the value zero and an enumerator  
6 for each predefined memory allocator in Table 2.10 on page 161;
- 7 • The type `omp_control_tool_result_t`;
- 8 • The type `omp_control_tool_t`;
- 9 • The type `omp_depend_t`;
- 10 • The type `omp_event_handle_t`, which must be an implementation-defined enum type;
- 11 • The type `omp_interop_t`;
- 12 • The type `omp_lock_hint_t` (deprecated);
- 13 • The type `omp_lock_t`;
- 14 • The type `omp_memspace_handle_t`, which must be an implementation-defined enum type  
15 with an enumerator for at least each predefined memory space in Table 2.8 on page 158;
- 16 • The type `omp_nest_lock_t`;
- 17 • The type `omp_pause_resource_t`;
- 18 • The type `omp_proc_bind_t`;
- 19 • The type `omp_sched_t`;
- 20 • The type `omp_sync_hint_t`; and
- 21 • The type `omp_uintptr_t`, which is an unsigned integer type capable of holding a pointer on  
22 any device;

## C / C++

### C++

23 The `omp.h` header file also defines a class template that models the `Allocator` concept in the  
24 `omp::allocator` namespace for each predefined memory allocator in Table 2.10 on page 161  
25 for which the name includes neither the `omp_` prefix nor the `_alloc` suffix.

### C++

▼----- Fortran -----▼

1 The OpenMP Fortran API runtime library routines are external procedures. The return values of  
2 these routines are of default kind, unless otherwise specified.

3 Interface declarations for the OpenMP Fortran runtime library routines described in this chapter  
4 shall be provided in the form of a Fortran **include** file named **omp\_lib.h** or a Fortran 90  
5 **module** named **omp\_lib**. It is implementation defined whether the **include** file or the  
6 **module** file (or both) is provided.

7 These files also define the following:

- 8 • The **integer parameter** **omp\_allocator\_handle\_kind**;
- 9 • An **integer parameter** of kind **omp\_allocator\_handle\_kind** for each predefined  
10 memory allocator in Table 2.10 on page 161;
- 11 • The **integer parameter** **omp\_alloctrail\_key\_kind**;
- 12 • The **integer parameter** **omp\_alloctrail\_val\_kind**;
- 13 • The **integer parameter** **omp\_control\_tool\_kind**;
- 14 • The **integer parameter** **omp\_control\_tool\_result\_kind**;
- 15 • The **integer parameter** **omp\_depend\_kind**;
- 16 • The **integer parameter** **omp\_event\_handle\_kind**;
- 17 • The **integer parameter** **omp\_lock\_hint\_kind** (deprecated);
- 18 • The **integer parameter** **omp\_lock\_kind**;
- 19 • The **integer parameter** **omp\_memspace\_handle\_kind**;
- 20 • An **integer parameter** of kind **omp\_memspace\_handle\_kind** for each predefined  
21 memory space in Table 2.8 on page 158;
- 22 • The **integer parameter** **omp\_nest\_lock\_kind**;
- 23 • The **integer parameter** **omp\_pause\_resource\_kind**;
- 24 • The **integer parameter** **omp\_proc\_bind\_kind**;
- 25 • The **integer parameter** **omp\_sched\_kind**;
- 26 • The **integer parameter** **omp\_sync\_hint\_kind**; and
- 27 • The **integer parameter** **openmp\_version** with a value *yyyymm* where *yyyy* and *mm* are  
28 the year and month designations of the version of the OpenMP Fortran API that the  
29 implementation supports; this value matches that of the C preprocessor macro **\_OPENMP**, when  
30 a macro preprocessor is supported (see Section 2.2 on page 48).

1 It is implementation defined whether any of the OpenMP runtime library routines that take an  
2 argument are extended with a generic interface so arguments of different **KIND** type can be  
3 accommodated.

Fortran

## 4 3.2 Execution Environment Routines

5 This section describes routines that affect and monitor threads, processors, and the parallel  
6 environment.

### 7 3.2.1 `omp_set_num_threads`

#### 8 Summary

9 The `omp_set_num_threads` routine affects the number of threads to be used for subsequent  
10 parallel regions that do not specify a `num_threads` clause, by setting the value of the first  
11 element of the *nthreads-var* ICV of the current task.

#### 12 Format

C / C++

```
13 void omp_set_num_threads(int num_threads);
```

C / C++

Fortran

```
14 subroutine omp_set_num_threads(num_threads)  
15 integer num_threads
```

Fortran

#### 16 Constraints on Arguments

17 The value of the argument passed to this routine must evaluate to a positive integer, or else the  
18 behavior of this routine is implementation defined.

#### 19 Binding

20 The binding task set for an `omp_set_num_threads` region is the generating task.

#### 21 Effect

22 The effect of this routine is to set the value of the first element of the *nthreads-var* ICV of the  
23 current task to the value specified in the argument.



## Cross References









- *nthreads-var* ICV, see Section 2.6 on page 66.
- **parallel** construct and **num\_threads** clause, see Section 2.8 on page 79.
- Determining the number of threads for a **parallel** region, see Section 2.8.1 on page 84.
- **omp\_get\_num\_threads** routine, see Section 3.2.2 on page 339.
- **omp\_get\_max\_threads** routine, see Section 3.2.3 on page 340.
- **OMP\_NUM\_THREADS** environment variable, see Section 6.2 on page 586.

## 3.2.2 omp\_get\_num\_threads

### Summary

The **omp\_get\_num\_threads** routine returns the number of threads in the current team.

### Format

		C / C++	
12	<code>int omp_get_num_threads(void);</code>		
		C / C++	
		Fortran	
13	<code>integer function omp_get_num_threads()</code>		
		Fortran	

### 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 that is 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

- *nthreads-var* ICV, see Section 2.6 on page 66.
- **parallel** construct and **num\_threads** clause, see Section 2.8 on page 79.
- Determining the number of threads for a **parallel** region, see Section 2.8.1 on page 84.
- **omp\_set\_num\_threads** routine, see Section 3.2.1 on page 338.
- **OMP\_NUM\_THREADS** environment variable, see Section 6.2 on page 586.

## 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++	▼
<code>int omp_get_max_threads(void);</code>	
C / C++	▲
Fortran	▼
<code>integer function omp_get_max_threads()</code>	
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 allocate sufficient storage dynamically for all threads in the team formed at the subsequent active `parallel` region.

### Cross References

- `nthreads-var` ICV, see Section 2.6 on page 66.
- `parallel` construct and `num_threads` clause, see Section 2.8 on page 79.
- Determining the number of threads for a `parallel` region, see Section 2.8.1 on page 84.
- `omp_set_num_threads` routine, see Section 3.2.1 on page 338.
- `omp_get_num_threads` routine, see Section 3.2.2 on page 339.
- `omp_get_thread_num` routine, see Section 3.2.4 on page 341.
- `OMP_NUM_THREADS` environment variable, see Section 6.2 on page 586.

## 3.2.4 `omp_get_thread_num`

### Summary

The `omp_get_thread_num` routine returns the thread number, within the current team, of the calling thread.

### Format

	C / C++
<code>int omp_get_thread_num(void);</code>	
	C / C++
	Fortran
<code>integer function omp_get_thread_num()</code>	
	Fortran

### Binding

The binding thread set for an `omp_get_thread_num` region is the current team. The binding region for an `omp_get_thread_num` region is the innermost enclosing `parallel` region.

### Effect

The `omp_get_thread_num` routine returns the thread number of the calling thread, within the team that is executing the `parallel` region to which the routine region binds. The thread number is an integer between 0 and one less than the value returned by `omp_get_num_threads`, inclusive. The thread number of the master thread of the team is 0. The routine returns 0 if it is called from the sequential part of a program.

Note – The thread number may change during the execution of an untied task. The value returned by `omp_get_thread_num` is not generally useful during the execution of such a task region.

### Cross References

- `nthreads-var` ICV, see Section 2.6 on page 66.
- `parallel` construct and `num_threads` clause, see Section 2.8 on page 79.
- Determining the number of threads for a `parallel` region, see Section 2.8.1 on page 84.
- `omp_set_num_threads` routine, see Section 3.2.1 on page 338.
- `omp_get_num_threads` routine, see Section 3.2.2 on page 339.
- `OMP_NUM_THREADS` environment variable, see Section 6.2 on page 586.

## 3.2.5 `omp_get_num_procs`

### Summary

The `omp_get_num_procs` routine returns the number of processors available to the device.

### Format

	C / C++
<code>int omp_get_num_procs(void);</code>	
	C / C++
	Fortran
<code>integer function omp_get_num_procs()</code>	
	Fortran

### Binding

The binding thread set for an `omp_get_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_num_procs` routine returns the number of processors that are available to the device at the time the routine is called. This value may change between the time that it is determined by the `omp_get_num_procs` routine and the time that it is read in the calling context due to system actions outside the control of the OpenMP implementation.

### Cross References

- `omp_get_num_places` routine, see Section [3.2.24](#) on page [359](#).
- `omp_get_place_num_procs` routine, see Section [3.2.25](#) on page [360](#).
- `omp_get_place_proc_ids` routine, see Section [3.2.26](#) on page [360](#).
- `omp_get_place_num` routine, see Section [3.2.27](#) on page [361](#).

## 3.2.6 `omp_in_parallel`

### Summary

The `omp_in_parallel` routine returns *true* if the *active-levels-var* ICV is greater than zero; otherwise, it returns *false*.

```

1      Format
2      int omp_in_parallel(void);
3      logical function omp_in_parallel()

```

C / C++  
C / C++  
Fortran  
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 2.6 on page 66.
- 12 • `parallel` construct, see Section 2.8 on page 79.
- 13 • `omp_get_num_threads` routine, see Section 3.2.2 on page 339.
- 14 • `omp_get_active_level` routine, see Section 3.2.21 on page 356.

### 15 **3.2.7 omp\_set\_dynamic**

#### 16 **Summary**

17 The `omp_set_dynamic` routine enables or disables dynamic adjustment of the number of  
18 threads available for the execution of subsequent `parallel` regions by setting the value of the  
19 *dyn-var* ICV.

```

20     Format
21     void omp_set_dynamic(int dynamic_threads);
22     subroutine omp_set_dynamic(dynamic_threads)
23     logical dynamic_threads

```

C / C++  
C / C++  
Fortran  
Fortran

## 1 Binding

2 The binding task set for an `omp_set_dynamic` region is the generating task.

## 3 Effect

4 For implementations that support dynamic adjustment of the number of threads, if the argument to  
5 `omp_set_dynamic` evaluates to *true*, dynamic adjustment is enabled for the current task;  
6 otherwise, dynamic adjustment is disabled for the current task. For implementations that do not  
7 support dynamic adjustment of the number of threads, this routine has no effect: the value of  
8 *dyn-var* remains *false*.

## 9 Cross References

- 10 • *dyn-var* ICV, see Section 2.6 on page 66.
- 11 • Determining the number of threads for a `parallel` region, see Section 2.8.1 on page 84.
- 12 • `omp_get_num_threads` routine, see Section 3.2.2 on page 339.
- 13 • `omp_get_dynamic` routine, see Section 3.2.8 on page 344.
- 14 • `OMP_DYNAMIC` environment variable, see Section 6.3 on page 587.

## 15 3.2.8 omp\_get\_dynamic

### 16 Summary

17 The `omp_get_dynamic` routine returns the value of the *dyn-var* ICV, which determines whether  
18 dynamic adjustment of the number of threads is enabled or disabled.

### 19 Format

20	<code>int omp_get_dynamic(void);</code>	C / C++
		C / C++
		Fortran
21	<code>logical function omp_get_dynamic()</code>	Fortran

### 22 Binding

23 The binding task set for an `omp_get_dynamic` region is the generating task.

### 24 Effect

25 This routine returns *true* if dynamic adjustment of the number of threads is enabled for the current  
26 task; it returns *false*, otherwise. If an implementation does not support dynamic adjustment of the  
27 number of threads, then this routine always returns *false*.

## Cross References

- *dyn-var* ICV, see Section 2.6 on page 66.
- Determining the number of threads for a **parallel** region, see Section 2.8.1 on page 84.
- **omp\_set\_dynamic** routine, see Section 3.2.7 on page 343.
- **OMP\_DYNAMIC** environment variable, see Section 6.3 on page 587.

## 3.2.9 omp\_get\_cancellation

### Summary

The **omp\_get\_cancellation** routine returns the value of the *cancel-var* ICV, which determines if cancellation is enabled or disabled.

### Format

		C / C++	
<b>int</b>	<b>omp_get_cancellation</b>	( <b>void</b> )	<b>;</b>
		C / C++	
		Fortran	
<b>logical function</b>	<b>omp_get_cancellation</b>	( <b>)</b>	
		Fortran	

### Binding

The binding task set for an **omp\_get\_cancellation** region is the whole program.

### Effect

This routine returns *true* if cancellation is enabled. It returns *false* otherwise.

## Cross References

- *cancel-var* ICV, see Section 2.6.1 on page 66.
- **cancel** construct, see Section 2.21.1 on page 265.
- **OMP\_CANCELLATION** environment variable, see Section 6.11 on page 594.

## 3.2.10 omp\_set\_nested

### Summary

The deprecated **omp\_set\_nested** routine enables or disables nested parallelism by setting the *max-active-levels-var* ICV.

## Format

C / C++  
`void omp_set_nested(int nested);`

C / C++  
Fortran

Fortran  
`subroutine omp_set_nested(nested)`  
`logical nested`

## Binding

The binding task set for an `omp_set_nested` region is the generating task.

## Effect

If the argument to `omp_set_nested` evaluates to *true*, the value of the *max-active-levels-var* ICV is set to the number of active levels of parallelism that the implementation supports; otherwise, if the value of *max-active-levels-var* is greater than 1 then it is set to 1. This routine has been deprecated.

## Cross References

- *max-active-levels-var* ICV, see Section 2.6 on page 66.
- Determining the number of threads for a `parallel` region, see Section 2.8.1 on page 84.
- `omp_get_nested` routine, see Section 3.2.11 on page 346.
- `omp_set_max_active_levels` routine, see Section 3.2.16 on page 352.
- `omp_get_max_active_levels` routine, see Section 3.2.17 on page 353.
- `OMP_NESTED` environment variable, see Section 6.9 on page 593.

## 3.2.11 omp\_get\_nested

### Summary

The deprecated `omp_get_nested` routine returns whether nested parallelism is enabled or disabled, according to the value of the *max-active-levels-var* ICV.

## Format

C / C++  
`int omp_get_nested(void);`

C / C++  
Fortran

Fortran  
`logical function omp_get_nested()`



1       **Binding**

2       The binding task set for an **omp\_get\_nested** region is the generating task.

3       **Effect**

4       This routine returns *true* if *max-active-levels-var* is greater than 1 for the current task; it returns  
5       *false*, otherwise. If an implementation does not support nested parallelism, this routine always  
6       returns *false*. This routine has been deprecated.

7       **Cross References**

- 8       • *max-active-levels-var* ICV, see Section 2.6 on page 66.
- 9       • Determining the number of threads for a **parallel** region, see Section 2.8.1 on page 84.
- 10      • **omp\_set\_nested** routine, see Section 3.2.10 on page 345.
- 11      • **omp\_set\_max\_active\_levels** routine, see Section 3.2.16 on page 352.
- 12      • **omp\_get\_max\_active\_levels** routine, see Section 3.2.17 on page 353.
- 13      • **OMP\_NESTED** environment variable, see Section 6.9 on page 593.

14      **3.2.12 omp\_set\_schedule**

15      **Summary**

16      The **omp\_set\_schedule** routine affects the schedule that is applied when **runtime** is used as  
17      schedule kind, by setting the value of the *run-sched-var* ICV.

18      **Format**

```

19      | void omp_set_schedule(omp_sched_t kind, int chunk_size);
20      | subroutine omp_set_schedule(kind, chunk_size)
21      | integer (kind=omp_sched_kind) kind
22      | integer chunk_size

```

C / C++  
C / C++  
Fortran  
Fortran

## Constraints on Arguments

The first argument passed to this routine can be one of the valid OpenMP schedule kinds (except for **runtime**) or any implementation specific schedule. 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, which can be extended with implementation specific values:

C / C++

```
7 typedef enum omp_sched_t {
8     // schedule kinds
9     omp_sched_static = 0x1,
10    omp_sched_dynamic = 0x2,
11    omp_sched_guided = 0x3,
12    omp_sched_auto = 0x4,
13
14    // schedule modifier
15    omp_sched_monotonic = 0x80000000u
16 } omp_sched_t;
```

C / C++

Fortran

```
17 ! schedule kinds
18 integer(kind=omp_sched_kind), &
19     parameter :: omp_sched_static = &
20         int(Z'1', kind=omp_sched_kind)
21 integer(kind=omp_sched_kind), &
22     parameter :: omp_sched_dynamic = &
23         int(Z'2', kind=omp_sched_kind)
24 integer(kind=omp_sched_kind), &
25     parameter :: omp_sched_guided = &
26         int(Z'3', kind=omp_sched_kind)
27 integer(kind=omp_sched_kind), &
28     parameter :: omp_sched_auto = &
29         int(Z'4', kind=omp_sched_kind)
30
31 ! schedule modifier
32 integer(kind=omp_sched_kind), &
33     parameter :: omp_sched_monotonic = &
34         int(Z'80000000', kind=omp_sched_kind)
```

Fortran

## Binding

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 kind that is specified by the first argument *kind*. It can be any of the standard schedule kinds or any other implementation specific one. For the schedule kinds **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 kind **auto** the second argument has no meaning; for implementation specific schedule kinds, the values and associated meanings of the second argument are implementation defined.

Each of the schedule kinds can be combined with the **omp\_sched\_monotonic** modifier by using the + or | operators in C/C++ or the + operator in Fortran. If the schedule kind is combined with the **omp\_sched\_monotonic** modifier, the schedule is modified as if the **monotonic** schedule modifier was specified. Otherwise, the schedule modifier is **nonmonotonic**.

## Cross References

- *run-sched-var* ICV, see Section 2.6 on page 66.
- Determining the schedule of a worksharing-loop, see Section 2.12.2.1 on page 114.
- **omp\_set\_schedule** routine, see Section 3.2.12 on page 347.
- **omp\_get\_schedule** routine, see Section 3.2.13 on page 349.
- **OMP\_SCHEDULE** environment variable, see Section 6.1 on page 585.

## 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
subroutine omp_get_schedule(kind, chunk_size)
integer (kind=omp_sched_kind) kind
integer chunk_size

Fortran
```

### Binding

The binding task set for an **omp\_get\_schedule** region is the generating task.

## Effect

This routine returns the *run-sched-var* ICV in the task to which the routine binds. The first argument *kind* returns the schedule to be used. It can be any of the standard schedule kinds as defined in Section 3.2.12 on page 347, or any implementation specific schedule kind. The second argument *chunk\_size* returns the chunk size to be used, or a value less than 1 if the default chunk size is to be used, if the returned schedule kind is **static**, **dynamic**, or **guided**. The value returned by the second argument is implementation defined for any other schedule kinds.

## Cross References

- *run-sched-var* ICV, see Section 2.6 on page 66.
- Determining the schedule of a worksharing-loop, see Section 2.12.2.1 on page 114.
- `omp_set_schedule` routine, see Section 3.2.12 on page 347.
- `OMP_SCHEDULE` environment variable, see Section 6.1 on page 585.

## 3.2.14 omp\_get\_thread\_limit

### Summary

The `omp_get_thread_limit` routine returns the maximum number of OpenMP threads available to participate in the current contention group.

### Format

	C / C++
<code>int omp_get_thread_limit(void);</code>	
	C / C++
	Fortran
<code>integer function omp_get_thread_limit()</code>	
	Fortran

### Binding

The binding thread set for an `omp_get_thread_limit` 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_thread_limit` routine returns the value of the *thread-limit-var* ICV.

## Cross References

- *thread-limit-var* ICV, see Section 2.6 on page 66.
- `omp_get_num_threads` routine, see Section 3.2.2 on page 339.
- `OMP_NUM_THREADS` environment variable, see Section 6.2 on page 586.
- `OMP_THREAD_LIMIT` environment variable, see Section 6.10 on page 593.

## 3.2.15 `omp_get_supported_active_levels`

### Summary

The `omp_get_supported_active_levels` routine returns the number of active levels of parallelism supported by the implementation.

### Format

	C / C++	
<code>int omp_get_supported_active_levels(void);</code>		
	C / C++	
	Fortran	
<code>integer function omp_get_supported_active_levels()</code>		
	Fortran	

### Binding

The binding task set for an `omp_get_supported_active_levels` region is the generating task.

### Effect

The `omp_get_supported_active_levels` routine returns the number of active levels of parallelism supported by the implementation. The *max-active-levels-var* ICV may not have a value that is greater than this number. The value returned by the `omp_get_supported_active_levels` routine is implementation defined, but it must be greater than 0.

## Cross References

- *max-active-levels-var* ICV, see Section 2.6 on page 66.
- `omp_set_max_active_levels` routine, see Section 3.2.16 on page 352.
- `omp_get_max_active_levels` routine, see Section 3.2.17 on page 353.

## 3.2.16 `omp_set_max_active_levels`

### Summary

The `omp_set_max_active_levels` routine limits the number of nested active parallel regions on the device, by setting the *max-active-levels-var* ICV

### Format

```
void omp_set_max_active_levels(int max_levels);
```

C / C++

```
subroutine omp_set_max_active_levels(max_levels)
integer max_levels
```

Fortran

### Constraints on Arguments

The value of the argument passed to this routine must evaluate to a non-negative integer, otherwise the behavior of this routine is implementation defined.

### 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 `parallel` or `teams` 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 active levels requested exceeds the number of active levels of parallelism supported by the implementation, the value of the *max-active-levels-var* ICV will be set to the number of active 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 a `parallel` or `teams` region, the effect of this routine is implementation defined.

### Cross References

- *max-active-levels-var* ICV, see Section 2.6 on page 66.
- `parallel` construct, see Section 2.8 on page 79.
- `omp_get_supported_active_levels` routine, see Section 3.2.15 on page 351.
- `omp_get_max_active_levels` routine, see Section 3.2.17 on page 353.
- `OMP_MAX_ACTIVE_LEVELS` environment variable, see Section 6.8 on page 592.

## 3.2.17 `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

	C / C++	
<code>int omp_get_max_active_levels(void);</code>		
	C / C++	
	Fortran	
<code>integer function omp_get_max_active_levels()</code>		
	Fortran	

### Binding

When called from a sequential part of the program, the binding thread set for an `omp_get_max_active_levels` region is the encountering thread. When called from within any `parallel` or `teams` region, the binding thread set (and binding region, if required) for the `omp_get_max_active_levels` region is implementation defined.

### Effect

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.

### Cross References

- *max-active-levels-var* ICV, see Section 2.6 on page 66.
- `parallel` construct, see Section 2.8 on page 79.
- `omp_get_supported_active_levels` routine, see Section 3.2.15 on page 351.
- `omp_set_max_active_levels` routine, see Section 3.2.16 on page 352.
- `OMP_MAX_ACTIVE_LEVELS` environment variable, see Section 6.8 on page 592.

## 3.2.18 `omp_get_level`

### Summary

The `omp_get_level` routine returns the value of the *levels-var* ICV.

1

### Format

2

```

int omp_get_level(void);

```

C / C++

3

```

integer function omp_get_level()

```

Fortran

4

### Binding

5

The binding task set for an `omp_get_level` region is the generating task.

6

### Effect

7

The effect of the `omp_get_level` routine is to return the number of nested `parallel` regions (whether active or inactive) that enclose the current task such that all of the `parallel` regions are enclosed by the outermost initial task region on the current device.

9

10

### Cross References

11

- `levels-var` ICV, see Section 2.6 on page 66.

12

- `parallel` construct, see Section 2.8 on page 79.

13

- `omp_get_active_level` routine, see Section 3.2.21 on page 356.

14

- `OMP_MAX_ACTIVE_LEVELS` environment variable, see Section 6.8 on page 592.

15

## 3.2.19 omp\_get\_ancestor\_thread\_num

16

### Summary

17

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.

18

19

### Format

20

```

int omp_get_ancestor_thread_num(int level);

```

C / C++

21

```

integer function omp_get_ancestor_thread_num(level)
integer level

```

Fortran

22



## Binding

The binding thread set for an `omp_get_ancestor_thread_num` region is the encountering thread. The binding region for an `omp_get_ancestor_thread_num` region is the innermost enclosing `parallel` region.

## Effect

The `omp_get_ancestor_thread_num` routine returns the thread number of the ancestor at a given nest level of the current thread or the thread number of the current thread. If the requested nest level is outside the range of 0 and the nest level of the current thread, as returned by the `omp_get_level` routine, the routine returns -1.

Note – When the `omp_get_ancestor_thread_num` routine is called with a value of `level=0`, the routine always returns 0. If `level=omp_get_level()`, the routine has the same effect as the `omp_get_thread_num` routine.

## Cross References

- `parallel` construct, see Section 2.8 on page 79.
- `omp_get_num_threads` routine, see Section 3.2.2 on page 339.
- `omp_get_thread_num` routine, see Section 3.2.4 on page 341.
- `omp_get_level` routine, see Section 3.2.18 on page 353.
- `omp_get_team_size` routine, see Section 3.2.20 on page 355.

## 3.2.20 omp\_get\_team\_size

### Summary

The `omp_get_team_size` routine returns, for a given nested level of the current thread, the size of the thread team to which the ancestor or the current thread belongs.

### Format

```
C / C++
int omp_get_team_size(int level);

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 339.
- `omp_get_level` routine, see Section 3.2.18 on page 353.
- `omp_get_ancestor_thread_num` routine, see Section 3.2.19 on page 354.

## 3.2.21 `omp_get_active_level`

### Summary

The `omp_get_active_level` routine returns the value of the *active-level-var* ICV..

### Format

	C / C++	
<code>int omp_get_active_level(void);</code>		
	C / C++	
	Fortran	
<code>integer function omp_get_active_level()</code>		
	Fortran	

### Binding

The binding task set for the an `omp_get_active_level` region is the generating task.

## Effect

The effect of the `omp_get_active_level` routine is to return the number of nested active `parallel` regions 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

- *active-levels-var* ICV, see Section 2.6 on page 66.
- `omp_set_max_active_levels` routine, see Section 3.2.16 on page 352.
- `omp_get_max_active_levels` routine, see Section 3.2.17 on page 353.
- `omp_get_level` routine, see Section 3.2.18 on page 353.
- `OMP_MAX_ACTIVE_LEVELS` environment variable, see Section 6.8 on page 592.

## 3.2.22 `omp_in_final`

### Summary

The `omp_in_final` routine returns *true* if the routine is executed in a final task region; otherwise, it returns *false*.

### Format

	C / C++
<code>int omp_in_final(void);</code>	
	C / C++
	Fortran
<code>logical function omp_in_final()</code>	
	Fortran

### Binding

The binding task set for an `omp_in_final` region is the generating task.

### Effect

`omp_in_final` returns *true* if the enclosing task region is final. Otherwise, it returns *false*.

## Cross References

- `task` construct, see Section 2.13.1 on page 142.

## 3.2.23 `omp_get_proc_bind`

### Summary

The `omp_get_proc_bind` routine returns the thread affinity policy to be used for the subsequent nested `parallel` regions that do not specify a `proc_bind` clause.

1 **Format**

2 `omp_proc_bind_t omp_get_proc_bind(void);`

C / C++

Fortran

3 `integer (kind=omp_proc_bind_kind) function omp_get_proc_bind()`

Fortran

4 **Constraints on Arguments**

5 The value returned by this routine must be one of the valid affinity policy kinds. The C/C++ header  
6 file (`omp.h`) and the Fortran include file (`omp_lib.h`) and/or Fortran 90 module file (`omp_lib`)  
7 define the valid constants. The valid constants must include the following:

C / C++

8 `typedef enum omp_proc_bind_t {`  
9  `omp_proc_bind_false = 0,`  
10  `omp_proc_bind_true = 1,`  
11  `omp_proc_bind_master = 2,`  
12  `omp_proc_bind_close = 3,`  
13  `omp_proc_bind_spread = 4`  
14 `} omp_proc_bind_t;`

C / C++

Fortran

15 `integer (kind=omp_proc_bind_kind), &`  
16  `parameter :: omp_proc_bind_false = 0`  
17 `integer (kind=omp_proc_bind_kind), &`  
18  `parameter :: omp_proc_bind_true = 1`  
19 `integer (kind=omp_proc_bind_kind), &`  
20  `parameter :: omp_proc_bind_master = 2`  
21 `integer (kind=omp_proc_bind_kind), &`  
22  `parameter :: omp_proc_bind_close = 3`  
23 `integer (kind=omp_proc_bind_kind), &`  
24  `parameter :: omp_proc_bind_spread = 4`

Fortran

25 **Binding**

26 The binding task set for an `omp_get_proc_bind` region is the generating task.

27 **Effect**

28 The effect of this routine is to return the value of the first element of the *bind-var* ICV of the current  
29 task. See Section 2.8.2 on page 85 for the rules that govern the thread affinity policy.

## Cross References

- *bind-var* ICV, see Section 2.6 on page 66.
- Controlling OpenMP thread affinity, see Section 2.8.2 on page 85.
- `omp_get_num_places` routine, see Section 3.2.24 on page 359.
- `OMP_PROC_BIND` environment variable, see Section 6.4 on page 588.
- `OMP_PLACES` environment variable, see Section 6.5 on page 588.

## 3.2.24 `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

	C / C++	
<code>int omp_get_num_places(void);</code>		
	C / C++	
	Fortran	
<code>integer function omp_get_num_places()</code>		
	Fortran	

### 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 2.6 on page 66.
- Controlling OpenMP thread affinity, see Section 2.8.2 on page 85.
- `omp_get_place_num` routine, see Section 3.2.27 on page 361.
- `OMP_PLACES` environment variable, see Section 6.5 on page 588.

## 3.2.25 `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++	
<code>int</code>	<code>omp_get_place_num_procs</code>	<code>(int place_num);</code>
	C / C++	
	Fortran	
<code>integer function</code>	<code>omp_get_place_num_procs</code>	<code>(place_num)</code>
<code>integer</code>		<code>place_num</code>
	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 greater than or equal to the value returned by `omp_get_num_places()`.

### Cross References

- *place-partition-var* ICV, see Section 2.6 on page 66.
- Controlling OpenMP thread affinity, see Section 2.8.2 on page 85.
- `omp_get_num_places` routine, see Section 3.2.24 on page 359.
- `omp_get_place_proc_ids` routine, see Section 3.2.26 on page 360.
- `OMP_PLACES` environment variable, see Section 6.5 on page 588.

## 3.2.26 `omp_get_place_proc_ids`

### Summary

The `omp_get_place_proc_ids` routine returns the numerical identifiers of the processors available to the execution environment in the specified place.

## Format

C / C++  
`void omp_get_place_proc_ids(int place_num, int *ids);`

C / C++

Fortran

Fortran  
`subroutine omp_get_place_proc_ids(place_num, ids)  
integer place_num  
integer ids(*)`

Fortran

## Binding

The binding thread set for an `omp_get_place_proc_ids` 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_proc_ids` routine returns the numerical identifiers of each processor associated with the place numbered `place_num`. The numerical identifiers are non-negative, and their meaning is implementation defined. The numerical identifiers are returned in the array `ids` and their order in the array is implementation defined. The array must be sufficiently large to contain `omp_get_place_num_procs(place_num)` integers; otherwise, the behavior is unspecified. The routine has no effect when `place_num` has a negative value, or a value greater than or equal to `omp_get_num_places()`.

## Cross References

- `place-partition-var` ICV, see Section 2.6 on page 66.
- Controlling OpenMP thread affinity, see Section 2.8.2 on page 85.
- `omp_get_num_places` routine, see Section 3.2.24 on page 359.
- `omp_get_place_num_procs` routine, see Section 3.2.25 on page 360.
- `OMP_PLACES` environment variable, see Section 6.5 on page 588.

## 3.2.27 `omp_get_place_num`

### Summary

The `omp_get_place_num` routine returns the place number of the place to which the encountering thread is bound.

1

### Format

2

	C / C++	
<code>int omp_get_place_num(void);</code>		

	C / C++	
--	---------	--

3

	Fortran	
<code>integer function omp_get_place_num()</code>		

	Fortran	
--	---------	--

4

### Binding

5

The binding thread set for an `omp_get_place_num` region is the encountering thread.

6

### Effect

7

When the encountering thread is bound to a place, the `omp_get_place_num` routine returns the place number associated with the thread. The returned value is between 0 and one less than the value returned by `omp_get_num_places()`, inclusive. When the encountering thread is not bound to a place, the routine returns -1.

10

11

### Cross References

12

- *place-partition-var* ICV, see Section 2.6 on page 66.
- Controlling OpenMP thread affinity, see Section 2.8.2 on page 85.
- `omp_get_num_places` routine, see Section 3.2.24 on page 359.
- `OMP_PLACES` environment variable, see Section 6.5 on page 588.

13

14

15

16

## 3.2.28 `omp_get_partition_num_places`

17

### Summary

18

The `omp_get_partition_num_places` routine returns the number of places in the place partition of the innermost implicit task.

19

20

### Format

21

	C / C++	
<code>int omp_get_partition_num_places(void);</code>		

	C / C++	
--	---------	--

22

	Fortran	
<code>integer function omp_get_partition_num_places()</code>		

	Fortran	
--	---------	--



## Binding

The binding task set for an `omp_get_partition_num_places` region is the encountering implicit task.

## Effect

The `omp_get_partition_num_places` routine returns the number of places in the *place-partition-var* ICV.

## Cross References

- *place-partition-var* ICV, see Section 2.6 on page 66.
- Controlling OpenMP thread affinity, see Section 2.8.2 on page 85.
- `omp_get_num_places` routine, see Section 3.2.24 on page 359.
- `OMP_PLACES` environment variable, see Section 6.5 on page 588.

## 3.2.29 `omp_get_partition_place_nums`

### Summary

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.

### Format

```
void omp_get_partition_place_nums (int *place_nums);
```

C / C++

```
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 that correspond 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 2.6 on page 66.
- Controlling OpenMP thread affinity, see Section 2.8.2 on page 85.
- `omp_get_partition_num_places` routine, see Section 3.2.28 on page 362.
- `OMP_PLACES` environment variable, see Section 6.5 on page 588.

## 3.2.30 `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.

### Format

	C / C++	
<code>void omp_set_affinity_format(const char *format);</code>		
	C / C++	
	Fortran	
<code>subroutine omp_set_affinity_format(format)</code>		
<code>character(len=*) , intent(in) :: format</code>		
	Fortran	

### Binding

When called from a sequential part of the program, the binding thread set for an `omp_set_affinity_format` region is the encountering thread. When called from within any `parallel` or `teams` region, the binding thread set (and binding region, if required) for the `omp_set_affinity_format` region is implementation defined.

### Effect

The effect of `omp_set_affinity_format` routine is to copy the character string specified by the *format* argument into the *affinity-format-var* ICV on the current device.

This routine has the described effect only when called from a sequential part of the program. When called from within a `parallel` or `teams` region, the effect of this routine is implementation defined.

## Cross References

- Controlling OpenMP thread affinity, see Section 2.8.2 on page 85.
- `omp_get_affinity_format` routine, see Section 3.2.31 on page 365.
- `omp_display_affinity` routine, see Section 3.2.32 on page 366.
- `omp_capture_affinity` routine, see Section 3.2.33 on page 367.
- `OMP_DISPLAY_AFFINITY` environment variable, see Section 6.13 on page 595.
- `OMP_AFFINITY_FORMAT` environment variable, see Section 6.14 on page 596.

## 3.2.31 `omp_get_affinity_format`

### Summary

The `omp_get_affinity_format` routine returns the value of the *affinity-format-var* ICV on the device.

### Format

**C / C++**  
`size_t omp_get_affinity_format(char *buffer, size_t size);`

**C / C++**

**Fortran**  
`integer function omp_get_affinity_format(buffer)  
character(len=*), intent(out) :: buffer`

**Fortran**

### Binding

When called from a sequential part of the program, the binding thread set for an `omp_get_affinity_format` region is the encountering thread. When called from within any `parallel` or `teams` region, the binding thread set (and binding region, if required) for the `omp_get_affinity_format` region is implementation defined.

## Effect

C / C++

The `omp_get_affinity_format` routine returns the number of characters in the *affinity-format-var* ICV on the current device, excluding the terminating null byte ('`\0`') and if *size* is non-zero, writes the value of the *affinity-format-var* ICV on the current device to *buffer* followed by a null byte. If the return value is larger or equal to *size*, the affinity format specification is 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

The `omp_get_affinity_format` routine returns the number of characters that are required to hold the *affinity-format-var* ICV on the current device and writes the value of the *affinity-format-var* ICV on the current device to *buffer*. If the return value is larger than `len(buffer)`, the affinity format specification is truncated.

Fortran

If the *buffer* argument does not conform to the specified format then the result is implementation defined.

## Cross References

- Controlling OpenMP thread affinity, see Section 2.8.2 on page 85.
- `omp_set_affinity_format` routine, see Section 3.2.30 on page 364.
- `omp_display_affinity` routine, see Section 3.2.32 on page 366.
- `omp_capture_affinity` routine, see Section 3.2.33 on page 367.
- `OMP_DISPLAY_AFFINITY` environment variable, see Section 6.13 on page 595.
- `OMP_AFFINITY_FORMAT` environment variable, see Section 6.14 on page 596.

## 3.2.32 `omp_display_affinity`

### Summary

The `omp_display_affinity` routine prints the OpenMP thread affinity information using the format specification provided.

## Format

C / C++  
`void omp_display_affinity(const char *format);`

C / C++

Fortran

Fortran  
`subroutine omp_display_affinity(format)  
character(len=*), intent(in) :: format`

Fortran

## Binding

The binding thread set for an `omp_display_affinity` region is the encountering thread.

## Effect

The `omp_display_affinity` routine prints the thread affinity information of the current thread in the format specified by the `format` argument, followed by a *new-line*. 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 `format` argument does not conform to the specified format then the result is implementation defined.

## Cross References

- Controlling OpenMP thread affinity, see Section 2.8.2 on page 85.
- `omp_set_affinity_format` routine, see Section 3.2.30 on page 364.
- `omp_get_affinity_format` routine, see Section 3.2.31 on page 365.
- `omp_capture_affinity` routine, see Section 3.2.33 on page 367.
- `OMP_DISPLAY_AFFINITY` environment variable, see Section 6.13 on page 595.
- `OMP_AFFINITY_FORMAT` environment variable, see Section 6.14 on page 596.

## 3.2.33 `omp_capture_affinity`

### Summary

The `omp_capture_affinity` routine prints the OpenMP thread affinity information into a buffer using the format specification provided.

## Format

C / C++

```
size_t omp_capture_affinity(  
    char *buffer,  
    size_t size,  
    const char *format  
);
```

C / C++

Fortran

```
integer function omp_capture_affinity(buffer,format)  
character(len=*) ,intent(out) :: buffer  
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 a 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 a zero-length string, 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

If the *format* argument does not conform to the specified format then the result is implementation defined.

## Cross References

- Controlling OpenMP thread affinity, see Section 2.8.2 on page 85.
- `omp_set_affinity_format` routine, see Section 3.2.30 on page 364.
- `omp_get_affinity_format` routine, see Section 3.2.31 on page 365.
- `omp_display_affinity` routine, see Section 3.2.32 on page 366.
- `OMP_DISPLAY_AFFINITY` environment variable, see Section 6.13 on page 595.
- `OMP_AFFINITY_FORMAT` environment variable, see Section 6.14 on page 596.

## 3.2.34 `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

```
void omp_set_default_device(int device_num);
```

C / C++

```
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 2.6 on page 66.
- **target** construct, see Section 2.15.5 on page 174.
- `omp_get_default_device`, see Section 3.2.35 on page 370.
- `OMP_DEFAULT_DEVICE` environment variable, see Section 6.15 on page 598.

### 3.2.35 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 2.6 on page 66.
- `target` construct, see Section 2.15.5 on page 174.
- `omp_set_default_device`, see Section 3.2.34 on page 369.
- `OMP_DEFAULT_DEVICE` environment variable, see Section 6.15 on page 598.

### 3.2.36 omp\_get\_num\_devices

#### Summary

The `omp_get_num_devices` routine returns the number of target devices.

#### Format

	C / C++	
<code>int omp_get_num_devices(void);</code>		
	C / C++	
	Fortran	
<code>integer function omp_get_num_devices()</code>		
	Fortran	



1       **Binding**

2       The binding task set for an `omp_get_num_devices` region is the generating task.

3       **Effect**

4       The `omp_get_num_devices` routine returns the number of available target devices. When  
5       called from within a `target` region the effect of this routine is unspecified.

6       **Cross References**

- 7       • `target` construct, see Section 2.15.5 on page 174.
- 8       • `omp_get_default_device`, see Section 3.2.35 on page 370.
- 9       • `omp_get_device_num`, see Section 3.2.37 on page 371.

10       **3.2.37 omp\_get\_device\_num**

11       **Summary**

12       The `omp_get_device_num` routine returns the device number of the device on which the  
13       calling thread is executing.

14       **Format**

15        C / C++  
16        Fortran

17       **Binding**

18       The binding task set for an `omp_get_device_num` region is the generating task.

19       **Effect**

20       The `omp_get_device_num` routine returns the device number of the device on which the  
21       calling thread is executing. When called on the host device, it will return the same value as the  
22       `omp_get_initial_device` routine.

23       **Cross References**

- 24       • `target` construct, see Section 2.15.5 on page 174.
- 25       • `omp_get_default_device`, see Section 3.2.35 on page 370.
- 26       • `omp_get_num_devices`, see Section 3.2.36 on page 370.
- 27       • `omp_get_initial_device` routine, see Section 3.2.41 on page 374.

### 3.2.38 omp\_get\_num\_teams

#### Summary

The `omp_get_num_teams` routine returns the number of initial teams in the current `teams` region.

#### Format

	C / C++
<code>int omp_get_num_teams(void);</code>	
	C / C++
	Fortran
<code>integer function omp_get_num_teams()</code>	
	Fortran

#### 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 2.9 on page 87.
- `target` construct, see Section 2.15.5 on page 174.
- `omp_get_team_num` routine, see Section 3.2.39 on page 372.

### 3.2.39 omp\_get\_team\_num

#### Summary

The `omp_get_team_num` routine returns the initial team number of the calling thread.

#### Format

	C / C++
<code>int omp_get_team_num(void);</code>	
	C / C++
	Fortran
<code>integer function omp_get_team_num()</code>	
	Fortran

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**

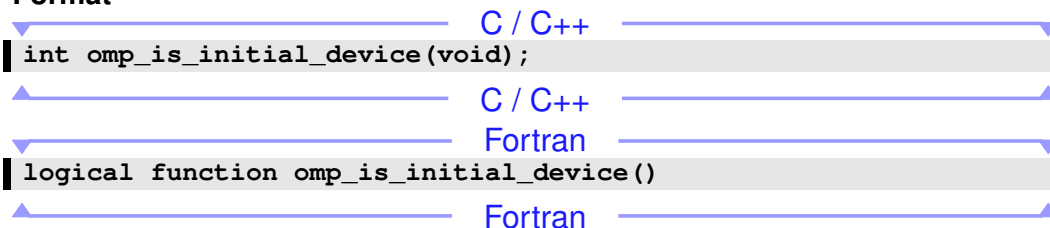
- 9       • `teams` construct, see Section 2.9 on page 87.
- 10      • `target` construct, see Section 2.15.5 on page 174.
- 11      • `omp_get_num_teams` routine, see Section 3.2.38 on page 372.

12      **3.2.40 omp\_is\_initial\_device**

13      **Summary**

14      The `omp_is_initial_device` routine returns *true* if the current task is executing on the host  
15      device; otherwise, it returns *false*.

16      **Format**

17      

```
int omp_is_initial_device(void);
logical function omp_is_initial_device()
```

19      **Binding**

20      The binding task set for an `omp_is_initial_device` region is the generating task.

21      **Effect**

22      The effect of this routine is to return *true* if the current task is executing on the host device;  
23      otherwise, it returns *false*.

24      **Cross References**

- 25      • `omp_get_initial_device` routine, see Section 3.2.41 on page 374.
- 26      • Device memory routines, see Section 3.6 on page 396.

### 3.2.41 `omp_get_initial_device`

#### Summary

The `omp_get_initial_device` routine returns a device number that represents the host device.

#### Format

	C / C++	
<code>int omp_get_initial_device(void);</code>		
	C / C++	
	Fortran	
<code>integer function omp_get_initial_device()</code>		
	Fortran	

#### Binding

The binding task set for an `omp_get_initial_device` region is the generating task.

#### Effect

The effect of this routine is to return the device number of the host device. The value of the device number is implementation defined. 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 174.
- `omp_is_initial_device` routine, see Section 3.2.40 on page 373.
- Device memory routines, see Section 3.6 on page 396.

### 3.2.42 `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++	
<code>int omp_get_max_task_priority(void);</code>		
	C / C++	
	Fortran	
<code>integer function omp_get_max_task_priority()</code>		
	Fortran	

1       **Binding**

2       The binding thread set for an `omp_get_max_task_priority` region is all threads on the  
3       device. The effect of executing this routine is not related to any specific region that corresponds to  
4       any construct or API routine.

5       **Effect**

6       The `omp_get_max_task_priority` routine returns the value of the *max-task-priority-var*  
7       ICV, which determines the maximum value that can be specified in the **priority** clause.

8       **Cross References**

- 9       • *max-task-priority-var*, see Section 2.6 on page 66.  
10      • **task** construct, see Section 2.13.1 on page 142.

11      **3.2.43 omp\_pause\_resource**

12      **Summary**

13      The `omp_pause_resource` routine allows the runtime to relinquish resources used by OpenMP  
14      on the specified device.

15      **Format**

C / C++

```
16      int omp_pause_resource(  
17          omp_pause_resource_t kind,  
18          int device_num  
19      );
```

C / C++

Fortran

```
20      integer function omp_pause_resource(kind, device_num)  
21      integer (kind=omp_pause_resource_kind) kind  
22      integer device_num
```

Fortran

## Constraints on Arguments

The first argument passed to this routine can be one of the valid OpenMP pause kind, or any implementation specific pause kind. 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, which can be extended with implementation specific values:

### Format

C / C++

```
typedef enum omp_pause_resource_t {  
    omp_pause_soft = 1,  
    omp_pause_hard = 2  
} omp_pause_resource_t;
```

C / C++

Fortran

```
integer (kind=omp_pause_resource_kind), parameter :: &  
    omp_pause_soft = 1  
integer (kind=omp_pause_resource_kind), parameter :: &  
    omp_pause_hard = 2
```

Fortran

The second argument passed to this routine indicates the device that will be paused. The `device_num` parameter must be greater than or equal to zero and less than the result of `omp_get_num_devices()` or equal to the result of `omp_get_initial_device()`.

### Binding

The binding task set for an `omp_pause_resource` region is the whole program.

### Effect

The `omp_pause_resource` routine allows the runtime to relinquish resources used by OpenMP on the specified device.

If successful, the `omp_pause_hard` value results in a hard pause for which the OpenMP state is not guaranteed to persist across the `omp_pause_resource` call. A hard pause may relinquish any data allocated by OpenMP on a given device, including data allocated by memory routines for that device as well as data present on the device as a result of a declare target directive or `target data` construct. A hard pause may also relinquish any data associated with a `threadprivate` directive. When relinquished and when applicable, base language appropriate deallocation/finalization is performed. When relinquished and when applicable, mapped data on a device will not be copied back from the device to the host.

If successful, the `omp_pause_soft` value results in a soft pause for which the OpenMP state is guaranteed to persist across the call, with the exception of any data associated with a `threadprivate` directive, which may be relinquished across the call. When relinquished and when applicable, base language appropriate deallocation/finalization is performed.

**Note** – A hard pause may relinquish more resources, but may resume processing OpenMP regions more slowly. A soft pause allows OpenMP regions to restart more quickly, but may relinquish fewer resources. An OpenMP implementation will reclaim resources as needed for OpenMP regions encountered after the `omp_pause_resource` region. Since a hard pause may unmap data on the specified device, appropriate data mapping is required before using data on the specified device after the `omp_pause_region` region.

The routine returns zero in case of success, and nonzero otherwise.

### Tool Callbacks

If the tool is not allowed to interact with the specified device after encountering this call, then the runtime must call the tool finalizer for that device.

### Restrictions

The `omp_pause_resource` routine has the following restrictions:

- The `omp_pause_resource` region may not be nested in any explicit OpenMP region.
- The routine may only be called when all explicit tasks have finalized execution. Calling the routine in any other circumstances may result in unspecified behavior.

### Cross References

- `target` construct, see Section 2.15.5 on page 174.
- Declare target directive, see Section 2.15.7 on page 185.
- `threadprivate` directives, see Section 2.22.2 on page 276.
- `omp_get_num_devices`, see Section 3.2.36 on page 370.
- `omp_get_initial_device` routine, see Section 3.2.41 on page 374.
- To pause resources on all devices at once, see Section 3.2.44 on page 377.

## 3.2.44 `omp_pause_resource_all`

### Summary

The `omp_pause_resource_all` routine allows the runtime to relinquish resources used by OpenMP on all devices.

### Format

```
int omp_pause_resource_all(omp_pause_resource_t kind);
```

## Fortran

```
1 integer function omp_pause_resource_all(kind)  
2 integer (kind=omp_pause_resource_kind) kind
```

## Fortran

### Binding

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 calling the `omp_pause_resource` routine once for each available device, including the host device.

The argument `kind` passed to this routine can be one of the valid OpenMP pause kind as defined in Section 3.2.43 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 restrictions:

- The `omp_pause_resource_all` region may not be nested in any explicit OpenMP region.
- The routine may only be called when all explicit tasks have finalized execution. Calling the routine in any other circumstances may result in unspecified behavior.

### Cross References

- `target` construct, see Section 2.15.5 on page 174.
- Declare target directive, see Section 2.15.7 on page 185.
- `omp_get_num_devices`, see Section 3.2.36 on page 370.
- `omp_get_initial_device` routine, see Section 3.2.41 on page 374.
- To pause resources on a specific device only, see Section 3.2.43 on page 375.

## 3.2.45 `omp_set_num_teams`

### Summary

The `omp_set_num_teams` routine affects the number of threads to be used for subsequent `teams` regions that do not specify a `num_teams` clause, by setting the value of the `nteams-var` ICV of the current task.



## Format

**void omp\_set\_num\_teams**(int *num\_teams*);

C / C++

**subroutine omp\_set\_num\_teams**(*num\_teams*)  
integer *num\_teams*

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\_teams** region is the generating task.

## Effect

The effect of this routine is to set the value of the *nteam*s-var ICV of the current task to the value specified in the argument.

## Restrictions

- The **omp\_set\_num\_teams** may not be called from within a parallel region that is not the implicit parallel region that surrounds the whole OpenMP program.

## Cross References

- *nteam*s-var ICV, see Section 2.6 on page 66.
- **teams** construct and **num\_teams** clause, see Section 2.9 on page 87.
- **omp\_get\_num\_teams** routine, see Section 3.2.38 on page 372.
- **omp\_get\_max\_teams** routine, see Section 3.2.46 on page 379.
- **OMP\_NUM\_TEAMS** environment variable, see Section 6.22 on page 601.

## 3.2.46 omp\_get\_max\_teams

### Summary

The **omp\_get\_max\_teams** routine returns an upper bound on the number of teams that could be created by a **teams** construct without a **num\_teams** clause that is encountered after execution returns from this routine.

## Format

C / C++  
`int omp_get_max_teams(void);`

C / C++

Fortran

Fortran  
`integer function omp_get_max_teams()`

## Binding

The binding task set for an `omp_get_max_teams` region is the generating task.

## Effect

The value returned by `omp_get_max_teams` is the value of the `nteamsvar` ICV of the current task. This value is also an upper bound on the number of teams that can be created by a `teams` construct without a `num_teams` clause that is encountered after execution returns from this routine.

## Cross References

- `nteamsvar` ICV, see Section 2.6 on page 66.
- `teams` construct and `num_teams` clause, see Section 2.9 on page 87.
- `omp_get_num_teams` routine, see Section 3.2.38 on page 372.
- `omp_set_num_teams` routine, see Section 3.2.45 on page 378.

## 3.2.47 `omp_set_teams_thread_limit`

### Summary

The `omp_set_teams_thread_limit` routine defines the maximum number of OpenMP threads that can participate in each contention group created by a `teams` construct.

## Format

C / C++  
`void omp_set_teams_thread_limit(int thread_limit);`

C / C++

Fortran  
`subroutine omp_set_teams_thread_limit(thread_limit)  
integer thread_limit`

Fortran

1       **Binding**  
2       The binding task set for an `omp_set_teams_thread_limit` region is the generating task.

3       **Effect**  
4       The `omp_set_teams_thread_limit` routine sets the value of the *teams-thread-limit-var*  
5       ICV to the value of the *thread\_limit* argument.

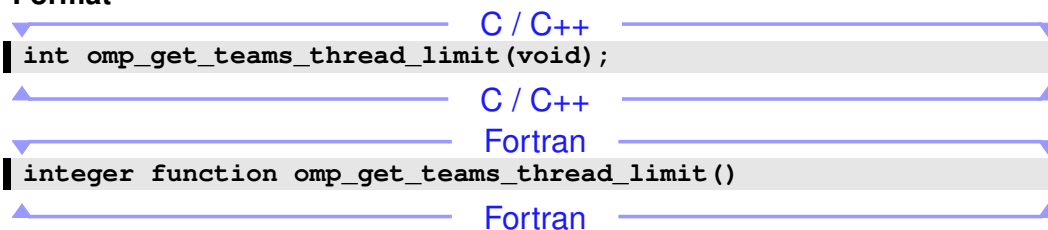
6       If the value of *thread\_limit* exceeds the number of OpenMP threads that an implementation  
7       supports for each contention group created by a **teams** construct, the value of the  
8       *teams-thread-limit-var* ICV will be set to the number that is supported by the implementation.

9       **Restrictions**  
10       • The `omp_set_teams_thread_limit` may not be called from within a parallel region that  
11       is not the implicit parallel region that surrounds the whole OpenMP program.

12       **Cross References**  
13       • *teams\_thread-limit-var* ICV, see Section 2.6 on page 66.  
14       • `omp_get_teams_thread_limit` routine, see Section 3.2.48 on page 381.  
15       • `OMP_TEAMS_THREAD_LIMIT` environment variable, see Section 6.23 on page 601.

## 16       3.2.48 `omp_get_teams_thread_limit`

17       **Summary**  
18       The `omp_get_teams_thread_limit` routine returns the maximum number of OpenMP  
19       threads available to participate in each contention group created by a **teams** construct.

20       **Format**  
21         
22       The diagram shows two code snippets. The first is for C/C++: `int omp_get_teams_thread_limit(void);`. The second is for Fortran: `integer function omp_get_teams_thread_limit()`. Blue arrows point from the text 'C / C++' and 'Fortran' to their respective code snippets.

23       **Binding**  
24       The binding task set for an `omp_get_teams_thread_limit` region is the generating task.

25       **Effect**  
26       The `omp_get_teams_thread_limit` routine returns the value of the *teams-thread-limit-var*  
27       ICV.

## Cross References

- *teams\_thread-limit-var* ICV, see Section 2.6 on page 66.
- `omp_set_teams_thread_limit` routine, see Section 3.2.47 on page 380.
- `OMP_TEAMS_THREAD_LIMIT` environment variable, see Section 6.23 on page 601.

## 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 only be passed to *simple lock* routines. *Nestable lock* variables are associated with *nestable locks* and can only be passed to *nestable lock* routines.

Each type of lock can also have a *synchronization hint* that contains information about the intended usage of the lock by the application code. The effect of the hint is implementation defined. An OpenMP implementation can use this hint to select a usage-specific lock, but hints do not change the mutual exclusion semantics of locks. A conforming implementation can safely ignore the hint.

Constraints on the state and ownership of the lock accessed by each of the lock routines are described with the routine. If these constraints are not met, the behavior of the routine is unspecified.

The OpenMP lock routines access a lock variable such that they always read and update the most current value of the lock variable. It is not necessary for an OpenMP program to include explicit **flush** directives to ensure that the lock variable's value is consistent among different tasks.

## Binding

The binding thread set for all lock routine regions is all threads in the contention group. As a consequence, for each OpenMP lock, the lock routine effects relate to all tasks that call the routines, without regard to which teams the threads in the contention group that are executing the tasks belong.

1  
2  
3  
4  
5  
6  
7  
8  
9  
10  
11  
12  
13  
14  
15  
16  
17  
18  
19

### Simple Lock Routines

C / C++

The type `omp_lock_t` represents a simple lock. For the following routines, a simple lock variable must be of `omp_lock_t` type. All simple lock routines require an argument that is a pointer to a variable of type `omp_lock_t`.

C / C++

Fortran

For the following routines, a simple lock variable must be an integer variable of `kind=omp_lock_kind`.

Fortran

The simple lock routines are as follows:

- The `omp_init_lock` routine initializes a simple lock;
- The `omp_init_lock_with_hint` routine initializes a simple lock and attaches a hint to it;
- The `omp_destroy_lock` routine uninitialized a simple lock;
- The `omp_set_lock` routine waits until a simple lock is available and then sets it;
- The `omp_unset_lock` routine unsets a simple lock; and
- 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

1 The nestable lock routines are as follows:

- 2 • The `omp_init_nest_lock` routine initializes a nestable lock;
- 3 • The `omp_init_nest_lock_with_hint` routine initializes a nestable lock and attaches a  
4 hint to it;
- 5 • The `omp_destroy_nest_lock` routine uninitialized a nestable lock;
- 6 • The `omp_set_nest_lock` routine waits until a nestable lock is available and then sets it;
- 7 • The `omp_unset_nest_lock` routine unsets a nestable lock; and
- 8 • The `omp_test_nest_lock` routine tests a nestable lock and sets it if it is available.

### 9 **Restrictions**

10 OpenMP lock routines have the following restriction:

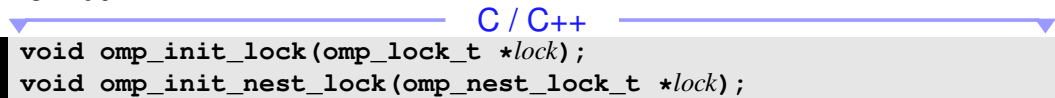
- 11 • The use of the same OpenMP lock in different contention groups results in unspecified behavior.

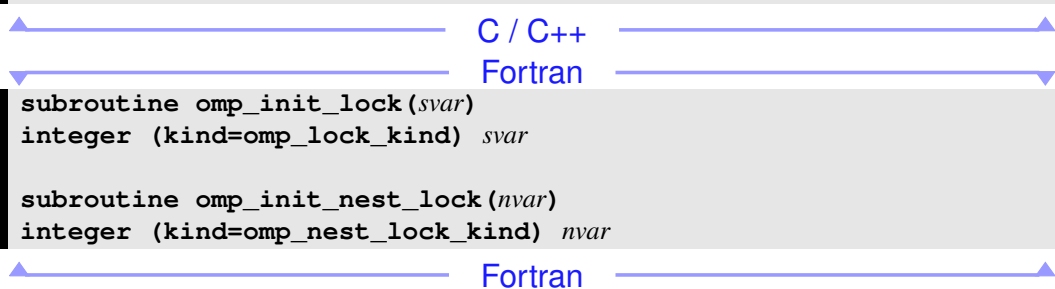
## 12 **3.3.1 `omp_init_lock` and `omp_init_nest_lock`**

### 13 **Summary**

14 These routines initialize an OpenMP lock without a hint.

### 15 **Format**

```
16  void omp_init_lock(omp_lock_t *lock);  
17 void omp_init_nest_lock(omp_nest_lock_t *lock);
```

```
18  subroutine omp_init_lock(svar)  
19 integer (kind=omp_lock_kind) svar  
20  
21 subroutine omp_init_nest_lock(nvar)  
22 integer (kind=omp_nest_lock_kind) nvar
```

### 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 **Effect**

27 The effect of these routines is to initialize the lock to the unlocked state; that is, no task owns the  
28 lock. In addition, the nesting count for a nestable lock is set to zero.

## Execution Model Events

The *lock-init* event occurs in a thread that executes an `omp_init_lock` region after initialization of the lock, but before it finishes the region. The *nest-lock-init* event occurs in a thread that executes an `omp_init_nest_lock` region after initialization of the lock, but before it finishes the region.

## Tool Callbacks

A thread dispatches a registered `ompt_callback_lock_init` callback with `omp_sync_hint_none` as the *hint* argument and `ompt_mutex_lock` as the *kind* argument for each occurrence of a *lock-init* event in that thread. Similarly, a thread dispatches a registered `ompt_callback_lock_init` callback with `omp_sync_hint_none` as the *hint* argument and `ompt_mutex_nest_lock` as the *kind* argument for each occurrence of a *nest-lock-init* event in that thread. These callbacks have the type signature `ompt_callback_mutex_acquire_t` and occur in the task that encounters the routine.

## Cross References

- `ompt_callback_mutex_acquire_t`, see Section 4.5.2.14 on page 477.

## 3.3.2 `omp_init_lock_with_hint` and `omp_init_nest_lock_with_hint`

### Summary

These routines initialize an OpenMP lock with a hint. The effect of the hint is implementation-defined. The OpenMP implementation can ignore the hint without changing program semantics.

### Format

```
C / C++
void omp_init_lock_with_hint(
    omp_lock_t *lock,
    omp_sync_hint_t hint
);
void omp_init_nest_lock_with_hint(
    omp_nest_lock_t *lock,
    omp_sync_hint_t hint
);
```

C / C++

## Fortran

```
1  subroutine omp_init_lock_with_hint(svar, hint)
2  integer (kind=omp_lock_kind) svar
3  integer (kind=omp_sync_hint_kind) hint
4
5  subroutine omp_init_nest_lock_with_hint(nvar, hint)
6  integer (kind=omp_nest_lock_kind) nvar
7  integer (kind=omp_sync_hint_kind) hint
```

## Fortran

### Constraints on Arguments

A program that accesses a lock that is not in the uninitialized state through either routine is non-conforming.

The second argument passed to these routines (*hint*) is a hint as described in Section 2.20.12 on page 262.

### 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* event occurs in a thread that executes an `omp_init_lock_with_hint` region after initialization of the lock, but before it finishes the region. The *nest-lock-init\_with\_hint* event occurs in a thread that executes an `omp_init_nest_lock` region after initialization of the lock, but before it finishes the region.

### Tool Callbacks

A thread dispatches a registered `ompt_callback_lock_init` callback with the same value for its *hint* argument as the *hint* argument of the call to `omp_init_lock_with_hint` and `ompt_mutex_lock` as the *kind* argument for each occurrence of a *lock-init* event in that thread. Similarly, a thread dispatches a registered `ompt_callback_lock_init` callback with the same value for its *hint* argument as the *hint* argument of the call to `omp_init_nest_lock_with_hint` and `ompt_mutex_nest_lock` as the *kind* argument for each occurrence of a *nest-lock-init* event in that thread. These callbacks have the type signature `ompt_callback_mutex_acquire_t` and occur in the task that encounters the routine.

### Cross References

- Synchronization Hints, see Section 2.20.12 on page 262.
- `ompt_callback_mutex_acquire_t`, see Section 4.5.2.14 on page 477.



### 3.3.3 `omp_destroy_lock` and `omp_destroy_nest_lock`

#### Summary

These routines ensure that the OpenMP lock is uninitialized.

#### Format

C / C++

```
void omp_destroy_lock(omp_lock_t *lock);  
void omp_destroy_nest_lock(omp_nest_lock_t *lock);
```

C / C++

Fortran

```
subroutine omp_destroy_lock(svar)  
integer (kind=omp_lock_kind) svar  
  
subroutine omp_destroy_nest_lock(nvar)  
integer (kind=omp_nest_lock_kind) nvar
```

Fortran

#### Constraints on Arguments

A program that accesses a lock that is not in the unlocked state through either routine is non-conforming.

#### Effect

The effect of these routines is to change the state of the lock to uninitialized.

#### Execution Model Events

The *lock-destroy* event occurs in a thread that executes an `omp_destroy_lock` region before it finishes the region. The *nest-lock-destroy\_with\_hint* event occurs in a thread that executes an `omp_destroy_nest_lock` region before it finishes the region.

#### Tool Callbacks

A thread dispatches a registered `ompt_callback_lock_destroy` callback with `ompt_mutex_lock` as the *kind* argument for each occurrence of a *lock-destroy* event in that thread. Similarly, a thread dispatches a registered `ompt_callback_lock_destroy` callback with `ompt_mutex_nest_lock` as the *kind* argument for each occurrence of a *nest-lock-destroy* event in that thread. These callbacks have the type signature `ompt_callback_mutex_acquire_t` and occur in the task that encounters the routine.

#### Cross References

- `ompt_callback_mutex_t`, see Section 4.5.2.15 on page 478.

### 3.3.4 omp\_set\_lock and omp\_set\_nest\_lock

#### Summary

These routines provide a means of setting an OpenMP lock. The calling task region behaves as if it 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 that is 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 that executes the routine.

A nestable lock is available if it is unlocked or if it is already owned by the task that executes the routine. The task that executes the routine is granted, or retains, ownership of the lock, and the nesting count for the lock is incremented.

## Execution Model Events

The *lock-acquire* event occurs in a thread that executes an `omp_set_lock` region before the associated lock is requested. The *nest-lock-acquire* event occurs in a thread that executes an `omp_set_nest_lock` region before the associated lock is requested.

The *lock-acquired* event occurs in a thread that executes an `omp_set_lock` region after it acquires the associated lock but before it finishes the region. The *nest-lock-acquired* event occurs in a thread that executes an `omp_set_nest_lock` region if the thread did not already own the lock, after it acquires the associated lock but before it finishes the region.

The *nest-lock-owned* event occurs in a thread when it already owns the lock and executes an `omp_set_nest_lock` region. The event occurs after the nesting count is incremented but before the thread finishes the region.

## Tool Callbacks

A thread dispatches a registered `ompt_callback_mutex_acquire` callback for each occurrence of a *lock-acquire* or *nest-lock-acquire* event in that thread. This callback has the type signature `ompt_callback_mutex_acquire_t`.

A thread dispatches a registered `ompt_callback_mutex_acquired` callback for each occurrence of a *lock-acquired* or *nest-lock-acquired* event in that thread. This callback has the type signature `ompt_callback_mutex_t`.

A thread dispatches a registered `ompt_callback_nest_lock` callback with `ompt_scope_begin` as its *endpoint* argument for each occurrence of a *nest-lock-owned* event in that thread. This callback has the type signature `ompt_callback_nest_lock_t`.

The above callbacks occur in the task that encounters the lock function. The *kind* argument of these callbacks is `ompt_mutex_lock` when the events arise from an `omp_set_lock` region while it is `ompt_mutex_nest_lock` when the events arise from an `omp_set_nest_lock` region.

## Cross References

- `ompt_callback_mutex_acquire_t`, see Section 4.5.2.14 on page 477.
- `ompt_callback_mutex_t`, see Section 4.5.2.15 on page 478.
- `ompt_callback_nest_lock_t`, see Section 4.5.2.16 on page 479.

## 3.3.5 `omp_unset_lock` and `omp_unset_nest_lock`

### Summary

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* event occurs in a thread that executes an `omp_unset_lock` region after it releases the associated lock but before it finishes the region. The *nest-lock-release* event occurs in a thread that executes an `omp_unset_nest_lock` region after it releases the associated lock but before it finishes the region.

The *nest-lock-held* event occurs in a thread that executes an `omp_unset_nest_lock` region before it finishes the region when the thread still owns the lock after the nesting count is decremented.

## 1 Tool Callbacks

2 A thread dispatches a registered `ompt_callback_mutex_released` callback with  
3 `ompt_mutex_lock` as the *kind* argument for each occurrence of a *lock-release* event in that  
4 thread. Similarly, a thread dispatches a registered `ompt_callback_mutex_released`  
5 callback with `ompt_mutex_nest_lock` as the *kind* argument for each occurrence of a  
6 *nest-lock-release* event in that thread. These callbacks have the type signature  
7 `ompt_callback_mutex_t` and occur in the task that encounters the routine.

8 A thread dispatches a registered `ompt_callback_nest_lock` callback with  
9 `ompt_scope_end` as its *endpoint* argument for each occurrence of a *nest-lock-held* event in that  
10 thread. This callback has the type signature `ompt_callback_nest_lock_t`.

## 11 Cross References

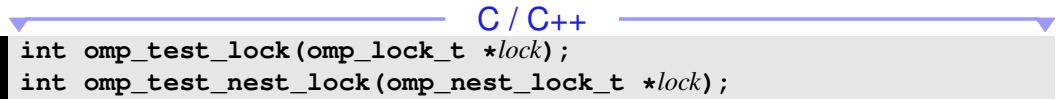
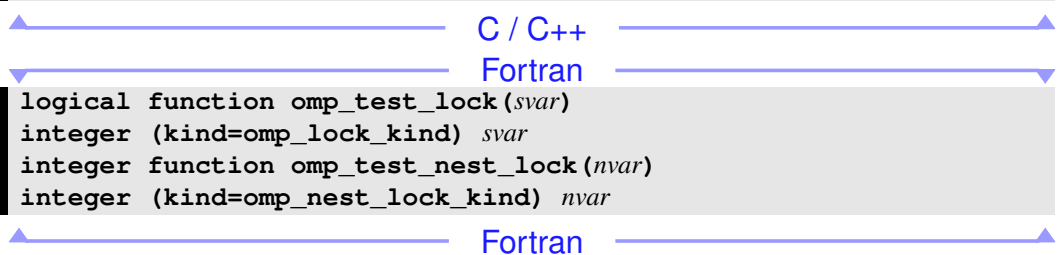
- 12 • `ompt_callback_mutex_t`, see Section 4.5.2.15 on page 478.
- 13 • `ompt_callback_nest_lock_t`, see Section 4.5.2.16 on page 479.

## 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 that  
17 executes 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  Fortran  
22 logical function omp_test_lock(svar)  
23 integer (kind=omp_lock_kind) svar  
24 integer function omp_test_nest_lock(nvar)  
integer (kind=omp_nest_lock_kind) nvar
```

### 25 Constraints on Arguments

26 A program that accesses a lock that is in the uninitialized state through either routine is  
27 non-conforming. The behavior is unspecified if a simple lock accessed by `omp_test_lock` is in  
28 the locked state and is owned by the task that contains the call.

## Effect

These routines attempt to set a lock in the same manner as `omp_set_lock` and `omp_set_nest_lock`, except that they do not suspend execution of the task that executes the routine.

For a simple lock, the `omp_test_lock` routine returns *true* if the lock is successfully set; otherwise, it returns *false*.

For a nestable lock, the `omp_test_nest_lock` routine returns the new nesting count if the lock is successfully set; otherwise, it returns zero.

## Execution Model Events

The *lock-test* event occurs in a thread that executes an `omp_test_lock` region before the associated lock is tested. The *nest-lock-test* event occurs in a thread that executes an `omp_test_nest_lock` region before the associated lock is tested.

The *lock-test-acquired* event occurs in a thread that executes an `omp_test_lock` region before it finishes the region if the associated lock was acquired. The *nest-lock-test-acquired* event occurs in a thread that executes an `omp_test_nest_lock` region before it finishes the region if the associated lock was acquired and the thread did not already own the lock.

The *nest-lock-owned* event occurs in a thread that executes an `omp_test_nest_lock` region before it finishes the region after the nesting count is incremented if the thread already owned the lock.

## Tool Callbacks

A thread dispatches a registered `ompt_callback_mutex_acquire` callback for each occurrence of a *lock-test* or *nest-lock-test* event in that thread. This callback has the type signature `ompt_callback_mutex_acquire_t`.

A thread dispatches a registered `ompt_callback_mutex_acquired` callback for each occurrence of a *lock-test-acquired* or *nest-lock-test-acquired* event in that thread. This callback has the type signature `ompt_callback_mutex_t`.

A thread dispatches a registered `ompt_callback_nest_lock` callback with `ompt_scope_begin` as its *endpoint* argument for each occurrence of a *nest-lock-owned* event in that thread. This callback has the type signature `ompt_callback_nest_lock_t`.

The above callbacks occur in the task that encounters the lock function. The *kind* argument of these callbacks is `ompt_mutex_test_lock` when the events arise from an `omp_test_lock` region while it is `ompt_mutex_test_nest_lock` when the events arise from an `omp_test_nest_lock` region.

## Cross References

- `omp_callback_mutex_acquire_t`, see Section 4.5.2.14 on page 477.
- `omp_callback_mutex_t`, see Section 4.5.2.15 on page 478.
- `omp_callback_nest_lock_t`, see Section 4.5.2.16 on page 479.

## 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++
<code>double omp_get_wtime(void);</code>	
	C / C++
	Fortran
<code>double precision function omp_get_wtime()</code>	
	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.

#### Effect

The `omp_get_wtime` routine returns a value equal to the elapsed wall clock time in seconds since some *time-in-the-past*. The actual *time-in-the-past* is arbitrary, but it is guaranteed not to change during the execution of the application program. The time returned is a *per-thread time*, so it is not required to be globally consistent across all threads that participate in an application.

Note – The routine is anticipated to be used to measure elapsed times as shown in the following example:

C / C++

```
1 double start;  
2 double end;  
3 start = omp_get_wtime();  
4 ... work to be timed ...  
5 end = omp_get_wtime();  
6 printf("Work took %f seconds\n", end - start);
```

C / C++

Fortran

```
7 DOUBLE PRECISION START, END  
8 START = omp_get_wtime()  
9 ... work to be timed ...  
10 END = omp_get_wtime()  
11 PRINT *, "Work took", END - START, "seconds"
```

Fortran

### 3.4.2 omp\_get\_wtick

#### Summary

The `omp_get_wtick` routine returns the precision of the timer used by `omp_get_wtime`.

#### Format

C / C++

```
double omp_get_wtick(void);
```

C / C++

Fortran

```
double precision function omp_get_wtick()
```

Fortran

#### Binding

The binding thread set for an `omp_get_wtick` region is the encountering thread. The routine's return value is not guaranteed to be consistent across any set of threads.

#### Effect

The `omp_get_wtick` routine returns a value equal to the number of seconds between successive clock ticks of the timer used by `omp_get_wtime`.



## 3.5 Event Routine

This section describes a routine that supports OpenMP event objects.

### Binding

The binding thread set for all event routine regions is the encountering thread.

### 3.5.1 `omp_fulfill_event`

#### Summary

This routine fulfills and destroys an OpenMP event.

#### Format

```
void omp_fulfill_event(omp_event_handle_t event);
subroutine omp_fulfill_event(event)
integer (kind=omp_event_handle_kind) event
```

#### Constraints on Arguments

A program that calls this routine on an event that was already fulfilled is non-conforming. A program that calls this routine with an event handle that was not created by the **detach** clause is non-conforming.

#### Effect

The effect of this routine is to fulfill the event associated with the event handle 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 handle becomes unassociated with any event.

#### Execution Model Events

The *task-fulfill* event occurs in a thread that executes an `omp_fulfill_event` region before the event is fulfilled if the OpenMP event object was created by a **detach** clause on a task.

#### Tool Callbacks

A thread dispatches a registered `ompt_callback_task_schedule` callback with **NULL** as its *next\_task\_data* argument while the argument *prior\_task\_data* binds to the detached task for each occurrence of a *task-fulfill* event. If the *task-fulfill* event occurs before the detached task finished the execution of the associated *structured-block*, the callback has `ompt_task_early_fulfill` as its *prior\_task\_status* argument; otherwise the callback has `ompt_task_late_fulfill` as its *prior\_task\_status* argument. This callback has type signature `ompt_callback_task_schedule_t`.

## Cross References

- **detach** clause, see Section 2.13.1 on page 142.
- **ompt\_callback\_task\_schedule\_t**, see Section 4.5.2.10 on page 472.

## 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);
```

C / C++

```
type(c_ptr) function omp_target_alloc(size, device_num) bind(c)
use, intrinsic :: iso_c_binding, only : c_ptr, c_size_t, c_int
integer(c_size_t), value :: size
integer(c_int), value :: device_num
```

Fortran

#### 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 `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` (or `C_NULL_PTR`, for Fortran) 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 174.

```
Unless unified_address clause appears on a requires directive in the compilation unit,
pointer arithmetic is not supported on the device address returned by omp_target_alloc.
```

C / C++

1 Freeing the storage returned by `omp_target_alloc` with any routine other than  
2 `omp_target_free` results in unspecified behavior.

### 3 Execution Model Events

4 The *target-data-allocation* event occurs when a thread allocates data on a target device.

### 5 Tool Callbacks

6 A thread invokes a registered `ompt_callback_target_data_op` callback for each  
7 occurrence of a *target-data-allocation* event in that thread. The callback occurs in the context of the  
8 target task and has type signature `ompt_callback_target_data_op_t`.

### 9 Cross References


- 10 • `target` construct, see Section 2.15.5 on page 174.
- 11 • `omp_get_num_devices` routine, see Section 3.2.36 on page 370.
- 12 • `omp_get_initial_device` routine, see Section 3.2.41 on page 374.
- 13 • `omp_target_free` routine, see Section 3.6.2 on page 397.
- 14 • `ompt_callback_target_data_op_t`, see Section 4.5.2.25 on page 488.

## 15 3.6.2 `omp_target_free`

### 16 Summary

17 The `omp_target_free` routine frees the device memory allocated by the  
18 `omp_target_alloc` routine.

### 19 Format

20   
`void omp_target_free(void *device_ptr, int device_num);`

21   
  
`subroutine omp_target_free(device_ptr, device_num) bind(c)`

22 `use, intrinsic :: iso_c_binding, only : c_ptr, c_int`  
23 `type(c_ptr), value :: device_ptr`  
24 `integer(c_int), value :: device_num`

  
`Fortran`

### 25 Constraints on Arguments

26 A program that calls `omp_target_free` with a non-null pointer that does not have a value  
27 returned from `omp_target_alloc` is non-conforming. The *device\_num* must be greater than or  
28 equal to zero and less than the result of `omp_get_num_devices()` or the result of  
29 `omp_get_initial_device()`.

## Effect

The `omp_target_free` routine frees the memory in the device data environment associated with `device_ptr`. If `device_ptr` is `NULL` (or `C_NULL_PTR`, for Fortran), the operation is ignored.

Synchronization must be inserted to ensure that all accesses to `device_ptr` are completed before the call to `omp_target_free`.

When called from within a `target` region the effect of this routine is unspecified.

## Execution Model Events

The `target-data-free` event occurs when a thread frees 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-free` event in that thread. The callback occurs in the context of the target task and has type signature `ompt_callback_target_data_op_t`.

## Cross References

- `target` construct, see Section 2.15.5 on page 174.
- `omp_get_num_devices` routine, see Section 3.2.36 on page 370.
- `omp_get_initial_device` routine, see Section 3.2.41 on page 374.
- `omp_target_alloc` routine, see Section 3.6.1 on page 396.
- `ompt_callback_target_data_op_t`, see Section 4.5.2.25 on page 488.

## 3.6.3 omp\_target\_is\_present

### Summary

The `omp_target_is_present` routine tests whether a host pointer has corresponding storage on a given device.

### Format

```

C / C++
int omp_target_is_present(const void *ptr, int device_num);

C / C++
integer(c_int) function omp_target_is_present(ptr, device_num) &
  bind(c)
use, intrinsic :: iso_c_binding, only : c_ptr, c_int
type(c_ptr), value :: ptr
integer(c_int), value :: device_num

Fortran
```

## Constraints on Arguments

The value of *ptr* must be a valid host pointer or **NULL** (or **C\_NULL\_PTR**, for Fortran). The *device\_num* must be greater than or equal to zero and less than the result of `omp_get_num_devices()` or the result of `omp_get_initial_device()`.

## Effect

This routine returns non-zero if the specified pointer would be found present on device *device\_num* by a **map** clause; otherwise, it returns zero.

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 174.
- **map** clause, see Section 2.22.7.1 on page 318.
- `omp_get_num_devices` routine, see Section 3.2.36 on page 370.
- `omp_get_initial_device` routine, see Section 3.2.41 on page 374.

## 3.6.4 omp\_target\_memcpy

### Summary

The `omp_target_memcpy` routine copies memory between any combination of host and device pointers.

### Format

C / C++

```
int omp_target_memcpy(  
    void *dst,  
    const void *src,  
    size_t length,  
    size_t dst_offset,  
    size_t src_offset,  
    int dst_device_num,  
    int src_device_num  
);
```

C / C++

Fortran

```
integer(c_int) function omp_target_memcpy(dst, src, length, &  
    dst_offset, src_offset, dst_device_num, src_device_num) bind(c)  
use, intrinsic :: iso_c_binding, only : c_ptr, c_int, c_size_t  
type(c_ptr), value :: dst, src  
integer(c_size_t), value :: length, dst_offset, src_offset  
integer(c_int), value :: dst_device_num, src_device_num
```

Fortran

## Constraints on Arguments

Each device pointer specified must be valid for the device on the same side of the copy. The *dst\_device\_num* and *src\_device\_num* arguments must be greater than or equal to zero and less than the result of `omp_get_num_devices()` or equal to the result of `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. 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 and has type signature `ompt_callback_target_data_op_t`.

## Cross References

- **target** construct, see Section [2.15.5](#) on page [174](#).
- `omp_get_initial_device` routine, see Section [3.2.41](#) on page [374](#).
- `ompt_callback_target_data_op_t`, see Section [4.5.2.25](#) on page [488](#).

## 3.6.5 omp\_target\_memcpy\_rect

### Summary

The `omp_target_memcpy_rect` routine copies a rectangular subvolume from a multi-dimensional array to another multi-dimensional array. The `omp_target_memcpy_rect` routine performs a copy between any combination of host and device pointers.

## Format

C / C++

```
1  int omp_target_memcpy_rect (  
2      void *dst,  
3      const void *src,  
4      size_t element_size,  
5      int num_dims,  
6      const size_t *volume,  
7      const size_t *dst_offsets,  
8      const size_t *src_offsets,  
9      const size_t *dst_dimensions,  
10     const size_t *src_dimensions,  
11     int dst_device_num,  
12     int src_device_num  
13 );  
14
```

C / C++

Fortran

```
15 integer(c_int) function omp_target_memcpy_rect (dst,src,element_size, &  
16     num_dims, volume, dst_offsets, src_offsets, dst_dimensions, &  
17     dst_device_num, src_device_num) bind(c)  
18 use, intrinsic :: iso_c_binding, only : c_ptr, c_int, c_size_t  
19 type(c_ptr), value :: dst, src  
20 integer(c_size_t), value :: element_size  
21 integer(c_int), value :: num_dims, dst_device_num, src_device_num  
22 integer(c_size_t), intent(in) :: volume(*), dst_offsets(*), &  
23     src_offsets(*), dst_dimensions(*), src_dimensions(*)
```

Fortran

## Constraints on Arguments

Each device pointer specified must be valid for the device on the same side of the copy. The *dst\_device\_num* and *src\_device\_num* arguments must be greater than or equal to zero and less than the result of `omp_get_num_devices()` or equal to the result of `omp_get_initial_device()`.

The length of the offset and dimension arrays must be at least the value of *num\_dims*.

The value of *num\_dims* must be between 1 and the implementation-defined limit, which must be at least three.

## Effect

This routine copies a rectangular subvolume of *src*, in the device data environment of device *src\_device\_num*, to *dst*, in the device data environment of device *dst\_device\_num*. The volume is specified in terms of the size of an element, number of dimensions, and constant arrays of length *num\_dims*. The maximum number of dimensions supported is at least three, support for higher dimensionality is implementation defined. The *volume* array specifies the length, in number of elements, to copy in each dimension from *src* to *dst*. The *dst\_offsets* (*src\_offsets*) parameter specifies number of elements from the origin of *dst* (*src*) in elements. The *dst\_dimensions* (*src\_dimensions*) parameter specifies the length of each dimension of *dst* (*src*).

The routine returns zero if successful. Otherwise, it returns a non-zero value. The routine contains a task scheduling point.

When called from within a **target** region the effect of this routine is unspecified.

An application can determine the inclusive number of dimensions supported by an implementation by passing **NULL** pointers for both *dst* and *src*. The routine returns the number of dimensions supported by the implementation for the specified device numbers. No copy operation is performed.

## 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 and has type signature **ompt\_callback\_target\_data\_op\_t**.

## Cross References

- **target** construct, see Section 2.15.5 on page 174.
- **omp\_get\_initial\_device** routine, see Section 3.2.41 on page 374.
- **ompt\_callback\_target\_data\_op\_t**, see Section 4.5.2.25 on page 488.

## 3.6.6 omp\_target\_memcpy\_async

### Summary

The **omp\_target\_memcpy\_async** routine asynchronously performs a copy between any combination of host and device pointers.



## Format

C / C++

```
1  int omp_target_memcpy_async(  
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      int depobj_count,  
10     omp_depend_t *depobj_list  
11 );  
12
```

C / C++

Fortran

```
13 integer(c_int) function omp_target_memcpy_async(dst, src, length, &  
14     dst_offset, src_offset, dst_device_num, src_device_num, &  
15     depobj_count, depobj_list) bind(c)  
16 use, intrinsic :: iso_c_binding, only : c_ptr, c_int, c_size_t  
17 type(c_ptr), value :: dst, src  
18 integer(c_size_t), value :: length, dst_offset, src_offset  
19 integer(c_int), value :: dst_device_num, src_device_num, depobj_count  
20 integer(omp_depend_kind), optional :: depobj_list(*)
```

Fortran

## Constraints on Arguments

Each device pointer specified must be valid for the device on the same side of the copy. The *dst\_device\_num* and *src\_device\_num* arguments must be greater than or equal to zero and less than the result of `omp_get_num_devices()` or equal to the result of `omp_get_initial_device()`.

## Effect

This routine performs an asynchronous memory copy where *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*. Logically the `omp_target_memcpy_async` routine generates a target task with an implicit nowait.

Task dependences are expressed with zero or more `omp_depend_t` objects. The dependences are specified by passing the number of `omp_depend_t` objects followed by an array of `omp_depend_t` objects. The generated target task is not a dependent task if the program passes in a count of zero for *depobj\_count*. *depobj\_list* is ignored if the value of *depobj\_count* is zero.

The routine returns zero if successful. Otherwise, it returns a non-zero value. The routine contains a task scheduling point.

## 1 Execution Model Events

2 The *target-data-op* event occurs when a thread transfers data on a target device.

## 3 Tool Callbacks

4 A thread invokes a registered `ompt_callback_target_data_op` callback for each  
5 occurrence of a *target-data-op* event in that thread. The callback occurs in the context of the target  
6 task and has type signature `ompt_callback_target_data_op_t`.

## 7 Cross References

- 8 • `target` construct, see Section 2.15.5 on page 174.
- 9 • `Depend` objects, see Section 2.20.10 on page 256.
- 10 • `omp_get_initial_device` routine, see Section 3.2.41 on page 374.
- 11 • `ompt_callback_target_data_op_t`, see Section 4.5.2.25 on page 488.

## 12 3.6.7 `omp_target_memcpy_rect_async`

### 13 Summary

14 The `omp_target_memcpy_rect_async` routine asynchronously performs a copy between  
15 any combination of host and device pointers.

### 16 Format

17 

```
18 int omp_target_memcpy_rect_async(  
19     void *dst,  
20     const void *src,  
21     size_t element_size,  
22     int num_dims,  
23     const size_t *volume,  
24     const size_t *dst_offsets,  
25     const size_t *src_offsets,  
26     const size_t *dst_dimensions,  
27     const size_t *src_dimensions,  
28     int dst_device_num,  
29     int src_device_num,  
30     int depobj_count,  
31     omp_depend_t *depobj_list  
);
```

18 

## Fortran

```
1 integer(c_int) function omp_target_memcpy_rect_async(dst, src, &
2     element_size, num_dims, volume, dst_offsets, src_offsets, &
3     dst_dimensions, src_dimensions, dst_device_num, src_device_num, &
4     depobj_count, depobj_list) bind(c)
5 use, intrinsic :: iso_c_binding, only : c_ptr, c_int, c_size_t
6 type(c_ptr), value :: dst, src
7 integer(c_size_t), value :: element_size
8 integer(c_int), value :: num_dims, dst_device_num, src_device_num, &
9     depobj_count
10 integer(c_size_t), intent(in) :: volume(*), dst_offsets(*), &
11     src_offsets(*), dst_dimensions(*), src_dimensions(*)
12 integer(omp_depobj_kind), optional :: depobj_list(*)
```

## Fortran

### Constraints on Arguments

Each device pointer specified must be valid for the device on the same side of the copy. The *dst\_device\_num* and *src\_device\_num* arguments must be greater than or equal to zero and less than the result of `omp_get_num_devices()` or equal to the result of `omp_get_initial_device()`.

The length of the offset and dimension arrays must be at least the value of *num\_dims*. The value of *num\_dims* must be between 1 and the implementation-defined limit, which must be at least three.

### Effect

This routine copies a rectangular subvolume of *src*, in the device data environment of device *src\_device\_num*, to *dst*, in the device data environment of device *dst\_device\_num*. The volume is specified in terms of the size of an element, number of dimensions, and constant arrays of length *num\_dims*. The maximum number of dimensions supported is at least three, support for higher dimensionality is implementation defined. The volume array specifies the length, in number of elements, to copy in each dimension from *src* to *dst*. The *dst\_offsets* (*src\_offsets*) parameter specifies number of elements from the origin of *dst* (*src*) in elements. The *dst\_dimensions* (*src\_dimensions*) parameter specifies the length of each dimension of *dst* (*src*). Logically the `omp_target_memcpy_rect_async` routine generates a target task with an implicit nowait.

Task dependences are expressed with zero or more `omp_depend_t` objects. The dependences are specified by passing the number of `omp_depend_t` objects followed by an array of `omp_depend_t` objects. The generated target task is not a dependent task if the program passes in a count of zero for *depobj\_count*. *depobj\_list* is ignored if the value of *depobj\_count* is zero.

The routine returns zero if successful. Otherwise, it returns a non-zero value. The routine contains a task scheduling point.

When called from within a `target` region the effect of this routine is unspecified.

1 An application can determine the number of inclusive dimensions supported by an implementation  
2 by passing **NULL** pointers (or **C\_NULL\_PTR**, for Fortran) for both *dst* and *src*. The routine returns  
3 the number of dimensions supported by the implementation for the specified device numbers. No  
4 copy operation is performed.

## 5 Execution Model Events

6 The *target-data-op* event occurs when a thread transfers data on a target device.

## 7 Tool Callbacks

8 A thread invokes a registered **ompt\_callback\_target\_data\_op** callback for each  
9 occurrence of a *target-data-op* event in that thread. The callback occurs in the context of the target  
10 task and has type signature **ompt\_callback\_target\_data\_op\_t**.

## 11 Cross References

- 12 • **target** construct, see Section 2.15.5 on page 174.
- 13 • **Depend objects**, see Section 2.20.10 on page 256.
- 14 • **omp\_get\_initial\_device** routine, see Section 3.2.41 on page 374.
- 15 • **ompt\_callback\_target\_data\_op\_t**, see Section 4.5.2.25 on page 488.

## 16 3.6.8 omp\_target\_associate\_ptr

### 17 Summary

18 The **omp\_target\_associate\_ptr** routine maps a device pointer, which may be returned  
19 from **omp\_target\_alloc** or implementation-defined runtime routines, to a host pointer.

### 20 Format

```
21 int omp_target_associate_ptr(  
22     const void *host_ptr,  
23     const void *device_ptr,  
24     size_t size,  
25     size_t device_offset,  
26     int device_num  
27 );
```

C / C++

## Fortran

```
1 integer(c_int) function omp_target_associate_ptr(host_ptr, &  
2 device_ptr, size, device_offset, device_num) bind(c)  
3 use, intrinsic :: iso_c_binding, only : c_ptr, c_size_t, c_int  
4 type(c_ptr), value :: host_ptr, device_ptr  
5 integer(c_size_t), value :: size, device_offset  
6 integer(c_int), value :: device_num
```

## Fortran

### Constraints on Arguments

The value of *device\_ptr* value must be a valid pointer to device memory for the device denoted by the value of *device\_num*. The *device\_num* argument must be greater than or equal to zero and less than the result of `omp_get_num_devices()` or equal to the result of `omp_get_initial_device()`.

### Effect

The `omp_target_associate_ptr` routine associates a device pointer in the device data environment of device *device\_num* with a host pointer such that when the host pointer appears in a subsequent `map` clause, the associated device pointer is used as the target for data motion associated with that host pointer. The *device\_offset* parameter specifies the offset into *device\_ptr* that is used as the base address for the device side of the mapping. The reference count of the resulting mapping will be infinite. After being successfully associated, the buffer to which the device pointer points is invalidated and accessing data directly through the device pointer results in unspecified behavior. The pointer can be retrieved for other uses by disassociating it. 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.

Only one device buffer can be associated with a given host pointer value and device number pair. Attempting to associate a second buffer will return non-zero. Associating the same pair of pointers on the same device with the same offset has no effect and returns zero. Associating pointers that share underlying storage will result in unspecified behavior. The `omp_target_is_present` function can be used to test whether a given host pointer has a corresponding variable in the device data environment.

### Execution Model Events

The *target-data-associate* event occurs when a thread associates 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-associate* event in that thread. The callback occurs in the context of the target task and has type signature `ompt_callback_target_data_op_t`.

## Cross References

- `target` construct, see Section 2.15.5 on page 174.
- `map` clause, see Section 2.22.7.1 on page 318.
- `omp_target_alloc` routine, see Section 3.6.1 on page 396.
- `omp_target_disassociate_ptr` routine, see Section 3.6.8 on page 406.
- `omp_get_mapped_ptr` routine, see Section 3.6.9 on page 408.
- `ompt_callback_target_data_op_t`, see Section 4.5.2.25 on page 488.

## 3.6.9 `omp_get_mapped_ptr`

### Summary

The `omp_get_mapped_ptr` routine returns the device pointer that is associated with a host pointer for a given device.

### Format

	C / C++	
<code>void * omp_get_mapped_ptr(const void *ptr, int device_num);</code>		
	C / C++	
	Fortran	
<code>type(c_ptr) function omp_get_mapped_ptr(ptr, &amp; device_num) bind(c) use, intrinsic :: iso_c_binding, only : c_ptr, c_int type(c_ptr), value :: ptr integer(c_int), value :: device_num</code>		
	Fortran	

### 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 the value returned by `omp_get_initial_device()`.

### Effect

The `omp_get_mapped_ptr` routine returns the associated device pointer on device `device_num`. A call to this routine for a pointer that is not `NULL` (or `C_NULL_PTR`, for Fortran) and does not have an associated pointer on the given device results in a `NULL` pointer.

This routine is intended to be called outside of any `target` region. When called from within a `target` region the effect of this routine is unspecified.

The routine returns `NULL` (or `C_NULL_PTR`, for Fortran) if unsuccessful. Otherwise it returns the device pointer, which is `ptr` if `device_num` is the value returned by `omp_get_initial_device()`.

## Execution Model Events

No events are associated with this routine.

## Cross References

- `omp_target_associate_ptr` routine, see Section 3.6.8 on page 406.
- `omp_target_disassociate_ptr` routine, see Section 3.6.10 on page 409.

## 3.6.10 `omp_target_disassociate_ptr`

### Summary

The `omp_target_disassociate_ptr` removes the associated pointer for a given device from a host pointer.

### Format

C / C++

```
int omp_target_disassociate_ptr(const void *ptr, int device_num);
```

C / C++

Fortran

```
integer(c_int) function omp_target_disassociate_ptr(ptr, &  
    device_num) bind(c)  
    use, intrinsic :: iso_c_binding, only : c_ptr, c_int  
    type(c_ptr), value :: ptr  
    integer(c_int), value :: device_num
```

Fortran

### 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 the value returned by `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` (or `C_NULL_PTR`, for Fortran) 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.

## 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 and has type signature `ompt_callback_target_data_op_t`.

## Cross References

- `target` construct, see Section 2.15.5 on page 174.
- `omp_target_associate_ptr` routine, see Section 3.6.8 on page 406.
- `ompt_callback_target_data_op_t`, see Section 4.5.2.25 on page 488.

# 3.7 Interoperability Routines

## 3.7.1 `omp_get_num_interop_properties`

### Summary

The `omp_get_num_interop_properties` routine retrieves the number of implementation-defined properties available from an `omp_interop_t` object.

### Format

```
int omp_get_num_interop_properties(omp_interop_t interop);
```

### Effect

The `omp_get_num_interop_properties` routine returns the number of implementation-defined properties available on *interop*, such that the returned value minus `OMP_INTEROP_FIRST` is the total number of properties available. This value allows iteration from `OMP_INTEROP_FIRST` to the result of calling `omp_get_num_interop_properties`.

## 3.7.2 `omp_get_interop_property`

### Summary

The `omp_get_interop_property` routine retrieves property data from an `omp_interop_t` object.



**TABLE 3.1:** Required values of the `omp_interop_property` enum

enum name	name	property
<code>OMP_INTEROP_TYPE = -1</code>	<code>type</code>	pointer to an <code>intptr_t</code> that represents the type of context
<code>OMP_INTEROP_TYPE_NAME = -2</code>	<code>type_name</code>	C string value that represents the type of context
<code>OMP_INTEROP_DEVICE = -3</code>	<code>device</code>	a foreign device handle
<code>OMP_INTEROP_DEVICE_CONTEXT = -4</code>	<code>device_context</code>	a handle to an instance of a foreign device context
<code>OMP_INTEROP_TASKSYNC = -5</code>	<code>tasksync</code>	a handle to a synchronization object of a foreign execution context
<code>OMP_INTEROP_FIRST = -5</code>		

## Format

```

int omp_get_interop_property(
    const omp_interop_t *interop,
    int property_id,
    void **data
);

```

## Effect

The `omp_get_interop_property` routine returns a status. The behavior of the routine is unspecified if an invalid object is provided. The returned status is zero if a valid `omp_interop_t` object is provided and the property is in the range between `OMP_INTEROP_FIRST` and `omp_get_num_interop_properties(\plc {interop})` minus one. The returned status is -1 if the object has the value `OMP_INTEROP_NONE`. The returned status is -2 if a valid `omp_interop_t` object is provided but the provided `property_id` is not in the range between `OMP_INTEROP_FIRST` and `omp_get_num_interop_properties(\plc {interop})` minus one. The value of the requested property is assigned into `*data` if and only if a valid `omp_interop_t` object is provided, `property_id` is in the range between `OMP_INTEROP_FIRST` and `omp_get_num_interop_properties(\plc {interop})` minus one, and `data` is not `NULL`.

All interop objects are required to support the values `OMP_INTEROP_FIRST` through -1 for the `property_id` argument.

OpenMP reserves all negative values of an `int` for properties; implementation defined properties may use zero and positive values. The value of `OMP_INTEROP_FIRST` is intended to match the

1 lowest property defined in a version of the specification and, thus, a program should not rely on the  
2 value as it is likely to change.

3 The value of *property\_id* must be in the range between **OMP\_INTEROP\_FIRST** and  
4 **omp\_get\_num\_interop\_properties (interop)** minus one.

### 5 **Restrictions**

- 6 • Memory referenced by **data** from the **omp\_get\_interop\_property** routine is managed  
7 by the OpenMP implementation and should not be freed or modified.

### 8 **Cross References**


- 9 • **interop** construct, see Section 2.15.8 on page 191.
- 10 • **omp\_get\_num\_interop\_properties** routine, see Section 3.7.1 on page 410.

## 11 **3.7.3 omp\_get\_interop\_property\_name**

### 12 **Summary**

13 The **omp\_get\_interop\_property\_name** routine retrieves the name of a property from the  
14 **omp\_interop\_t** object that is specified by *interop*.

### 15 **Format**

```
16  C / C++  
17 const char* omp_get_interop_property_name (  
18     omp_interop_t interop,  
19     int property_id  
20 );
```

### 20 **Effect**

21 The **omp\_get\_interop\_property\_name** routine returns the name of a property retrieved  
22 from an **omp\_interop\_t** object.

23 If *property\_id* is not in the range between **OMP\_INTEROP\_FIRST** and  
24 **omp\_get\_num\_interop\_properties (\plc {interop})** minus one or if *interop* has  
25 the value **OMP\_INTEROP\_NONE**, **NULL** will be returned.

### 26 **Restrictions**

- 27 • Memory returned from the **omp\_get\_interop\_property\_name** routine is managed by  
28 the OpenMP implementation and should not be freed or modified.

## 3.8 Memory Management Routines

This section describes routines that support memory management on the current device.

Instances of memory management types must be accessed only through the routines described in this section; programs that otherwise access instances of these types are non-conforming.

### 3.8.1 Memory Management Types

The following type definitions are used by the memory management routines:

C/C++

```
typedef enum omp_alloctrail_key_t {
    omp_atk_sync_hint = 1,
    omp_atk_alignment = 2,
    omp_atk_access = 3,
    omp_atk_pool_size = 4,
    omp_atk_fallback = 5,
    omp_atk_fb_data = 6,
    omp_atk_pinned = 7,
    omp_atk_partition = 8
} omp_alloctrail_key_t;

typedef enum omp_alloctrail_value_t {
    omp_atv_false = 0,
    omp_atv_true = 1,
    omp_atv_default = 2,
    omp_atv_contended = 3,
    omp_atv_uncontended = 4,
    omp_atv_sequential = 5,
    omp_atv_private = 6,
    omp_atv_all = 7,
    omp_atv_thread = 8,
    omp_atv_pteam = 9,
    omp_atv_cgroup = 10,
    omp_atv_default_mem_fb = 11,
    omp_atv_null_fb = 12,
    omp_atv_abort_fb = 13,
    omp_atv_allocator_fb = 14,
    omp_atv_environment = 15,
    omp_atv_nearest = 16,
    omp_atv_blocked = 17,
    omp_atv_interleaved = 18
} omp_alloctrail_value_t;
```

```

1
2 typedef struct omp_alloctrail_t {
3     omp_alloctrail_key_t key;
4     omp_uintptr_t value;
5 } omp_alloctrail_t;

```

← C / C++ →

← Fortran →

```

6
7 integer(kind=omp_alloctrail_key_kind), &
8     parameter :: omp_atk_sync_hint = 1
9 integer(kind=omp_alloctrail_key_kind), &
10    parameter :: omp_atk_alignment = 2
11 integer(kind=omp_alloctrail_key_kind), &
12    parameter :: omp_atk_access = 3
13 integer(kind=omp_alloctrail_key_kind), &
14    parameter :: omp_atk_pool_size = 4
15 integer(kind=omp_alloctrail_key_kind), &
16    parameter :: omp_atk_fallback = 5
17 integer(kind=omp_alloctrail_key_kind), &
18    parameter :: omp_atk_fb_data = 6
19 integer(kind=omp_alloctrail_key_kind), &
20    parameter :: omp_atk_pinned = 7
21 integer(kind=omp_alloctrail_key_kind), &
22    parameter :: omp_atk_partition = 8
23
24 integer(kind=omp_alloctrail_val_kind), &
25    parameter :: omp_atv_false = 0
26 integer(kind=omp_alloctrail_val_kind), &
27    parameter :: omp_atv_true = 1
28 integer(kind=omp_alloctrail_val_kind), &
29    parameter :: omp_atv_default = 2
30 integer(kind=omp_alloctrail_val_kind), &
31    parameter :: omp_atv_contended = 3
32 integer(kind=omp_alloctrail_val_kind), &
33    parameter :: omp_atv_uncontended = 4
34 integer(kind=omp_alloctrail_val_kind), &
35    parameter :: omp_atv_sequential = 5
36 integer(kind=omp_alloctrail_val_kind), &
37    parameter :: omp_atv_private = 6
38 integer(kind=omp_alloctrail_val_kind), &
39    parameter :: omp_atv_all = 7
40 integer(kind=omp_alloctrail_val_kind), &
41    parameter :: omp_atv_thread = 8

```

```

1 integer(kind=omp_alloctrail_val_kind), &
2   parameter :: omp_atv_pteam = 9
3 integer(kind=omp_alloctrail_val_kind), &
4   parameter :: omp_atv_cgroup = 10
5 integer(kind=omp_alloctrail_val_kind), &
6   parameter :: omp_atv_default_mem_fb = 11
7 integer(kind=omp_alloctrail_val_kind), &
8   parameter :: omp_atv_null_fb = 12
9 integer(kind=omp_alloctrail_val_kind), &
10  parameter :: omp_atv_abort_fb = 13
11 integer(kind=omp_alloctrail_val_kind), &
12  parameter :: omp_atv_allocator_fb = 14
13 integer(kind=omp_alloctrail_val_kind), &
14  parameter :: omp_atv_environment = 15
15 integer(kind=omp_alloctrail_val_kind), &
16  parameter :: omp_atv_nearest = 16
17 integer(kind=omp_alloctrail_val_kind), &
18  parameter :: omp_atv_blocked = 17
19 integer(kind=omp_alloctrail_val_kind), &
20  parameter :: omp_atv_interleaved = 18
21
22 type omp_alloctrail
23   integer(kind=omp_alloctrail_key_kind) key
24   integer(kind=omp_alloctrail_val_kind) value
25 end type omp_alloctrail
26
27 integer(kind=omp_allocator_handle_kind), &
28   parameter :: omp_null_allocator = 0

```

Fortran

## 29 3.8.2 omp\_init\_allocator

### 30 Summary

31 The `omp_init_allocator` routine initializes an allocator and associates it with a memory  
32 space.

### 33 Format

C / C++

```

34 omp_allocator_handle_t omp_init_allocator (
35   omp_memspace_handle_t memspace,
36   int ntraits,
37   const omp_alloctrail_t traits[]
38 );

```

C / C++

Fortran

```
1 integer(kind=omp_allocator_handle_kind) &  
2 function omp_init_allocator ( memspace, ntraits, traits )  
3 integer(kind=omp_memspace_handle_kind),intent(in) :: memspace  
4 integer,intent(in) :: ntraits  
5 type(omp_alloctrait),intent(in) :: traits(*)
```

Fortran

## Constraints on Arguments

The *memspace* argument must be one of the predefined memory spaces defined in Table 2.8.

If the *ntraits* argument is greater than zero then the *traits* argument must specify at least that many traits. If it specifies fewer than *ntraits* traits the behavior is unspecified.

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\_init\_allocator** region is all threads on a device. The effect of executing this routine is not related to any specific region that corresponds to any construct or API routine.

## Effect

The **omp\_init\_allocator** routine creates a new allocator that is associated with the *memspace* memory space and returns a handle to it. All allocations 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 2.9 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 **omp\_null\_allocator** handle 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, see Section 2.14.1 on page 158.
- Memory Allocators, see Section 2.14.2 on page 159.

### 3.8.3 `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_handle_t allocator);`

C / C++  
Fortran  
`subroutine omp_destroy_allocator ( allocator )  
integer (kind=omp_allocator_handle_kind), intent (in) :: allocator`

Fortran

#### Constraints on Arguments

The *allocator* argument must not represent 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 that corresponds 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 unspecified behavior.

If *allocator* is `omp_null_allocator` then this routine will have no effect.

#### Cross References

- Memory Allocators, see Section 2.14.2 on page 159.

### 3.8.4 `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

C / C++

```
void omp_set_default_allocator (omp_allocator_handle_t allocator);
```

C / C++

Fortran

```
subroutine omp_set_default_allocator ( allocator )  
integer(kind=omp_allocator_handle_kind),intent(in) :: allocator
```

Fortran

## Constraints on Arguments

The *allocator* argument must be a valid memory allocator handle.

## 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 2.6 on page 66.
- Memory Allocators, see Section 2.14.2 on page 159.
- **omp\_alloc** routine, see Section 3.8.6 on page 419.

## 3.8.5 omp\_get\_default\_allocator

### Summary

The **omp\_get\_default\_allocator** routine returns a handle to the memory allocator to be used by allocation calls, **allocate** directives and **allocate** clauses that do not specify an allocator.

## Format

C / C++

```
omp_allocator_handle_t omp_get_default_allocator (void);
```

C / C++

Fortran

```
integer(kind=omp_allocator_handle_kind) &  
function omp_get_default_allocator ()
```

Fortran



1       **Binding**

2       The binding task set for an `omp_get_default_allocator` region is the binding implicit task.

3       **Effect**

4       The effect of this routine is to return the value of the *def-allocator-var* ICV of the binding implicit  
5       task.

6       **Cross References**

- 7       • *def-allocator-var* ICV, see Section 2.6 on page 66.
- 8       • Memory Allocators, see Section 2.14.2 on page 159.
- 9       • `omp_alloc` routine, see Section 3.8.6 on page 419.

▼                   C / C++                   ▼

10       **3.8.6 omp\_alloc**

11       **Summary**

12       The `omp_alloc` routine requests a memory allocation from a memory allocator.

13       **Format**

▼                   C                   ▼

14       | `void *omp_alloc (size_t size, omp_allocator_handle_t allocator);`

▲                   C                   ▲

▼                   C++                  ▼

15       | `void *omp_alloc(  
16       |     size_t size,  
17       |     omp_allocator_handle_t allocator=omp_null_allocator  
18       | );`

▲                   C++                  ▲

19       **Constraints on Arguments**

20       Unless `dynamic_allocators` appears on a `requires` directive in the same compilation unit,  
21       `omp_alloc` invocations that appear in `target` regions must not pass `omp_null_allocator`  
22       as the *allocator* argument, which must be a constant expression that evaluates to one of the  
23       predefined memory allocator values.

## 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.

Allocated memory will be byte aligned to at least the alignment required by `malloc`.

## Cross References

- Memory allocators, see Section 2.14.2 on page 159.

## 3.8.7 `omp_free`

### Summary

The `omp_free` routine deallocates previously allocated memory.

### Format

C

```
void omp_free (void *ptr, omp_allocator_handle_t allocator);
```

C

C++

```
void omp_free(  
    void *ptr,  
    omp_allocator_handle_t allocator=omp_null_allocator  
);
```

C++

### Effect

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 determine that value automatically. 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.

### Cross References

- Memory allocators, see Section 2.14.2 on page 159.

C / C++

## 3.9 Tool Control Routine

### Summary

The `omp_control_tool` routine enables a program to pass commands to an active tool.

### Format

C / C++  
`int omp_control_tool(int command, int modifier, void *arg);`

C / C++  
Fortran  
`integer function omp_control_tool(command, modifier)  
integer (kind=omp_control_tool_kind) command  
integer modifier`

### Description

An OpenMP program may use `omp_control_tool` to pass commands to a tool. An application can use `omp_control_tool` to request that a tool starts or restarts data collection when a code region of interest is encountered, that a tool pauses data collection when leaving the region of interest, that a tool flushes any data that it has collected so far, or that a tool ends data collection. Additionally, `omp_control_tool` can be used to pass tool-specific commands to a particular tool.

The following types correspond to return values from `omp_control_tool`:

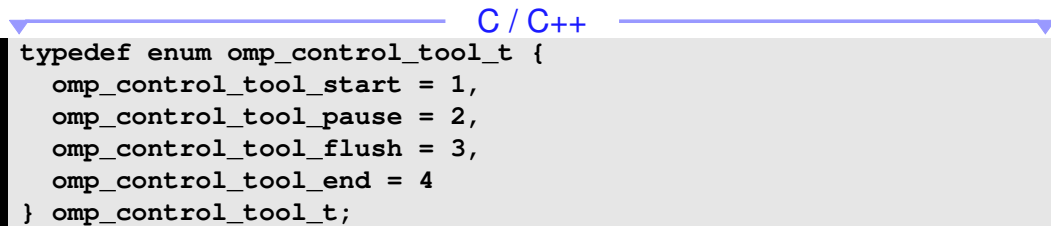
C / C++  
`typedef enum omp_control_tool_result_t {  
 omp_control_tool_notool = -2,  
 omp_control_tool_nocallback = -1,  
 omp_control_tool_success = 0,  
 omp_control_tool_ignored = 1  
} omp_control_tool_result_t;`

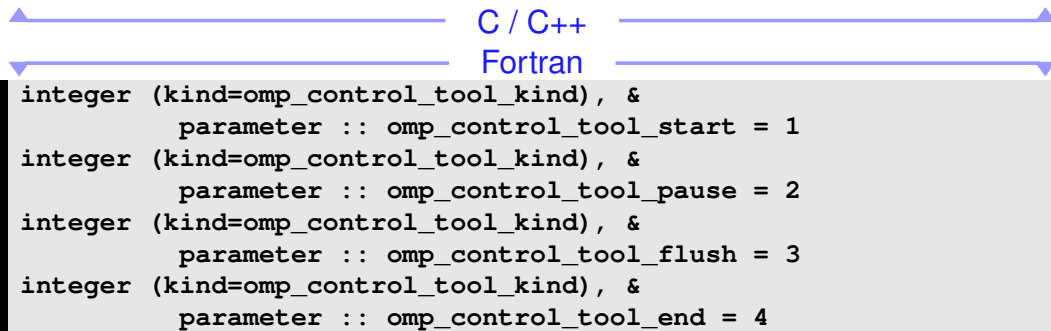
C / C++  
Fortran  
`integer (kind=omp_control_tool_result_kind), &  
 parameter :: omp_control_tool_notool = -2  
integer (kind=omp_control_tool_result_kind), &  
 parameter :: omp_control_tool_nocallback = -1  
integer (kind=omp_control_tool_result_kind), &  
 parameter :: omp_control_tool_success = 0  
integer (kind=omp_control_tool_result_kind), &  
 parameter :: omp_control_tool_ignored = 1`

1 If the OMPT interface state is inactive, the OpenMP implementation returns  
2 **omp\_control\_tool\_notool**. If the OMPT interface state is active, but no callback is  
3 registered for the *tool-control* event, the OpenMP implementation returns  
4 **omp\_control\_tool\_nocallback**. An OpenMP implementation may return other  
5 implementation-defined negative values strictly smaller than -64; an application may assume that  
6 any negative return value indicates that a tool has not received the command. A return value of  
7 **omp\_control\_tool\_success** indicates that the tool has performed the specified command. A  
8 return value of **omp\_control\_tool\_ignored** indicates that the tool has ignored the specified  
9 command. A tool may return other positive values strictly greater than 64 that are tool-defined.

## 10 Constraints on Arguments

11 The following enumeration type defines four standard commands. Table 3.2 describes the actions  
12 that these commands request from a tool.

13  C / C++  
14 `typedef enum omp_control_tool_t {`  
15  `omp_control_tool_start = 1,`  
16  `omp_control_tool_pause = 2,`  
17  `omp_control_tool_flush = 3,`  
18  `omp_control_tool_end = 4`  
`} omp_control_tool_t;`

19  C / C++  
Fortran  
20 `integer (kind=omp_control_tool_kind), &`  
21  `parameter :: omp_control_tool_start = 1`  
22 `integer (kind=omp_control_tool_kind), &`  
23  `parameter :: omp_control_tool_pause = 2`  
24 `integer (kind=omp_control_tool_kind), &`  
25  `parameter :: omp_control_tool_flush = 3`  
26 `integer (kind=omp_control_tool_kind), &`  
 `parameter :: omp_control_tool_end = 4`

27 Tool-specific values for *command* must be greater or equal to 64. Tools must ignore *command*  
28 values that they are not explicitly designed to handle. Other values accepted by a tool for *command*,  
29 and any values for *modifier* and *arg* are tool-defined.

**TABLE 3.2:** 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.

### 1 Execution Model Events

2 The *tool-control* event occurs in the thread that encounters a call to `omp_control_tool` at a  
3 point inside its corresponding OpenMP region.

### 4 Tool Callbacks

5 A thread dispatches a registered `ompt_callback_control_tool` callback for each  
6 occurrence of a *tool-control* event. The callback executes in the context of the call that occurs in the  
7 user program and has type signature `ompt_callback_control_tool_t`. The callback may  
8 return any non-negative value, which will be returned to the application by the OpenMP  
9 implementation as the return value of the `omp_control_tool` call that triggered the callback.

10 Arguments passed to the callback are those passed by the user to `omp_control_tool`. If the  
11 call is made in Fortran, the tool will be passed `NULL` as the third argument to the callback. If any of  
12 the four standard commands is presented to a tool, the tool will ignore the *modifier* and *arg*  
13 argument values.

### 14 Restrictions

15 Access to the state of an OpenMP first-party tool has the following restriction:

- 16 • An application may access the tool state modified by an OMPT callback only by using  
17 `omp_control_tool`.

### 18 Cross References

- 19 • OMPT Interface, see Chapter 4 on page 425
- 20 • `ompt_callback_control_tool_t`, see Section 4.5.2.29 on page 495.

*This page intentionally left blank*

# 4 OMPT Interface

This chapter describes OMPT, which is an interface for *first-party* tools. *First-party* tools are linked or loaded directly into the OpenMP program. OMPT defines mechanisms to initialize a tool, to examine OpenMP state associated with an OpenMP thread, to interpret the call stack of an OpenMP thread, to receive notification about OpenMP *events*, to trace activity on OpenMP target devices, to assess implementation-dependent details of an OpenMP implementation (such as supported states and mutual exclusion implementations), and to control a tool from an OpenMP application.

## 4.1 OMPT Interfaces Definitions

▼ C / C++ ▼

A compliant implementation must supply a set of definitions for the OMPT runtime entry points, OMPT callback signatures, and the special data types of their parameters and return values. These definitions, which are listed throughout this chapter, and their associated declarations shall be provided in a header file named `ompt-tools.h`. In addition, the set of definitions may specify other implementation-specific values.

The `ompt_start_tool` function is an external function with C linkage.

▲ C / C++ ▲

## 4.2 Activating a First-Party Tool

To activate a tool, an OpenMP implementation first determines whether the tool should be initialized. If so, the OpenMP implementation invokes the initializer of the tool, which enables the tool to prepare to monitor execution on the host. The tool may then also arrange to monitor computation that executes on target devices. This section explains how the tool and an OpenMP implementation interact to accomplish these tasks.

### 4.2.1 `ompt_start_tool`

#### Summary

In order to use the OMPT interface provided by an OpenMP implementation, a tool must implement the `ompt_start_tool` function, through which the OpenMP implementation initializes the tool.

## Format

```
1
2 ompt_start_tool_result_t *ompt_start_tool(  
3     unsigned int omp_version,  
4     const char *runtime_version  
5 );
```

## Description

For a tool to use the OMPT interface that an OpenMP implementation provides, the tool must define a globally-visible implementation of the function `ompt_start_tool`. The tool indicates that it will use the OMPT interface that an OpenMP implementation provides by returning a non-null pointer to an `ompt_start_tool_result_t` structure from the `ompt_start_tool` implementation that it provides. The `ompt_start_tool_result_t` structure contains pointers to tool initialization and finalization callbacks as well as a tool data word that an OpenMP implementation must pass by reference to these callbacks. A tool may return `NULL` from `ompt_start_tool` to indicate that it will not use the OMPT interface in a particular execution.

A tool may use the `omp_version` argument to determine if it is compatible with the OMPT interface that the OpenMP implementation provides.

## Description of Arguments

The argument `omp_version` is the value of the `_OPENMP` version macro associated with the OpenMP API implementation. This value identifies the OpenMP API version that an OpenMP implementation supports, which specifies the version of the OMPT interface that it supports.

The argument `runtime_version` is a version string that unambiguously identifies the OpenMP implementation.

## Constraints on Arguments

The argument `runtime_version` must be an immutable string that is defined for the lifetime of a program execution.

## Effect

If a tool returns a non-null pointer to an `ompt_start_tool_result_t` structure, an OpenMP implementation will call the tool initializer specified by the `initialize` field in this structure before beginning execution of any OpenMP construct or completing execution of any environment routine invocation; the OpenMP implementation will call the tool finalizer specified by the `finalize` field in this structure when the OpenMP implementation shuts down.

## Cross References

- `ompt_start_tool_result_t`, see Section 4.4.1 on page 438.



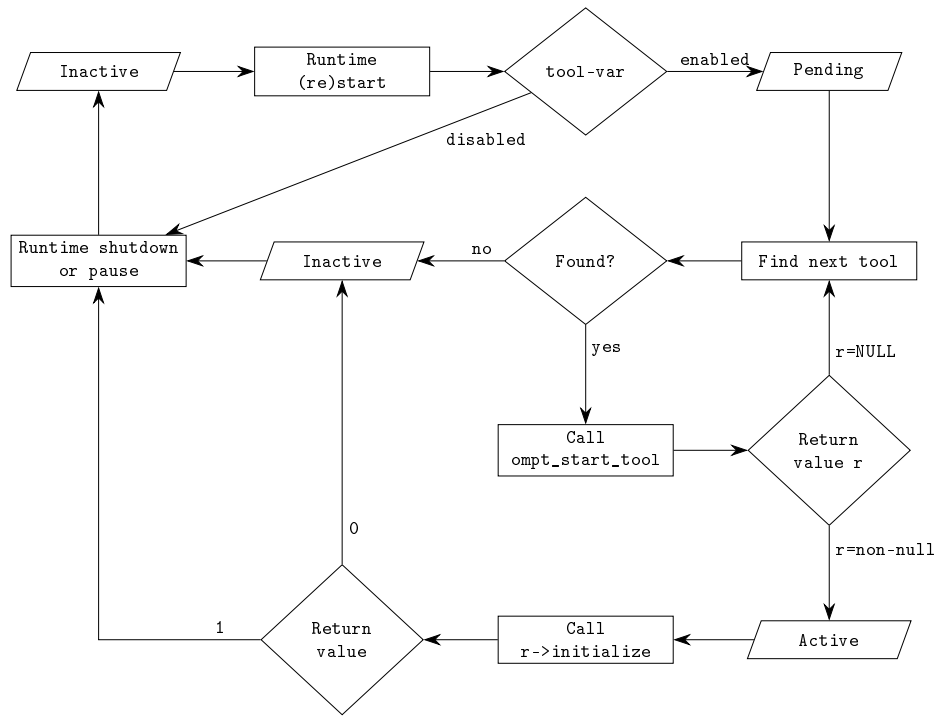


FIGURE 4.1: First-Party Tool Activation Flow Chart

## 4.2.2 Determining Whether a First-Party Tool Should be Initialized

An OpenMP implementation examines the *tool-var* ICV as one of its first initialization steps. If the value of *tool-var* is *disabled*, the initialization continues without a check for the presence of a tool and the functionality of the OMPT interface will be unavailable as the program executes. In this case, the OMPT interface state remains *inactive*.

Otherwise, the OMPT interface state changes to *pending* and the OpenMP implementation activates any first-party tool that it finds. A tool can provide a definition of `ompt_start_tool` to an OpenMP implementation in three ways:

- By statically-linking its definition of `ompt_start_tool` into an OpenMP application;
- By introducing a dynamically-linked library that includes its definition of `ompt_start_tool` into the application's address space; or
- By providing, in the *tool-libraries-var* ICV, the name of a dynamically-linked library that is appropriate for the architecture and operating system used by the application and that includes a

1 definition of `ompt_start_tool`.

2 If the value of `tool-var` is *enabled*, the OpenMP implementation must check if a tool has provided  
3 an implementation of `ompt_start_tool`. The OpenMP implementation first checks if a  
4 tool-provided implementation of `ompt_start_tool` is available in the address space, either  
5 statically-linked into the application or in a dynamically-linked library loaded in the address space.  
6 If multiple implementations of `ompt_start_tool` are available, the OpenMP implementation  
7 will use the first tool-provided implementation of `ompt_start_tool` that it finds.

8 If the implementation does not find a tool-provided implementation of `ompt_start_tool` in the  
9 address space, it consults the `tool-libraries-var` ICV, which contains a (possibly empty) list of  
10 dynamically-linked libraries. As described in detail in Section 6.19 on page 600, the libraries in  
11 `tool-libraries-var` are then searched for the first usable implementation of `ompt_start_tool`  
12 that one of the libraries in the list provides.

13 If the implementation finds a tool-provided definition of `ompt_start_tool`, it invokes that  
14 method; if a `NULL` pointer is returned, the OMPT interface state remains *pending* and the  
15 implementation continues to look for implementations of `ompt_start_tool`; otherwise a  
16 non-null pointer to an `ompt_start_tool_result_t` structure is returned, the OMPT  
17 interface state changes to *active* and the OpenMP implementation makes the OMPT interface  
18 available as the program executes. In this case, as the OpenMP implementation completes its  
19 initialization, it initializes the OMPT interface.

20 If no tool can be found, the OMPT interface state changes to *inactive*.

## 21 Cross References

- 22 • `tool-libraries-var` ICV, see Section 2.6 on page 66.
- 23 • `tool-var` ICV, see Section 2.6 on page 66.
- 24 • `ompt_start_tool` function, see Section 4.2.1 on page 425.
- 25 • `ompt_start_tool_result_t` type, see Section 4.4.1 on page 438.

## 26 4.2.3 Initializing a First-Party Tool

27 To initialize the OMPT interface, the OpenMP implementation invokes the tool initializer that is  
28 specified in the `ompt_start_tool_result_t` structure that is indicated by the non-null  
29 pointer that `ompt_start_tool` returns. The initializer is invoked prior to the occurrence of any  
30 OpenMP *event*.

31 A tool initializer, described in Section 4.5.1.1 on page 460, uses the function specified in its *lookup*  
32 argument to look up pointers to OMPT interface runtime entry points that the OpenMP  
33 implementation provides; this process is described in Section 4.2.3.1 on page 429. Typically, a tool  
34 initializer obtains a pointer to the `ompt_set_callback` runtime entry point with type signature  
35 `ompt_set_callback_t` and then uses this runtime entry point to register tool callbacks for  
36 OpenMP events, as described in Section 4.2.4 on page 430.

1 A tool initializer may use the `ompt_enumerate_states` runtime entry point, which has type  
2 signature `ompt_enumerate_states_t`, to determine the thread states that an OpenMP  
3 implementation employs. Similarly, it may use the `ompt_enumerate_mutex_impls` runtime  
4 entry point, which has type signature `ompt_enumerate_mutex_impls_t`, to determine the  
5 mutual exclusion implementations that the OpenMP implementation employs.

6 If a tool initializer returns a non-zero value, the OMPT interface state remains *active* for the  
7 execution; otherwise, the OMPT interface state changes to *inactive*.

## 8 Cross References

- 9 • `ompt_start_tool` function, see Section 4.2.1 on page 425.
- 10 • `ompt_start_tool_result_t` type, see Section 4.4.1 on page 438.
- 11 • `ompt_initialize_t` type, see Section 4.5.1.1 on page 460.
- 12 • `ompt_callback_thread_begin_t` type, see Section 4.5.2.1 on page 462.
- 13 • `ompt_enumerate_states_t` type, see Section 4.6.1.1 on page 497.
- 14 • `ompt_enumerate_mutex_impls_t` type, see Section 4.6.1.2 on page 498.
- 15 • `ompt_set_callback_t` type, see Section 4.6.1.3 on page 499.
- 16 • `ompt_function_lookup_t` type, see Section 4.6.3 on page 526.

### 17 4.2.3.1 Binding Entry Points in the OMPT Callback Interface

18 Functions that an OpenMP implementation provides to support the OMPT interface are not defined  
19 as global function symbols. Instead, they are defined as runtime entry points that a tool can only  
20 identify through the *lookup* function that is provided as an argument with type signature  
21 `ompt_function_lookup_t` to the tool initializer. A tool can use this function to obtain a  
22 pointer to each of the runtime entry points that an OpenMP implementation provides to support the  
23 OMPT interface. Once a tool has obtained a *lookup* function, it may employ it at any point in the  
24 future.

25 For each runtime entry point in the OMPT interface for the host device, Table 4.1 provides the  
26 string name by which it is known and its associated type signature. Implementations can provide  
27 additional implementation-specific names and corresponding entry points. Any names that begin  
28 with `ompt_` are reserved names.

29 During initialization, a tool should look up each runtime entry point in the OMPT interface by  
30 name and bind a pointer maintained by the tool that can later be used to invoke the entry point. The  
31 entry points described in Table 4.1 enable a tool to assess the thread states and mutual exclusion  
32 implementations that an OpenMP implementation supports, to register tool callbacks, to inspect  
33 registered callbacks, to introspect OpenMP state associated with threads, and to use tracing to  
34 monitor computations that execute on target devices.

1 Detailed information about each runtime entry point listed in Table 4.1 is included as part of the  
2 description of its type signature.

### 3 **Cross References**

- 4 • `ompt_enumerate_states_t` type, see Section 4.6.1.1 on page 497.
- 5 • `ompt_enumerate_mutex_impls_t` type, see Section 4.6.1.2 on page 498.
- 6 • `ompt_set_callback_t` type, see Section 4.6.1.3 on page 499.
- 7 • `ompt_get_callback_t` type, see Section 4.6.1.4 on page 501.
- 8 • `ompt_get_thread_data_t` type, see Section 4.6.1.5 on page 501.
- 9 • `ompt_get_num_procs_t` type, see Section 4.6.1.6 on page 502.
- 10 • `ompt_get_num_places_t` type, see Section 4.6.1.7 on page 503.
- 11 • `ompt_get_place_proc_ids_t` type, see Section 4.6.1.8 on page 503.
- 12 • `ompt_get_place_num_t` type, see Section 4.6.1.9 on page 504.
- 13 • `ompt_get_partition_place_nums_t` type, see Section 4.6.1.10 on page 505.
- 14 • `ompt_get_proc_id_t` type, see Section 4.6.1.11 on page 506.
- 15 • `ompt_get_state_t` type, see Section 4.6.1.12 on page 506.
- 16 • `ompt_get_parallel_info_t` type, see Section 4.6.1.13 on page 508.
- 17 • `ompt_get_task_info_t` type, see Section 4.6.1.14 on page 509.
- 18 • `ompt_get_task_memory_t` type, see Section 4.6.1.15 on page 511.
- 19 • `ompt_get_target_info_t` type, see Section 4.6.1.16 on page 512.
- 20 • `ompt_get_num_devices_t` type, see Section 4.6.1.17 on page 513.
- 21 • `ompt_get_unique_id_t` type, see Section 4.6.1.18 on page 514.
- 22 • `ompt_finalize_tool_t` type, see Section 4.6.1.19 on page 514.
- 23 • `ompt_function_lookup_t` type, see Section 4.6.3 on page 526.

## 24 **4.2.4 Monitoring Activity on the Host with OMPT**

25 To monitor the execution of an OpenMP program on the host device, a tool initializer must register  
26 to receive notification of events that occur as an OpenMP program executes. A tool can use the  
27 `ompt_set_callback` runtime entry point to register callbacks for OpenMP events. The return  
28 codes for `ompt_set_callback` use the `ompt_set_result_t` enumeration type. If the  
29 `ompt_set_callback` runtime entry point is called outside a tool initializer, registration of  
30 supported callbacks may fail with a return value of `ompt_set_error`.

**TABLE 4.1:** OMPT Callback Interface Runtime Entry Point Names and Their Type Signatures

Entry Point String Name	Type signature
<code>"ompt_enumerate_states"</code>	<code>ompt_enumerate_states_t</code>
<code>"ompt_enumerate_mutex_impls"</code>	<code>ompt_enumerate_mutex_impls_t</code>
<code>"ompt_set_callback"</code>	<code>ompt_set_callback_t</code>
<code>"ompt_get_callback"</code>	<code>ompt_get_callback_t</code>
<code>"ompt_get_thread_data"</code>	<code>ompt_get_thread_data_t</code>
<code>"ompt_get_num_places"</code>	<code>ompt_get_num_places_t</code>
<code>"ompt_get_place_proc_ids"</code>	<code>ompt_get_place_proc_ids_t</code>
<code>"ompt_get_place_num"</code>	<code>ompt_get_place_num_t</code>
<code>"ompt_get_partition_place_nums"</code>	<code>ompt_get_partition_place_nums_t</code>
<code>"ompt_get_proc_id"</code>	<code>ompt_get_proc_id_t</code>
<code>"ompt_get_state"</code>	<code>ompt_get_state_t</code>
<code>"ompt_get_parallel_info"</code>	<code>ompt_get_parallel_info_t</code>
<code>"ompt_get_task_info"</code>	<code>ompt_get_task_info_t</code>
<code>"ompt_get_task_memory"</code>	<code>ompt_get_task_memory_t</code>
<code>"ompt_get_num_devices"</code>	<code>ompt_get_num_devices_t</code>
<code>"ompt_get_num_procs"</code>	<code>ompt_get_num_procs_t</code>
<code>"ompt_get_target_info"</code>	<code>ompt_get_target_info_t</code>
<code>"ompt_get_unique_id"</code>	<code>ompt_get_unique_id_t</code>
<code>"ompt_finalize_tool"</code>	<code>ompt_finalize_tool_t</code>

1 All callbacks registered with `ompt_set_callback` or returned by `ompt_get_callback` use  
2 the dummy type signature `ompt_callback_t`.

3 Table 4.2 shows the valid registration return codes of the `ompt_set_callback` runtime entry  
4 point with specific values of its *event* argument. For callbacks for which `ompt_set_always` is  
5 the only registration return code that is allowed, an OpenMP implementation must guarantee that  
6 the callback will be invoked every time that a runtime event that is associated with it occurs.  
7 Support for such callbacks is required in a minimal implementation of the OMPT interface. For  
8 callbacks for which the `ompt_set_callback` runtime entry may return values other than  
9 `ompt_set_always`, whether an OpenMP implementation invokes a registered callback never,  
10 sometimes, or always is implementation-defined. If registration for a callback allows a return code  
11 of `ompt_set_never`, support for invoking such a callback may not be present in a minimal  
12 implementation of the OMPT interface. The return code from registering a callback indicates the  
13 implementation-defined level of support for the callback.

14 Two techniques reduce the size of the OMPT interface. First, in cases where events are naturally  
15 paired, for example, the beginning and end of a region, and the arguments needed by the callback at  
16 each endpoint are identical, a tool registers a single callback for the pair of events, with  
17 `ompt_scope_begin` or `ompt_scope_end` provided as an argument to identify for which  
18 endpoint the callback is invoked. Second, when a class of events is amenable to uniform treatment,  
19 OMPT provides a single callback for that class of events, for example, an  
20 `ompt_callback_sync_region_wait` callback is used for multiple kinds of synchronization  
21 regions, such as barrier, taskwait, and taskgroup regions. Some events, for example,  
22 `ompt_callback_sync_region_wait`, use both techniques.

## 23 Cross References

- 24 • `ompt_set_result_t` type, see Section 4.4.4.2 on page 443.
- 25 • `ompt_set_callback_t` type, see Section 4.6.1.3 on page 499.
- 26 • `ompt_get_callback_t` type, see Section 4.6.1.4 on page 501.

## 27 4.2.5 Tracing Activity on Target Devices with OMPT

28 A target device may or may not initialize a full OpenMP runtime system. Unless it does, it may not  
29 be possible to monitor activity on a device using a tool interface based on callbacks. To  
30 accommodate such cases, the OMPT interface defines a monitoring interface for tracing activity on  
31 target devices. Tracing activity on a target device involves the following steps:

- 32 • To prepare to trace activity on a target device, a tool must register for an  
33 `ompt_callback_device_initialize` callback. A tool may also register for an  
34 `ompt_callback_device_load` callback to be notified when code is loaded onto a target  
35 device or an `ompt_callback_device_unload` callback to be notified when code is  
36 unloaded from a target device. A tool may also optionally register an  
37 `ompt_callback_device_finalize` callback.

**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_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 = `ompt_set_never`                      S = `ompt_set_sometimes`  
P = `ompt_set_sometimes_paired`            A = `ompt_set_always`

- 1 • When an OpenMP implementation initializes a target device, the OpenMP implementation  
2 dispatches the device initialization callback of the tool on the host device. If the OpenMP  
3 implementation or target device does not support tracing, the OpenMP implementation passes  
4 **NULL** to the device initializer of the tool for its *lookup* argument; otherwise, the OpenMP  
5 implementation passes a pointer to a device-specific runtime entry point with type signature  
6 **ompt\_function\_lookup\_t** to the device initializer of the tool.
- 7 • If a non-null *lookup* pointer is provided to the device initializer of the tool, the tool may use it to  
8 determine the runtime entry points in the tracing interface that are available for the device and  
9 may bind the returned function pointers to tool variables. Table 4.3 indicates the names of  
10 runtime entry points that may be available for a device; an implementations may provide  
11 additional implementation-defined names and corresponding entry points. The driver for the  
12 device provides the runtime entry points that enable a tool to control the trace collection interface  
13 of the device. The *native* trace format that the interface uses may be device specific and the  
14 available kinds of trace records are implementation-defined. Some devices may allow a tool to  
15 collect traces of records in a standard format known as OMPT trace records. Each OMPT trace  
16 record serves as a substitute for an OMPT callback that cannot be made on the device. The fields  
17 in each trace record type are defined in the description of the callback that the record represents.  
18 If this type of record is provided then the *lookup* function returns values for the runtime entry  
19 points **ompt\_set\_trace\_ompt** and **ompt\_get\_record\_ompt**, which support collecting  
20 and decoding OMPT traces. If the native tracing format for a device is the OMPT format then  
21 tracing can be controlled using the runtime entry points for native or OMPT tracing.
- 22 • The tool uses the **ompt\_set\_trace\_native** and/or the **ompt\_set\_trace\_ompt**  
23 runtime entry point to specify what types of events or activities to monitor on the device. The  
24 return codes for **ompt\_set\_trace\_ompt** and **ompt\_set\_trace\_native** use the  
25 **ompt\_set\_result\_t** enumeration type. If the **ompt\_set\_trace\_native** /or the  
26 **ompt\_set\_trace\_ompt** runtime entry point is called outside a device initializer, registration  
27 of supported callbacks may fail with a return code of **ompt\_set\_error**.
- 28 • The tool initiates tracing on the device by invoking **ompt\_start\_trace**. Arguments to  
29 **ompt\_start\_trace** include two tool callbacks through which the OpenMP implementation  
30 can manage traces associated with the device. One allocates a buffer in which the device can  
31 deposit trace events. The second callback processes a buffer of trace events from the device.
- 32 • If the device requires a trace buffer, the OpenMP implementation invokes the tool-supplied  
33 callback function on the host device to request a new buffer.
- 34 • The OpenMP implementation monitors the execution of OpenMP constructs on the device and  
35 records a trace of events or activities into a trace buffer. If possible, device trace records are  
36 marked with a *host\_op\_id*—an identifier that associates device activities with the target operation  
37 that the host initiated to cause these activities. To correlate activities on the host with activities  
38 on a device, a tool can register a **ompt\_callback\_target\_submit** callback. Before the  
39 host initiates each distinct activity associated with a structured block for a **target** construct on  
40 a device, the OpenMP implementation dispatches the **ompt\_callback\_target\_submit**  
41 callback on the host in the thread that is executing the task that encounters the **target** construct.



**TABLE 4.3:** OMPT Tracing Interface Runtime Entry Point Names and Their Type Signatures

Entry Point String Name	Type Signature
<code>"ompt_get_device_num_procs"</code>	<code>ompt_get_device_num_procs_t</code>
<code>"ompt_get_device_time"</code>	<code>ompt_get_device_time_t</code>
<code>"ompt_translate_time"</code>	<code>ompt_translate_time_t</code>
<code>"ompt_set_trace_ompt"</code>	<code>ompt_set_trace_ompt_t</code>
<code>"ompt_set_trace_native"</code>	<code>ompt_set_trace_native_t</code>
<code>"ompt_start_trace"</code>	<code>ompt_start_trace_t</code>
<code>"ompt_pause_trace"</code>	<code>ompt_pause_trace_t</code>
<code>"ompt_flush_trace"</code>	<code>ompt_flush_trace_t</code>
<code>"ompt_stop_trace"</code>	<code>ompt_stop_trace_t</code>
<code>"ompt_advance_buffer_cursor"</code>	<code>ompt_advance_buffer_cursor_t</code>
<code>"ompt_get_record_type"</code>	<code>ompt_get_record_type_t</code>
<code>"ompt_get_record_ompt"</code>	<code>ompt_get_record_ompt_t</code>
<code>"ompt_get_record_native"</code>	<code>ompt_get_record_native_t</code>
<code>"ompt_get_record_abstract"</code>	<code>ompt_get_record_abstract_t</code>

1 Examples of activities that could cause an `ompt_callback_target_submit` callback to  
2 be dispatched include an explicit data copy between a host and target device or execution of a  
3 computation. This callback provides the tool with a pair of identifiers: one that identifies the  
4 target region and a second that uniquely identifies an activity associated with that region. These  
5 identifiers help the tool correlate activities on the target device with their target region.

- 6 • When appropriate, for example, when a trace buffer fills or needs to be flushed, the OpenMP  
7 implementation invokes the tool-supplied buffer completion callback to process a non-empty  
8 sequence of records in a trace buffer that is associated with the device.
- 9 • The tool-supplied buffer completion callback may return immediately, ignoring records in the  
10 trace buffer, or it may iterate through them using the `ompt_advance_buffer_cursor`  
11 entry point to inspect each record. A tool may use the `ompt_get_record_type` runtime  
12 entry point to inspect the type of the record at the current cursor position. Three runtime entry  
13 points (`ompt_get_record_ompt`, `ompt_get_record_native`, and  
14 `ompt_get_record_abstract`) allow tools to inspect the contents of some or all records in  
15 a trace buffer. The `ompt_get_record_native` runtime entry point uses the native trace  
16 format of the device. The `ompt_get_record_abstract` runtime entry point decodes the  
17 contents of a native trace record and summarizes them as an `ompt_record_abstract_t`  
18 record. The `ompt_get_record_ompt` runtime entry point can only be used to retrieve  
19 records in OMPT format.
- 20 • Once tracing has been started on a device, a tool may pause or resume tracing on the device at  
21 any time by invoking `ompt_pause_trace` with an appropriate flag value as an argument.
- 22 • A tool may invoke the `ompt_flush_trace` runtime entry point for a device at any time  
23 between device initialization and finalization to cause the device to flush pending trace records.
- 24 • At any time, a tool may use the `ompt_start_trace` runtime entry point to start tracing or the  
25 `ompt_stop_trace` runtime entry point to stop tracing on a device. When tracing is stopped  
26 on a device, the OpenMP implementation eventually gathers all trace records already collected  
27 on the device and presents them to the tool using the buffer completion callback.
- 28 • An OpenMP implementation can be shut down while device tracing is in progress.
- 29 • When an OpenMP implementation is shut down, it finalizes each device. Device finalization  
30 occurs in three steps. First, the OpenMP implementation halts any tracing in progress for the  
31 device. Second, the OpenMP implementation flushes all trace records collected for the device  
32 and uses the buffer completion callback associated with that device to present them to the tool.  
33 Finally, the OpenMP implementation dispatches any `ompt_callback_device_finalize`  
34 callback registered for the device.

## 35 Restrictions

36 Tracing activity on devices has the following restriction:

- 37 • Implementation-defined names must not start with the prefix `ompt_`, which is reserved for the  
38 OpenMP specification.

## Cross References

- `ompt_callback_device_initialize_t` callback type, see Section 4.5.2.19 on page 482.
- `ompt_callback_device_finalize_t` callback type, see Section 4.5.2.20 on page 484.
- `ompt_get_device_num_procs` runtime entry point, see Section 4.6.2.1 on page 515.
- `ompt_get_device_time` runtime entry point, see Section 4.6.2.2 on page 515.
- `ompt_translate_time` runtime entry point, see Section 4.6.2.3 on page 516.
- `ompt_set_trace_ompt` runtime entry point, see Section 4.6.2.4 on page 517.
- `ompt_set_trace_native` runtime entry point, see Section 4.6.2.5 on page 518.
- `ompt_start_trace` runtime entry point, see Section 4.6.2.6 on page 519.
- `ompt_pause_trace` runtime entry point, see Section 4.6.2.7 on page 520.
- `ompt_flush_trace` runtime entry point, see Section 4.6.2.8 on page 521.
- `ompt_stop_trace` runtime entry point, see Section 4.6.2.9 on page 521.
- `ompt_advance_buffer_cursor` runtime entry point, see Section 4.6.2.10 on page 522.
- `ompt_get_record_type` runtime entry point, see Section 4.6.2.11 on page 523.
- `ompt_get_record_ompt` runtime entry point, see Section 4.6.2.12 on page 524.
- `ompt_get_record_native` runtime entry point, see Section 4.6.2.13 on page 525.
- `ompt_get_record_abstract` runtime entry point, see Section 4.6.2.14 on page 526.

## 4.3 Finalizing a First-Party Tool

If the OMPT interface state is active, the tool finalizer, which has type signature `ompt_finalize_t` and is specified by the *finalize* field in the `ompt_start_tool_result_t` structure returned from the `ompt_start_tool` function, is called when the OpenMP implementation shuts down.

### Cross References

- `ompt_finalize_t` callback type, see Section 4.5.1.2 on page 461

## 4.4 OMPT Data Types

The C/C++ header file (`omp-tools.h`) provides the definitions of the types that are specified throughout this subsection.

## 4.4.1 Tool Initialization and Finalization

### Summary

A tool's implementation of `ompt_start_tool` returns a pointer to an `ompt_start_tool_result_t` structure, which contains pointers to the tool's initialization and finalization callbacks as well as an `ompt_data_t` object for use by the tool.

### Format

```
C / C++
typedef struct ompt_start_tool_result_t {
    ompt_initialize_t initialize;
    ompt_finalize_t finalize;
    ompt_data_t tool_data;
} ompt_start_tool_result_t;
C / C++
```

### Restrictions

The `ompt_start_tool_result_t` type has the following restriction:

- The *initialize* and *finalize* callback pointer values in an `ompt_start_tool_result_t` structure that `ompt_start_tool` returns must be non-null.

### Cross References

- `ompt_start_tool` function, see Section 4.2.1 on page 425.
- `ompt_data_t` type, see Section 4.4.4.4 on page 444.
- `ompt_initialize_t` callback type, see Section 4.5.1.1 on page 460.
- `ompt_finalize_t` callback type, see Section 4.5.1.2 on page 461.

## 4.4.2 Callbacks

### Summary

The `ompt_callbacks_t` enumeration type indicates the integer codes used to identify OpenMP callbacks when registering or querying them.

## Format

C/C++

```
1
2 typedef enum ompt_callbacks_t {
3     ompt_callback_thread_begin           = 1,
4     ompt_callback_thread_end            = 2,
5     ompt_callback_parallel_begin        = 3,
6     ompt_callback_parallel_end         = 4,
7     ompt_callback_task_create           = 5,
8     ompt_callback_task_schedule         = 6,
9     ompt_callback_implicit_task        = 7,
10    ompt_callback_target                 = 8,
11    ompt_callback_target_data_op        = 9,
12    ompt_callback_target_submit         = 10,
13    ompt_callback_control_tool          = 11,
14    ompt_callback_device_initialize     = 12,
15    ompt_callback_device_finalize       = 13,
16    ompt_callback_device_load           = 14,
17    ompt_callback_device_unload         = 15,
18    ompt_callback_sync_region_wait     = 16,
19    ompt_callback_mutex_released        = 17,
20    ompt_callback_dependences           = 18,
21    ompt_callback_task_dependence       = 19,
22    ompt_callback_work                   = 20,
23    ompt_callback_master                 = 21,
24    ompt_callback_target_map            = 22,
25    ompt_callback_sync_region           = 23,
26    ompt_callback_lock_init             = 24,
27    ompt_callback_lock_destroy          = 25,
28    ompt_callback_mutex_acquire         = 26,
29    ompt_callback_mutex_acquired        = 27,
30    ompt_callback_nest_lock             = 28,
31    ompt_callback_flush                  = 29,
32    ompt_callback_cancel                 = 30,
33    ompt_callback_reduction             = 31,
34    ompt_callback_dispatch               = 32
35 } ompt_callbacks_t;
```

C/C++

### 4.4.3 Tracing

OpenMP provides type definitions that support tracing with OMPT.

1  
2  
3  
4  
5  
6  
7  
8  
9  
10  
11  
12  
13  
14  
15  
16  
17  
18  
19  
20  
21  
22  
23

### 4.4.3.1 Record Type

#### Summary

The `ompt_record_t` enumeration type indicates the integer codes used to identify OpenMP trace record formats.

#### Format

C / C++

```
typedef enum ompt_record_t {  
    ompt_record_ompt           = 1,  
    ompt_record_native        = 2,  
    ompt_record_invalid       = 3  
} ompt_record_t;
```

C / C++

### 4.4.3.2 Native Record Kind

#### Summary

The `ompt_record_native_t` enumeration type indicates the integer codes used to identify OpenMP native trace record contents.

#### Format

C / C++

```
typedef enum ompt_record_native_t {  
    ompt_record_native_info = 1,  
    ompt_record_native_event = 2  
} ompt_record_native_t;
```

C / C++

### 4.4.3.3 Native Record Abstract Type

#### Summary

The `ompt_record_abstract_t` type provides an abstract trace record format that is used to summarize native device trace records.

## Format

C / C++

```
typedef struct ompt_record_abstract_t {
    ompt_record_native_t rclass;
    const char *type;
    ompt_device_time_t start_time;
    ompt_device_time_t end_time;
    ompt_hwid_t hwid;
} ompt_record_abstract_t;
```

C / C++

## Description

An `ompt_record_abstract_t` record contains information that a tool can use to process a native record that it may not fully understand. The `rclass` field indicates that the record is informational or that it represents an event; this information 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 can 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 does not have an associated `start_time` (`end_time`), the value of the `start_time` (`end_time`) field is `ompt_time_none`. The hardware identifier field, `hwid`, indicates the location on the device where the event occurred. A `hwid` may represent a hardware abstraction such as a core or a hardware thread identifier. The meaning of a `hwid` value for a device is implementation defined. If no hardware abstraction is associated with the record then the value of `hwid` is `ompt_hwid_none`.

### 4.4.3.4 Record Type

#### Summary

The `ompt_record_ompt_t` type provides an standard complete trace record format.

## Format

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_parallel_begin_t parallel_begin;
        ompt_record_parallel_end_t parallel_end;
        ompt_record_work_t work;
        ompt_record_dispatch_t dispatch;
        ompt_record_task_create_t task_create;
    };
};
```

```

1  ompt_record_dependences_t dependences;
2  ompt_record_task_dependence_t task_dependence;
3  ompt_record_task_schedule_t task_schedule;
4  ompt_record_implicit_task_t implicit_task;
5  ompt_record_master_t master;
6  ompt_record_sync_region_t sync_region;
7  ompt_record_mutex_acquire_t mutex_acquire;
8  ompt_record_mutex_t mutex;
9  ompt_record_nest_lock_t nest_lock;
10 ompt_record_flush_t flush;
11 ompt_record_cancel_t cancel;
12 ompt_record_target_t target;
13 ompt_record_target_data_op_t target_data_op;
14 ompt_record_target_map_t target_map;
15 ompt_record_target_kernel_t target_kernel;
16 ompt_record_control_tool_t control_tool;
17 } record;
18 } ompt_record_ompt_t;

```

C / C++

## Description

The field *type* specifies the type of record provided by this structure. According to the type, event specific information is stored in the matching *record* entry.

## Restrictions

The `ompt_record_ompt_t` type has the following restriction:

- If *type* is set to `ompt_callback_thread_end_t` then the value of *record* is undefined.

## 4.4.4 Miscellaneous Type Definitions

This section describes miscellaneous types and enumerations used by the tool interface.

### 4.4.4.1 ompt\_callback\_t

#### Summary

Pointers to tool callback functions with different type signatures are passed to the `ompt_set_callback` runtime entry point and returned by the `ompt_get_callback` runtime entry point. For convenience, these runtime entry points expect all type signatures to be cast to a dummy type `ompt_callback_t`.



## Format

```
typedef void (*ompt_callback_t) (void);
```

C / C++

### 4.4.4.2 ompt\_set\_result\_t

#### Summary

The `ompt_result_t` enumeration type corresponds to values that the `ompt_set_callback`, `ompt_set_trace_ompt` and `ompt_set_trace_native` runtime entry points return.

## Format

```
typedef enum ompt_set_result_t {
    ompt_set_error          = 0,
    ompt_set_never          = 1,
    ompt_set_impossible     = 2,
    ompt_set_sometimes      = 3,
    ompt_set_sometimes_paired = 4,
    ompt_set_always         = 5
} ompt_set_result_t;
```

C / C++

## Description

Values of `ompt_set_result_t`, may indicate several possible outcomes. The `ompt_set_error` value indicates that the associated call failed. Otherwise, the value indicates when an event may occur and, when appropriate, *dispatching* a callback event leads to the invocation of the callback. The `ompt_set_never` value indicates that the event will never occur or that the callback will never be invoked at runtime. The `ompt_set_impossible` value indicates that the event may occur but that tracing of it is not possible. The `ompt_set_sometimes` value indicates that the event may occur and, for an implementation-defined subset of associated event occurrences, will be traced or the callback will be invoked at runtime. The `ompt_set_sometimes_paired` value indicates the same result as `ompt_set_sometimes` and, in addition, that a callback with an *endpoint* value of `ompt_scope_begin` will be invoked if and only if the same callback with an *endpoint* value of `ompt_scope_end` will also be invoked sometime in the future. The `ompt_set_always` value indicates that, whenever an associated event occurs, it will be traced or the callback will be invoked.

## Cross References

- Monitoring activity on the host with OMPT, see Section 4.2.4 on page 430.
- Tracing activity on target devices with OMPT, see Section 4.2.5 on page 432.
- `ompt_set_callback` runtime entry point, see Section 4.6.1.3 on page 499.
- `ompt_set_trace_ompt` runtime entry point, see Section 4.6.2.4 on page 517.
- `ompt_set_trace_native` runtime entry point, see Section 4.6.2.5 on page 518.

### 4.4.4.3 `ompt_id_t`

#### Summary

The `ompt_id_t` type is used to provide various identifiers to tools.

#### Format

```
▼ _____ C / C++ _____ ▼  
| typedef uint64_t ompt_id_t; |  
▲ _____ C / C++ _____ ▲
```

#### Description

When tracing asynchronous activity on devices, identifiers enable tools to correlate target regions and operations that the host initiates with associated activities on a target device. In addition, OMPT provides identifiers to refer to parallel regions and tasks that execute on a device. These various identifiers are of type `ompt_id_t`.

`ompt_id_none` is defined as an instance of type `ompt_id_t` with the value 0.

#### Restrictions

The `ompt_id_t` type has the following restriction:

- Identifiers created on each device must be unique from the time an OpenMP implementation is initialized until it is shut down. Identifiers for each target region and target operation instance that the host device initiates must be unique over time on the host. Identifiers for parallel and task region instances that execute on a device must be unique over time within that device.

### 4.4.4.4 `ompt_data_t`

#### Summary

The `ompt_data_t` type represents data associated with threads and with parallel and task regions.

## Format

C / C++

```
typedef union ompt_data_t {  
    uint64_t value;  
    void *ptr;  
} ompt_data_t;
```

C / C++

## Description

The `ompt_data_t` type represents data that is reserved for tool use and that is related to a thread or to a parallel or task region. When an OpenMP implementation creates a thread or an instance of a parallel or task region, it initializes the associated `ompt_data_t` object with the value `ompt_data_none`, which is an instance of the type with the data and pointer fields equal to 0.

### 4.4.4.5 `ompt_device_t`

#### Summary

The `ompt_device_t` opaque object type represents a device.

## Format

C / C++

```
typedef void ompt_device_t;
```

C / C++

### 4.4.4.6 `ompt_device_time_t`

#### Summary

The `ompt_device_time_t` type represents raw device time values.

## Format

C / C++

```
typedef uint64_t ompt_device_time_t;
```

C / C++

## Description

The `ompt_device_time_t` opaque object type represents raw device time values. `ompt_time_none` refers to an unknown or unspecified time and is defined as an instance of type `ompt_device_time_t` with the value 0.

#### 4.4.4.7 `ompt_buffer_t`

##### Summary

The `ompt_buffer_t` opaque object type is a handle for a target buffer.

##### Format

C / C++

```
typedef void ompt_buffer_t;
```

C / C++

#### 4.4.4.8 `ompt_buffer_cursor_t`

##### Summary

The `ompt_buffer_cursor_t` opaque type is a handle for a position in a target buffer.

##### Format

C / C++

```
typedef uint64_t ompt_buffer_cursor_t;
```

C / C++

#### 4.4.4.9 `ompt_dependence_t`

##### Summary

The `ompt_dependence_t` type represents a task dependence.

##### Format

C / C++

```
typedef struct ompt_dependence_t {  
    ompt_data_t variable;  
    ompt_dependence_type_t dependence_type;  
} ompt_dependence_t;
```

C / C++

##### Description

The `ompt_dependence_t` type is a structure that holds information about a depend clause. For task dependences, the *variable* field points to the storage location of the dependence. For *doacross* dependences, the *variable* field contains the value of a vector element that describes the dependence. The *dependence\_type* field indicates the type of the dependence.

##### Cross References

- `ompt_dependence_type_t` type, see Section 4.4.4.23 on page 453.

#### 1        4.4.4.10 `ompt_thread_t`

##### 2        **Summary**

3        The `ompt_thread_t` enumeration type defines the valid thread type values.

##### 4        **Format**

C / C++

```
5        typedef enum ompt_thread_t {  
6            ompt_thread_initial                         = 1,  
7            ompt_thread_worker                         = 2,  
8            ompt_thread_other                           = 3,  
9            ompt_thread_unknown                        = 4  
10        } ompt_thread_t;
```

C / C++

##### 11       **Description**

12       Any *initial thread* has thread type `ompt_thread_initial`. All *OpenMP threads* that are not  
13       initial threads have thread type `ompt_thread_worker`. A thread that an OpenMP  
14       implementation uses but that does not execute user code has thread type `ompt_thread_other`.  
15       Any thread that is created outside an OpenMP implementation and that is not an *initial thread* has  
16       thread type `ompt_thread_unknown`.

#### 17       4.4.4.11 `ompt_scope_endpoint_t`

##### 18       **Summary**

19       The `ompt_scope_endpoint_t` enumeration type defines valid scope endpoint values.

##### 20       **Format**

C / C++

```
21       typedef enum ompt_scope_endpoint_t {  
22            ompt_scope_begin                            = 1,  
23            ompt_scope_end                              = 2  
24        } ompt_scope_endpoint_t;
```

C / C++

#### 25       4.4.4.12 `ompt_dispatch_t`

##### 26       **Summary**

27       The `ompt_dispatch_t` enumeration type defines the valid dispatch kind values.

## Format

C / C++

```
typedef enum ompt_dispatch_t {
    ompt_dispatch_iteration          = 1,
    ompt_dispatch_section            = 2
} ompt_dispatch_t;
```

C / C++

### 4.4.4.13 ompt\_sync\_region\_t

#### Summary

The `ompt_sync_region_t` enumeration type defines the valid synchronization region kind values.

## Format

C / C++

```
typedef enum ompt_sync_region_t {
    ompt_sync_region_barrier          = 1,
    ompt_sync_region_barrier_implicit = 2,
    ompt_sync_region_barrier_explicit = 3,
    ompt_sync_region_barrier_implementation = 4,
    ompt_sync_region_taskwait         = 5,
    ompt_sync_region_taskgroup        = 6,
    ompt_sync_region_reduction        = 7
} ompt_sync_region_t;
```

C / C++

### 4.4.4.14 ompt\_target\_data\_op\_t

#### Summary

The `ompt_target_data_op_t` enumeration type defines the valid target data operation values.

## Format

C / C++

```
typedef enum ompt_target_data_op_t {
    ompt_target_data_alloc          = 1,
    ompt_target_data_transfer_to_device = 2,
    ompt_target_data_transfer_from_device = 3,
    ompt_target_data_delete         = 4,
    ompt_target_data_associate      = 5,
    ompt_target_data_disassociate   = 6
} ompt_target_data_op_t;
```

C / C++

#### 1 4.4.4.15 ompt\_work\_t

##### 2 Summary

3 The `ompt_work_t` enumeration type defines the valid work type values.

##### 4 Format

C / C++

```
5 typedef enum ompt_work_t {  
6     ompt_work_loop           = 1,  
7     ompt_work_sections      = 2,  
8     ompt_work_single_executor = 3,  
9     ompt_work_single_other  = 4,  
10    ompt_work_workshare     = 5,  
11    ompt_work_distribute    = 6,  
12    ompt_work_taskloop     = 7,  
13    ompt_work_scope        = 8  
14 } ompt_work_t;
```

C / C++

#### 15 4.4.4.16 ompt\_mutex\_t

##### 16 Summary

17 The `ompt_mutex_t` enumeration type defines the valid mutex kind values.

##### 18 Format

C / C++

```
19 typedef enum ompt_mutex_t {  
20     ompt_mutex_lock           = 1,  
21     ompt_mutex_test_lock     = 2,  
22     ompt_mutex_nest_lock     = 3,  
23     ompt_mutex_test_nest_lock = 4,  
24     ompt_mutex_critical      = 5,  
25     ompt_mutex_atomic        = 6,  
26     ompt_mutex_ordered       = 7  
27 } ompt_mutex_t;
```

C / C++

#### 28 4.4.4.17 ompt\_native\_mon\_flag\_t

##### 29 Summary

30 The `ompt_native_mon_flag_t` enumeration type defines the valid native monitoring flag  
31 values.

## Format

C / C++

```
typedef enum ompt_native_mon_flag_t {
    ompt_native_data_motion_explicit      = 0x01,
    ompt_native_data_motion_implicit     = 0x02,
    ompt_native_kernel_invocation        = 0x04,
    ompt_native_kernel_execution         = 0x08,
    ompt_native_driver                   = 0x10,
    ompt_native_runtime                   = 0x20,
    ompt_native_overhead                  = 0x40,
    ompt_native_idleness                  = 0x80
} ompt_native_mon_flag_t;
```

C / C++

## 4.4.4.18 ompt\_task\_flag\_t

### Summary

The `ompt_task_flag_t` enumeration type defines valid task types.

## Format

C / C++

```
typedef enum ompt_task_flag_t {
    ompt_task_initial                    = 0x00000001,
    ompt_task_implicit                   = 0x00000002,
    ompt_task_explicit                   = 0x00000004,
    ompt_task_target                     = 0x00000008,
    ompt_task_undeferrred                = 0x08000000,
    ompt_task_untied                     = 0x10000000,
    ompt_task_final                      = 0x20000000,
    ompt_task_mergeable                  = 0x40000000,
    ompt_task_merged                     = 0x80000000
} ompt_task_flag_t;
```

C / C++

### Description

The `ompt_task_flag_t` enumeration type defines valid task type values. The least significant byte provides information about the general classification of the task. The other bits represent properties of the task.



#### 4.4.4.19 ompt\_task\_status\_t

##### Summary

The `ompt_task_status_t` enumeration type indicates the reason that a task was switched when it reached a task scheduling point.

##### Format

C / C++

```
typedef enum ompt_task_status_t {  
    ompt_task_complete      = 1,  
    ompt_task_yield         = 2,  
    ompt_task_cancel        = 3,  
    ompt_task_detach        = 4,  
    ompt_task_early_fulfill = 5,  
    ompt_task_late_fulfill  = 6,  
    ompt_task_switch        = 7  
} ompt_task_status_t;
```

C / C++

##### Description

The value `ompt_task_complete` of the `ompt_task_status_t` type indicates that the task that encountered the task scheduling point completed execution of the associated *structured-block* and an associated *allow-completion-event* was fulfilled. The value `ompt_task_yield` indicates that the task encountered a `taskyield` construct. The value `ompt_task_cancel` indicates that the task was canceled when it encountered an active cancellation point. The value `ompt_task_detach` indicates that a task with `detach` clause completed execution of the associated *structured-block* and is waiting for an *allow-completion-event* to be fulfilled. The value `ompt_task_early_fulfill` indicates that the *allow-completion-event* of the task is fulfilled before the task completed execution of the associated *structured-block*. The value `ompt_task_late_fulfill` indicates that the *allow-completion-event* of the task is fulfilled after the task completed execution of the associated *structured-block*. The value `ompt_task_switch` is used for all other cases that a task was switched.

#### 4.4.4.20 ompt\_target\_t

##### Summary

The `ompt_target_t` enumeration type defines the valid target type values.

## Format

C / C++

```
typedef enum ompt_target_t {
    ompt_target                = 1,
    ompt_target_enter_data     = 2,
    ompt_target_exit_data      = 3,
    ompt_target_update         = 4
} ompt_target_t;
```

C / C++

### 4.4.4.21 ompt\_parallel\_flag\_t

#### Summary

The `ompt_parallel_flag_t` enumeration type defines valid invoker values.

#### Format

C / C++

```
typedef enum ompt_parallel_flag_t {
    ompt_parallel_invoker_program = 0x00000001,
    ompt_parallel_invoker_runtime = 0x00000002,
    ompt_parallel_league         = 0x40000000,
    ompt_parallel_team           = 0x80000000
} ompt_parallel_flag_t;
```

C / C++

#### Description

The `ompt_parallel_flag_t` enumeration type defines valid invoker values, which indicate how an outlined function is invoked.

The value `ompt_parallel_invoker_program` indicates that the outlined function associated with implicit tasks for the region is invoked directly by the application on the master thread for a parallel region.

The value `ompt_parallel_invoker_runtime` indicates that the outlined function associated with implicit tasks for the region is invoked by the runtime on the master thread for a parallel region.

The value `ompt_parallel_league` indicates that the callback is invoked due to the creation of a league of teams by a `teams` construct.

The value `ompt_parallel_team` indicates that the callback is invoked due to the creation of a team of threads by a `parallel` construct.

#### 1 4.4.4.22 ompt\_target\_map\_flag\_t

##### 2 Summary

3 The `ompt_target_map_flag_t` enumeration type defines the valid target map flag values.

##### 4 Format

C / C++

```
5 typedef enum ompt_target_map_flag_t {  
6     ompt_target_map_flag_to           = 0x01,  
7     ompt_target_map_flag_from        = 0x02,  
8     ompt_target_map_flag_alloc       = 0x04,  
9     ompt_target_map_flag_release     = 0x08,  
10    ompt_target_map_flag_delete      = 0x10,  
11    ompt_target_map_flag_implicit    = 0x20  
12 } ompt_target_map_flag_t;
```

C / C++

#### 13 4.4.4.23 ompt\_dependence\_type\_t

##### 14 Summary

15 The `ompt_dependence_type_t` enumeration type defines the valid task dependence type  
16 values.

##### 17 Format

C / C++

```
18 typedef enum ompt_dependence_type_t {  
19     ompt_dependence_type_in          = 1,  
20     ompt_dependence_type_out        = 2,  
21     ompt_dependence_type_inout      = 3,  
22     ompt_dependence_type_mutexinoutset = 4,  
23     ompt_dependence_type_source     = 5,  
24     ompt_dependence_type_sink       = 6,  
25     ompt_dependence_type_inoutset   = 7  
26 } ompt_dependence_type_t;
```

C / C++

#### 27 4.4.4.24 ompt\_cancel\_flag\_t

##### 28 Summary

29 The `ompt_cancel_flag_t` enumeration type defines the valid cancel flag values.

## Format

C / C++

```
typedef enum ompt_cancel_flag_t {
    ompt_cancel_parallel      = 0x01,
    ompt_cancel_sections     = 0x02,
    ompt_cancel_loop         = 0x04,
    ompt_cancel_taskgroup    = 0x08,
    ompt_cancel_activated    = 0x10,
    ompt_cancel_detected     = 0x20,
    ompt_cancel_discarded_task = 0x40
} ompt_cancel_flag_t;
```

C / C++

### 4.4.4.25 ompt\_hwid\_t

#### Summary

The `ompt_hwid_t` opaque type is a handle for a hardware identifier for a target device.

#### Format

C / C++

```
typedef uint64_t ompt_hwid_t;
```

C / C++

#### Description

The `ompt_hwid_t` opaque type is a handle for a hardware identifier for a target device. `ompt_hwid_none` is an instance of the type that refers to an unknown or unspecified hardware identifier and that has the value 0. If no `hwid` is associated with an `ompt_record_abstract_t` then the value of `hwid` is `ompt_hwid_none`.

#### Cross References

- `ompt_record_abstract_t` type, see Section 4.4.3.3 on page 440.

### 4.4.4.26 ompt\_state\_t

#### Summary

If the OMPT interface is in the *active* state then an OpenMP implementation must maintain *thread state* information for each thread. The thread state maintained is an approximation of the instantaneous state of a thread.

## Format

C/C++

A thread state must be one of the values of the enumeration type `ompt_state_t` or an implementation-defined state value of 512 or higher.

```
typedef enum ompt_state_t {
    ompt_state_work_serial          = 0x000,
    ompt_state_work_parallel        = 0x001,
    ompt_state_work_reduction       = 0x002,

    ompt_state_wait_barrier         = 0x010,
    ompt_state_wait_barrier_implicit_parallel = 0x011,
    ompt_state_wait_barrier_implicit_workshare = 0x012,
    ompt_state_wait_barrier_implicit = 0x013,
    ompt_state_wait_barrier_explicit = 0x014,

    ompt_state_wait_taskwait        = 0x020,
    ompt_state_wait_taskgroup        = 0x021,

    ompt_state_wait_mutex            = 0x040,
    ompt_state_wait_lock              = 0x041,
    ompt_state_wait_critical          = 0x042,
    ompt_state_wait_atomic            = 0x043,
    ompt_state_wait_ordered           = 0x044,

    ompt_state_wait_target           = 0x080,
    ompt_state_wait_target_map        = 0x081,
    ompt_state_wait_target_update     = 0x082,

    ompt_state_idle                  = 0x100,
    ompt_state_overhead               = 0x101,
    ompt_state_undefined              = 0x102
} ompt_state_t;
```

C/C++

## Description

A tool can query the OpenMP state of a thread at any time. If a tool queries the state of a thread that is not associated with OpenMP then the implementation reports the state as `ompt_state_undefined`.

The value `ompt_state_work_serial` indicates that the thread is executing code outside all `parallel` regions.

The value `ompt_state_work_parallel` indicates that the thread is executing code within the scope of a `parallel` region.

1 The value `ompt_state_work_reduction` indicates that the thread is combining partial  
2 reduction results from threads in its team. An OpenMP implementation may never report a thread  
3 in this state; a thread that is combining partial reduction results may have its state reported as  
4 `ompt_state_work_parallel` or `ompt_state_overhead`.

5 The value `ompt_state_wait_barrier` indicates that the thread is waiting at either an  
6 implicit or explicit barrier. An implementation may never report a thread in this state; instead, a  
7 thread may have its state reported as `ompt_state_wait_barrier_implicit` or  
8 `ompt_state_wait_barrier_explicit`, as appropriate.

9 The value `ompt_state_wait_barrier_implicit` indicates that the thread is waiting at an  
10 implicit barrier in a `parallel` region. An OpenMP implementation may report  
11 `ompt_state_wait_barrier` for implicit barriers.

12 The value `ompt_state_wait_barrier_implicit_parallel` indicates that the thread is  
13 waiting at an implicit barrier at the end of a `parallel` region. An OpenMP implementation may  
14 report `ompt_state_wait_barrier` or `ompt_state_wait_barrier_implicit` for  
15 these barriers.

16 The value `ompt_state_wait_barrier_implicit_workshare` indicates that the thread  
17 is waiting at an implicit barrier at the end of a worksharing construct. An OpenMP implementation  
18 may report `ompt_state_wait_barrier` or `ompt_state_wait_barrier_implicit`  
19 for these barriers.

20 The value `ompt_state_wait_barrier_explicit` indicates that the thread is waiting in a  
21 `barrier` region. An OpenMP implementation may report `ompt_state_wait_barrier` for  
22 these barriers.

23 The value `ompt_state_wait_taskwait` indicates that the thread is waiting at a `taskwait`  
24 construct.

25 The value `ompt_state_wait_taskgroup` indicates that the thread is waiting at the end of a  
26 `taskgroup` construct.

27 The value `ompt_state_wait_mutex` indicates that the thread is waiting for a mutex of an  
28 unspecified type.

29 The value `ompt_state_wait_lock` indicates that the thread is waiting for a lock or nestable  
30 lock.

31 The value `ompt_state_wait_critical` indicates that the thread is waiting to enter a  
32 `critical` region.

33 The value `ompt_state_wait_atomic` indicates that the thread is waiting to enter an `atomic`  
34 region.

35 The value `ompt_state_wait_ordered` indicates that the thread is waiting to enter an  
36 `ordered` region.

1 The value `ompt_state_wait_target` indicates that the thread is waiting for a **target**  
2 region to complete.

3 The value `ompt_state_wait_target_map` indicates that the thread is waiting for a target  
4 data mapping operation to complete. An implementation may report  
5 `ompt_state_wait_target` for **target data** constructs.

6 The value `ompt_state_wait_target_update` indicates that the thread is waiting for a  
7 **target update** operation to complete. An implementation may report  
8 `ompt_state_wait_target` for **target update** constructs.

9 The value `ompt_state_idle` indicates that the thread is idle, that is, it is not part of an  
10 OpenMP team.

11 The value `ompt_state_overhead` indicates that the thread is in the overhead state at any point  
12 while executing within the OpenMP runtime, except while waiting at a synchronization point.

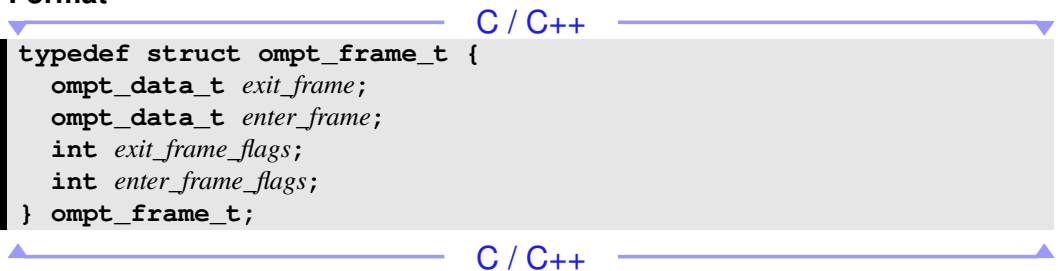
13 The value `ompt_state_undefined` indicates that the native thread is not created by the  
14 OpenMP implementation.

#### 15 4.4.4.27 `ompt_frame_t`

##### 16 Summary

17 The `ompt_frame_t` type describes procedure frame information for an OpenMP task.

##### 18 Format

19  C / C++

```
20 typedef struct ompt_frame_t {  
21     ompt_data_t exit_frame;  
22     ompt_data_t enter_frame;  
23     int exit_frame_flags;  
24     int enter_frame_flags;  
25 } ompt_frame_t;
```

##### 25 Description

26 Each `ompt_frame_t` object is associated with the task to which the procedure frames belong.  
27 Each non-merged initial, implicit, explicit, or target task with one or more frames on the stack of a  
28 native thread has an associated `ompt_frame_t` object.

29 The `exit_frame` field of an `ompt_frame_t` object contains information to identify the first  
30 procedure frame executing the task region. The `exit_frame` for the `ompt_frame_t` object  
31 associated with the *initial task* that is not nested inside any OpenMP construct is **NULL**.

32 The `enter_frame` field of an `ompt_frame_t` object contains information to identify the latest still  
33 active procedure frame executing the task region before entering the OpenMP runtime

1 implementation or before executing a different task. If a task with frames on the stack has not been  
2 suspended, the value of *enter\_frame* for the `ompt_frame_t` object associated with the task may  
3 contain `NULL`.

4 For *exit\_frame*, the *exit\_frame\_flags* and, for *enter\_frame*, the *enter\_frame\_flags* field indicates that  
5 the provided frame information points to a runtime or an application frame address. The same  
6 fields also specify the kind of information that is provided to identify the frame, These fields are a  
7 disjunction of values in the `ompt_frame_flag_t` enumeration type.

8 The lifetime of an `ompt_frame_t` object begins when a task is created and ends when the task is  
9 destroyed. Tools should not assume that a frame structure remains at a constant location in memory  
10 throughout the lifetime of the task. A pointer to an `ompt_frame_t` object is passed to some  
11 callbacks; a pointer to the `ompt_frame_t` object of a task can also be retrieved by a tool at any  
12 time, including in a signal handler, by invoking the `ompt_get_task_info` runtime entry point  
13 (described in Section 4.6.1.14). A pointer to an `ompt_frame_t` object that a tool retrieved is  
14 valid as long as the tool does not pass back control to the OpenMP implementation.

15  
16 **Note** – A monitoring tool that uses asynchronous sampling can observe values of *exit\_frame* and  
17 *enter\_frame* at inconvenient times. Tools must be prepared to handle `ompt_frame_t` objects  
18 observed just prior to when their field values will be set or cleared.  
19

#### 20 4.4.4.28 `ompt_frame_flag_t`

##### 21 **Summary**

22 The `ompt_frame_flag_t` enumeration type defines valid frame information flags.

##### 23 **Format**

24 `C / C++`  
25 

```
typedef enum ompt_frame_flag_t {  
26     ompt_frame_runtime      = 0x00,  
27     ompt_frame_application  = 0x01,  
28     ompt_frame_cfa         = 0x10,  
29     ompt_frame_framepointer = 0x20,  
30     ompt_frame_stackaddress = 0x30  
} ompt_frame_flag_t;
```

26 `C / C++`



## Description

The value `ompt_frame_runtime` of the `ompt_frame_flag_t` type indicates that a frame address is a procedure frame in the OpenMP runtime implementation. The value `ompt_frame_application` of the `ompt_frame_flag_t` type indicates that an exit frame address is a procedure frame in the OpenMP application.

Higher order bits indicate the kind of provided information that is unique for the particular frame pointer. The value `ompt_frame_cfa` indicates that a frame address specifies a *canonical frame address*. The value `ompt_frame_framepointer` indicates that a frame address provides the value of the frame pointer register. The value `ompt_frame_stackaddress` indicates that a frame address specifies a pointer address that is contained in the current stack frame.

### 4.4.4.29 `ompt_wait_id_t`

#### Summary

The `ompt_wait_id_t` type describes wait identifiers for an OpenMP thread.

#### Format

```
typedef uint64_t ompt_wait_id_t;
```

#### Description

Each thread maintains a *wait identifier* of type `ompt_wait_id_t`. When a task that a thread executes is waiting for mutual exclusion, the wait identifier of the thread indicates the reason that the thread is waiting. A wait identifier may represent a critical section *name*, a lock, a program variable accessed in an atomic region, or a synchronization object that is internal to an OpenMP implementation. When a thread is not in a wait state then the value of the wait identifier of the thread is undefined.

`ompt_wait_id_none` is defined as an instance of type `ompt_wait_id_t` with the value 0.

## 4.5 OMPT Tool Callback Signatures and Trace Records

The C/C++ header file (`omp-tools.h`) provides the definitions of the types that are specified throughout this subsection.

#### Restrictions

- Tool callbacks may not use OpenMP directives or call any runtime library routines described in Section 3.
- Tool callbacks must exit by either returning to the caller or aborting.

## 4.5.1 Initialization and Finalization Callback Signature

### 4.5.1.1 `ompt_initialize_t`

#### Summary

A callback with type signature `ompt_initialize_t` initializes use of the OMPT interface.

#### Format

```
C / C++  
typedef int (*ompt_initialize_t) (  
    ompt_function_lookup_t lookup,  
    int initial_device_num,  
    ompt_data_t *tool_data  
);
```

C / C++

#### Description

To use the OMPT interface, an 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 initializer with type signature `ompt_initialize_t`. An OpenMP implementation will call the initializer after fully initializing itself but before beginning execution of any OpenMP construct or completing execution of any environment routine invocation.

The initializer returns a non-zero value if it succeeds.

#### Description of Arguments

The `lookup` argument is a callback to an OpenMP runtime routine that must be used to obtain a pointer to each runtime entry point in the OMPT interface. The `initial_device_num` argument provides the value of `omp_get_initial_device()`. The `tool_data` argument is a pointer to the `tool_data` field in the `ompt_start_tool_result_t` structure that `ompt_start_tool` returned. The expected actions of an initializer are described in Section 4.2.3.

#### Cross References

- `omp_get_initial_device` routine, see Section 3.2.41 on page 374.
- `ompt_start_tool` function, see Section 4.2.1 on page 425.
- `ompt_start_tool_result_t` type, see Section 4.4.1 on page 438.
- `ompt_data_t` type, see Section 4.4.4.4 on page 444.
- `ompt_function_lookup_t` type, see Section 4.6.3 on page 526.

## 4.5.1.2 `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

To use the OMPT interface, an 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 `tool_data` argument 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` function, see Section 4.2.1 on page 425.
- `ompt_start_tool_result_t` type, see Section 4.4.1 on page 438.
- `ompt_data_t` type, see Section 4.4.4.4 on page 444.

## 4.5.2 Event Callback Signatures and Trace Records

This section describes the signatures of tool callback functions that an OMPT tool may register and that are called during runtime of an OpenMP program. An implementation may also provide a trace of events per device. Along with the callbacks, the following defines standard trace records. For the trace records, tool data arguments are replaced by an ID, which must be initialized by the OpenMP implementation. Each of `parallel_id`, `task_id`, and `thread_id` must be unique per target region. Tool implementations of callbacks are not required to be *async signal safe*.

### Cross References

- `ompt_id_t` type, see Section 4.4.4.3 on page 444.
- `ompt_data_t` type, see Section 4.4.4.4 on page 444.

## 4.5.2.1 `ompt_callback_thread_begin_t`

### Summary

The `ompt_callback_thread_begin_t` type is used for callbacks that are dispatched when native threads are created.

### Format

```
typedef void (*ompt_callback_thread_begin_t) (  
    ompt_thread_t thread_type,  
    ompt_data_t *thread_data  
);
```

### Trace Record

```
typedef struct ompt_record_thread_begin_t {  
    ompt_thread_t thread_type;  
} ompt_record_thread_begin_t;
```

### Description of Arguments

The *thread\_type* argument indicates the type of the new thread: initial, worker, or other. The binding of the *thread\_data* argument is the new thread.

### Cross References

- `parallel` construct, see Section 2.8 on page 79.
- `teams` construct, see Section 2.9 on page 87.
- Initial task, see Section 2.13.5 on page 155.
- `ompt_data_t` type, see Section 4.4.4.4 on page 444.
- `ompt_thread_t` type, see Section 4.4.4.10 on page 447.

## 4.5.2.2 `ompt_callback_thread_end_t`

### Summary

The `ompt_callback_thread_end_t` type is used for callbacks that are dispatched when native threads are destroyed.

## Format

C / C++

```
typedef void (*ompt_callback_thread_end_t) (  
    ompt_data_t *thread_data  
);
```

C / C++

## Description of Arguments

The binding of the *thread\_data* argument is the thread that will be destroyed.

## Cross References

- **parallel** construct, see Section 2.8 on page 79.
- **teams** construct, see Section 2.9 on page 87.
- Initial task, see Section 2.13.5 on page 155.
- **ompt\_record\_ompt\_t** type, see Section 4.4.3.4 on page 441.
- **ompt\_data\_t** type, see Section 4.4.4.4 on page 444.

### 4.5.2.3 ompt\_callback\_parallel\_begin\_t

#### Summary

The **ompt\_callback\_parallel\_begin\_t** type is used for callbacks that are dispatched when **parallel** and **teams** regions start.

## Format

C / C++

```
typedef void (*ompt_callback_parallel_begin_t) (  
    ompt_data_t *encountering_task_data,  
    const ompt_frame_t *encountering_task_frame,  
    ompt_data_t *parallel_data,  
    unsigned int requested_parallelism,  
    int flags,  
    const void *codeptr_ra  
);
```

C / C++

## Trace Record

C / C++

```
typedef struct ompt_record_parallel_begin_t {
    ompt_id_t encountering_task_id;
    ompt_id_t parallel_id;
    unsigned int requested_parallelism;
    int flags;
    const void *codeptr_ra;
} ompt_record_parallel_begin_t;
```

C / C++

## Description of Arguments

The binding of the *encountering\_task\_data* argument is the encountering task.

The *encountering\_task\_frame* argument points to the frame object that is associated with the encountering task.

The binding of the *parallel\_data* argument is the **parallel** or **teams** region that is beginning.

The *requested\_parallelism* argument indicates the number of threads or teams that the user requested.

The *flags* argument indicates whether the code for the region is inlined into the application or invoked by the runtime and also whether the region is a **parallel** or **teams** region. Valid values for *flags* are a disjunction of elements in the enum **ompt\_parallel\_flag\_t**.

The *codeptr\_ra* argument relates the implementation of an OpenMP region to its source code. If a runtime routine implements the region associated with a callback that has type signature **ompt\_callback\_parallel\_begin\_t** then *codeptr\_ra* contains the return address of the call to that runtime routine. If the implementation the region is inlined then *codeptr\_ra* contains the return address of the invocation of the callback. If attribution to source code is impossible or inappropriate, *codeptr\_ra* may be **NULL**.

## Cross References

- **parallel** construct, see Section 2.8 on page 79.
- **teams** construct, see Section 2.9 on page 87.
- **ompt\_data\_t** type, see Section 4.4.4.4 on page 444.
- **ompt\_parallel\_flag\_t** type, see Section 4.4.4.21 on page 452.
- **ompt\_frame\_t** type, see Section 4.4.4.27 on page 457.

## 4.5.2.4 `ompt_callback_parallel_end_t`

### Summary

The `ompt_callback_parallel_end_t` type is used for callbacks that are dispatched when `parallel` and `teams` regions ends.

### Format

```
C / C++  
typedef void (*ompt_callback_parallel_end_t) (  
    ompt_data_t *parallel_data,  
    ompt_data_t *encountering_task_data,  
    int flags,  
    const void *codeptr_ra  
);
```

### Trace Record

```
C / C++  
typedef struct ompt_record_parallel_end_t {  
    ompt_id_t parallel_id;  
    ompt_id_t encountering_task_id;  
    int flags;  
    const void *codeptr_ra;  
} ompt_record_parallel_end_t;
```

### Description of Arguments

The binding of the `parallel_data` argument is the `parallel` or `teams` region that is ending.

The binding of the `encountering_task_data` argument is the encountering task.

The `flags` argument indicates whether the execution of the region is inlined into the application or invoked by the runtime and also whether it is a `parallel` or `teams` region. Values for `flags` are a disjunction of elements in the enum `ompt_parallel_flag_t`.

The `codeptr_ra` argument relates the implementation of an OpenMP region to its source code. If a runtime routine implements the region associated with a callback that has type signature `ompt_callback_parallel_end_t` then `codeptr_ra` contains the return address of the call to that runtime routine. If the implementation of the region is inlined then `codeptr_ra` contains the return address of the invocation of the callback. If attribution to source code is impossible or inappropriate, `codeptr_ra` may be `NULL`.

## Cross References

- `parallel` construct, see Section 2.8 on page 79.
- `teams` construct, see Section 2.9 on page 87.
- `ompt_data_t` type, see Section 4.4.4.4 on page 444.
- `ompt_parallel_flag_t` type, see Section 4.4.4.21 on page 452.

### 4.5.2.5 `ompt_callback_work_t`

#### Summary

The `ompt_callback_work_t` type is used for callbacks that are dispatched when worksharing regions, loop-related regions, `taskloop` regions and `scope` regions begin and end.

#### Format

```
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++

#### Trace Record

```
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;
```

C / C++



## Description of Arguments

The *wstype* argument indicates the kind of region.

The *endpoint* argument indicates that the callback signals the beginning of a scope or the end of a scope.

The binding of the *parallel\_data* argument is the current parallel region.

The binding of the *task\_data* argument is the current task.

The *count* argument is a measure of the quantity of work involved in the 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 signals the end of a scope, a *count* value of 0 indicates that the actual *count* value is not available.

The *codeptr\_ra* argument relates the implementation of an OpenMP region to its source code. If a runtime routine implements the region associated with a callback that has type signature **ompt\_callback\_work\_t** then *codeptr\_ra* contains the return address of the call to that runtime routine. If the implementation of the region is inlined then *codeptr\_ra* contains the return address of the invocation of the callback. If attribution to source code is impossible or inappropriate, *codeptr\_ra* may be **NULL**.

## Cross References

- Worksharing constructs, see Section [2.11](#) on page [93](#).
- Loop-related constructs, see Section [2.12](#) on page [101](#).
- Worksharing-Loop construct, see Section [2.12.2](#) on page [106](#).
- **taskloop** construct, see Section [2.13.2](#) on page [147](#).
- **ompt\_data\_t** type, see Section [4.4.4.4](#) on page [444](#).
- **ompt\_scope\_endpoint\_t** type, see Section [4.4.4.11](#) on page [447](#).
- **ompt\_work\_t** type, see Section [4.4.4.15](#) on page [449](#).

### 4.5.2.6 **ompt\_callback\_dispatch\_t**

#### Summary

The **ompt\_callback\_dispatch\_t** type is used for callbacks that are dispatched when a thread begins to execute a section or loop iteration.

## Format

C / C++

```
typedef void (*ompt_callback_dispatch_t) (  
    ompt_data_t *parallel_data,  
    ompt_data_t *task_data,  
    ompt_dispatch_t kind,  
    ompt_data_t instance  
);
```

C / C++

## Trace Record

C / C++

```
typedef struct ompt_record_dispatch_t {  
    ompt_id_t parallel_id;  
    ompt_id_t task_id;  
    ompt_dispatch_t kind;  
    ompt_data_t instance;  
} ompt_record_dispatch_t;
```

C / C++

## Description of Arguments

The binding of the *parallel\_data* argument is the current parallel region.

The binding of the *task\_data* argument is the implicit task that executes the structured block of the parallel region.

The *kind* argument indicates whether a loop iteration or a section is being dispatched.

For a loop iteration, the *instance.value* argument contains the iteration variable value. For a structured block in the **sections** construct, *instance.ptr* contains a code address that identifies the structured block. In cases where a runtime routine implements the structured block associated with this callback, *instance.ptr* contains the return address of the call to the runtime routine. In cases where the implementation of the structured block is inlined, *instance.ptr* contains the return address of the invocation of this callback.

## Cross References

- **sections** and **section** constructs, see Section 2.11.1 on page 93.
- Worksharing-loop construct, see Section 2.12.2 on page 106.
- **taskloop** construct, see Section 2.13.2 on page 147.
- **ompt\_data\_t** type, see Section 4.4.4.4 on page 444.
- **ompt\_dispatch\_t** type, see Section 4.4.4.12 on page 447.

## 4.5.2.7 `ompt_callback_task_create_t`

### Summary

The `ompt_callback_task_create_t` type is used for callbacks that are dispatched when `task` regions or initial tasks are generated.

### Format

```
C / C++
typedef void (*ompt_callback_task_create_t) (
    ompt_data_t *encountering_task_data,
    const ompt_frame_t *encountering_task_frame,
    ompt_data_t *new_task_data,
    int flags,
    int has_dependences,
    const void *codeptr_ra
);
```

### Trace Record

```
C / C++
typedef struct ompt_record_task_create_t {
    ompt_id_t encountering_task_id;
    ompt_id_t new_task_id;
    int flags;
    int has_dependences;
    const void *codeptr_ra;
} ompt_record_task_create_t;
```

### Description of Arguments

The binding of the `encountering_task_data` argument is the encountering task. This argument is `NULL` for an initial task.

The `encountering_task_frame` argument points to the frame object associated with the encountering task. This argument is `NULL` for an initial task.

The binding of the `new_task_data` argument is the generated task.

The `flags` argument indicates the kind of the task (initial, explicit, or target) that is generated. Values for `flags` are a disjunction of elements in the `ompt_task_flag_t` enumeration type.

The `has_dependences` argument is `true` if the generated task has dependences and `false` otherwise.

The `codeptr_ra` argument relates the implementation of an OpenMP region to its source code. If a runtime routine implements the region associated with a callback that has type signature `ompt_callback_task_create_t` then `codeptr_ra` contains the return address of the call to

1 that runtime routine. If the implementation of the region is inlined then *codeptr\_ra* contains the  
2 return address of the invocation of the callback. If attribution to source code is impossible or  
3 inappropriate, *codeptr\_ra* may be **NULL**.

## 4 **Cross References**

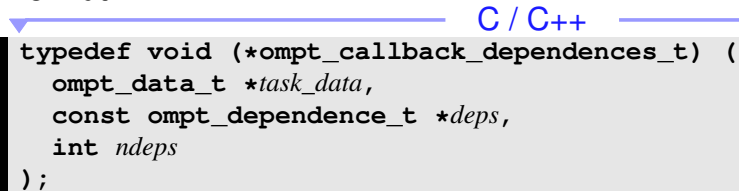
- 5 • **task** construct, see Section 2.13.1 on page 142.
- 6 • Initial task, see Section 2.13.5 on page 155.
- 7 • **ompt\_data\_t** type, see Section 4.4.4.4 on page 444.
- 8 • **ompt\_task\_flag\_t** type, see Section 4.4.4.18 on page 450.
- 9 • **ompt\_frame\_t** type, see Section 4.4.4.27 on page 457.

## 10 **4.5.2.8 ompt\_callback\_dependences\_t**

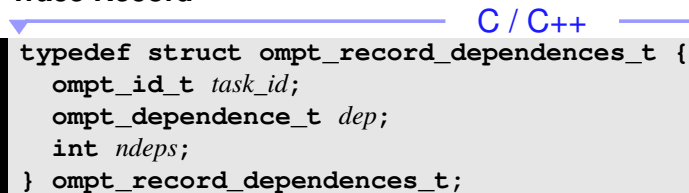
### 11 **Summary**

12 The **ompt\_callback\_dependences\_t** type is used for callbacks that are related to  
13 dependences and that are dispatched when new tasks are generated and when **ordered** constructs  
14 are encountered.

### 15 **Format**

16   
17 `typedef void (*ompt_callback_dependences_t) (  
18 ompt_data_t *task_data,  
19 const ompt_dependence_t *deps,  
20 int ndeps  
);`

### 21 **Trace Record**

22   
23 `typedef struct ompt_record_dependences_t {  
24 ompt_id_t task_id;  
25 ompt_dependence_t dep;  
26 int ndeps;  
} ompt_record_dependences_t;`

## Description of Arguments

The binding of the *task\_data* argument is the generated task.

The *deps* argument lists dependences of the new task or the dependence vector of the ordered construct.

The *ndeps* argument specifies the length of the list passed by the *deps* argument. The memory for *deps* is owned by the caller; the tool cannot rely on the data after the callback returns.

The performance monitor interface for tracing activity on target devices provides one record per dependence.

## Cross References

- **ordered** construct, see Section 2.20.9 on page 252.
- **depend** clause, see Section 2.20.11 on page 257.
- **ompt\_data\_t** type, see Section 4.4.4.4 on page 444.
- **ompt\_dependence\_t** type, see Section 4.4.4.9 on page 446.

## 4.5.2.9 ompt\_callback\_task\_dependence\_t

### Summary

The **ompt\_callback\_task\_dependence\_t** type is used for callbacks that are dispatched when unfulfilled task dependences are encountered.

### Format

```
C / C++  
typedef void (*ompt_callback_task_dependence_t) (  
    ompt_data_t *src_task_data,  
    ompt_data_t *sink_task_data  
);
```

### Trace Record

```
C / C++  
typedef struct ompt_record_task_dependence_t {  
    ompt_id_t src_task_id;  
    ompt_id_t sink_task_id;  
} ompt_record_task_dependence_t;
```

## Description of Arguments

The binding of the *src\_task\_data* argument is a running task with an outgoing dependence.

The binding of the *sink\_task\_data* argument is a task with an unsatisfied incoming dependence.

## Cross References

- **depend** clause, see Section 2.20.11 on page 257.
- **ompt\_data\_t** type, see Section 4.4.4.4 on page 444.

### 4.5.2.10 ompt\_callback\_task\_schedule\_t

#### Summary

The **ompt\_callback\_task\_schedule\_t** type is used for callbacks that are dispatched when task scheduling decisions are made.

#### Format

```
typedef void (*ompt_callback_task_schedule_t) (  
    ompt_data_t *prior_task_data,  
    ompt_task_status_t prior_task_status,  
    ompt_data_t *next_task_data  
);
```

#### Trace Record

```
typedef struct ompt_record_task_schedule_t {  
    ompt_id_t prior_task_id;  
    ompt_task_status_t prior_task_status;  
    ompt_id_t next_task_id;  
} ompt_record_task_schedule_t;
```

## Description of Arguments

The *prior\_task\_status* argument indicates the status of the task that arrived at a task scheduling point.

The binding of the *prior\_task\_data* argument is the task that arrived at the scheduling point.

The binding of the *next\_task\_data* argument is the task that is resumed at the scheduling point. This argument is **NULL** if the callback is dispatched for a *task-fulfill* event.

## Cross References

- Task scheduling, see Section 2.13.6 on page 156.
- `ompt_data_t` type, see Section 4.4.4.4 on page 444.
- `ompt_task_status_t` type, see Section 4.4.4.19 on page 451.

### 4.5.2.11 `ompt_callback_implicit_task_t`

#### Summary

The `ompt_callback_implicit_task_t` type is used for callbacks that are dispatched when initial tasks and implicit tasks are generated and completed.

#### Format

C / C++

```
typedef void (*ompt_callback_implicit_task_t) (  
    ompt_scope_endpoint_t endpoint,  
    ompt_data_t *parallel_data,  
    ompt_data_t *task_data,  
    unsigned int actual_parallelism,  
    unsigned int index,  
    int flags  
);
```

C / C++

#### Trace Record

C / C++

```
typedef struct ompt_record_implicit_task_t {  
    ompt_scope_endpoint_t endpoint;  
    ompt_id_t parallel_id;  
    ompt_id_t task_id;  
    unsigned int actual_parallelism;  
    unsigned int index;  
    int flags;  
} ompt_record_implicit_task_t;
```

C / C++

## Description of Arguments

The *endpoint* argument indicates that the callback signals the beginning of a scope or the end of a scope.

The binding of the *parallel\_data* argument is the current parallel region. For the *implicit-task-end* event, this argument is **NULL**.

The binding of the *task\_data* argument is the implicit task that executes the structured block of the parallel region.

The *actual\_parallelism* argument indicates the number of threads in the **parallel** region or the number of teams in the **teams** region. For initial tasks, that are not closely nested in a **teams** construct, this argument is **1**. For the *implicit-task-end* and the *initial-task-end* events, this argument is **0**.

The *index* argument indicates the thread number or team number of the calling thread, within the team or league that is executing the parallel or **teams** region to which the implicit task region binds. For initial tasks, that are not created by a **teams** construct, this argument is **1**.

The *flags* argument indicates the kind of the task (initial or implicit).

## Cross References

- **parallel** construct, see Section 2.8 on page 79.
- **teams** construct, see Section 2.9 on page 87.
- **ompt\_data\_t** type, see Section 4.4.4.4 on page 444.
- **ompt\_scope\_endpoint\_t** enumeration type, see Section 4.4.4.11 on page 447.

### 4.5.2.12 ompt\_callback\_master\_t

#### Summary

The **ompt\_callback\_master\_t** type is used for callbacks that are dispatched when **master** regions start and end.

#### Format

```
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++

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 *endpoint* argument indicates that the callback signals the beginning of a scope or the end of a scope.

The binding of the *parallel\_data* argument is the current parallel region.

The binding of the *task\_data* argument is the encountering task.

The *codeptr\_ra* argument relates the implementation of an OpenMP region to its source code. If a runtime routine implements the region associated with a callback that has type signature **ompt\_callback\_master\_t** then *codeptr\_ra* contains the return address of the call to that runtime routine. If the implementation of the region is inlined then *codeptr\_ra* contains the return address of the invocation of the callback. If attribution to source code is impossible or inappropriate, *codeptr\_ra* may be **NULL**.

## Cross References

- **master** construct, see Section 2.19 on page 225.
- **ompt\_data\_t** type, see Section 4.4.4.4 on page 444.
- **ompt\_scope\_endpoint\_t** type, see Section 4.4.4.11 on page 447.

### 4.5.2.13 ompt\_callback\_sync\_region\_t

#### Summary

The **ompt\_callback\_sync\_region\_t** type is used for callbacks that are dispatched when barrier regions, **taskwait** regions, and **taskgroup** regions begin and end and when waiting begins and ends for them as well as for when reductions are performed.

## Format

C / C++

```
typedef void (*ompt_callback_sync_region_t) (  
    ompt_sync_region_t kind,  
    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_sync_region_t {  
    ompt_sync_region_t kind;  
    ompt_scope_endpoint_t endpoint;  
    ompt_id_t parallel_id;  
    ompt_id_t task_id;  
    const void *codeptr_ra;  
} ompt_record_sync_region_t;
```

C / C++

## Description of Arguments

The *kind* argument indicates the kind of synchronization.

The *endpoint* argument indicates that the callback signals the beginning of a scope or the end of a scope.

The binding of the *parallel\_data* argument is the current parallel region. For the *barrier-end* event at the end of a parallel region this argument is **NULL**.

The binding of the *task\_data* argument is the current task.

The *codeptr\_ra* argument relates the implementation of an OpenMP region to its source code. If a runtime routine implements the region associated with a callback that has type signature **ompt\_callback\_sync\_region\_t** then *codeptr\_ra* contains the return address of the call to that runtime routine. If the implementation of the region is inlined then *codeptr\_ra* contains the return address of the invocation of the callback. If attribution to source code is impossible or inappropriate, *codeptr\_ra* may be **NULL**.

## Cross References

- **barrier** construct, see Section 2.20.2 on page 229.
- Implicit barriers, see Section 2.20.3 on page 231.
- **taskwait** construct, see Section 2.20.5 on page 233.
- **taskgroup** construct, see Section 2.20.6 on page 235.
- Properties common to all reduction clauses, see Section 2.22.5.1 on page 295.
- **ompt\_data\_t** type, see Section 4.4.4.4 on page 444.
- **ompt\_scope\_endpoint\_t** type, see Section 4.4.4.11 on page 447.
- **ompt\_sync\_region\_t** type, see Section 4.4.4.13 on page 448.

### 4.5.2.14 ompt\_callback\_mutex\_acquire\_t

#### Summary

The **ompt\_callback\_mutex\_acquire\_t** type is used for callbacks that are dispatched when locks are initialized, acquired and tested and when **critical** regions, **atomic** regions, and **ordered** regions are begun.

#### Format

C / C++

```
typedef void (*ompt_callback_mutex_acquire_t) (  
    ompt_mutex_t kind,  
    unsigned int hint,  
    unsigned int impl,  
    ompt_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;  
    ompt_wait_id_t wait_id;  
    const void *codeptr_ra;  
} ompt_record_mutex_acquire_t;
```

C / C++

## Description of Arguments

The *kind* argument indicates the kind of the lock involved.

The *hint* argument indicates the hint that was 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 *impl* argument indicates the mechanism chosen by the runtime to implement the mutual exclusion.

The *wait\_id* argument indicates the object being awaited.

The *codeptr\_ra* argument relates the implementation of an OpenMP region to its source code. If a runtime routine implements the region associated with a callback that has type signature `ompt_callback_mutex_acquire_t` then *codeptr\_ra* contains the return address of the call to that runtime routine. If the implementation of the region is inlined then *codeptr\_ra* contains the return address of the invocation of the callback. If attribution to source code is impossible or inappropriate, *codeptr\_ra* may be `NULL`.

## Cross References

- `critical` construct, see Section 2.20.1 on page 226.
- `atomic` construct, see Section 2.20.7 on page 237.
- `ordered` construct, see Section 2.20.9 on page 252.
- `omp_init_lock` and `omp_init_nest_lock` routines, see Section 3.3.1 on page 384.
- `ompt_mutex_t` type, see Section 4.4.4.16 on page 449.
- `ompt_wait_id_t` type, see Section 4.4.4.29 on page 459.

### 4.5.2.15 `ompt_callback_mutex_t`

#### Summary

The `ompt_callback_mutex_t` type is used for callbacks that indicate important synchronization events.

#### Format

```
typedef void (*ompt_callback_mutex_t) (  
    ompt_mutex_t kind,  
    ompt_wait_id_t wait_id,  
    const void *codeptr_ra  
);
```

C / C++

## Trace Record

C / C++

```
typedef struct ompt_record_mutex_t {  
    ompt_mutex_t kind;  
    ompt_wait_id_t wait_id;  
    const void *codeptr_ra;  
} ompt_record_mutex_t;
```

C / C++

## Description of Arguments

The *kind* argument indicates the kind of mutual exclusion event.

The *wait\_id* argument indicates the object being awaited.

The *codeptr\_ra* argument relates the implementation of an OpenMP region to its source code. If a runtime routine implements the region associated with a callback that has type signature **ompt\_callback\_mutex\_t** then *codeptr\_ra* contains the return address of the call to that runtime routine. If the implementation of the region is inlined then *codeptr\_ra* contains the return address of the invocation of the callback. If attribution to source code is impossible or inappropriate, *codeptr\_ra* may be **NULL**.

## Cross References

- **critical** construct, see Section 2.20.1 on page 226.
- **atomic** construct, see Section 2.20.7 on page 237.
- **ordered** construct, see Section 2.20.9 on page 252.
- **omp\_destroy\_lock** and **omp\_destroy\_nest\_lock** routines, see Section 3.3.3 on page 387.
- **omp\_set\_lock** and **omp\_set\_nest\_lock** routines, see Section 3.3.4 on page 388.
- **omp\_unset\_lock** and **omp\_unset\_nest\_lock** routines, see Section 3.3.5 on page 389.
- **omp\_test\_lock** and **omp\_test\_nest\_lock** routines, see Section 3.3.6 on page 391.
- **ompt\_mutex\_t** type, see Section 4.4.4.16 on page 449.
- **ompt\_wait\_id\_t** type, see Section 4.4.4.29 on page 459.

### 4.5.2.16 ompt\_callback\_nest\_lock\_t

#### Summary

The **ompt\_callback\_nest\_lock\_t** type is used for callbacks that indicate that a thread that owns a nested lock has performed an action related to the lock but has not relinquished ownership of it.

## Format

C / C++

```
typedef void (*ompt_callback_nest_lock_t) (  
    ompt_scope_endpoint_t endpoint,  
    ompt_wait_id_t wait_id,  
    const void *codeptr_ra  
);
```

C / C++

## 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 *endpoint* argument indicates that the callback signals the beginning of a scope or the end of a scope.

The *wait\_id* argument indicates the object being awaited.

The *codeptr\_ra* argument relates the implementation of an OpenMP region to its source code. If a runtime routine implements the region associated with a callback that has type signature **ompt\_callback\_nest\_lock\_t** then *codeptr\_ra* contains the return address of the call to that runtime routine. If the implementation of the region is inlined then *codeptr\_ra* contains the return address of the invocation of the callback. If attribution to source code is impossible or inappropriate, *codeptr\_ra* may be **NULL**.

## Cross References

- **ompt\_set\_nest\_lock** routine, see Section 3.3.4 on page 388.
- **ompt\_unset\_nest\_lock** routine, see Section 3.3.5 on page 389.
- **ompt\_test\_nest\_lock** routine, see Section 3.3.6 on page 391.
- **ompt\_scope\_endpoint\_t** type, see Section 4.4.4.11 on page 447.
- **ompt\_wait\_id\_t** type, see Section 4.4.4.29 on page 459.

### 4.5.2.17 `ompt_callback_flush_t`

#### Summary

The `ompt_callback_flush_t` type is used for callbacks that are dispatched when `flush` constructs are encountered.

#### 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 binding of the `thread_data` argument is the executing thread.

The `codeptr_ra` argument relates the implementation of an OpenMP region to its source code. If a runtime routine implements the region associated with a callback that has type signature `ompt_callback_flush_t` then `codeptr_ra` contains the return address of the call to that runtime routine. If the implementation of the region is inlined then `codeptr_ra` contains the return address of the invocation of the callback. If attribution to source code is impossible or inappropriate, `codeptr_ra` may be `NULL`.

#### Cross References

- `flush` construct, see Section 2.20.8 on page 245.
- `ompt_data_t` type, see Section 4.4.4.4 on page 444.

### 4.5.2.18 `ompt_callback_cancel_t`

#### Summary

The `ompt_callback_cancel_t` type is used for callbacks that are dispatched for `cancellation`, `cancel` and `discarded-task` events.

## Format

C / C++

```
typedef void (*ompt_callback_cancel_t) (  
    ompt_data_t *task_data,  
    int flags,  
    const void *codeptr_ra  
);
```

C / C++

## Trace Record

C / C++

```
typedef struct ompt_record_cancel_t {  
    ompt_id_t task_id;  
    int flags;  
    const void *codeptr_ra;  
} ompt_record_cancel_t;
```

C / C++

## Description of Arguments

The binding of the *task\_data* argument is the task that encounters a **cancel** construct, a **cancellation point** construct, or a construct defined as having an implicit cancellation point.

The *flags* argument, defined by the **ompt\_cancel\_flag\_t** enumeration type, indicates whether cancellation is activated by the current task, or detected as being activated by another task. The construct that is being canceled is also described in the *flags* argument. When several constructs are detected as being concurrently canceled, each corresponding bit in the argument will be set.

The *codeptr\_ra* argument relates the implementation of an OpenMP region to its source code. If a runtime routine implements the region associated with a callback that has type signature **ompt\_callback\_cancel\_t** then *codeptr\_ra* contains the return address of the call to that runtime routine. If the implementation of the region is inlined then *codeptr\_ra* contains the return address of the invocation of the callback. If attribution to source code is impossible or inappropriate, *codeptr\_ra* may be **NULL**.

## Cross References

- **ompt\_cancel\_flag\_t** enumeration type, see Section 4.4.4.24 on page 453.

### 4.5.2.19 ompt\_callback\_device\_initialize\_t

#### Summary

The **ompt\_callback\_device\_initialize\_t** type is used for callbacks that initialize device tracing interfaces.



## Format

C / C++

```
typedef void (*ompt_callback_device_initialize_t) (  
    int device_num,  
    const char *type,  
    ompt_device_t *device,  
    ompt_function_lookup_t lookup,  
    const char *documentation  
);
```

C / C++

## Description

Registration of a callback with type signature `ompt_callback_device_initialize_t` for the `ompt_callback_device_initialize` event enables asynchronous collection of a trace for a device. The OpenMP implementation invokes this callback after OpenMP is initialized for the device but before execution of any OpenMP construct is started on the device.

## Description of Arguments

The `device_num` argument identifies the logical device that is being initialized.

The `type` argument is a character string that indicates the type of the device. A device type string is a semicolon separated character string that includes at a minimum the vendor and model name of the device. These names may be followed by a semicolon-separated sequence of properties that describe the hardware or software of the device.

The `device` argument is a pointer to an opaque object that represents the target device instance. Functions in the device tracing interface use this pointer to identify the device that is being addressed.

The `lookup` argument points to a runtime callback that a tool must use to obtain pointers to runtime entry points in the device's OMPT tracing interface. If a device does not support tracing then `lookup` is **NULL**.

The `documentation` argument is a string that describes how to use any device-specific runtime entry points that can be obtained through the `lookup` argument. This documentation string may be a pointer to external documentation, or it may be inline descriptions that include names and type signatures for any device-specific interfaces that are available through the `lookup` argument along with descriptions of how to use these interface functions to control monitoring and analysis of device traces.

## Constraints on Arguments

The `type` and `documentation` arguments must be immutable strings that are defined for the lifetime of a program execution.

## Effect

A device initializer must fulfill several duties. First, the *type* argument should be used to determine if any special knowledge about the hardware and/or software of a device is employed. Second, the *lookup* argument should be used to look up pointers to runtime entry points in the OMPT tracing interface for the device. Finally, these runtime entry points should be used to set up tracing for the device.

Initialization of tracing for a target device is described in Section [4.2.5](#) on page [432](#).

## Cross References

- `ompt_function_lookup_t` type, see Section [4.6.3](#) on page [526](#).

### 4.5.2.20 `ompt_callback_device_finalize_t`

#### Summary

The `ompt_callback_device_initialize_t` type is used for callbacks that finalize device tracing interfaces.

#### Format

```
typedef void (*ompt_callback_device_finalize_t) (  
    int device_num  
);
```

C / C++

C / C++

#### Description of Arguments

The *device\_num* argument identifies the logical device that is being finalized.

#### Description

A registered callback with type signature `ompt_callback_device_finalize_t` is dispatched for a device immediately prior to finalizing the device. Prior to dispatching a finalization callback for a device on which tracing is active, the OpenMP implementation stops tracing on the device and synchronously flushes all trace records for the device that have not yet been reported. These trace records are flushed through one or more buffer completion callbacks with type signature `ompt_callback_buffer_complete_t` as needed prior to the dispatch of the callback with type signature `ompt_callback_device_finalize_t`.

## Cross References

- `ompt_callback_buffer_complete_t` callback type, see Section [4.5.2.24](#) on page [487](#).

## 4.5.2.21 `ompt_callback_device_load_t`

### Summary

The `ompt_callback_device_load_t` type is used for callbacks that the OpenMP runtime invokes to indicate that it has just loaded code onto the specified device.

### Format

```
C / C++
typedef void (*ompt_callback_device_load_t) (
    int device_num,
    const char *filename,
    int64_t offset_in_file,
    void *vma_in_file,
    size_t bytes,
    void *host_addr,
    void *device_addr,
    uint64_t module_id
);
C / C++
```

### Description of Arguments

The `device_num` argument specifies the device.

The `filename` argument indicates the name of a file in which the device code can be found. A NULL `filename` indicates that the code is not available in a file in the file system.

The `offset_in_file` argument indicates an offset into `filename` at which the code can be found. A value of -1 indicates that no offset is provided.

`ompt_addr_none` is defined as a pointer with the value ~0.

The `vma_in_file` argument indicates an virtual address in `filename` at which the code can be found. A value of `ompt_addr_none` indicates that a virtual address in the file is not available.

The `bytes` argument indicates the size of the device code object in bytes.

The `host_addr` argument indicates the address at which a copy of the device code is available in host memory. A value of `ompt_addr_none` indicates that a host code address is not available.

The `device_addr` argument indicates the address at which the device code has been loaded in device memory. A value of `ompt_addr_none` indicates that a device code address is not available.

The `module_id` argument is an identifier that is associated with the device code object.

### Cross References

- Device directives, see Section 2.15 on page 165.

### 4.5.2.22 `ompt_callback_device_unload_t`

#### Summary

The `ompt_callback_device_unload_t` type is used for callbacks that the OpenMP runtime invokes to indicate that it is about to unload code from the specified device.

#### Format

```
typedef void (*ompt_callback_device_unload_t) (  
    int device_num,  
    uint64_t module_id  
);
```

#### Description of Arguments

The `device_num` argument specifies the device.

The `module_id` argument is an identifier that is associated with the device code object.

#### Cross References

- Device directives, see Section 2.15 on page 165.

### 4.5.2.23 `ompt_callback_buffer_request_t`

#### Summary

The `ompt_callback_buffer_request_t` type is used for callbacks that are dispatched when a buffer to store event records for a device is requested.

#### Format

```
typedef void (*ompt_callback_buffer_request_t) (  
    int device_num,  
    ompt_buffer_t **buffer,  
    size_t *bytes  
);
```

#### Description

A callback with type signature `ompt_callback_buffer_request_t` requests a buffer to store trace records for the specified device. A buffer request callback may set `*bytes` to 0 if it does not provide a buffer. If a callback sets `*bytes` to 0, further recording of events for the device is disabled until the next invocation of `ompt_start_trace`. This action causes the device to drop future trace records until recording is restarted.

## Description of Arguments

The *device\_num* argument specifies the device.

The *\*buffer* argument points to a buffer where device events may be recorded. The *\*bytes* argument indicates the length of that buffer.

## Cross References

- `ompt_buffer_t` type, see Section 4.4.4.7 on page 446.

### 4.5.2.24 `ompt_callback_buffer_complete_t`

#### Summary

The `ompt_callback_buffer_complete_t` type is used for callbacks that are dispatched when devices will not record any more trace records in an event buffer and all records written to the buffer are valid.

#### Format

```
C / C++
typedef void (*ompt_callback_buffer_complete_t) (
    int device_num,
    ompt_buffer_t *buffer,
    size_t bytes,
    ompt_buffer_cursor_t begin,
    int buffer_owned
);
```

C / C++

#### Description

A callback with type signature `ompt_callback_buffer_complete_t` provides a buffer that contains trace records for the specified device. Typically, a tool will iterate through the records in the buffer and process them.

The OpenMP implementation makes these callbacks on a thread that is not an OpenMP master or worker thread.

The callee may not delete the buffer if the *buffer\_owned* argument is 0.

The buffer completion callback is not required to be *async signal safe*.

## Description of Arguments

The *device\_num* argument indicates the device which the buffer contains events.

The *buffer* argument is the address of a buffer that was previously allocated by a *buffer request* callback.

The *bytes* argument indicates the full size of the buffer.

The *begin* argument is an opaque cursor that indicates the position of the beginning of the first record in the buffer.

The *buffer\_owned* argument is 1 if the data to which the buffer points can be deleted by the callback and 0 otherwise. If multiple devices accumulate trace events into a single buffer, this callback may be invoked with a pointer to one or more trace records in a shared buffer with *buffer\_owned* = 0. In this case, the callback may not delete the buffer.

## Cross References

- `ompt_buffer_t` type, see Section 4.4.4.7 on page 446.
- `ompt_buffer_cursor_t` type, see Section 4.4.4.8 on page 446.

### 4.5.2.25 `ompt_callback_target_data_op_t`

#### Summary

The `ompt_callback_target_data_op_t` type is used for callbacks that are dispatched when a thread maps data to a device.

#### Format

```
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  
);
```

C / C++

## Trace Record

C / C++

```
typedef struct ompt_record_target_data_op_t {
    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;
    ompt_device_time_t end_time;
    const void *codeptr_ra;
} ompt_record_target_data_op_t;
```

C / C++

## Description

A registered **ompt\_callback\_target\_data\_op** callback is dispatched when device memory is allocated or freed, as well as when data is copied to or from a device.

**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. Thus, callbacks may not be dispatched as separate data operations on each variable.

## Description of Arguments

The *host\_op\_id* argument is a unique identifier for a data operations on a target device.

The *optype* argument indicates the kind of data mapping.

The *src\_addr* argument indicates the data address before the operation, where applicable.

The *src\_device\_num* argument indicates the source device number for the data operation, where applicable.

The *dest\_addr* argument indicates the data address after the operation.

The *dest\_device\_num* argument indicates the destination device number for the data operation.

It is implementation defined whether in some operations *src\_addr* or *dest\_addr* may point to an intermediate buffer.

The *bytes* argument indicates the size of data.

1 The *codeptr\_ra* argument relates the implementation of an OpenMP region to its source code. If a  
2 runtime routine implements the region associated with a callback that has type signature  
3 **ompt\_callback\_target\_data\_op\_t** then *codeptr\_ra* contains the return address of the call  
4 to that runtime routine. If the implementation of the region is inlined then *codeptr\_ra* contains the  
5 return address of the invocation of the callback. If attribution to source code is impossible or  
6 inappropriate, *codeptr\_ra* may be **NULL**.

## 7 Cross References

- 8 • **map** clause, see Section 2.22.7.1 on page 318.
- 9 • **ompt\_id\_t** type, see Section 4.4.4.3 on page 444.
- 10 • **ompt\_target\_data\_op\_t** type, see Section 4.4.4.14 on page 448.

### 11 4.5.2.26 ompt\_callback\_target\_t

#### 12 Summary

13 The **ompt\_callback\_target\_t** type is used for callbacks that are dispatched when a thread  
14 begins to execute a device construct.

#### 15 Format

C / C++

```
16 typedef void (*ompt_callback_target_t) (  
17     ompt_target_t kind,  
18     ompt_scope_endpoint_t endpoint,  
19     int device_num,  
20     ompt_data_t *task_data,  
21     ompt_id_t target_id,  
22     const void *codeptr_ra  
23 );
```

C / C++

#### 24 Trace Record

C / C++

```
25 typedef struct ompt_record_target_t {  
26     ompt_target_t kind;  
27     ompt_scope_endpoint_t endpoint;  
28     int device_num;  
29     ompt_id_t task_id;  
30     ompt_id_t target_id;  
31     const void *codeptr_ra;  
32 } ompt_record_target_t;
```

C / C++



## Description of Arguments

The *kind* argument indicates the kind of target region.

The *endpoint* argument indicates that the callback signals the beginning of a scope or the end of a scope.

The *device\_num* argument indicates the id of the device that will execute the target region.

The binding of the *task\_data* argument is the generating task.

The binding of the *target\_id* argument is the target region.

The *codeptr\_ra* argument relates the implementation of an OpenMP region to its source code. If a runtime routine implements the region associated with a callback that has type signature **ompt\_callback\_target\_t** then *codeptr\_ra* contains the return address of the call to that runtime routine. If the implementation of the region is inlined then *codeptr\_ra* contains the return address of the invocation of the callback. If attribution to source code is impossible or inappropriate, *codeptr\_ra* may be **NULL**.

## Cross References

- **target data** construct, see Section 2.15.2 on page 166.
- **target enter data** construct, see Section 2.15.3 on page 169.
- **target exit data** construct, see Section 2.15.4 on page 172.
- **target** construct, see Section 2.15.5 on page 174.
- **target update** construct, see Section 2.15.6 on page 181.
- **ompt\_id\_t** type, see Section 4.4.4.3 on page 444.
- **ompt\_data\_t** type, see Section 4.4.4.4 on page 444.
- **ompt\_scope\_endpoint\_t** type, see Section 4.4.4.11 on page 447.
- **ompt\_target\_t** type, see Section 4.4.4.20 on page 451.

### 4.5.2.27 ompt\_callback\_target\_map\_t

#### Summary

The **ompt\_callback\_target\_map\_t** type is used for callbacks that are dispatched to indicate data mapping relationships.

## Format

C / C++

```
typedef void (*ompt_callback_target_map_t) (  
    ompt_id_t target_id,  
    unsigned int nitems,  
    void **host_addr,  
    void **device_addr,  
    size_t *bytes,  
    unsigned int *mapping_flags,  
    const void *codeptr_ra  
);
```

C / C++

## Trace Record

C / C++

```
typedef struct ompt_record_target_map_t {  
    ompt_id_t target_id;  
    unsigned int nitems;  
    void **host_addr;  
    void **device_addr;  
    size_t *bytes;  
    unsigned int *mapping_flags;  
    const void *codeptr_ra;  
} ompt_record_target_map_t;
```

C / C++

## Description

An instance of a **target**, **target data**, **target enter data**, or **target exit data** construct may contain one or more **map** clauses. An OpenMP implementation may report the set of mappings associated with **map** clauses for a construct with a single **ompt\_callback\_target\_map** callback to report the effect of all mappings or multiple **ompt\_callback\_target\_map** callbacks with each reporting a subset of the mappings. Furthermore, an OpenMP implementation may omit mappings that it determines are unnecessary. If an OpenMP implementation issues multiple **ompt\_callback\_target\_map** callbacks, these callbacks may be interleaved with **ompt\_callback\_target\_data\_op** callbacks used to report data operations associated with the mappings.

## Description of Arguments

The binding of the *target\_id* argument is the target region.

The *nitems* argument indicates the number of data mappings that this callback reports.

The *host\_addr* argument indicates an array of host data addresses.

The *device\_addr* argument indicates an array of device data addresses.

The *bytes* argument indicates an array of size of data.

The *mapping\_flags* argument indicates the kind of data mapping. Flags for a mapping include one or more values specified by the `ompt_target_map_flag_t` type.

The *codeptr\_ra* argument relates the implementation of an OpenMP region to its source code. If a runtime routine implements the region associated with a callback that has type signature `ompt_callback_target_map_t` then *codeptr\_ra* contains the return address of the call to that runtime routine. If the implementation of the region is inlined then *codeptr\_ra* contains the return address of the invocation of the callback. If attribution to source code is impossible or inappropriate, *codeptr\_ra* may be `NULL`.

## Cross References

- `target data` construct, see Section 2.15.2 on page 166.
- `target enter data` construct, see Section 2.15.3 on page 169.
- `target exit data` construct, see Section 2.15.4 on page 172.
- `target` construct, see Section 2.15.5 on page 174.
- `ompt_id_t` type, see Section 4.4.4.3 on page 444.
- `ompt_target_map_flag_t` type, see Section 4.4.4.22 on page 453.
- `ompt_callback_target_data_op_t` callback type, see Section 4.5.2.25 on page 488.

### 4.5.2.28 `ompt_callback_target_submit_t`

#### Summary

The `ompt_callback_target_submit_t` type is used for callbacks that are dispatched when an initial task is created on a device.

#### Format

```
C / C++  
typedef void (*ompt_callback_target_submit_t) (  
    ompt_id_t target_id,  
    ompt_id_t host_op_id,  
    unsigned int requested_num_teams  
);  
C / C++
```

## 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;
```

C / C++

## 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 `target_id` argument is a unique identifier for the associated target region.

The `host_op_id` argument is a unique identifier for the initial task on the target device.

The `requested_num_teams` argument is the number of teams that the host requested to execute the kernel. The actual number of teams that execute the kernel may be smaller and generally will not be known until the kernel begins to execute on the device.

If `ompt_set_trace_ompt` has configured the device to trace kernel execution then 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 can be used 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; and
- The time when the initial task completed execution on the device is recorded in the `end_time` field.

## Cross References

- `target` construct, see Section [2.15.5](#) on page [174](#).
- `ompt_id_t` type, see Section [4.4.4.3](#) on page [444](#).

## 4.5.2.29 `ompt_callback_control_tool_t`

### Summary

The `ompt_callback_control_tool_t` type is used for callbacks that dispatch *tool-control* events.

### Format

C / C++

```
typedef int (*ompt_callback_control_tool_t) (  
    uint64_t command,  
    uint64_t modifier,  
    void *arg,  
    const void *codeptr_ra  
);
```

C / C++

### Trace Record

C / C++

```
typedef struct ompt_record_control_tool_t {  
    uint64_t command;  
    uint64_t modifier;  
    const void *codeptr_ra;  
} ompt_record_control_tool_t;
```

C / C++

### Description

Callbacks with type signature `ompt_callback_control_tool_t` may return any non-negative value, which will be returned to the application as the return value of the `ompt_control_tool` call that triggered the callback.

## Description of Arguments

The *command* argument passes a command from an application to a tool. Standard values for *command* are defined by `omp_control_tool_t` in Section 3.9 on page 421.

The *modifier* argument passes a command modifier from an application to a tool.

The *command* and *modifier* arguments may have tool-specific values. Tools must ignore *command* values that they are not designed to handle.

The *arg* argument is a void pointer that enables a tool and an application to exchange arbitrary state. The *arg* argument may be `NULL`.

The *codeptr\_ra* argument relates the implementation of an OpenMP region to its source code. If a runtime routine implements the region associated with a callback that has type signature `omp_callback_control_tool_t` then *codeptr\_ra* contains the return address of the call to that runtime routine. If the implementation of the region is inlined then *codeptr\_ra* contains the return address of the invocation of the callback. If attribution to source code is impossible or inappropriate, *codeptr\_ra* may be `NULL`.

## Constraints on Arguments

Tool-specific values for *command* must be  $\geq 64$ .

## Cross References

- `omp_control_tool_t` enumeration type, see Section 3.9 on page 421.

## 4.6 OMPT Runtime Entry Points for Tools

OMPT supports two principal sets of runtime entry points for tools. One set of runtime entry points enables a tool to register callbacks for OpenMP events and to inspect the state of an OpenMP thread while executing in a tool callback or a signal handler. The second set of runtime entry points enables a tool to trace activities on a device. When directed by the tracing interface, an OpenMP implementation will trace activities on a device, collect buffers of trace records, and invoke callbacks on the host to process these records. OMPT runtime entry points should not be global symbols since tools cannot rely on the visibility of such symbols.

OMPT also supports runtime entry points for two classes of lookup routines. The first class of lookup routines contains a single member: a routine that returns runtime entry points in the OMPT callback interface. The second class of lookup routines includes a unique lookup routine for each kind of device that can return runtime entry points in a device's OMPT tracing interface.

The C/C++ header file (`omp-tools.h`) provides the definitions of the types that are specified throughout this subsection.

## Restrictions

OMPT runtime entry points have the following restrictions:

- OMPT runtime entry points must not be called from a signal handler on a native thread before a *native-thread-begin* or after a *native-thread-end* event.
- OMPT device runtime entry points must not be called after a *device-finalize* event for that device.

## 4.6.1 Entry Points in the OMPT Callback Interface

Entry points in the OMPT callback interface enable a tool to register callbacks for OpenMP events and to inspect the state of an OpenMP thread while executing in a tool callback or a signal handler. Pointers to these runtime entry points are obtained through the lookup function that is provided through the OMPT initializer.

### 4.6.1.1 `ompt_enumerate_states_t`

#### Summary

The `ompt_enumerate_states_t` type is the type signature of the `ompt_enumerate_states` runtime entry point, which enumerates the thread states that an OpenMP implementation supports.

#### Format

```
C / C++
typedef int (*ompt_enumerate_states_t) (
    int current_state,
    int *next_state,
    const char **next_state_name
);
```

C / C++

#### Description

An OpenMP implementation may support only a subset of the states defined by the `ompt_state_t` enumeration type. An OpenMP implementation may also support implementation-specific states. The `ompt_enumerate_states` runtime entry point, which has type signature `ompt_enumerate_states_t`, enables a tool to enumerate the supported thread states.

When a supported thread state is passed as `current_state`, the runtime entry point assigns the next thread state in the enumeration to the variable passed by reference in `next_state` and assigns the name associated with that state to the character pointer passed by reference in `next_state_name`.

Whenever one or more states are left in the enumeration, the `ompt_enumerate_states` runtime entry point returns 1. When the last state in the enumeration is passed as `current_state`, `ompt_enumerate_states` returns 0, which indicates that the enumeration is complete.

## Description of Arguments

The *current\_state* argument must be a thread state that the OpenMP implementation supports. To begin enumerating the supported states, a tool should pass `ompt_state_undefined` as *current\_state*. Subsequent invocations of `ompt_enumerate_states` should pass the value assigned to the variable passed by reference in *next\_state* to the previous call.

The value `ompt_state_undefined` is reserved to indicate an invalid thread state. `ompt_state_undefined` is defined as an integer with the value 0.

The *next\_state* argument is a pointer to an integer in which `ompt_enumerate_states` returns the value of the next state in the enumeration.

The *next\_state\_name* argument is a pointer to a character string pointer through which `ompt_enumerate_states` returns a string that describes the next state.

## Constraints on Arguments

Any string returned through the *next\_state\_name* argument must be immutable and defined for the lifetime of a program execution.

## Cross References

- `ompt_state_t` type, see Section 4.4.4.26 on page 454.

### 4.6.1.2 `ompt_enumerate_mutex_impls_t`

#### Summary

The `ompt_enumerate_mutex_impls_t` type is the type signature of the `ompt_enumerate_mutex_impls` runtime entry point, which enumerates the kinds of mutual exclusion implementations that an OpenMP implementation employs.

#### Format

```
typedef int (*ompt_enumerate_mutex_impls_t) (  
    int current_impl,  
    int *next_impl,  
    const char **next_impl_name  
);
```

C / C++



## Description

Mutual exclusion for locks, **critical** sections, and **atomic** regions may be implemented in several ways. The **ompt\_enumerate\_mutex\_impls** runtime entry point, which has type signature **ompt\_enumerate\_mutex\_impls\_t**, enables a tool to enumerate the supported mutual exclusion implementations.

When a supported mutex implementation is passed as *current\_impl*, the runtime entry point assigns the next mutex implementation in the enumeration to the variable passed by reference in *next\_impl* and assigns the name associated with that mutex implementation to the character pointer passed by reference in *next\_impl\_name*.

Whenever one or more mutex implementations are left in the enumeration, the **ompt\_enumerate\_mutex\_impls** runtime entry point returns 1. When the last mutex implementation in the enumeration is passed as *current\_impl*, the runtime entry point returns 0, which indicates that the enumeration is complete.

## Description of Arguments

The *current\_impl* argument must be a mutex implementation that an OpenMP implementation supports. To begin enumerating the supported mutex implementations, a tool should pass **ompt\_mutex\_impl\_none** as *current\_impl*. Subsequent invocations of **ompt\_enumerate\_mutex\_impls** should pass the value assigned to the variable passed in *next\_impl* to the previous call.

The value **ompt\_mutex\_impl\_none** is reserved to indicate an invalid mutex implementation. **ompt\_mutex\_impl\_none** is defined as an integer with the value 0.

The *next\_impl* argument is a pointer to an integer in which **ompt\_enumerate\_mutex\_impls** returns the value of the next mutex implementation in the enumeration.

The *next\_impl\_name* argument is a pointer to a character string pointer in which **ompt\_enumerate\_mutex\_impls** returns a string that describes the next mutex implementation.

## Constraints on Arguments

Any string returned through the *next\_impl\_name* argument must be immutable and defined for the lifetime of a program execution.

## Cross References

- **ompt\_mutex\_t** type, see Section [4.4.4.16](#) on page [449](#).

### 4.6.1.3 ompt\_set\_callback\_t

#### Summary

The **ompt\_set\_callback\_t** type is the type signature of the **ompt\_set\_callback** runtime entry point, which registers a pointer to a tool callback that an OpenMP implementation invokes when a host OpenMP event occurs.

## Format

C / C++

```
typedef ompt_set_result_t (*ompt_set_callback_t) (  
    ompt_callbacks_t event,  
    ompt_callback_t callback  
);
```

C / C++

## Description

OpenMP implementations can use callbacks to indicate the occurrence of events during the execution of an OpenMP program. The **ompt\_set\_callback** runtime entry point, which has type signature **ompt\_set\_callback\_t**, registers a callback for an OpenMP event on the current device. The return value of **ompt\_set\_callback** indicates the outcome of registering the callback.

## Description of Arguments

The *event* argument indicates the event for which the callback is being registered.

The *callback* argument is a tool callback function. If *callback* is **NULL** then callbacks associated with *event* are disabled. If callbacks are successfully disabled then **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.

## Restrictions

The **ompt\_set\_callback** runtime entry point has the following restriction:

- The entry point must not return **ompt\_set\_impossible**.

## Cross References

- Monitoring activity on the host with OMPT, see Section 4.2.4 on page 430.
- **ompt\_callbacks\_t** enumeration type, see Section 4.4.2 on page 438.
- **ompt\_callback\_t** type, see Section 4.4.4.1 on page 442.
- **ompt\_set\_result\_t** type, see Section 4.4.4.2 on page 443.
- **ompt\_get\_callback\_t** host callback type signature, see Section 4.6.1.4 on page 501.

## 4.6.1.4 `ompt_get_callback_t`

### Summary

The `ompt_get_callback_t` type is the type signature of the `ompt_get_callback` runtime entry point, which retrieves a pointer to a registered tool callback routine (if any) that an OpenMP implementation invokes when a host OpenMP event occurs.

### Format

```
C / C++  
typedef int (*ompt_get_callback_t) (  
    ompt_callbacks_t event,  
    ompt_callback_t *callback  
);
```

C / C++

### Description

The `ompt_get_callback` runtime entry point, which has type signature `ompt_get_callback_t`, retrieves a pointer to the tool callback that an OpenMP implementation may 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 is assigned to the variable passed by reference in `callback` and `ompt_get_callback` returns 1; otherwise, it returns 0. If `ompt_get_callback` returns 0, the value of the variable passed by reference as `callback` is undefined.

### Description of Arguments

The `event` argument indicates the event for which the callback would be invoked.

The `callback` argument returns a pointer to the callback associated with `event`.

### Constraints on Arguments

The `callback` argument must be a reference to a variable of specified type.

### Cross References

- `ompt_callbacks_t` enumeration type, see Section 4.4.2 on page 438.
- `ompt_callback_t` type, see Section 4.4.4.1 on page 442.
- `ompt_set_callback_t` type signature, see Section 4.6.1.3 on page 499.

## 4.6.1.5 `ompt_get_thread_data_t`

### Summary

The `ompt_get_thread_data_t` type is the type signature of the `ompt_get_thread_data` runtime entry point, which returns the address of the thread data object for the current thread.

## Format

C / C++

```
typedef omp_t_data_t *(*omp_t_get_thread_data_t) (void);
```

C / C++

## Binding

The binding thread for the **omp\_t\_get\_thread\_data** runtime entry point is the current thread.

## Description

Each OpenMP thread can have an associated thread data object of type **omp\_t\_data\_t**. The **omp\_t\_get\_thread\_data** runtime entry point, which has type signature **omp\_t\_get\_thread\_data\_t**, retrieves a pointer to the thread data object, if any, that is associated with the current thread. A tool may use a pointer to an OpenMP thread's data object that **omp\_t\_get\_thread\_data** retrieves to inspect or to modify the value of the data object. When an OpenMP thread is created, its data object is initialized with value **omp\_t\_data\_none**.

This runtime entry point is *async signal safe*.

## Cross References

- **omp\_t\_data\_t** type, see Section 4.4.4.4 on page 444.

### 4.6.1.6 omp\_t\_get\_num\_procs\_t

#### Summary

The **omp\_t\_get\_num\_procs\_t** type is the type signature of the **omp\_t\_get\_num\_procs** runtime entry point, which returns the number of processors currently available to the execution environment on the host device.

## Format

C / C++

```
typedef int (*omp_t_get_num_procs_t) (void);
```

C / C++

## Binding

The binding thread set for the **omp\_t\_get\_num\_procs** runtime entry point is all threads on the host device.

## Description

The `ompt_get_num_procs` runtime entry point, which has type signature `ompt_get_num_procs_t`, 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*.

### 4.6.1.7 `ompt_get_num_places_t`

#### Summary

The `ompt_get_num_places_t` type is the type signature of the `ompt_get_num_places` runtime entry point, which returns the number of places currently 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 `ompt_get_num_places` runtime entry point is all threads on a device.

#### Description

The `ompt_get_num_places` runtime entry point, which has type signature `ompt_get_num_places_t`, 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 2.6 on page 66.
- `OMP_PLACES` environment variable, see Section 6.5 on page 588.

### 4.6.1.8 `ompt_get_place_proc_ids_t`

#### Summary

The `ompt_get_place_procs_ids_t` type is the type signature of the `ompt_get_num_place_procs_ids` runtime entry point, which returns the numerical identifiers of the processors that are available to the execution environment in the specified place.

## Format

C / C++

```
typedef int (*ompt_get_place_proc_ids_t) (  
    int place_num,  
    int ids_size,  
    int *ids  
);
```

C / C++

## Binding

The binding thread set for the `ompt_get_place_proc_ids` runtime entry point is all threads on a device.

## Description

The `ompt_get_place_proc_ids` runtime entry point, which has type signature `ompt_get_place_proc_ids_t`, returns the numerical identifiers of each processor that is associated with the specified place. These numerical identifiers are non-negative and their meaning is implementation defined.

## Description of Arguments

The `place_num` argument specifies the place that is being queried.

The `ids` argument is an array in which the routine can return a vector of processor identifiers in the specified place.

The `ids_size` argument indicates the size of the result array that is specified by `ids`.

## Effect

If the `ids` array of size `ids_size` is large enough to contain all identifiers then they are returned in `ids` and their order in the array is implementation defined. Otherwise, if the `ids` array is too small the values in `ids` when the function returns are unspecified. The routine always returns the number of numerical identifiers of the processors that are available to the execution environment in the specified place.

### 4.6.1.9 `ompt_get_place_num_t`

## Summary

The `ompt_get_place_num_t` type is the type signature of the `ompt_get_place_num` runtime entry point, which returns the place number of the place to which the current thread is bound.

## Format

C / C++

```
typedef int (*ompt_get_place_num_t) (void);
```

C / C++

## Binding

The binding thread set of the `ompt_get_place_num` runtime entry point is the current thread.

## Description

When the current thread is bound to a place, `ompt_get_place_num` returns the place number associated with the thread. The returned value is between 0 and one less than the value returned by `ompt_get_num_places`, inclusive. When the current thread is not bound to a place, the routine returns -1.

This runtime entry point is *async signal safe*.

### 4.6.1.10 `ompt_get_partition_place_nums_t`

## Summary

The `ompt_get_partition_place_nums_t` type is the type signature of the `ompt_get_partition_place_nums` runtime entry point, which returns a list of place numbers that correspond to the places in the *place-partition-var* ICV of the innermost implicit task.

## Format

C / C++

```
typedef int (*ompt_get_partition_place_nums_t) (  
    int place_nums_size,  
    int *place_nums  
);
```

C / C++

## Binding

The binding task set for the `ompt_get_partition_place_nums` runtime entry point is the current implicit task.

## Description

The `ompt_get_partition_place_nums` runtime entry point, which has type signature `ompt_get_partition_place_nums_t`, returns a list of place numbers that correspond to the places in the *place-partition-var* ICV of the innermost implicit task.

This runtime entry point is *async signal safe*.

## Description of Arguments

The `place_nums` argument is an array in which the routine can return a vector of place identifiers.

The `place_nums_size` argument indicates the size of the result array that the `place_nums` argument specifies.

## Effect

If the `place_nums` array of size `place_nums_size` is large enough to contain all identifiers then 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. The routine always returns the number of places in the `place-partition-var` ICV of the innermost implicit task.

## Cross References

- `place-partition-var` ICV, see Section 2.6 on page 66.
- `OMP_PLACES` environment variable, see Section 6.5 on page 588.

### 4.6.1.11 `ompt_get_proc_id_t`

#### Summary

The `ompt_get_proc_id_t` type is the type signature of the `ompt_get_proc_id` runtime entry point, which 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 `ompt_get_proc_id` runtime entry point is the current thread.

#### Description

The `ompt_get_proc_id` runtime entry point, which has type signature `ompt_get_proc_id_t`, 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*.

### 4.6.1.12 `ompt_get_state_t`

#### Summary

The `ompt_get_state_t` type is the type signature of the `ompt_get_state` runtime entry point, which returns the state and the wait identifier of the current thread.



## Format

```
typedef int (*ompt_get_state_t) (  
    ompt_wait_id_t *wait_id  
);
```

C / C++

## Binding

The binding thread for the `ompt_get_state` runtime entry point is the current thread.

## Description

Each OpenMP thread has an associated state and a wait identifier. If a thread's state indicates that the thread is waiting for mutual exclusion then its wait identifier contains an opaque handle that indicates the data object upon which the thread is waiting. The `ompt_get_state` runtime entry point, which has type signature `ompt_get_state_t`, retrieves the state and wait identifier of the current thread. The returned value may be any one of the states predefined by `ompt_state_t` or a value that represents any implementation specific state. The tool may obtain a string representation for each state with the `ompt_enumerate_states` function.

If the returned state indicates that the thread is waiting for a lock, nest lock, critical section, atomic region, or ordered region then the value of the thread's wait identifier is assigned to a non-null wait identifier passed as the `wait_id` argument.

This runtime entry point is *async signal safe*.

## Description of Arguments

The `wait_id` argument is a pointer to an opaque handle that is available to receive the value of the thread's wait identifier. If `wait_id` is not `NULL` then the entry point assigns the value of the thread's wait identifier to the object to which `wait_id` points. If the returned state is not one of the specified wait states then the value of opaque object to which `wait_id` points is undefined after the call.

## Constraints on Arguments

The argument passed to the entry point must be a reference to a variable of the specified type or `NULL`.

## Cross References

- `ompt_state_t` type, see Section [4.4.4.26](#) on page [454](#).
- `ompt_wait_id_t` type, see Section [4.4.4.29](#) on page [459](#).
- `ompt_enumerate_states_t` type, see Section [4.6.1.1](#) on page [497](#).

### 4.6.1.13 `ompt_get_parallel_info_t`

#### Summary

The `ompt_get_parallel_info_t` type is the type signature of the `ompt_get_parallel_info` runtime entry point, which returns information about the parallel region, if any, at the specified ancestor level for the current execution context.

#### Format

```
C / C++
typedef int (*ompt_get_parallel_info_t) (
    int ancestor_level,
    ompt_data_t **parallel_data,
    int *team_size
);
```

C / C++

#### Description

During execution, an OpenMP program may employ nested parallel regions. The `ompt_get_parallel_info` runtime entry point known, which has type signature `ompt_get_parallel_info_t`, retrieves information, about the current parallel region 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 to 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*.

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 `ompt_state_wait_barrier_implicit_parallel` then 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.

## Description of Arguments

The *ancestor\_level* argument specifies the parallel region of interest by its ancestor level. Ancestor level 0 refers to the innermost parallel region; information about enclosing parallel regions may be obtained using larger values for *ancestor\_level*.

The *parallel\_data* argument returns the parallel data if the argument is not **NULL**.

The *team\_size* argument 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, **ompt\_get\_parallel\_info** has the following effects:

- If a non-null value was passed for *parallel\_data*, the value returned in *parallel\_data* is a pointer to a data word that is associated with the parallel region at the specified level; and
- If a non-null value was passed for *team\_size*, the value returned in the integer to which *team\_size* point is the number of threads in the team that is 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**.

## Cross References

- **ompt\_data\_t** type, see Section 4.4.4.4 on page 444.

### 4.6.1.14 ompt\_get\_task\_info\_t

#### Summary

The **ompt\_get\_task\_info\_t** type is the type signature of the **ompt\_get\_task\_info** runtime entry point, which returns information about the task, if any, at the specified ancestor level in the current execution context.

#### Format

```
typedef int (*ompt_get_task_info_t) (  
    int ancestor_level,  
    int *flags,  
    ompt_data_t **task_data,  
    ompt_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 that are 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 `ompt_get_task_info` runtime entry point, which has 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 then information is 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 then the values of variables passed by reference to the entry point are undefined when `ompt_get_task_info` returns.

A tool may use a pointer to a data object for a task or parallel region that it obtains from `ompt_get_task_info` to inspect or to 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 *ancestor\_level* argument specifies the task region of interest 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 *flags* argument returns the task type if the argument is not **NULL**.

The *task\_data* argument returns the task data if the argument is not **NULL**.

The *task\_frame* argument returns the task frame pointer if the argument is not **NULL**.

The *parallel\_data* argument returns the parallel data if the argument is not **NULL**.

The *thread\_num* argument 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, `ompt_get_task_info` has the following effects:

- If a non-null value was passed for *flags* then the value returned in the integer to which *flags* points represents the type of the task at the specified level; possible task types include initial, implicit, explicit, and target tasks;

- 1 • If a non-null value was passed for *task\_data* then the value that is returned in the object to which  
2 it points is a pointer to a data word that is associated with the task at the specified level;
- 3 • If a non-null value was passed for *task\_frame* then the value that is returned in the object to  
4 which *task\_frame* points is a pointer to the **ompt\_frame\_t** structure that is associated with the  
5 task at the specified level;
- 6 • If a non-null value was passed for *parallel\_data* then the value that is returned in the object to  
7 which *parallel\_data* points is a pointer to a data word that is associated with the parallel region  
8 that contains the task at the specified level or, if the task at the specified level is an initial task,  
9 **NULL**; and
- 10 • If a non-null value was passed for *thread\_num* then the value that is returned in the object to  
11 which *thread\_num* points indicates the number of the thread in the parallel region that is  
12 executing the task at the specified level.

### 13 Constraints on Arguments

14 While argument *ancestor\_level* is passed by value, all other arguments to  
15 **ompt\_get\_task\_info** must be pointers to variables of the specified types or **NULL**.

### 16 Cross References

- 17 • **ompt\_data\_t** type, see Section 4.4.4.4 on page 444.
- 18 • **ompt\_task\_flag\_t** type, see Section 4.4.4.18 on page 450.
- 19 • **ompt\_frame\_t** type, see Section 4.4.4.27 on page 457.

### 20 4.6.1.15 ompt\_get\_task\_memory\_t

#### 21 Summary

22 The **ompt\_get\_task\_memory\_t** type is the type signature of the  
23 **ompt\_get\_task\_memory** runtime entry point, which returns information about memory ranges  
24 that are associated with the task.

#### 25 Format

```
26 typedef int (*ompt_get_task_memory_t) (  
27     void **addr,  
28     size_t *size,  
29     int block  
30 );
```

C / C++

## Description

During execution, an OpenMP thread may be executing an OpenMP task. The OpenMP implementation must preserve the data environment from the creation of the task for the execution of the task. The `ompt_get_task_memory` runtime entry point, which has type signature `ompt_get_task_memory_t`, provides information about the memory ranges used to store the data environment for the current task.

Multiple memory ranges may be used to store these data. The *block* argument supports iteration 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 no memory is 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 *addr* argument is a pointer to a void pointer return value to provide the start address of a memory block.

The *size* argument is a pointer to a size type return value to provide the size of the memory block.

The *block* argument is an integer value to specify the memory block of interest.

### 4.6.1.16 `ompt_get_target_info_t`

#### Summary

The `ompt_get_target_info_t` type is the type signature of the `ompt_get_target_info` runtime entry point, which returns identifiers that specify a thread's current `target` region and target operation ID, if any.

#### Format

```
typedef int (*ompt_get_target_info_t) (  
    uint64_t *device_num,  
    ompt_id_t *target_id,  
    ompt_id_t *host_op_id  
);
```

## Description

The `ompt_get_target_info` entry point, which has type signature `ompt_get_target_info_t`, returns 1 if the current thread is in a **target** region and 0 otherwise. If the entry point returns 0 then the values of the variables passed by reference as its arguments are undefined.

If the current thread is in a **target** region then `ompt_get_target_info` returns 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

The `device_num` argument returns the device number if the current thread is in a **target** region.

The `target_id` argument returns the **target** region identifier if the current thread is in a **target** region.

If the current thread is in the process of initiating an operation on a target device (for example, copying data to or from an accelerator or launching a kernel) then `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 4.4.4.3 on page 444.

### 4.6.1.17 `ompt_get_num_devices_t`

#### Summary

The `ompt_get_num_devices_t` type is the type signature of the `ompt_get_num_devices` runtime entry point, which returns the number of available devices.

#### Format

```

C / C++
typedef int (*ompt_get_num_devices_t) (void);
C / C++
```

#### Description

The `ompt_get_num_devices` runtime entry point, which has type signature `ompt_get_num_devices_t`, returns the number of devices available to an OpenMP program.

This runtime entry point is *async signal safe*.

### 4.6.1.18 `ompt_get_unique_id_t`

#### Summary

The `ompt_get_unique_id_t` type is the type signature of the `ompt_get_unique_id` runtime entry point, which returns a unique number.

#### Format

```
typedef uint64_t (*ompt_get_unique_id_t) (void);
```

#### Description

The `ompt_get_unique_id` runtime entry point, which has type signature `ompt_get_unique_id_t`, returns a number that is unique for the duration of an OpenMP program. Successive invocations may not result in consecutive or even increasing numbers.

This runtime entry point is *async signal safe*.

### 4.6.1.19 `ompt_finalize_tool_t`

#### Summary

The `ompt_finalize_tool_t` type is the type signature of the `ompt_finalize_tool` runtime entry point, which enables a tool to finalize itself.

#### Format

```
typedef void (*ompt_finalize_tool_t) (void);
```

#### Description

A tool may detect that the execution of an OpenMP program is ending before the OpenMP implementation does. To facilitate clean termination of the tool, the tool may invoke the `ompt_finalize_tool` runtime entry point, which has type signature `ompt_finalize_tool_t`. Upon completion of `ompt_finalize_tool`, no OMPT callbacks are dispatched.

#### Effect

The `ompt_finalize_tool` routine detaches the tool from the runtime, unregisters all callbacks and invalidates all OMPT entry points passed to the tool in the *lookup-function*. Upon completion of `ompt_finalize_tool`, no further callbacks will be issued on any thread.

Before the callbacks are unregistered, the OpenMP runtime should attempt to dispatch all outstanding registered callbacks as well as the callbacks that would be encountered during shutdown of the runtime, if possible in the current execution context.



## 4.6.2 Entry Points in the OMPT Device Tracing Interface

The runtime entry points with type signatures of the types that are specified in this section enable a tool to trace activities on a device.

### 4.6.2.1 `ompt_get_device_num_procs_t`

#### Summary

The `ompt_get_device_num_procs_t` type is the type signature of the `ompt_get_device_num_procs` runtime entry point, which returns the number of processors currently available to the execution environment on the specified device.

#### Format

```
typedef int (*ompt_get_device_num_procs_t) (  
    ompt_device_t *device  
);
```

#### Description

The `ompt_get_device_num_procs` runtime entry point, which has type signature `ompt_get_device_num_procs_t`, returns the number of processors that are available on the 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.

#### Description of Arguments

The `device` argument 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` type, see Section 4.4.4.5 on page 445.

### 4.6.2.2 `ompt_get_device_time_t`

#### Summary

The `ompt_get_device_time_t` type is the type signature of the `ompt_get_device_time` runtime entry point, which returns the current time on the specified device.

## Format

C / C++

```
typedef ompt_device_time_t (*ompt_get_device_time_t) (  
    ompt_device_t *device  
);
```

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, a common time base for ordering host-side and device-side events may not be available.

The `ompt_get_device_time` runtime entry point, which has 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 `device` argument 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` type, see Section 4.4.4.5 on page 445.
- `ompt_device_time_t` type, see Section 4.4.4.6 on page 445.

### 4.6.2.3 `ompt_translate_time_t`

## Summary

The `ompt_translate_time_t` type is the type signature of the `ompt_translate_time` runtime entry point, which translates a time value that is 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

The `ompt_translate_time` runtime entry point, which has 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 `ompt_get_wtime`.

**Note** – The accuracy of time translations may degrade if they are not performed promptly after a device time value is received and if either the host or device vary their clock speeds. Prompt translation of device times to host times is recommended.

## Description of Arguments

The *device* argument 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 *time* argument is a time from the specified device.

## Cross References

- `ompt_get_wtime` routine, see Section 3.4.1 on page 393.
- `ompt_device_t` type, see Section 4.4.4.5 on page 445.
- `ompt_device_time_t` type, see Section 4.4.4.6 on page 445.

## 4.6.2.4 `ompt_set_trace_ompt_t`

### Summary

The `ompt_set_trace_ompt_t` type is the type signature of the `ompt_set_trace_ompt` runtime entry point, which enables or disables the recording of trace records for one or more types of OMPT events.

### Format

```
C / C++
typedef ompt_set_result_t (*ompt_set_trace_ompt_t) (
    ompt_device_t *device,
    unsigned int enable,
    unsigned int etype
);
C / C++
```

## Description of Arguments

The *device* argument points to an opaque object that represents the target device instance. Functions in the device tracing interface use this pointer to identify the device that is being addressed.

The *etype* argument indicates the events to which the invocation of `ompt_set_trace_ompt` applies. If the value of *etype* is 0 then the invocation applies to all events. If *etype* is positive then it applies to the event in `ompt_callbacks_t` that matches that value.

The *enable* argument indicates whether tracing should be enabled or disabled for the event or events that the *etype* argument specifies. A positive value for *enable* indicates that recording should be enabled; a value of 0 for *enable* indicates that recording should be disabled.

## Restrictions

The `ompt_set_trace_ompt` runtime entry point has the following restriction:

- The entry point must not return `ompt_set_sometimes_paired`.

## Cross References

- Tracing activity on target devices with OMPT, see Section 4.2.5 on page 432.
- `ompt_callbacks_t` type, see Section 4.4.2 on page 438.
- `ompt_set_result_t` type, see Section 4.4.4.2 on page 443.
- `ompt_device_t` type, see Section 4.4.4.5 on page 445.

### 4.6.2.5 `ompt_set_trace_native_t`

#### Summary

The `ompt_set_trace_native_t` type is the type signature of the `ompt_set_trace_native` runtime entry point, which enables or disables the recording of native trace records for a device.

#### Format

```
C / C++
typedef ompt_set_result_t (*ompt_set_trace_native_t) (
    ompt_device_t *device,
    int enable,
    int flags
);
C / C++
```

## Description

This interface is designed for use by a tool that cannot directly use native control functions for the device. If a tool can directly use the native control functions then it can invoke native control functions directly using pointers that the *lookup* function associated with the device provides and that are described in the *documentation* string that is provided to the device initializer callback.

## Description of Arguments

The *device* argument points to an opaque object that represents the target device instance. Functions in the device tracing interface use this pointer to identify the device that is being addressed.

The *enable* argument indicates whether this invocation should enable or disable recording of events.

The *flags* argument specifies the kinds of native device monitoring to enable or to 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`.

To start, to pause, to flush, or to stop tracing for a specific target device associated with *device*, a tool invokes the `ompt_start_trace`, `ompt_pause_trace`, `ompt_flush_trace`, or `ompt_stop_trace` runtime entry point for the device.

## Restrictions

The `ompt_set_trace_native` runtime entry point has the following restriction:

- The entry point must not return `ompt_set_sometimes_paired`.

## Cross References

- Tracing activity on target devices with OMPT, see Section 4.2.5 on page 432.
- `ompt_set_result_t` type, see Section 4.4.4.2 on page 443.
- `ompt_device_t` type, see Section 4.4.4.5 on page 445.

### 4.6.2.6 `ompt_start_trace_t`

#### Summary

The `ompt_start_trace_t` type is the type signature of the `ompt_start_trace` runtime entry point, which starts tracing of activity on a specific device.

#### Format

```
typedef int (*ompt_start_trace_t) (  
    ompt_device_t *device,  
    ompt_callback_buffer_request_t request,  
    ompt_callback_buffer_complete_t complete  
);
```

## Description

A device's `ompt_start_trace` runtime entry point, which has type signature `ompt_start_trace_t`, initiates tracing on the device. Under normal operating conditions, every event buffer provided to a device by a tool callback is 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 succeeds and 0 otherwise.

## Description of Arguments

The *device* argument points to an opaque object that represents the target device instance. Functions in the device tracing interface use this pointer to identify the device that is being addressed.

The *request* argument specifies a tool callback that supplies a device with a buffer to deposit events.

The *complete* argument specifies a tool callback that is invoked by the OpenMP implementation to empty a buffer that contains event records.

## Cross References

- `ompt_device_t` type, see Section 4.4.4.5 on page 445.
- `ompt_callback_buffer_request_t` callback type, see Section 4.5.2.23 on page 486.
- `ompt_callback_buffer_complete_t` callback type, see Section 4.5.2.24 on page 487.

### 4.6.2.7 `ompt_pause_trace_t`

#### Summary

The `ompt_pause_trace_t` type is the type signature of the `ompt_pause_trace` runtime entry point, which pauses or restarts activity tracing on a specific device.

#### Format

```
typedef int (*ompt_pause_trace_t) (  
    ompt_device_t *device,  
    int begin_pause  
);
```

C / C++

C / C++

#### Description

A device's `ompt_pause_trace` runtime entry point, which has type signature `ompt_pause_trace_t`, pauses or resumes tracing on a device. An invocation of `ompt_pause_trace` returns 1 if the command succeeds and 0 otherwise. Redundant pause or resume commands are idempotent and will return the same value as the prior command.

## Description of Arguments

The *device* argument points to an opaque object that represents the target device instance. Functions in the device tracing interface use this pointer to identify the device that is being addressed.

The *begin\_pause* argument indicates whether to pause or to resume tracing. To resume tracing, zero should be supplied for *begin\_pause*; To pause tracing, any other value should be supplied.

## Cross References

- `ompt_device_t` type, see Section [4.4.4.5](#) on page [445](#).

## 4.6.2.8 `ompt_flush_trace_t`

### Summary

The `ompt_flush_trace_t` type is the type signature of the `ompt_flush_trace` runtime entry point, which causes all pending trace records for the specified device to be delivered.

### Format

```
typedef int (*ompt_flush_trace_t) (  
    ompt_device_t *device  
);
```

C / C++

C / C++

### Description

A device's `ompt_flush_trace` runtime entry point, which has type signature `ompt_flush_trace_t`, causes the OpenMP implementation to issue a sequence of zero or more buffer completion callbacks to deliver all trace records that have been collected prior to the flush. An invocation of `ompt_flush_trace` returns 1 if the command succeeds and 0 otherwise.

## Description of Arguments

The *device* argument points to an opaque object that represents the target device instance. Functions in the device tracing interface use this pointer to identify the device that is being addressed.

## Cross References

- `ompt_device_t` type, see Section [4.4.4.5](#) on page [445](#).

## 4.6.2.9 `ompt_stop_trace_t`

### Summary

The `ompt_stop_trace_t` type is the type signature of the `ompt_stop_trace` runtime entry point, which stops tracing for a device.

## Format

C / C++

```
typedef int (*ompt_stop_trace_t) (  
    ompt_device_t *device  
);
```

C / C++

## Description

A device's **ompt\_stop\_trace** runtime entry point, which has type signature **ompt\_stop\_trace\_t**, halts tracing on the device and requests that any pending trace records are flushed. An invocation of **ompt\_stop\_trace** returns 1 if the command succeeds and 0 otherwise.

## Description of Arguments

The *device* argument points to an opaque object that represents the target device instance. Functions in the device tracing interface use this pointer to identify the device that is being addressed.

## Cross References

- **ompt\_device\_t** type, see Section 4.4.4.5 on page 445.

### 4.6.2.10 ompt\_advance\_buffer\_cursor\_t

#### Summary

The **ompt\_advance\_buffer\_cursor\_t** type is the type signature of the **ompt\_advance\_buffer\_cursor** runtime entry point, which advances a trace buffer cursor to the next record.

## Format

C / C++

```
typedef int (*ompt_advance_buffer_cursor_t) (  
    ompt_device_t *device,  
    ompt_buffer_t *buffer,  
    size_t size,  
    ompt_buffer_cursor_t current,  
    ompt_buffer_cursor_t *next  
);
```

C / C++

## Description

A device's **ompt\_advance\_buffer\_cursor** runtime entry point, which has type signature **ompt\_advance\_buffer\_cursor\_t**, advances a trace buffer pointer to the next trace record. An invocation of **ompt\_advance\_buffer\_cursor** returns *true* if the advance is successful and the next position in the buffer is valid.



## Description of Arguments

The *device* argument points to an opaque object that represents the target device instance. Functions in the device tracing interface use this pointer to identify the device that is being addressed.

The *buffer* argument indicates a trace buffer that is associated with the cursors.

The argument *size* indicates the size of *buffer* in bytes.

The *current* argument is an opaque buffer cursor.

The *next* argument returns the next value of an opaque buffer cursor.

## Cross References

- `ompt_device_t` type, see Section 4.4.4.5 on page 445.
- `ompt_buffer_cursor_t` type, see Section 4.4.4.8 on page 446.

### 4.6.2.11 `ompt_get_record_type_t`

#### Summary

The `ompt_get_record_type_t` type is the type signature of the `ompt_get_record_type` runtime entry point, which inspects the type of a trace record.

#### Format

```
C / C++
typedef ompt_record_t (*ompt_get_record_type_t) (
    ompt_buffer_t *buffer,
    ompt_buffer_cursor_t current
);
```

C / C++

#### Description

Trace records for a device may be in one of two forms: *native* record format, which may be device-specific, or *OMPT* record format, in which each trace record corresponds to an OpenMP *event* and most fields in the record structure are the arguments that would be passed to the OMPT callback for the event.

A device's `ompt_get_record_type` runtime entry point, which has 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 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 *buffer* argument indicates a trace buffer.

The *current* argument is an opaque buffer cursor.

## Cross References

- `ompt_record_t` type, see Section 4.4.3.1 on page 440.
- `ompt_buffer_t` type, see Section 4.4.4.7 on page 446.
- `ompt_buffer_cursor_t` type, see Section 4.4.4.8 on page 446.

### 4.6.2.12 `ompt_get_record_ompt_t`

#### Summary

The `ompt_get_record_ompt_t` type is the type signature of the `ompt_get_record_ompt` runtime entry point, which 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++
```

#### Description

A device's `ompt_get_record_ompt` runtime entry point, which has type signature `ompt_get_record_ompt_t`, returns a pointer that may point to a record in the trace buffer, or it may point to a record in thread local storage in which the information extracted from a record was assembled. The information available for an event depends upon its type.

The return value of the `ompt_record_ompt_t` type includes a field of a union type that can represent information for any OMPT event record type. Another call to the runtime entry point may overwrite the contents of the fields in a record returned by a prior invocation.

#### Description of Arguments

The *buffer* argument indicates a trace buffer.

The *current* argument is an opaque buffer cursor.

#### Cross References

- `ompt_record_ompt_t` type, see Section 4.4.3.4 on page 441.
- `ompt_device_t` type, see Section 4.4.4.5 on page 445.
- `ompt_buffer_cursor_t` type, see Section 4.4.4.8 on page 446.

### 4.6.2.13 `ompt_get_record_native_t`

#### Summary

The `ompt_get_record_native_t` type is the type signature of the `ompt_get_record_native` runtime entry point, which obtains a pointer to a native trace record from a trace buffer associated with a device.

#### Format

```
C / C++
typedef void *(*ompt_get_record_native_t) (
    ompt_buffer_t *buffer,
    ompt_buffer_cursor_t current,
    ompt_id_t *host_op_id
);
```

#### Description

A device's `ompt_get_record_native` runtime entry point, which has type signature `ompt_get_record_native_t`, returns a pointer that may point into the specified trace buffer, or into thread local storage in which 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 the object to which `host_op_id` points to a host-side identifier for the operation that is 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 *buffer* argument indicates a trace buffer.

The *current* argument is an opaque buffer cursor.

The *host\_op\_id* argument is a pointer to an identifier that is returned by the function. The entry point sets the identifier to which *host\_op\_id* points 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` type, see Section 4.4.4.3 on page 444.
- `ompt_buffer_t` type, see Section 4.4.4.7 on page 446.
- `ompt_buffer_cursor_t` type, see Section 4.4.4.8 on page 446.

#### 4.6.2.14 `ompt_get_record_abstract_t`

##### Summary

The `ompt_get_record_abstract_t` type is the type signature of the `ompt_get_record_abstract` runtime entry point, which 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++

##### Description

An OpenMP implementation may execute on a device that logs trace records in a native (device-specific) format that a tool cannot interpret directly. A device's `ompt_get_record_abstract` runtime entry point, which has type signature `ompt_get_record_abstract_t`, translates a native trace record into a standard form.

##### Description of Arguments

The *native\_record* argument is a pointer to a native trace record.

##### Cross References

- `ompt_record_abstract_t` type, see Section 4.4.3.3 on page 440.

### 4.6.3 Lookup Entry Points: `ompt_function_lookup_t`

##### Summary

The `ompt_function_lookup_t` type is the type signature of the lookup runtime entry points that provide pointers to runtime entry points that are part of the OMPT interface.

##### Format

C / C++

```
typedef void (*ompt_interface_fn_t) (void);  
  
typedef ompt_interface_fn_t (*ompt_function_lookup_t) (  
    const char *interface_function_name  
);
```

C / C++

## Description

An OpenMP implementation provides a pointer to a lookup routine that provides pointers to OMPT runtime entry points. When the implementation invokes a tool initializer to configure the OMPT callback interface, it provides a lookup function that provides pointers to runtime entry points that implement routines that are part of the OMPT callback interface. Alternatively, when it invokes a tool initializer to configure the OMPT tracing interface for a device, it provides a lookup function that provides pointers to runtime entry points that implement tracing control routines appropriate for that device.

## Description of Arguments

The *interface\_function\_name* argument is a C string that represents the name of a runtime entry point.

## Cross References

- Tool initializer for a device's OMPT tracing interface, see Section [4.2.5](#) on page [432](#).
- Tool initializer for the OMPT callback interface, see Section [4.5.1.1](#) on page [460](#).
- Entry points in the OMPT callback interface, see Table [4.1](#) on page [431](#) for a list and Section [4.6.1](#) on page [497](#) for detailed definitions.
- Entry points in the OMPT tracing interface, see Table [4.3](#) on page [435](#) for a list and Section [4.6.2](#) on page [515](#) for detailed definitions.

*This page intentionally left blank*

## 5 OMPD Interface

This chapter describes OMPD, which is an interface for *third-party* tools. *Third-party* tools exist in separate processes from the OpenMP program. To provide OMPD support, an OpenMP implementation must provide an OMPD library to be loaded by the *third-party* tool. An OpenMP implementation does not need to maintain any extra information to support OMPD inquiries from *third-party* tools *unless* it is explicitly instructed to do so.

OMP allows *third-party tools* such as a debuggers to inspect the OpenMP state of a live program or core file in an implementation-agnostic manner. That is, a tool that uses OMPD should work with any conforming OpenMP implementation. An OpenMP implementor provides a library for OMPD that a third-party tool can dynamically load. Using the interface exported by the OMPD library, the external tool can inspect the OpenMP state of a program. In order to satisfy requests from the third-party tool, the OMPD library may need to read data from, or to find the addresses of symbols in the OpenMP program. The OMPD library provides this functionality through a callback interface that the third-party tool must instantiate for the OMPD library.

To use OMPD, the third-party tool loads the OMPD library. The OMPD library exports the API that is defined throughout this section and that the tool uses to determine OpenMP information about the OpenMP program. The OMPD library must look up the symbols and read data out of the program. It does not perform these operations directly, but instead it uses the callback interface that the tool exports to cause the tool to perform them.

The OMPD architecture insulates tools from the internal structure of the OpenMP runtime while the OMPD library is insulated from the details of how to access the OpenMP program. This decoupled design allows for flexibility in how the OpenMP program and tool are deployed, so that, for example, the tool and the OpenMP program are not required to execute on the same machine.

Generally the tool does not interact directly with the OpenMP runtime and, instead, interacts with it through the OMPD library. However, a few cases require the tool to access the OpenMP runtime directly. These cases fall into two broad categories. The first is during initialization, where the tool must look up symbols and read variables in the OpenMP runtime in order to identify the OMPD library that it should use, which is discussed in Section 5.2.2 on page 530 and Section 5.2.3 on page 531. The second category relates to arranging for the tool to be notified when certain events occur during the execution of the OpenMP program. For this purpose, the OpenMP implementation must define certain symbols in the runtime code, as is discussed in Section 5.6 on page 580. Each of these symbols corresponds to an event type. The runtime must ensure that control passes through the appropriate named location when events occur. If the tool requires notification of an event, it can plant a breakpoint at the matching location. The location can, but may not, be a function. It can, for example, simply be a label. However, the names of the locations must have external C linkage.

## 5.1 OMPD Interfaces Definitions

C / C++

A compliant implementation must supply a set of definitions for the OMPD runtime entry points, OMPD tool callback signatures, OMPD tool interface routines, and the special data types of their parameters and return values. These definitions, which are listed throughout this chapter, and their associated declarations shall be provided in a header file named `omp-tools.h`. In addition, the set of definitions may specify other implementation-specific values.

The `ompd_dll_locations` variable, all OMPD tool interface functions, and all OMPD runtime entry points are external symbols with `C` linkage.

C / C++

## 5.2 Activating an OMPD Tool

The tool and the OpenMP program exist as separate processes. Thus, coordination is required between the OpenMP runtime and the external tool for OMPD.

### 5.2.1 Enabling the Runtime for OMPD

In order to support third-party tools, the OpenMP runtime may need to collect and to maintain information that it might not otherwise. The OpenMP runtime collects whatever information is necessary to support OMPD if the environment variable `OMP_DEBUG` is set to *enabled*.

#### Cross References

- Activating an OMPT Tool, Section 4.2 on page 425
- `OMP_DEBUG`, Section 6.20 on page 600

### 5.2.2 `ompd_dll_locations`

#### Summary

The `ompd_dll_locations` global variable indicates the location of OMPD libraries that are compatible with the OpenMP implementation.

#### Format

```
const char **ompd_dll_locations;
```

C

C



## Description

An OpenMP runtime may have more than one OMPD library. The tool must be able to locate the right library to use for the OpenMP program that it is examining. The OpenMP runtime system must provide a public variable `ompd_dll_locations`, which is an `argv`-style vector of filename string pointers that provides the name(s) of any compatible OMPD library. This variable must have `C` linkage. The tool uses the name of the variable verbatim and, in particular, does not apply any name mangling before performing the look up.

The programming model or architecture of the tool and, thus, that of OMPD does not have to match that of the OpenMP program that is being examined. The tool must interpret the contents of `ompd_dll_locations` to find a suitable OMPD that matches its own architectural characteristics. On platforms that support different programming models (for example, 32-bit vs 64-bit), OpenMP implementations are encouraged to provide OMPD libraries for all models, and that can handle OpenMP programs of any model. Thus, for example, a 32-bit debugger that uses OMPD should be able to debug a 64-bit OpenMP program by loading a 32-bit OMPD implementation that can manage a 64-bit OpenMP runtime.

`ompd_dll_locations` points to a NULL-terminated vector of zero or more NULL-terminated pathname strings that do not have any filename conventions. This vector must be fully initialized *before* `ompd_dll_locations` is set to a non-null value, such that if a tool, such as a debugger, stops execution of the OpenMP program at any point at which `ompd_dll_locations` is non-null, then the vector of strings to which it points is valid and complete.

## Cross References

- `ompd_dll_locations_valid`, see Section 5.2.3 on page 531

## 5.2.3 `ompd_dll_locations_valid`

### Summary

The OpenMP runtime notifies third-party tools that `ompd_dll_locations` is valid by allowing execution to pass through a location that the symbol `ompd_dll_locations_valid` identifies.

### Format

```
void ompd_dll_locations_valid(void);
```

### Description

Since `ompd_dll_locations` may not be a static variable, it may require runtime initialization. The OpenMP runtime notifies third-party tools that `ompd_dll_locations` is valid by having execution pass through a location that the symbol `ompd_dll_locations_valid` identifies. If `ompd_dll_locations` is NULL, a third-party tool can place a breakpoint at `ompd_dll_locations_valid` to be notified that `ompd_dll_locations` is initialized. In practice, the symbol `ompd_dll_locations_valid` may not be a function; instead, it may be a labeled machine instruction through which execution passes once the vector is valid.

## 5.3 OMPD Data Types

This section defines the OMPD types.

### 5.3.1 Size Type

#### Summary

The `ompd_size_t` type specifies the number of bytes in opaque data objects that are passed across the OMPD API.

#### Format

C / C++

```
typedef uint64_t ompd_size_t;
```

C / C++

### 5.3.2 Wait ID Type

#### Summary

This `ompd_wait_id_t` type identifies the object on which a thread.

#### Format

C / C++

```
typedef uint64_t ompd_wait_id_t;
```

C / C++

### 5.3.3 Basic Value Types

#### Summary

These definitions represent a word, address, and segment value types.

#### Format

C / C++

```
typedef uint64_t ompd_addr_t;  
typedef int64_t ompd_word_t;  
typedef uint64_t ompd_seg_t;
```

C / C++

#### Description

The `ompd_addr_t` type represents an unsigned integer address in an OpenMP process. The `ompd_word_t` type represents a signed version of `ompd_addr_t` to hold a signed integer of the OpenMP process. The `ompd_seg_t` type represents an unsigned integer segment value.

## 5.3.4 Address Type

### Summary

The `ompd_address_t` type is used to specify device addresses.

### Format

C / C++

```
typedef struct ompd_address_t {  
    ompd_seg_t segment;  
    ompd_addr_t address;  
} ompd_address_t;
```

C / C++

### Description

The `ompd_address_t` type is a structure that OMPD uses to specify device addresses, which may or may not be segmented. For non-segmented architectures, `ompd_segment_none` is used in the `segment` field of `ompd_address_t`; it is an instance of the `ompd_seg_t` type that has the value 0.

## 5.3.5 Frame Information Type

### Summary

The `ompd_frame_info_t` type is used to specify frame information.

### Format

C / C++

```
typedef struct ompd_frame_info_t {  
    ompd_address_t frame_address;  
    ompd_word_t frame_flag;  
} ompd_frame_info_t;
```

C / C++

### Description

The `ompd_frame_info_t` type is a structure that OMPD uses to specify frame information. The `frame_address` field of `ompd_frame_info_t` identifies a frame. The `frame_flag` field of `ompd_frame_info_t` indicates what type of information is provided in `frame_address`. The values and meaning is the same as defined for the `ompt_frame_t` enumeration type.

### Cross References

- `ompt_frame_t`, see Section 4.4.4.27 on page 457

## 5.3.6 System Device Identifiers

### Summary

The `ompd_device_t` type provides information about OpenMP devices.

### Format

```
▼ C / C++ ▼  
| typedef uint64_t ompd_device_t;  
▲ C / C++ ▲
```

### Description

Different OpenMP runtimes may utilize different underlying devices. The Device identifiers can vary in size and format and, thus, are not explicitly represented in OMPD. Instead, device identifiers are passed across the interface via the `ompd_device_t` type, which is a pointer to where the device identifier is stored, and the size of the device identifier in bytes. The OMPD library and a tool that uses it must agree on the format of the object that is passed. Each different kind of device identifier uses a unique unsigned 64-bit integer value.

Recommended values of `ompd_device_t` are defined in the `ompd-types.h` header file, which is available on <http://www.openmp.org/>.

## 5.3.7 Native Thread Identifiers

### Summary

The `ompd_thread_id_t` type provides information about native threads.

### Format

```
▼ C / C++ ▼  
| typedef uint64_t ompd_thread_id_t;  
▲ C / C++ ▲
```

### Description

Different OpenMP runtimes may use different native thread implementations. Native thread identifiers can vary in size and format and, thus, are not explicitly represented in the OMPD API. Instead, native thread identifiers are passed across the interface via the `ompd_thread_id_t` type, which is a pointer to where the native thread identifier is stored, and the size of the native thread identifier in bytes. The OMPD library and a tool that uses it must agree on the format of the object that is passed. Each different kind of native thread identifier uses a unique unsigned 64-bit integer value.

Recommended values of `ompd_thread_id_t` are defined in the `ompd-types.h` header file, which is available on <http://www.openmp.org/>.

## 5.3.8 OMPD Handle Types

### Summary

OMPd handle types are opaque types.

### Format

C / C++

```
typedef struct _ompd_aspace_handle ompd_address_space_handle_t;  
typedef struct _ompd_thread_handle ompd_thread_handle_t;  
typedef struct _ompd_parallel_handle ompd_parallel_handle_t;  
typedef struct _ompd_task_handle ompd_task_handle_t;
```

C / C++

### Description

OMPd uses handles for address spaces (`ompd_address_space_handle_t`), threads (`ompd_thread_handle_t`), parallel regions (`ompd_parallel_handle_t`), and tasks (`ompd_task_handle_t`). Each operation of the OMPd interface that applies to a particular address space, thread, parallel region, or task must explicitly specify a corresponding handle. A handle for an entity is constant while the entity itself is alive. Handles are defined by the OMPd library, and are opaque to the tool.

Defining externally visible type names in this way introduces type safety to the interface, and helps to catch instances where incorrect handles are passed by the tool to the OMPd library. The structures do not need to be defined; instead, the OMPd library must cast incoming (pointers to) handles to the appropriate internal, private types.

## 5.3.9 OMPD Scope Types

### Summary

The `ompd_scope_t` type identifies OMPd scopes.

### Format

C / C++

```
typedef enum ompd_scope_t {  
    ompd_scope_global = 1,  
    ompd_scope_address_space = 2,  
    ompd_scope_thread = 3,  
    ompd_scope_parallel = 4,  
    ompd_scope_implicit_task = 5,  
    ompd_scope_task = 6  
} ompd_scope_t;
```

C / C++

## Description

The `ompd_scope_t` type identifies OpenMP scopes, including those related to parallel regions and tasks. When used in an OMPD interface function call, the scope type and the ompd handle must match according to Table 5.1.

**TABLE 5.1:** Mapping of Scope Type and OMPD Handles

Scope types	Handles
<code>ompd_scope_global</code>	Address space handle for the host device
<code>ompd_scope_address_space</code>	Any address space handle
<code>ompd_scope_thread</code>	Any thread handle
<code>ompd_scope_parallel</code>	Any parallel handle
<code>ompd_scope_implicit_task</code>	Task handle for an implicit task
<code>ompd_scope_task</code>	Any task handle

## 5.3.10 ICV ID Type

### Summary

The `ompd_icv_id_t` type identifies an OpenMP implementation ICV.

### Format

```
typedef uint64_t ompd_icv_id_t;
```

The `ompd_icv_id_t` type identifies OpenMP implementation ICVs. `ompd_icv_undefined` is an instance of this type with the value 0.

## 5.3.11 Tool Context Types

### Summary

A third-party tool uses contexts to uniquely identify abstractions. These contexts are opaque to the OMPD library and are defined as follows:

### Format

```
typedef struct _ompd_aspace_cont ompd_address_space_context_t;  
typedef struct _ompd_thread_cont ompd_thread_context_t;
```

## 5.3.12 Return Code Types

### Summary

The `ompd_rc_t` type is the return code type of OMPD operations

### Format

C / C++

```
typedef enum ompd_rc_t {
    ompd_rc_ok = 0,
    ompd_rc_unavailable = 1,
    ompd_rc_stale_handle = 2,
    ompd_rc_bad_input = 3,
    ompd_rc_error = 4,
    ompd_rc_unsupported = 5,
    ompd_rc_needs_state_tracking = 6,
    ompd_rc_incompatible = 7,
    ompd_rc_device_read_error = 8,
    ompd_rc_device_write_error = 9,
    ompd_rc_nomem = 10,
} ompd_rc_t;
```

C / C++

### Description

The `ompd_rc_t` type is used for the return codes of OMPD operations. The return code types and their semantics are defined as follows:

- `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_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 on the device;
- `ompd_rc_incompatible` is returned when this OMPD library is incompatible with, or is not capable of handling, the OpenMP program; and
- `ompd_rc_nomem` is returned when a memory allocation fails.

## 5.3.13 Primitive Type Sizes

### Summary

The `ompd_device_type_sizes_t` type provides the “sizeof” of primitive types in the OpenMP architecture address space.

### Format

```
typedef struct ompd_device_type_sizes_t {  
    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;
```

### Description

The `ompd_device_type_sizes_t` type is used in operations through which the OMPD library can interrogate the tool about the “sizeof” of primitive types in the OpenMP architecture address space. 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 `uint8_t`.

### Cross References

- `ompd_callback_sizeof_fn_t`, see Section 5.4.2.2 on page 542

## 5.4 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 on which the tool is operating 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, to look up the addresses of symbols, and to read and to write memory in the device. The OMPD library uses these callbacks to implement its interface operations. The OMPD library only invokes the callback functions in direct response to calls made by the tool to the OMPD library.



## 5.4.1 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 5.4.1.1 on page 539) and `ompd_callback_memory_free_fn_t` (see Section 5.4.1.2 on page 540), which are provided by the tool to obtain and to release heap memory. This mechanism ensures 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 that the tool provides. The OMPD library must be coded so that any of its definitions of `new` or `delete` do not interfere with any that the tool defines.

In some cases, the OMPD library must allocate memory to return results to the tool. The tool then owns this memory and has the responsibility to release it. Thus, the OMPD library and the tool must use the same memory manager.

The OMPD library creates OMPD handles, which are opaque to the tool and may have a complex internal structure. The tool cannot determine if the handle pointers that the API returns correspond to discrete heap allocations. Thus, the tool must not simply deallocate a handle by passing an address that it receives from the OMPD library to its own memory manager. Instead, the API includes functions that the tool must use when it no longer needs a handle.

A tool creates contexts and passes them to the OMPD library. The OMPD library does not release contexts; instead the tool release them after it releases any handles that may reference the contexts.

### 5.4.1.1 `ompd_callback_memory_alloc_fn_t`

#### Summary

The `ompd_callback_memory_alloc_fn_t` type is the type signature of the callback routine that the tool provides to the OMPD library to allocate memory.

#### Format

```
typedef ompd_rc_t (*ompd_callback_memory_alloc_fn_t) (  
    ompd_size_t nbytes,  
    void **ptr  
);
```

#### Description

The `ompd_callback_memory_alloc_fn_t` type is the type signature of the memory allocation callback routine that the tool provides. The OMPD library may call the `ompd_callback_memory_alloc_fn_t` callback function to allocate memory.

## Description of Arguments

The *nbytes* argument is the size in bytes of the block of memory to allocate.

The address of the newly allocated block of memory is returned in the location to which the *ptr* argument points. The newly allocated block is suitably aligned for any type of variable, and is not guaranteed to be zeroed.

## Cross References

- `ompd_size_t`, see Section 5.3.1 on page 532.
- `ompd_rc_t`, see Section 5.3.12 on page 537.

### 5.4.1.2 `ompd_callback_memory_free_fn_t`

#### Summary

The `ompd_callback_memory_free_fn_t` type is the type signature of the callback routine that the tool provides to the OMPD library to deallocate memory.

#### Format

```
typedef ompd_rc_t (*ompd_callback_memory_free_fn_t) (  
    void *ptr  
);
```

#### Description

The `ompd_callback_memory_free_fn_t` type is the type signature of the memory deallocation callback routine that the tool provides. The OMPD library may call the `ompd_callback_memory_free_fn_t` callback function to deallocate memory that was obtained from a prior call to the `ompd_callback_memory_alloc_fn_t` callback function.

#### Description of Arguments

The *ptr* argument is the address of the block to be deallocated.

#### Cross References

- `ompd_rc_t`, see Section 5.3.12 on page 537.
- `ompd_callback_memory_alloc_fn_t`, see Section 5.4.1.1 on page 539.
- `ompd_callbacks_t`, see Section 5.4.6 on page 548.

## 5.4.2 Context Management and Navigation

#### Summary

The tool provides the OMPD library with callbacks to manage and to navigate context relationships.

## 5.4.2.1 `ompd_callback_get_thread_context_for_thread_id_fn_t`

### Summary

The `ompd_callback_get_thread_context_for_thread_id_fn_t` is the type signature of the callback routine that the tool provides to the OMPD library to map a thread identifier to a tool thread context.

### Format

```
typedef ompd_rc_t  
(*ompd_callback_get_thread_context_for_thread_id_fn_t) (  
    ompd_address_space_context_t *address_space_context,  
    ompd_thread_id_t kind,  
    ompd_size_t sizeof_thread_id,  
    const void *thread_id,  
    ompd_thread_context_t **thread_context  
);
```

### Description

The `ompd_callback_get_thread_context_for_thread_id_fn_t` is the type signature of the context mapping callback routine that the tool provides. This callback maps a thread identifier to a tool thread context. The thread identifier is within the address space that `address_space_context` identifies. The OMPD library can use the thread context, for example, to access thread local storage.

### Description of Arguments

The `address_space_context` argument is an opaque handle that the tool provides to reference an address space. The `kind`, `sizeof_thread_id`, and `thread_id` arguments represent a native thread identifier. On return, the `thread_context` argument provides an opaque handle that maps a native thread identifier to a tool thread context.

### Restrictions

Routines that use `ompd_callback_get_thread_context_for_thread_id_fn_t` have the following restriction:

- The provided `thread_context` must be valid until the OMPD library returns from the OMPD tool interface routine.

## Cross References

- `ompd_size_t`, see Section 5.3.1 on page 532.
- `ompd_thread_id_t`, see Section 5.3.7 on page 534.
- `ompd_address_space_context_t`, see Section 5.3.11 on page 536.
- `ompd_thread_context_t`, see Section 5.3.11 on page 536.
- `ompd_rc_t`, see Section 5.3.12 on page 537.

## 5.4.2.2 `ompd_callback_sizeof_fn_t`

### Summary

The `ompd_callback_sizeof_fn_t` type is the type signature of the callback routine that the tool provides to the OMPD library to determine the sizes of the primitive types in an address space.

### Format

```
typedef ompd_rc_t (*ompd_callback_sizeof_fn_t) (  
    ompd_address_space_context_t *address_space_context,  
    ompd_device_type_sizes_t *sizes  
);
```

### Description

The `ompd_callback_sizeof_fn_t` is the type signature of the type-size query callback routine that the tool provides. This callback provides 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 that the `address_space_context` argument specifies in the location to which the `sizes` argument points.

### Cross References

- `ompd_address_space_context_t`, see Section 5.3.11 on page 536.
- `ompd_rc_t`, see Section 5.3.12 on page 537.
- `ompd_device_type_sizes_t`, see Section 5.3.13 on page 538.
- `ompd_callbacks_t`, see Section 5.4.6 on page 548.

## 5.4.3 Accessing Memory in the OpenMP Program or Runtime

The OMPD library may need to read from or to write to the OpenMP program. It cannot do this directly. Instead the OMPD library must use callbacks that the tool provides so that the tool performs the operation.

### 5.4.3.1 `ompd_callback_symbol_addr_fn_t`

#### Summary

The `ompd_callback_symbol_addr_fn_t` type is the type signature of the callback that the tool provides to look up the addresses of symbols in an OpenMP program.

#### Format

```
typedef ompd_rc_t (ompd_callback_symbol_addr_fn_t) (  
    ompd_address_space_context_t *address_space_context,  
    ompd_thread_context_t *thread_context,  
    const char *symbol_name,  
    ompd_address_t *symbol_addr,  
    const char *file_name  
);
```

#### Description

The `ompd_callback_symbol_addr_fn_t` is the type signature of the symbol-address query callback routine that the tool provides. This callback looks up addresses of symbols within a specified address space.

#### Description of Arguments

This callback looks up the symbol provided in the *symbol\_name* argument.

The *address\_space\_context* argument is the tool's representation of the address space of the process, core file, or device.

The *thread\_context* argument is NULL for global memory access. If *thread\_context* is not NULL, *thread\_context* gives the thread specific context for the symbol lookup, for the purpose of calculating thread local storage addresses. If *thread\_context* is non-null then the thread to which *thread\_context* refers must be associated with either the process or the device that corresponds to the *address\_space\_context* argument.

The tool uses the *symbol\_name* argument that the OMPD library supplies verbatim. In particular, no name mangling, demangling or other transformations are performed prior to the lookup. The *symbol\_name* parameter must correspond to a statically allocated symbol within the specified

1 address space. The symbol can correspond to any type of object, such as a variable, thread local  
2 storage variable, function, or untyped label. The symbol can have a local, global, or weak binding.

3 The *file\_name* argument is an optional input parameter that indicates the name of the shared library  
4 in which the symbol is defined, and is intended to help the third party tool disambiguate symbols  
5 that are defined multiple times across the executable or shared library files. The shared library  
6 name may not be an exact match for the name seen by the tool. If *file\_name* is NULL then the tool  
7 first tries to find the symbol in the executable file, and, if the symbol is not found, the tool tries to  
8 find the symbol in the shared libraries in the order in which the shared libraries are loaded into the  
9 address space. If *file\_name* is non-null then the tool first tries to find the symbol in the libraries that  
10 match the name in the *file\_name* argument and, if the symbol is not found, the tool then uses the  
11 same procedure as when *file\_name* is NULL.

12 The callback does not support finding symbols that are dynamically allocated on the call stack, or  
13 statically allocated symbols that are defined within the scope of a function or subroutine.

14 The callback returns the symbol's address in the location to which *symbol\_addr* points.

## 15 Restrictions

16 Routines that use the `ompd_callback_symbol_addr_fn_t` type have the following  
17 restrictions:

- 18 • The *address\_space\_context* argument must be non-null.
- 19 • The symbol that the *symbol\_name* argument specifies must be defined.

## 20 Cross References

- 21 • `ompd_address_t`, see Section [5.3.4](#) on page [533](#).
- 22 • `ompd_address_space_context_t`, see Section [5.3.11](#) on page [536](#).
- 23 • `ompd_thread_context_t`, see Section [5.3.11](#) on page [536](#).
- 24 • `ompd_rc_t`, see Section [5.3.12](#) on page [537](#).
- 25 • `ompd_callbacks_t`, see Section [5.4.6](#) on page [548](#).

### 26 5.4.3.2 `ompd_callback_memory_read_fn_t`

#### 27 Summary

28 The `ompd_callback_memory_read_fn_t` type is the type signature of the callback that the  
29 tool provides to read data from an OpenMP program.

## Format

```
typedef ompd_rc_t (*ompd_callback_memory_read_fn_t) (  
    ompd_address_space_context_t *address_space_context,  
    ompd_thread_context_t *thread_context,  
    const ompd_address_t *addr,  
    ompd_size_t nbytes,  
    void *buffer  
);
```

## Description

The `ompd_callback_memory_read_fn_t` is the type signature of the read callback routines that the tool provides.

The `read_memory` callback copies a block of data from `addr` within the address space to the tool `buffer`.

The `read_string` callback copies a string to which `addr` points, including the terminating null byte (`'\0'`), to the tool `buffer`. At most `nbytes` bytes are copied. If a null byte is not among the first `nbytes` bytes, the string placed in `buffer` is not null-terminated.

## Description of Arguments

The address from which the data are to be read from the OpenMP program specified by `address_space_context` is given by `addr`. while `nbytes` gives the number of bytes to be transferred. The `thread_context` argument is optional for global memory access, and in this case should be NULL. If it is non-null, `thread_context` identifies the thread specific context for the memory access for the purpose of accessing thread local storage.

The data are returned through `buffer`, which is allocated and owned by the OMPD library. The contents of the buffer are unstructured, raw bytes. The OMPD library must arrange for any transformations such as byte-swapping that may be necessary (see Section 5.4.4 on page 547) to interpret the data.

## Cross References

- `ompd_size_t`, see Section 5.3.1 on page 532.
- `ompd_address_t`, see Section 5.3.4 on page 533.
- `ompd_address_space_context_t`, see Section 5.3.11 on page 536.
- `ompd_thread_context_t`, see Section 5.3.11 on page 536.
- `ompd_rc_t`, see Section 5.3.12 on page 537.
- `ompd_callback_device_host_fn_t`, see Section 5.4.4 on page 547.
- `ompd_callbacks_t`, see Section 5.4.6 on page 548.

### 5.4.3.3 `ompd_callback_memory_write_fn_t`

#### Summary

The `ompd_callback_memory_write_fn_t` type is the type signature of the callback that the tool provides to write data to an OpenMP program.

#### Format

```
typedef ompd_rc_t (*ompd_callback_memory_write_fn_t) (  
    ompd_address_space_context_t *address_space_context,  
    ompd_thread_context_t *thread_context,  
    const ompd_address_t *addr,  
    ompd_size_t nbytes,  
    const void *buffer  
);
```

#### Description

The `ompd_callback_memory_write_fn_t` is the type signature of the write callback routine that the tool provides. The OMPD library may call this callback to have the tool write a block of data to a location within an address space from a provided buffer.

#### Description of Arguments

The address to which the data are to be written in the OpenMP program that `address_space_context` specifies is given by `addr`. The `nbytes` argument is the number of bytes to be transferred. The `thread_context` argument is optional for global memory access, and, in this case, should be NULL. If it is non-null then `thread_context` identifies the thread-specific context for the memory access for the purpose of accessing thread local storage.

The data to be written are passed through `buffer`, which is allocated and owned by the OMPD library. The contents of the buffer are unstructured, raw bytes. The OMPD library must arrange for any transformations such as byte-swapping that may be necessary (see Section 5.4.4 on page 547) to render the data into a form that is compatible with the OpenMP runtime.

#### Cross References

- `ompd_size_t`, see Section 5.3.1 on page 532.
- `ompd_address_t`, see Section 5.3.4 on page 533.
- `ompd_address_space_context_t`, see Section 5.3.11 on page 536.
- `ompd_thread_context_t`, see Section 5.3.11 on page 536.
- `ompd_rc_t`, see Section 5.3.12 on page 537.
- `ompd_callback_device_host_fn_t`, see Section 5.4.4 on page 547.
- `ompd_callbacks_t`, see Section 5.4.6 on page 548.



## 5.4.4 Data Format Conversion: `ompd_callback_device_host_fn_t`

### Summary

The `ompd_callback_device_host_fn_t` type is the type signature of the callback that the tool provides to convert data between the formats that the tool and the OMPD library use and that the OpenMP program uses.

### Format

```
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 architecture and/or programming-model of the tool and the OMPD library may be different from that of the OpenMP program that is being examined. Thus, the conventions for representing data may differ. The callback interface includes operations to convert between the conventions, such as the byte order (endianness), that the tool and OMPD library use and the one that the OpenMP program uses. The callback with the `ompd_callback_device_host_fn_t` type signature convert data between formats

### Description of Arguments

The `address_space_context` argument specifies the OpenMP address space that is associated with the data. The `input` argument is the source buffer and the `output` argument is the destination buffer. The `unit_size` argument is the size of each of the elements to be converted. The `count` argument is the number of elements to be transformed.

The OMPD library allocates and owns the input and output buffers. It must ensure that the buffers have the correct size, and are eventually deallocated when they are no longer needed.

### Cross References

- `ompd_size_t`, see Section 5.3.1 on page 532.
- `ompd_address_space_context_t`, see Section 5.3.11 on page 536.
- `ompd_rc_t`, see Section 5.3.12 on page 537.
- `ompd_callbacks_t`, see Section 5.4.6 on page 548.

## 5.4.5 Output: `ompd_callback_print_string_fn_t`

### Summary

The `ompd_callback_print_string_fn_t` type is the type signature of the callback that tool provides so that the OMPD library can emit output.

### Format

```
typedef ompd_rc_t (*ompd_callback_print_string_fn_t) (  
    const char *string,  
    int category  
);
```

### Description

The OMPD library may call the `ompd_callback_print_string_fn_t` callback function to emit output, such as logging or debug information. The tool may set the `ompd_callback_print_string_fn_t` callback function to NULL to prevent the OMPD library from emitting output; the OMPD may not write to file descriptors that it did not open.

### Description of Arguments

The *string* argument is the null-terminated string to be printed. No conversion or formatting is performed on the string.

The *category* argument is the implementation-defined category of the string to be printed.

### Cross References

- `ompd_rc_t`, see Section [5.3.12](#) on page [537](#).
- `ompd_callbacks_t`, see Section [5.4.6](#) on page [548](#).

## 5.4.6 The Callback Interface

### Summary

All OMPD library interactions with the OpenMP program must be through a set of callbacks that the tool provides. These callbacks must also be used for allocating or releasing resources, such as memory, that the library needs.

## Format

```
typedef struct ompd_callbacks_t {
    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 that the OMPD library must 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 5.5.1.1 on page 550). Each field points to a function that the OMPD library must use to interact with the OpenMP program or for memory operations.

The `alloc_memory` and `free_memory` fields are pointers to functions the OMPD library uses to allocate and to release dynamic memory.

`print_string` points to a function that prints a string.

The architectures or programming models of the OMPD library and third party tool may be different from that of the OpenMP program that is being examined. `sizeof_type` points to function that allows the OMPD library to determine the sizes of the basic integer and pointer types that the OpenMP program uses. 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 that translates data from the conventions that the OpenMP program uses to those that the tool and OMPD library use. The reverse operation is performed by the function to which the `host_to_device` field points.

The `symbol_addr_lookup` field points to a callback that the OMPD library can use to find the address of a global or thread local storage symbol. The `read_memory`, `read_string`, and `write_memory` fields are pointers to functions for reading from and writing to global memory or thread local storage in the OpenMP program.

The `get_thread_context_for_thread_id` field is a pointer to a function that the OMPD library can use to obtain a thread context that corresponds to a native thread identifier.

## Cross References

- `ompd_callback_memory_alloc_fn_t`, see Section 5.4.1.1 on page 539.
- `ompd_callback_memory_free_fn_t`, see Section 5.4.1.2 on page 540.
- `ompd_callback_get_thread_context_for_thread_id_fn_t`, see Section 5.4.2.1 on page 541.
- `ompd_callback_sizeof_fn_t`, see Section 5.4.2.2 on page 542.
- `ompd_callback_symbol_addr_fn_t`, see Section 5.4.3.1 on page 543.
- `ompd_callback_memory_read_fn_t`, see Section 5.4.3.2 on page 544.
- `ompd_callback_memory_write_fn_t`, see Section 5.4.3.3 on page 546.
- `ompd_callback_device_host_fn_t`, see Section 5.4.4 on page 547.
- `ompd_callback_print_string_fn_t`, see Section 5.4.5 on page 548

## 5.5 OMPD Tool Interface Routines

### 5.5.1 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 that the library supports by calling `ompd_get_api_version` (see Section 5.5.1.2 on page 551). If the tool supports the version that `ompd_get_api_version` returns, the tool starts the initialization by calling `ompd_initialize` (see Section 5.5.1.1 on page 550) using the version of the OMPD API that the library supports. If the tool does not support the version that `ompd_get_api_version` returns, it may attempt to call `ompd_initialize` with a different version.

#### 5.5.1.1 `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  
);
```

C

C

## Description

A tool that uses OMPD calls `ompd_initialize` to initialize each OMPD library that it loads. More than one library may be present in a third-party tool, such as a debugger, because the tool may control multiple devices, which may use different runtime systems that require different OMPD libraries. This initialization must be performed exactly once before the tool can begin to operate on an OpenMP process or core file.

## Description of Arguments

The `api_version` argument is the OMPD API version that the tool requests to use. The tool may call `ompd_get_api_version` to obtain the latest OMPD API version that the OMPD library supports.

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 to 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, and to read and to write memory in the device.

## Cross References

- `ompd_rc_t` type, see Section 5.3.12 on page 537.
- `ompd_callbacks_t` type, see Section 5.4.6 on page 548.
- `ompd_get_api_version` call, see Section 5.5.1.2 on page 551.

### 5.5.1.2 `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 the `ompd_get_api_version` function to obtain the latest OMPD API version number of the OMPD library. The OMPD API version number is equal to the value of the `_OPENMP` macro defined in the associated OpenMP implementation, if the C preprocessor is supported. If the associated OpenMP implementation compiles Fortran codes without the use of a C preprocessor, the OMPD API version number is equal to the value of the Fortran integer parameter `openmp_version`.

#### Description of Arguments

The latest version number is returned into the location to which the `version` argument points.

## Cross References

- `ompd_rc_t` type, see Section 5.3.12 on page 537.

### 5.5.1.3 `ompd_get_version_string`

#### Summary

The `ompd_get_version_string` function returns a descriptive string for the OMPD library version.

#### Format

```
ompd_rc_t ompd_get_version_string(const char **string);
```

#### Description

The tool may call this function to obtain a pointer to a descriptive version string of the OMPD library vendor, implementation, internal version, date, or any other information that may be useful to a tool user or vendor. An implementation should provide a different string for every change to its source code or build that could be visible to the interface user.

#### Description of Arguments

A pointer to a descriptive version string is placed into the location to which *string* output argument points. The OMPD library owns the string that the OMPD library returns; the tool must not modify or release this string. The string remains valid for as long as the library is loaded. The `ompd_get_version_string` function may be called before `ompd_initialize` (see Section 5.5.1.1 on page 550). Accordingly, the OMPD library must not use heap or stack memory for the string.

The signatures of `ompd_get_api_version` (see Section 5.5.1.2 on page 551) and `ompd_get_version_string` are guaranteed not to change in future versions of the API. In contrast, the type definitions and prototypes in the rest of the API do not carry the same guarantee. Therefore a tool that uses OMPD should check the version of the API of the loaded OMPD library before it calls any other function of the API.

## Cross References

- `ompd_rc_t` type, see Section 5.3.12 on page 537.

### 5.5.1.4 `ompd_finalize`

#### Summary

When the tool is finished with the OMPD library it should call `ompd_finalize` before it unloads the library.

## Format

```
ompd_rc_t ompd_finalize(void);
```

## Description

The call to `ompd_finalize` must be the last OMPD call that the tool makes before it unloads the library. This call allows the OMPD library to free any resources that it may be holding.

The OMPD library may implement a *finalizer* section, which executes as the library is unloaded and therefore after the call to `ompd_finalize`. During finalization, the OMPD library may use the callbacks that the tool earlier provided after the call to `ompd_initialize`.

## Cross References

- `ompd_rc_t` type, see Section 5.3.12 on page 537.

## 5.5.2 Per OpenMP Process Initialization and Finalization

### 5.5.2.1 `ompd_process_initialize`

#### Summary

A tool calls `ompd_process_initialize` to obtain an address space handle when it initializes a session on a live process or core file.

## Format

```
ompd_rc_t ompd_process_initialize(  
    ompd_address_space_context_t *context,  
    ompd_address_space_handle_t **handle  
);
```

## Description

A tool calls `ompd_process_initialize` to obtain an address space handle when it initializes a session on a live process or core file. On return from `ompd_process_initialize`, the tool owns the address space handle, which it must release with

`ompd_rel_address_space_handle`. The initialization function must be called before any OMPD operations are performed on the OpenMP process. This call allows the OMPD library to confirm that it can handle the OpenMP process or core file that the *context* identifies.

Incompatibility is signaled by a return value of `ompd_rc_incompatible`.

## Description of Arguments

The *context* argument is an opaque handle that the tool provides to address an address space. On return, the *handle* argument provides an opaque handle to the tool for this address space, which the tool must release when it is no longer needed.

## Cross References

- `ompd_address_space_handle_t` type, see Section 5.3.8 on page 535.
- `ompd_address_space_context_t` type, see Section 5.3.11 on page 536.
- `ompd_rc_t` type, see Section 5.3.12 on page 537.
- `ompd_rel_address_space_handle` type, see Section 5.5.2.3 on page 555.

### 5.5.2.2 `ompd_device_initialize`

#### Summary

A tool calls `ompd_device_initialize` to obtain an address space handle for a device that has at least one active target region.

#### Format

```
ompd_rc_t ompd_device_initialize(  
    ompd_address_space_handle_t *process_handle,  
    ompd_address_space_context_t *device_context,  
    ompd_device_t kind,  
    ompd_size_t sizeof_id,  
    void *id,  
    ompd_address_space_handle_t **device_handle  
);
```

#### Description

A tool calls `ompd_device_initialize` to obtain an address space handle for a device that has at least one active target region. On return from `ompd_device_initialize`, the tool owns the address space handle.

#### Description of Arguments

The *process\_handle* argument is an opaque handle that the tool provides to reference the address space of the OpenMP process. The *device\_context* argument is an opaque handle that the tool provides to reference a device address space. The *kind*, *sizeof\_id*, and *id* arguments represent a device identifier. On return the *device\_handle* argument provides an opaque handle to the tool for this address space.



## Cross References

- `ompd_size_t` type, see Section 5.3.1 on page 532.
- `ompd_device_t` type, see Section 5.3.6 on page 534.
- `ompd_address_space_handle_t` type, see Section 5.3.8 on page 535.
- `ompd_address_space_context_t` type, see Section 5.3.11 on page 536.
- `ompd_rc_t` type, see Section 5.3.12 on page 537.

## 5.5.2.3 `ompd_rel_address_space_handle`

### Summary

A tool calls `ompd_rel_address_space_handle` to release an address space handle.

### Format

```
ompd_rc_t ompd_rel_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_rel_address_space_handle` to release the handle, which allows the OMPD library to release any resources that it has related to the address space.

### Description of Arguments

The *handle* argument is an opaque handle for the address space to be released.

### Restrictions

The `ompd_rel_address_space_handle` has the following restriction:

- An address space context must not be used after the corresponding address space handle is released.

## Cross References

- `ompd_address_space_handle_t` type, see Section 5.3.8 on page 535.
- `ompd_rc_t` type, see Section 5.3.12 on page 537.

## 5.5.3 Thread and Signal Safety

The OMPD library does not need to be reentrant. The tool must ensure that only one thread enters the OMPD library at a time. The OMPD library must not install signal handlers or otherwise interfere with the tool's signal configuration.

## 5.5.4 Address Space Information

### 5.5.4.1 `ompd_get_omp_version`

#### Summary

The tool may call the `ompd_get_omp_version` function to obtain the version of the OpenMP API that is associated with an address space.

#### Format

```
ompd_rc_t ompd_get_omp_version(  
    ompd_address_space_handle_t *address_space,  
    ompd_word_t *omp_version  
);
```

#### Description

The tool may call the `ompd_get_omp_version` function to obtain the version of the OpenMP API that is associated with the address space.

#### Description of Arguments

The `address_space` argument is an opaque handle that the tool provides to reference the address space of the OpenMP process or device.

Upon return, the `omp_version` argument contains the version of the OpenMP runtime in the `_OPENMP` version macro format.

#### Cross References

- `ompd_address_space_handle_t` type, see Section 5.3.8 on page 535.
- `ompd_rc_t` type, see Section 5.3.12 on page 537.

### 5.5.4.2 `ompd_get_omp_version_string`

#### Summary

The `ompd_get_omp_version_string` function returns a descriptive string for the OpenMP API version that is associated with an address space.

## Format

```
ompd_rc_t ompd_get_omp_version_string(  
    ompd_address_space_handle_t *address_space,  
    const char **string  
);
```

## Description

After initialization, the tool may call the `ompd_get_omp_version_string` function to obtain the version of the OpenMP API that is associated with an address space.

## Description of Arguments

The `address_space` argument is an opaque handle that the tool provides to reference the address space of the OpenMP process or device. A pointer to a descriptive version string is placed into the location to which the `string` output argument points. After returning from the call, the tool owns the string. The OMPD library must use the memory allocation callback that the tool provides to allocate the string storage. The tool is responsible for releasing the memory.

## Cross References

- `ompd_address_space_handle_t` type, see Section 5.3.8 on page 535.
- `ompd_rc_t` type, see Section 5.3.12 on page 537.

## 5.5.5 Thread Handles

### 5.5.5.1 `ompd_get_thread_in_parallel`

#### Summary

The `ompd_get_thread_in_parallel` function enables a tool to obtain handles for OpenMP threads that are 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 the location to which `thread_handle` points. This call yields meaningful results only if all OpenMP threads in the parallel region are stopped.

## Description of Arguments

The *parallel\_handle* argument is an opaque handle for a parallel region and selects the parallel region on which to operate. The *thread\_num* argument selects the thread of the team to be returned. On return, the *thread\_handle* argument is an opaque handle for the selected thread.

## Restrictions

The `ompd_get_thread_in_parallel` function has the following restriction:

- The value of *thread\_num* must be a non-negative integer smaller than the team size that was provided as the *ompd-team-size-var* from `ompd_get_icv_from_scope`.

## Cross References

- `ompd_parallel_handle_t` type, see Section 5.3.8 on page 535.
- `ompd_thread_handle_t` type, see Section 5.3.8 on page 535.
- `ompd_rc_t` type, see Section 5.3.12 on page 537.
- `ompd_get_icv_from_scope` call, see Section 5.5.9.2 on page 576.

## 5.5.5.2 ompd\_get\_thread\_handle

### Summary

The `ompd_get_thread_handle` function maps 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

The `ompd_get_thread_handle` function determines if the native thread identifier to which *thread\_id* points represents an OpenMP thread. If so, the function returns `ompd_rc_ok` and the location to which *thread\_handle* points is set to the thread handle for the OpenMP thread.

## Description of Arguments

The *handle* argument is an opaque handle that the tool provides to reference an address space. The *kind*, *sizeof\_thread\_id*, and *thread\_id* arguments represent a native thread identifier. On return, the *thread\_handle* argument provides an opaque handle to the thread within the provided address space.

The native thread identifier to which *thread\_id* points is guaranteed to be valid for the duration of the call. If the OMPD library must retain the native thread identifier, it must copy it.

## Cross References

- `ompd_size_t` type, see Section 5.3.1 on page 532.
- `ompd_thread_id_t` type, see Section 5.3.7 on page 534.
- `ompd_address_space_handle_t` type, see Section 5.3.8 on page 535.
- `ompd_thread_handle_t` type, see Section 5.3.8 on page 535.
- `ompd_rc_t` type, see Section 5.3.12 on page 537.

### 5.5.5.3 `ompd_rel_thread_handle`

#### Summary

The `ompd_rel_thread_handle` function releases a thread handle.

#### Format

```
ompd_rc_t ompd_rel_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 `ompd_rel_thread_handle` for disposal.

#### Description of Arguments

The *thread\_handle* argument is an opaque handle for a thread to be released.

#### Cross References

- `ompd_thread_handle_t` type, see Section 5.3.8 on page 535.
- `ompd_rc_t` type, see Section 5.3.12 on page 537.

#### 5.5.5.4 ompd\_thread\_handle\_compare

##### Summary

The `ompd_thread_handle_compare` function allows tools to compare two thread handles.

##### Format

```
ompd_rc_t ompd_thread_handle_compare(  
    ompd_thread_handle_t *thread_handle_1,  
    ompd_thread_handle_t *thread_handle_2,  
    int *cmp_value  
);
```

##### 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. The `ompd_thread_handle_compare` function compares thread handles.

On success, `ompd_thread_handle_compare` returns in the location to which `cmp_value` points 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 `thread_handle_1` and `thread_handle_2` arguments are opaque handles for threads. On return the `cmp_value` argument is set to a signed integer value.

##### Cross References

- `ompd_thread_handle_t` type, see Section 5.3.8 on page 535.
- `ompd_rc_t` type, see Section 5.3.12 on page 537.

#### 5.5.5.5 ompd\_get\_thread\_id

##### Summary

The `ompd_get_thread_id` maps 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

The `ompd_get_thread_id` function maps an OMPD thread handle to a native thread identifier.

## Description of Arguments

The `thread_handle` argument is an opaque thread handle. The `kind` argument represents the native thread identifier. The `sizeof_thread_id` argument represents the size of the native thread identifier. On return, the `thread_id` argument is a buffer that represents a native thread identifier.

## Cross References

- `ompd_size_t` type, see Section 5.3.1 on page 532.
- `ompd_thread_id_t` type, see Section 5.3.7 on page 534.
- `ompd_thread_handle_t` type, see Section 5.3.8 on page 535.
- `ompd_rc_t` type, see Section 5.3.12 on page 537.

## 5.5.6 Parallel Region Handles

### 5.5.6.1 `ompd_get_curr_parallel_handle`

#### Summary

The `ompd_get_curr_parallel_handle` function obtains a pointer to the parallel handle for an OpenMP thread's current parallel region.

## Format

```
ompd_rc_t ompd_get_curr_parallel_handle(  
    ompd_thread_handle_t *thread_handle,  
    ompd_parallel_handle_t **parallel_handle  
);
```

## Description

The `ompd_get_curr_parallel_handle` function enables the tool to obtain a pointer to the parallel handle for the current parallel region that is associated with an OpenMP thread. This call is meaningful only if the associated thread is stopped. The parallel handle must be released by calling `ompd_rel_parallel_handle`.

## Description of Arguments

The `thread_handle` argument is an opaque handle for a thread and selects the thread on which to operate. On return, the `parallel_handle` argument is set to a handle for the parallel region that the associated thread is currently executing, if any.

## Cross References

- `ompd_thread_handle_t` type, see Section 5.3.8 on page 535.
- `ompd_parallel_handle_t` type, see Section 5.3.8 on page 535.
- `ompd_rc_t` type, see Section 5.3.12 on page 537.
- `ompd_rel_parallel_handle` call, see Section 5.5.6.4 on page 564.

### 5.5.6.2 `ompd_get_enclosing_parallel_handle`

#### Summary

The `ompd_get_enclosing_parallel_handle` function obtains a pointer to the parallel handle for an enclosing parallel region.

#### Format

```
ompd_rc_t ompd_get_enclosing_parallel_handle(  
    ompd_parallel_handle_t *parallel_handle,  
    ompd_parallel_handle_t **enclosing_parallel_handle  
);
```

#### Description

The `ompd_get_enclosing_parallel_handle` function enables a tool to obtain a pointer to the parallel handle for the parallel region that encloses the parallel region that `parallel_handle` specifies. 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 the location to which `enclosing_parallel_handle` points. After the call, the tool owns the handle; the tool must release the handle with `ompd_rel_parallel_handle` when it is no longer required.



## Description of Arguments

The *parallel\_handle* argument is an opaque handle for a parallel region that selects the parallel region on which to operate. On return, the *enclosing\_parallel\_handle* argument is set to a handle for the parallel region that encloses the selected parallel region.

## Cross References

- `ompd_parallel_handle_t` type, see Section 5.3.8 on page 535.
- `ompd_rc_t` type, see Section 5.3.12 on page 537.
- `ompd_rel_parallel_handle` call, see Section 5.5.6.4 on page 564.

### 5.5.6.3 `ompd_get_task_parallel_handle`

#### Summary

The `ompd_get_task_parallel_handle` function obtains a pointer to the parallel handle for the parallel region that encloses a task region.

#### Format

```
ompd_rc_t ompd_get_task_parallel_handle(  
    ompd_task_handle_t *task_handle,  
    ompd_parallel_handle_t **task_parallel_handle  
);
```

#### Description

The `ompd_get_task_parallel_handle` function enables a tool to obtain a pointer to the parallel handle for the parallel region that encloses the task region that *task\_handle* specifies. 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 the location to which *task\_parallel\_handle* points. The tool owns that parallel handle, which it must release with `ompd_rel_parallel_handle`.

#### Description of Arguments

The *task\_handle* argument is an opaque handle that selects the task on which to operate. On return, the *parallel\_handle* argument is set to a handle for the parallel region that encloses the selected task.

## Cross References

- `ompd_task_handle_t` type, see Section 5.3.8 on page 535.
- `ompd_parallel_handle_t` type, see Section 5.3.8 on page 535.
- `ompd_rc_t` type, see Section 5.3.12 on page 537.
- `ompd_rel_parallel_handle` call, see Section 5.5.6.4 on page 564.

### 5.5.6.4 ompd\_rel\_parallel\_handle

#### Summary

The `ompd_rel_parallel_handle` function releases a parallel region handle.

#### Format

```
ompd_rc_t ompd_rel_parallel_handle(  
    ompd_parallel_handle_t *parallel_handle  
);
```

#### Description

Parallel region handles are opaque so tools cannot release them directly. Instead, a tool must pass a parallel region handle to the `ompd_rel_parallel_handle` function for disposal when finished with it.

#### Description of Arguments

The `parallel_handle` argument is an opaque handle to be released.

#### Cross References

- `ompd_parallel_handle_t` type, see Section 5.3.8 on page 535.
- `ompd_rc_t` type, see Section 5.3.12 on page 537.

### 5.5.6.5 ompd\_parallel\_handle\_compare

#### Summary

The `ompd_parallel_handle_compare` function compares two parallel 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 tools. While tools can easily compare pointers to parallel region handles, they cannot determine whether handles at two different addresses refer to the same underlying parallel region and, instead must use the `ompd_parallel_handle_compare` function.

On success, `ompd_parallel_handle_compare` returns a signed integer value in the location to which `cmp_value` points 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`. This function is provided since the means by which parallel region handles are ordered is implementation defined.

## Description of Arguments

The `parallel_handle_1` and `parallel_handle_2` arguments are opaque handles that correspond to parallel regions. On return the `cmp_value` argument points to a signed integer value that indicates how the underlying parallel regions compare.

## Cross References

- `ompd_parallel_handle_t` type, see Section 5.3.8 on page 535.
- `ompd_rc_t` type, see Section 5.3.12 on page 537.

## 5.5.7 Task Handles

### 5.5.7.1 `ompd_get_curr_task_handle`

#### Summary

The `ompd_get_curr_task_handle` function obtains a pointer to the task handle for the current task region that is associated with an OpenMP thread.

#### Format

```
ompd_rc_t ompd_get_curr_task_handle(  
    ompd_thread_handle_t *thread_handle,  
    ompd_task_handle_t **task_handle  
);
```

#### Description

The `ompd_get_curr_task_handle` function obtains a pointer to the task handle for the current task region that is associated with an OpenMP thread. This call is meaningful only if the thread for which the handle is provided is stopped. The task handle must be released with `ompd_rel_task_handle`.

## Description of Arguments

The *thread\_handle* argument is an opaque handle that selects the thread on which to operate. On return, the *task\_handle* argument points to a location that points to a handle for the task that the thread is currently executing.

## Cross References

- `ompd_thread_handle_t` type, see Section 5.3.8 on page 535.
- `ompd_task_handle_t` type, see Section 5.3.8 on page 535.
- `ompd_rc_t` type, see Section 5.3.12 on page 537.
- `ompd_rel_task_handle` call, see Section 5.5.7.5 on page 568.

### 5.5.7.2 `ompd_get_generating_task_handle`

#### Summary

The `ompd_get_generating_task_handle` function obtains a pointer to the task handle of the generating task region.

#### Format

```
ompd_rc_t ompd_get_generating_task_handle(  
    ompd_task_handle_t *task_handle,  
    ompd_task_handle_t **generating_task_handle  
);
```

#### Description

The `ompd_get_generating_task_handle` function obtains a pointer to the task handle for the task that encountered the OpenMP task construct that generated the task represented by *task\_handle*. The generating task is the OpenMP task that was active when the task specified by *task\_handle* was created. This call is meaningful only if the thread that is executing the task that *task\_handle* specifies is stopped. The generating task handle must be released with `ompd_rel_task_handle`.

#### Description of Arguments

The *task\_handle* argument is an opaque handle that selects the task on which to operate. On return, the *generating\_task\_handle* argument points to a location that points to a handle for the generating task.

## Cross References

- `ompd_task_handle_t` type, see Section 5.3.8 on page 535.
- `ompd_rc_t` type, see Section 5.3.12 on page 537.
- `ompd_rel_task_handle` call, see Section 5.5.7.5 on page 568.

### 5.5.7.3 `ompd_get_scheduling_task_handle`

#### Summary

The `ompd_get_scheduling_task_handle` function obtains a task handle for the task that was active at a task scheduling point.

#### Format

```
ompd_rc_t ompd_get_scheduling_task_handle(  
    ompd_task_handle_t *task_handle,  
    ompd_task_handle_t **scheduling_task_handle  
);
```

#### Description

The `ompd_get_scheduling_task_handle` function obtains a task handle for the task that was active when the task that `task_handle` represents was scheduled. This call is meaningful only if the thread that is executing the task that `task_handle` specifies is stopped. The scheduling task handle must be released with `ompd_rel_task_handle`.

#### Description of Arguments

The `task_handle` argument is an opaque handle for a task and selects the task on which to operate. On return, the `scheduling_task_handle` argument points to a location that points to a handle for the task that is still on the stack of execution on the same thread and was deferred in favor of executing the selected task.

#### Cross References

- `ompd_task_handle_t` type, see Section 5.3.8 on page 535.
- `ompd_rc_t` type, see Section 5.3.12 on page 537.
- `ompd_rel_task_handle` call, see Section 5.5.7.5 on page 568.

### 5.5.7.4 `ompd_get_task_in_parallel`

#### Summary

The `ompd_get_task_in_parallel` function obtains handles for the implicit tasks that are associated with a parallel region.

## Format

```
ompd_rc_t ompd_get_task_in_parallel(  
    ompd_parallel_handle_t *parallel_handle,  
    int thread_num,  
    ompd_task_handle_t **task_handle  
);
```

## Description

The `ompd_get_task_in_parallel` function obtains handles for the implicit tasks that are associated with a parallel region. A successful invocation of `ompd_get_task_in_parallel` returns a pointer to a task handle in the location to which `task_handle` points. This call yields meaningful results only if all OpenMP threads in the parallel region are stopped.

## Description of Arguments

The `parallel_handle` argument is an opaque handle that selects the parallel region on which to operate. The `thread_num` argument selects the implicit task of the team that is returned. The selected implicit task would return `thread_num` from a call of the `omp_get_thread_num()` routine. On return, the `task_handle` argument points to a location that points to an opaque handle for the selected implicit task.

## Restrictions

The following restriction applies to the `ompd_get_task_in_parallel` function:

- The value of `thread_num` must be a non-negative integer that is smaller than the size of the team size that is the value of the `ompd-team-size-var` that `ompd_get_icv_from_scope` returns.

## Cross References

- `ompd_parallel_handle_t` type, see Section 5.3.8 on page 535.
- `ompd_task_handle_t` type, see Section 5.3.8 on page 535.
- `ompd_rc_t` type, see Section 5.3.12 on page 537.
- `ompd_get_icv_from_scope` call, see Section 5.5.9.2 on page 576.

### 5.5.7.5 `ompd_rel_task_handle`

#### Summary

This `ompd_rel_task_handle` function releases a task handle.

## Format

```
ompd_rc_t ompd_rel_task_handle(  
    ompd_task_handle_t *task_handle  
);
```

## Description

Task handles are opaque so tools cannot release them directly. Instead, when a tool is finished with a task handle it must use the `ompd_rel_task_handle` function to release it.

## Description of Arguments

The `task_handle` argument is an opaque task handle to be released.

## Cross References

- `ompd_task_handle_t` type, see Section 5.3.8 on page 535.
- `ompd_rc_t` type, see Section 5.3.12 on page 537.

## 5.5.7.6 ompd\_task\_handle\_compare

### Summary

The `ompd_task_handle_compare` function compares 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 so tools cannot directly determine if handles at two different addresses refer to the same underlying task. The `ompd_task_handle_compare` function compares task handles. After a successful call to `ompd_task_handle_compare`, the value of the location to which `cmp_value` points is a signed integer that indicates how the underlying tasks compare: a value less than, equal to, or greater than 0 indicates that the task that corresponds to `task_handle_1` is, respectively, less than, equal to, or greater than the task that corresponds to `task_handle_2`. The means by which task handles are ordered is implementation defined.

## Description of Arguments

The *task\_handle\_1* and *task\_handle\_2* arguments are opaque handles that correspond to tasks. On return, the *cmp\_value* argument points to a location in which a signed integer value indicates how the underlying tasks compare.

## Cross References

- `ompd_task_handle_t` type, see Section 5.3.8 on page 535.
- `ompd_rc_t` type, see Section 5.3.12 on page 537.

## 5.5.7.7 `ompd_get_task_function`

### Summary

This `ompd_get_task_function` function returns the entry point of the code that corresponds to the body of a task.

### Format

```
ompd_rc_t ompd_get_task_function (  
    ompd_task_handle_t *task_handle,  
    ompd_address_t *entry_point  
);
```

### Description

The `ompd_get_task_function` function returns the entry point of the code that corresponds to the body of code that the task executes.

### Description of Arguments

The *task\_handle* argument is an opaque handle that selects the task on which to operate. On return, the *entry\_point* argument is set to an address that describes the beginning of application code that executes the task region.

## Cross References

- `ompd_address_t` type, see Section 5.3.4 on page 533.
- `ompd_task_handle_t` type, see Section 5.3.8 on page 535.
- `ompd_rc_t` type, see Section 5.3.12 on page 537.

## 5.5.7.8 `ompd_get_task_frame`

### Summary

The `ompd_get_task_frame` function extracts the frame pointers of a task.



## Format

```
ompd_rc_t ompd_get_task_frame (  
    ompd_task_handle_t *task_handle,  
    ompd_frame_info_t *exit_frame,  
    ompd_frame_info_t *enter_frame  
);
```

## Description

An OpenMP implementation maintains an **ompd\_frame\_t** object for every implicit or explicit task. The **ompd\_get\_task\_frame** function extracts the *enter\_frame* and *exit\_frame* fields of the **ompd\_frame\_t** object of the task that *task\_handle* identifies.

## Description of Arguments

The *task\_handle* argument specifies an OpenMP task. On return, the *exit\_frame* argument points to an **ompd\_frame\_info\_t** object that has the frame information with the same semantics as the *exit\_frame* field in the **ompd\_frame\_t** object that is associated with the specified task. On return, the *enter\_frame* argument points to an **ompd\_frame\_info\_t** object that has the frame information with the same semantics as the *enter\_frame* field in the **ompd\_frame\_t** object that is associated with the specified task.

## Cross References

- **ompd\_frame\_t** type, see Section 4.4.4.27 on page 457.
- **ompd\_address\_t** type, see Section 5.3.4 on page 533.
- **ompd\_frame\_info\_t** type, see Section 5.3.5 on page 533.
- **ompd\_task\_handle\_t** type, see Section 5.3.8 on page 535.
- **ompd\_rc\_t** type, see Section 5.3.12 on page 537.

## 5.5.7.9 ompd\_enumerate\_states

### Summary

The **ompd\_enumerate\_states** function enumerates thread states that an OpenMP implementation supports.

## 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 that the `ompt_state_t` enumeration type defines. In addition, an OpenMP implementation may support implementation-specific states. The `ompd_enumerate_states` call enables a tool to enumerate the thread states that an OpenMP implementation supports.

When the `current_state` argument is a thread state that an OpenMP implementation supports, the call assigns the value and string name of the next thread state in the enumeration to the locations to which the `next_state` and `next_state_name` arguments point.

On return, the third-party tool owns the `next_state_name` string. The OMPD library allocates storage for the string with the memory allocation callback that the tool provides. The tool is responsible for releasing the memory.

On return, the location to which the `more_enums` argument points has the value 1 whenever one or more states are left in the enumeration. On return, the location to which the `more_enums` argument points has the value 0 when `current_state` is the last state in the enumeration.

## Description of Arguments

The `address_space_handle` argument identifies the address space. The `current_state` argument must be a thread state that the OpenMP implementation supports. To begin enumerating the supported states, a tool should pass `ompt_state_undefined` as the value of `current_state`. Subsequent calls to `ompd_enumerate_states` by the tool should pass the value that the call returned in the `next_state` argument. On return, the `next_state` argument points to an integer with the value of the next state in the enumeration. On return, the `next_state_name` argument points to a character string that describes the next state. On return, the `more_enums` argument points to an integer with a value of 1 when more states are left to enumerate and a value of 0 when no more states are left.

## Cross References

- `ompt_state_t` type, see Section 4.4.4.26 on page 454.
- `ompd_address_space_handle_t` type, see Section 5.3.8 on page 535.
- `ompd_rc_t` type, see Section 5.3.12 on page 537.

## 5.5.7.10 ompd\_get\_state

### Summary

The `ompd_get_state` function obtains the state of a thread.

### Format

```
ompd_rc_t ompd_get_state (  
    ompd_thread_handle_t *thread_handle,  
    ompd_word_t *state,  
    ompd_wait_id_t *wait_id  
);
```

### Description

The `ompd_get_state` function returns the state of an OpenMP thread.

### Description of Arguments

The `thread_handle` argument identifies the thread. The `state` argument represents the state of that thread as represented by a value that `ompd_enumerate_states` returns. On return, if the `wait_id` argument is non-null then it points to a handle that corresponds to the `wait_id` wait identifier of the thread. If the thread state is not one of the specified wait states, the value to which `wait_id` points is undefined.

### Cross References

- `ompd_wait_id_t` type, see Section 5.3.2 on page 532.
- `ompd_thread_handle_t` type, see Section 5.3.8 on page 535.
- `ompd_rc_t` type, see Section 5.3.12 on page 537.
- `ompd_enumerate_states` call, see Section 5.5.7.9 on page 571.

## 5.5.8 Display Control Variables

### 5.5.8.1 ompd\_get\_display\_control\_vars

#### Summary

The `ompd_get_display_control_vars` function returns a list of name/value pairs for OpenMP control variables.

## Format

```
ompd_rc_t ompd_get_display_control_vars (  
    ompd_address_space_handle_t *address_space_handle,  
    const char * const **control_vars  
);
```

## Description

The `ompd_get_display_control_vars` function returns a NULL-terminated vector of NULL-terminated strings of name/value pairs of control variables that have user controllable settings and are important to the operation or performance of an OpenMP runtime system. The control variables that this interface exposes include all 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**.

On return, the third-party tool owns the vector and the strings. The OMP library must satisfy the termination constraints; it may use static or dynamic memory for the vector and/or the strings and is unconstrained in how it arranges them in memory. If it uses dynamic memory then the OMPD library must use the `allocate` callback that the tool provides to `ompd_initialize`. The tool must use `ompd_rel_display_control_vars()` to release the vector and the strings.

## Description of Arguments

The `address_space_handle` argument identifies the address space. On return, the `control_vars` argument points to the vector of display control variables.

## Cross References

- `ompd_address_space_handle_t` type, see Section 5.3.8 on page 535.
- `ompd_rc_t` type, see Section 5.3.12 on page 537.
- `ompd_initialize` call, see Section 5.5.1.1 on page 550.
- `ompd_rel_display_control_vars` type, see Section 5.5.8.2 on page 574.

## 5.5.8.2 ompd\_rel\_display\_control\_vars

### Summary

The `ompd_rel_display_control_vars` releases a list of name/value pairs of OpenMP control variables previously acquired with `ompd_get_display_control_vars`.

## Format

```
ompd_rc_t ompd_rel_display_control_vars (  
    const char * const **control_vars  
);
```

## Description

The third-party tool owns the vector and strings that `ompd_get_display_control_vars` returns. The tool must call `ompd_rel_display_control_vars` to release the vector and the strings.

## Description of Arguments

The `control_vars` argument is the vector of display control variables to be released.

## Cross References

- `ompd_rc_t` type, see Section 5.3.12 on page 537.
- `ompd_get_display_control_vars` call, see Section 5.5.8.1 on page 573.

## 5.5.9 Accessing Scope-Specific Information

### 5.5.9.1 `ompd_enumerate_icvs`

#### Summary

The `ompd_enumerate_icvs` function enumerates ICVs.

## Format

```
ompd_rc_t ompd_enumerate_icvs (  
    ompd_address_space_handle_t *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

An OpenMP implementation must support all ICVs listed in Section 2.6.1. An OpenMP implementation may support additional implementation specific variables. An implementation may store ICVs in a different scope than Table 2.3 indicates. The `ompd_enumerate_icvs` function enables a tool to enumerate the ICVs that an OpenMP implementation supports and their related scopes. The ICVs `num-procs-var`, `thread-num-var`, `final-task-var`, `implicit-task-var`, and `team-size-var` will also be available with an `ompd-` prefix.

When the `current` argument is set to the identifier of a supported ICV, `ompd_enumerate_icvs` assigns the value, string name, and scope of the next ICV in the enumeration to the locations to which the `next_id`, `next_icv_name`, and `next_scope` arguments point. On return, the third-party tool owns the `next_icv_name` string. The OMPD library uses the memory allocation callback that the tool provides to allocate the string storage; the tool is responsible for releasing the memory.

On return, the location to which the `more` argument points has the value of 1 whenever one or more ICV are left in the enumeration. On return, that location has the value 0 when `current` is the last ICV in the enumeration.

## Description of Arguments

The `address_space_handle` argument identifies the address space. The `current` argument must be an ICV that the OpenMP implementation supports. To begin enumerating the ICVs, a tool should pass `ompd_icv_undefined` as the value of `current`. Subsequent calls to `ompd_enumerate_icvs` should pass the value returned by the call in the `next_id` output argument. On return, the `next_id` argument points to an integer with the value of the ID of the next ICV in the enumeration. On return, the `next_icv_name` argument points to a character string with the name of the next ICV. On return, the `next_scope` argument points to the scope enum value of the scope of the next ICV. On return, the `more_enums` argument points to an integer with the value of 1 when more ICVs are left to enumerate and the value of 0 when no more ICVs are left.

## Cross References

- `ompd_address_space_handle_t` type, see Section 5.3.8 on page 535.
- `ompd_scope_t` type, see Section 5.3.9 on page 535.
- `ompd_icv_id_t` type, see Section 5.3.10 on page 536.
- `ompd_rc_t` type, see Section 5.3.12 on page 537.

### 5.5.9.2 `ompd_get_icv_from_scope`

#### Summary

The `ompd_get_icv_from_scope` function returns the value of an ICV.

## 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 `ompd_get_icv_from_scope` function provides access to the ICVs that `ompd_enumerate_icvs` identifies.

## Description of Arguments

The *handle* argument provides an OpenMP scope handle. The *scope* argument specifies the kind of scope provided in *handle*. The *icv\_id* argument specifies the ID of the requested ICV. On return, the *icv\_value* argument points to a location with the value of the requested ICV.

## Constraints on Arguments

If the ICV cannot be represented by an integer type value then the function returns `ompd_rc_incompatible`.

The provided *handle* must match the *scope* as defined in Section 5.3.10 on page 536.

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 5.3.8 on page 535.
- `ompd_thread_handle_t` type, see Section 5.3.8 on page 535.
- `ompd_parallel_handle_t` type, see Section 5.3.8 on page 535.
- `ompd_task_handle_t` type, see Section 5.3.8 on page 535.
- `ompd_scope_t` type, see Section 5.3.9 on page 535.
- `ompd_icv_id_t` type, see Section 5.3.10 on page 536.
- `ompd_rc_t` type, see Section 5.3.12 on page 537.
- `ompd_enumerate_icvs`, see Section 5.5.9.1 on page 575.

### 5.5.9.3 `ompd_get_icv_string_from_scope`

#### Summary

The `ompd_get_icv_string_from_scope` function returns the value of an ICV.

## 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 `ompd_get_icv_string_from_scope` function provides access to the ICVs that `ompd_enumerate_icvs` identifies.

## Description of Arguments

The *handle* argument provides an OpenMP scope handle. The *scope* argument specifies the kind of scope provided in *handle*. The *icv\_id* argument specifies the ID of the requested ICV. On return, the *icv\_string* argument points to a string representation of the requested ICV.

On return, the third-party tool owns the *icv\_string* string. The OMPD library allocates the string storage with the memory allocation callback that the tool provides. The tool is responsible for releasing the memory.

## Constraints on Arguments

The provided *handle* must match the *scope* as defined in Section 5.3.10 on page 536.

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 5.3.8 on page 535.
- `ompd_thread_handle_t` type, see Section 5.3.8 on page 535.
- `ompd_parallel_handle_t` type, see Section 5.3.8 on page 535.
- `ompd_task_handle_t` type, see Section 5.3.8 on page 535.
- `ompd_scope_t` type, see Section 5.3.9 on page 535.
- `ompd_icv_id_t` type, see Section 5.3.10 on page 536.
- `ompd_rc_t` type, see Section 5.3.12 on page 537.
- `ompd_enumerate_icvs`, see Section 5.5.9.1 on page 575.



## 5.5.9.4 ompd\_get\_tool\_data

### Summary

The `ompd_get_tool_data` function provides access to the OMPT data variable stored for each OpenMP scope.

### Format

```
ompd_rc_t ompd_get_tool_data(  
    void* handle,  
    ompd_scope_t scope,  
    ompd_word_t *value,  
    ompd_address_t *ptr  
);
```

### Description

The `ompd_get_tool_data` function provides access to the OMPT tool data stored for each scope. If the runtime library does not support OMPT then the function returns `ompd_rc_unsupported`.

### Description of Arguments

The `handle` argument provides an OpenMP scope handle. The `scope` argument specifies the kind of scope provided in `handle`. On return, the `value` argument points to the `value` field of the `ompt_data_t` union stored for the selected scope. On return, the `ptr` argument points to the `ptr` field of the `ompt_data_t` union stored for the selected scope.

### Cross References

- `ompt_data_t` type, see Section 4.4.4.4 on page 444.
- `ompd_address_space_handle_t` type, see Section 5.3.8 on page 535.
- `ompd_thread_handle_t` type, see Section 5.3.8 on page 535.
- `ompd_parallel_handle_t` type, see Section 5.3.8 on page 535.
- `ompd_task_handle_t` type, see Section 5.3.8 on page 535.
- `ompd_scope_t` type, see Section 5.3.9 on page 535.
- `ompd_rc_t` type, see Section 5.3.12 on page 537.

## 5.6 Runtime Entry Points for OMPD

The OpenMP implementation must define several entry point symbols through which execution must pass when particular events occur *and* data collection for OMPD is enabled. A tool can enable notification of an event by setting a breakpoint at the address of the entry point symbol.

Entry point symbols have external C linkage and do not require demangling or other transformations to look up their names to obtain the address in the OpenMP program. While each entry point symbol conceptually has a function type signature, it may not be a function. It may be a labeled location

### 5.6.1 Beginning Parallel Regions

#### Summary

Before starting the execution of an OpenMP parallel region, the implementation executes `ompd_bp_parallel_begin`.

#### Format

```
void ompd_bp_parallel_begin(void);
```

#### Description

The OpenMP implementation must execute `ompd_bp_parallel_begin` at every *parallel-begin* event. At the point that the implementation reaches `ompd_bp_parallel_begin`, the binding for `ompd_get_curr_parallel_handle` is the parallel region that is beginning and the binding for `ompd_get_curr_task_handle` is the task that encountered the `parallel` construct.

#### Cross References

- `parallel` construct, see Section 2.8 on page 79.
- `ompd_get_curr_parallel_handle`, see Section 5.5.6.1 on page 561.
- `ompd_get_curr_task_handle`, see Section 5.5.7.1 on page 565.

### 5.6.2 Ending Parallel Regions

#### Summary

After finishing the execution of an OpenMP parallel region, the implementation executes `ompd_bp_parallel_end`.

## Format

```
void ompd_bp_parallel_end(void);
```

## Description

The OpenMP implementation must execute `ompd_bp_parallel_end` at every *parallel-end* event. At the point that the implementation reaches `ompd_bp_parallel_end`, the binding for `ompd_get_curr_parallel_handle` is the `parallel` region that is ending and the binding for `ompd_get_curr_task_handle` is the task that encountered the `parallel` construct. After execution of `ompd_bp_parallel_end`, any *parallel\_handle* that was acquired for the `parallel` region is invalid and should be released.

## Cross References

- `parallel` construct, see Section 2.8 on page 79.
- `ompd_get_curr_parallel_handle`, see Section 5.5.6.1 on page 561.
- `ompd_rel_parallel_handle`, see Section 5.5.6.4 on page 564.
- `ompd_get_curr_task_handle`, see Section 5.5.7.1 on page 565.

## 5.6.3 Beginning Task Regions

### Summary

Before starting the execution of an OpenMP task region, the implementation executes `ompd_bp_task_begin`.

## Format

```
void ompd_bp_task_begin(void);
```

## Description

The OpenMP implementation must execute `ompd_bp_task_begin` immediately before starting execution of a *structured-block* that is associated with a non-merged task. At the point that the implementation reaches `ompd_bp_task_begin`, the binding for `ompd_get_curr_task_handle` is the task that is scheduled to execute.

## Cross References

- `ompd_get_curr_task_handle`, see Section 5.5.7.1 on page 565.

## 5.6.4 Ending Task Regions

### Summary

After finishing the execution of an OpenMP task region, the implementation executes `ompd_bp_task_end`.

### Format

```
void ompd_bp_task_end(void);
```

### Description

The OpenMP implementation must execute `ompd_bp_task_end` immediately after completion of a *structured-block* that is associated with a non-merged task. At the point that the implementation reaches `ompd_bp_task_end`, the binding for `ompd_get_curr_task_handle` is the task that finished execution. After execution of `ompd_bp_task_end`, any *task\_handle* that was acquired for the task region is invalid and should be released.

### Cross References

- `ompd_get_curr_task_handle`, see Section 5.5.7.1 on page 565.
- `ompd_rel_task_handle`, see Section 5.5.7.5 on page 568.

## 5.6.5 Beginning OpenMP Threads

### Summary

When starting an OpenMP thread, the implementation executes `ompd_bp_thread_begin`.

### Format

```
void ompd_bp_thread_begin(void);
```

### Description

The OpenMP implementation must execute `ompd_bp_thread_begin` at every *native-thread-begin* and *initial-thread-begin* event. This execution occurs before the thread starts the execution of any OpenMP region.

### Cross References

- `parallel` construct, see Section 2.8 on page 79.
- Initial task, see Section 2.13.5 on page 155.

## 5.6.6 Ending OpenMP Threads

### Summary

When terminating an OpenMP thread, the implementation executes `ompd_bp_thread_end`.

### Format

```
void ompd_bp_thread_end(void);
```

### Description

The OpenMP implementation must execute `ompd_bp_thread_end` at every *native-thread-end* and the *initial-thread-end* event. This execution occurs after the thread completes the execution of all OpenMP regions. After executing `ompd_bp_thread_end`, any *thread\_handle* that was acquired for this thread is invalid and should be released.

### Cross References

- `parallel` construct, see Section 2.8 on page 79.
- Initial task, see Section 2.13.5 on page 155.
- `ompd_rel_thread_handle`, see Section 5.5.5.3 on page 559.

## 5.6.7 Initializing OpenMP Devices

### Summary

The OpenMP implementation must execute `ompd_bp_device_begin` at every *device-initialize* event.

### Format

```
void ompd_bp_device_begin(void);
```

### Description

When initializing a device for execution of a **target** region, the implementation must execute `ompd_bp_device_begin`. This execution occurs before the work associated with any OpenMP region executes on the device.

### Cross References



- Device Initialization, see Section 2.15.1 on page 165.

## 5.6.8 Finalizing OpenMP Devices

### Summary

When terminating an OpenMP thread, the implementation executes `ompd_bp_device_end`.

### Format

 `void ompd_bp_device_end(void);`  


### Description

The OpenMP implementation must execute `ompd_bp_device_end` at every *device-finalize* event. This execution occurs after the thread executes all OpenMP regions. After execution of `ompd_bp_device_end`, any *address\_space\_handle* that was acquired for this device is invalid and should be released.

### Cross References

- Device Initialization, see Section [2.15.1](#) on page [165](#).
- `ompd_rel_address_space_handle`, see Section [5.5.2.3](#) on page [555](#).

# 6 Environment Variables

This chapter describes the OpenMP environment variables that specify the settings of the ICVs that affect the execution of OpenMP programs (see Section 2.6 on page 66). The names of the environment variables must be upper case. The values assigned to the environment variables are case insensitive and may have leading and trailing white space. Modifications to the environment variables after the program has started, even if modified by the program itself, are ignored by the OpenMP implementation. However, the settings of some of the ICVs can be modified during the execution of the OpenMP program by the use of the appropriate directive clauses or OpenMP API routines.

The following examples demonstrate how the OpenMP environment variables can be set in different environments:

- csh-like shells:

```
setenv OMP_SCHEDULE "dynamic"
```

- bash-like shells:

```
export OMP_SCHEDULE="dynamic"
```

- Windows Command Line:

```
set OMP_SCHEDULE=dynamic
```

## 6.1 OMP\_SCHEDULE

The **OMP\_SCHEDULE** environment variable controls the schedule kind and chunk size of all loop directives that have the schedule kind **runtime**, by setting the value of the *run-sched-var* ICV.

The value of this environment variable takes the form:

```
[modifier:]kind[, chunk]
```

where

- *modifier* is one of **monotonic** or **nonmonotonic**;
- *kind* is one of **static**, **dynamic**, **guided**, or **auto**;
- *chunk* is an optional positive integer that specifies the chunk size.

1 If the *modifier* is not present, the *modifier* is set to **monotonic** if *kind* is **static**; for any other  
2 *kind* it is set to **nonmonotonic**.

3 If *chunk* is present, white space may be on either side of the “,”. See Section 2.12.2 on page 106  
4 for a detailed description of the schedule kinds.

5 The behavior of the program is implementation defined if the value of **OMP\_SCHEDULE** does not  
6 conform to the above format.

7 Implementation specific schedules cannot be specified in **OMP\_SCHEDULE**. They can only be  
8 specified by calling **omp\_set\_schedule**, described in Section 3.2.12 on page 347.

9 Examples:

```
10 setenv OMP_SCHEDULE "guided, 4"  
11 setenv OMP_SCHEDULE "dynamic"  
12 setenv OMP_SCHEDULE "nonmonotonic:dynamic, 4"
```

## 13 Cross References

- 14 • *run-sched-var* ICV, see Section 2.6 on page 66.
- 15 • Worksharing-Loop construct, see Section 2.12.2 on page 106.
- 16 • Parallel worksharing-loop construct, see Section 2.16.1 on page 193.
- 17 • **omp\_set\_schedule** routine, see Section 3.2.12 on page 347.
- 18 • **omp\_get\_schedule** routine, see Section 3.2.13 on page 349.

## 19 6.2 OMP\_NUM\_THREADS

20 The **OMP\_NUM\_THREADS** environment variable sets the number of threads to use for **parallel**  
21 regions by setting the initial value of the *nthreads-var* ICV. See Section 2.6 on page 66 for a  
22 comprehensive set of rules about the interaction between the **OMP\_NUM\_THREADS** environment  
23 variable, the **num\_threads** clause, the **omp\_set\_num\_threads** library routine and dynamic  
24 adjustment of threads, and Section 2.8.1 on page 84 for a complete algorithm that describes how the  
25 number of threads for a **parallel** region is determined.

26 The value of this environment variable must be a list of positive integer values. The values of the  
27 list set the number of threads to use for **parallel** regions at the corresponding nested levels.

28 The behavior of the program is implementation defined if any value of the list specified in the  
29 **OMP\_NUM\_THREADS** environment variable leads to a number of threads that is greater than an  
30 implementation can support, or if any value is not a positive integer.

31 The **OMP\_NUM\_THREADS** environment variable sets the *max-active-levels-var* ICV to the number  
32 of active levels of parallelism that the implementation supports if the **OMP\_NUM\_THREADS**  
33 environment variable is set to a comma-separated list of more than one value. The value of the



1 *max-active-level-var* ICV may be overridden by setting **OMP\_MAX\_ACTIVE\_LEVELS** or  
2 **OMP\_NESTED**. See Section 6.8 on page 592 and Section 6.9 on page 593 for details.

3 Example:

```
4 | setenv OMP_NUM_THREADS 4,3,2
```

## 5 **Cross References**

- 6 • *nthreads-var* ICV, see Section 2.6 on page 66.
- 7 • **num\_threads** clause, see Section 2.8 on page 79.
- 8 • **omp\_set\_num\_threads** routine, see Section 3.2.1 on page 338.
- 9 • **omp\_get\_num\_threads** routine, see Section 3.2.2 on page 339.
- 10 • **omp\_get\_max\_threads** routine, see Section 3.2.3 on page 340.
- 11 • **omp\_get\_team\_size** routine, see Section 3.2.20 on page 355.

## 12 **6.3 OMP\_DYNAMIC**

13 The **OMP\_DYNAMIC** environment variable controls dynamic adjustment of the number of threads  
14 to use for executing **parallel** regions by setting the initial value of the *dyn-var* ICV.

15 The value of this environment variable must be one of the following:

16 **true** | **false**

17 If the environment variable is set to **true**, the OpenMP implementation may adjust the number of  
18 threads to use for executing **parallel** regions in order to optimize the use of system resources. If  
19 the environment variable is set to **false**, the dynamic adjustment of the number of threads is  
20 disabled. The behavior of the program is implementation defined if the value of **OMP\_DYNAMIC** is  
21 neither **true** nor **false**.

22 Example:

```
23 | setenv OMP_DYNAMIC true
```

## 24 **Cross References**

- 25 • *dyn-var* ICV, see Section 2.6 on page 66.
- 26 • **omp\_set\_dynamic** routine, see Section 3.2.7 on page 343.
- 27 • **omp\_get\_dynamic** routine, see Section 3.2.8 on page 344.

## 6.4 OMP\_PROC\_BIND

The `OMP_PROC_BIND` environment variable sets the initial value of the *bind-var* ICV. The value of this environment variable is either `true`, `false`, or a comma separated list of `master`, `close`, or `spread`. The values of the list set the thread affinity policy to be used for parallel regions at the corresponding nested level.

If the environment variable is set to `false`, the execution environment may move OpenMP threads between OpenMP places, thread affinity is disabled, and `proc_bind` clauses on `parallel` constructs are ignored.

Otherwise, the execution environment should not move OpenMP threads between OpenMP places, thread affinity is enabled, and the initial thread is bound to the first place in the OpenMP place list prior to the first active parallel region.

The behavior of the program is implementation defined if the value in the `OMP_PROC_BIND` environment variable is not `true`, `false`, or a comma separated list of `master`, `close`, or `spread`. The behavior is also implementation defined if an initial thread cannot be bound to the first place in the OpenMP place list.

The `OMP_PROC_BIND` environment variable sets the *max-active-levels-var* ICV to the number of active levels of parallelism that the implementation supports if the `OMP_PROC_BIND` environment variable is set to a comma-separated list of more than one element. The value of the *max-active-level-var* ICV may be overridden by setting `OMP_MAX_ACTIVE_LEVELS` or `OMP_NESTED`. See Section 6.8 on page 592 and Section 6.9 on page 593 for details.

Examples:

```
setenv OMP_PROC_BIND false
setenv OMP_PROC_BIND "spread, spread, close"
```

### Cross References

- *bind-var* ICV, see Section 2.6 on page 66.
- `proc_bind` clause, see Section 2.8.2 on page 85.
- `omp_get_proc_bind` routine, see Section 3.2.23 on page 357.

## 6.5 OMP\_PLACES

A list of places can be specified in the `OMP_PLACES` environment variable. The *place-partition-var* ICV obtains its initial value from the `OMP_PLACES` value, and makes the list available to the execution environment. The value of `OMP_PLACES` can be one of two types of values: either an abstract name that describes a set of places or an explicit list of places described by non-negative numbers.

1 The **OMP\_PLACES** environment variable can be defined using an explicit ordered list of  
2 comma-separated places. A place is defined by an unordered set of comma-separated non-negative  
3 numbers enclosed by braces. The meaning of the numbers and how the numbering is done are  
4 implementation defined. Generally, the numbers represent the smallest unit of execution exposed by  
5 the execution environment, typically a hardware thread.

6 Intervals may also be used to define places. Intervals can be specified using the *<lower-bound> :*  
7 *<length> : <stride>* notation to represent the following list of numbers: “*<lower-bound>*,  
8 *<lower-bound> + <stride>*, ..., *<lower-bound> + (<length> - 1)\*<stride>*.” When *<stride>* is  
9 omitted, a unit stride is assumed. Intervals can specify numbers within a place as well as sequences  
10 of places.

11 An exclusion operator “!” can also be used to exclude the number or place immediately following  
12 the operator.

13 Alternatively, the abstract names listed in Table 6.1 should be understood by the execution and  
14 runtime environment. The precise definitions of the abstract names are implementation defined. An  
15 implementation may also add abstract names as appropriate for the target platform.

16 The abstract name may be appended by a positive number in parentheses to denote the length of the  
17 place list to be created, that is *abstract\_name(num-places)*. When requesting fewer places than  
18 available on the system, the determination of which resources of type *abstract\_name* are to be  
19 included in the place list is implementation defined. When requesting more resources than  
20 available, the length of the place list is implementation defined.

**TABLE 6.1:** Defined Abstract Names for **OMP\_PLACES**

Abstract Name	Meaning
<b>threads</b>	Each place corresponds to a single hardware thread on the device.
<b>cores</b>	Each place corresponds to a single core (having one or more hardware threads) on the device.
<b>ll_caches</b>	Each place corresponds to a set of cores that share the last level cache on the device.
<b>numa_domains</b>	Each place corresponds to a set of cores for which their closest memory on the device is: <ul style="list-style-type: none"><li>• the same memory; and</li><li>• at a similar distance from the cores.</li></ul>

*table continued on next page*

Abstract Name	Meaning
<b>sockets</b>	Each place corresponds to a single socket (consisting of one or more cores) on the device.

1 The behavior of the program is implementation defined when the execution environment cannot  
 2 map a numerical value (either explicitly defined or implicitly derived from an interval) within the  
 3 **OMP\_PLACES** list to a processor on the target platform, or if it maps to an unavailable processor.  
 4 The behavior is also implementation defined when the **OMP\_PLACES** environment variable is  
 5 defined using an abstract name.  
 6 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 | ll_caches | numa_domains | threads
              | <implementation-defined abstract name>
    <res>    = non-negative integer
    <num-places> = positive integer
    <stride>  = integer
    <len>     = positive integer
  
```

7 Examples:

```

8 setenv OMP_PLACES threads
9 setenv OMP_PLACES "threads (4) "
10 setenv OMP_PLACES
11     "{0,1,2,3},{4,5,6,7},{8,9,10,11},{12,13,14,15}"
12 setenv OMP_PLACES "{0:4},{4:4},{8:4},{12:4}"
13 setenv OMP_PLACES "{0:4}:4:4"
  
```

14 where each of the last three definitions corresponds to the same 4 places including the smallest

1 units of execution exposed by the execution environment numbered, in turn, 0 to 3, 4 to 7, 8 to 11,  
2 and 12 to 15.

## 3 **Cross References**

- 4 • *place-partition-var*, see Section 2.6 on page 66.
- 5 • Controlling OpenMP thread affinity, see Section 2.8.2 on page 85.
- 6 • `omp_get_num_places` routine, see Section 3.2.24 on page 359.
- 7 • `omp_get_place_num_procs` routine, see Section 3.2.25 on page 360.
- 8 • `omp_get_place_proc_ids` routine, see Section 3.2.26 on page 360.
- 9 • `omp_get_place_num` routine, see Section 3.2.27 on page 361.
- 10 • `omp_get_partition_num_places` routine, see Section 3.2.28 on page 362.
- 11 • `omp_get_partition_place_nums` routine, see Section 3.2.29 on page 363.

## 12 **6.6 OMP\_STACKSIZE**

13 The `OMP_STACKSIZE` environment variable controls the size of the stack for threads created by  
14 the OpenMP implementation, by setting the value of the *stacksize-var* ICV. The environment  
15 variable does not control the size of the stack for an initial thread.

16 The value of this environment variable takes the form:

17 *size* | *size***B** | *size***K** | *size***M** | *size***G**

18 where:

- 19 • *size* is a positive integer that specifies the size of the stack for threads that are created by the  
20 OpenMP implementation.
- 21 • **B**, **K**, **M**, and **G** are letters that specify whether the given size is in Bytes, Kilobytes (1024 Bytes),  
22 Megabytes (1024 Kilobytes), or Gigabytes (1024 Megabytes), respectively. If one of these letters  
23 is present, there may be white space between *size* and the letter.

24 If only *size* is specified and none of **B**, **K**, **M**, or **G** is specified, then *size* is assumed to be in Kilobytes.

25 The behavior of the program is implementation defined if `OMP_STACKSIZE` does not conform to  
26 the above format, or if the implementation cannot provide a stack with the requested size.

27 Examples:

```
28 setenv OMP_STACKSIZE 2000500B  
29 setenv OMP_STACKSIZE "3000 k "  
30 setenv OMP_STACKSIZE 10M  
31 setenv OMP_STACKSIZE " 10 M "  
32 setenv OMP_STACKSIZE "20 m "
```

```
1 setenv OMP_STACKSIZE " 1G"  
2 setenv OMP_STACKSIZE 20000
```

### Cross References

- *stacksize-var* ICV, see Section 2.6 on page 66.

## 6.7 OMP\_WAIT\_POLICY

The **OMP\_WAIT\_POLICY** environment variable provides a hint to an OpenMP implementation about the desired behavior of waiting threads by setting the *wait-policy-var* ICV. A compliant OpenMP implementation may or may not abide by the setting of the environment variable.

The value of this environment variable must be one of the following:

**ACTIVE** | **PASSIVE**

The **ACTIVE** value specifies that waiting threads should mostly be active, consuming processor cycles, while waiting. An OpenMP implementation may, for example, make waiting threads spin.

The **PASSIVE** value specifies that waiting threads should mostly be passive, not consuming processor cycles, while waiting. For example, an OpenMP implementation may make waiting threads yield the processor to other threads or go to sleep.

The details of the **ACTIVE** and **PASSIVE** behaviors are implementation defined.

The behavior of the program is implementation defined if the value of **OMP\_WAIT\_POLICY** is neither **ACTIVE** nor **PASSIVE**.

Examples:

```
20 setenv OMP_WAIT_POLICY ACTIVE  
21 setenv OMP_WAIT_POLICY active  
22 setenv OMP_WAIT_POLICY PASSIVE  
23 setenv OMP_WAIT_POLICY passive
```

### Cross References

- *wait-policy-var* ICV, see Section 2.6 on page 66.

## 6.8 OMP\_MAX\_ACTIVE\_LEVELS

The **OMP\_MAX\_ACTIVE\_LEVELS** environment variable controls the maximum number of nested active **parallel** regions by setting the initial value of the *max-active-levels-var* ICV.

The value of this environment variable must be a non-negative integer. The behavior of the program is implementation defined if the requested value of **OMP\_MAX\_ACTIVE\_LEVELS** is

greater than the maximum number of nested active parallel levels an implementation can support, or if the value is not a non-negative integer.

## Cross References

- *max-active-levels-var* ICV, see Section 2.6 on page 66.
- `omp_set_max_active_levels` routine, see Section 3.2.16 on page 352.
- `omp_get_max_active_levels` routine, see Section 3.2.17 on page 353.

## 6.9 OMP\_NESTED

The `OMP_NESTED` environment variable controls nested parallelism by setting the initial value of the *max-active-levels-var* ICV. If the environment variable is set to `true`, the initial value of *max-active-levels-var* is set to the number of active levels of parallelism supported by the implementation. If the environment variable is set to `false`, the initial value of *max-active-levels-var* is set to 1. The behavior of the program is implementation defined if the value of `OMP_NESTED` is neither `true` nor `false`.

If both the `OMP_NESTED` and `OMP_MAX_ACTIVE_LEVELS` environment variables are set, the value of `OMP_NESTED` is `false`, and the value of `OMP_MAX_ACTIVE_LEVELS` is greater than 1, the behavior is implementation defined. Otherwise, if both environment variables are set then the `OMP_NESTED` environment variable has no effect.

The `OMP_NESTED` environment variable has been deprecated.

Example:

```
setenv OMP_NESTED false
```

## Cross References

- *max-active-levels-var* ICV, see Section 2.6 on page 66.
- `omp_set_nested` routine, see Section 3.2.10 on page 345.
- `omp_get_team_size` routine, see Section 3.2.20 on page 355.
- `OMP_MAX_ACTIVE_LEVELS` environment variable, see Section 6.8 on page 592.

## 6.10 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 2.6 on page 66.
- `omp_get_thread_limit` routine, see Section 3.2.14 on page 350.

## 6.11 OMP\_CANCELLATION

The `OMP_CANCELLATION` environment variable sets the initial value of the *cancel-var* ICV.

The value of this environment variable must be one of the following:

`true` | `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. The behavior of the program is implementation defined if `OMP_CANCELLATION` is set to neither `true` nor `false`.

## Cross References

- *cancel-var*, see Section 2.6.1 on page 66.
- `cancel` construct, see Section 2.21.1 on page 265.
- `cancellation point` construct, see Section 2.21.2 on page 269.
- `omp_get_cancellation` routine, see Section 3.2.9 on page 345.

## 6.12 OMP\_DISPLAY\_ENV

The `OMP_DISPLAY_ENV` environment variable instructs the runtime to display the OpenMP version number and the value of the ICVs associated with the environment variables described in Chapter 6, as *name = value* pairs. The runtime displays this information once, after processing the environment variables and before any user calls to change the ICV values by runtime routines defined in Chapter 3.

The value of the `OMP_DISPLAY_ENV` environment variable may be set to one of these values:

`TRUE` | `FALSE` | `VERBOSE`

The `TRUE` value instructs the runtime to display the OpenMP version number defined by the `_OPENMP` version macro (or the `openmp_version` Fortran parameter) value and the initial ICV values for the environment variables listed in Chapter 6. The `VERBOSE` value indicates that the runtime may also display the values of runtime variables that may be modified by vendor-specific environment variables. The runtime does not display any information when the `OMP_DISPLAY_ENV` environment variable is `FALSE` or undefined. For all values of the



1 environment variable other than **TRUE**, **FALSE**, and **VERBOSE**, the displayed information is  
2 unspecified.

3 The display begins with "OPENMP DISPLAY ENVIRONMENT BEGIN", followed by the  
4 `_OPENMP` version macro (or the `openmp_version` Fortran parameter) value and ICV values, in  
5 the format `NAME '=' VALUE`. `NAME` corresponds to the macro or environment variable name,  
6 optionally prepended by a bracketed *device-type*. `VALUE` corresponds to the value of the macro or  
7 ICV associated with this environment variable. Values are enclosed in single quotes. The display is  
8 terminated with "OPENMP DISPLAY ENVIRONMENT END".

9 For the **OMP\_NESTED** environment variable, the printed value is *true* if the *max-active-levels-var*  
10 ICV is initialized to a value greater than 1; otherwise the printed value is *false*.

11 Example:

```
12 % setenv OMP_DISPLAY_ENV TRUE
```

13 The above example causes an OpenMP implementation to generate output of the following form:

```
14 OPENMP DISPLAY ENVIRONMENT BEGIN  
15   _OPENMP=' 201811'  
16   [host] OMP_SCHEDULE=' GUIDED, 4'  
17   [host] OMP_NUM_THREADS=' 4, 3, 2'  
18   [device] OMP_NUM_THREADS=' 2'  
19   [host,device] OMP_DYNAMIC=' TRUE'  
20   [host] OMP_PLACES=' {0:4}, {4:4}, {8:4}, {12:4}'  
21   ...  
22 OPENMP DISPLAY ENVIRONMENT END
```

## 23 6.13 OMP\_DISPLAY\_AFFINITY

24 The **OMP\_DISPLAY\_AFFINITY** environment variable instructs the runtime to display formatted  
25 affinity information for all OpenMP threads in the parallel region upon entering the first parallel  
26 region and when any change occurs in the information accessible by the format specifiers listed in  
27 Table 6.2. If affinity of any thread in a parallel region changes then thread affinity information for  
28 all threads in that region is displayed. If the thread affinity for each respective parallel region at  
29 each nesting level has already been displayed and the thread affinity has not changed, then the  
30 information is not displayed again. There is no specific order in displaying thread affinity  
31 information for all threads in the same parallel region.

32 The value of the **OMP\_DISPLAY\_AFFINITY** environment variable may be set to one of these  
33 values:

34 **TRUE** | **FALSE**

35 The **TRUE** value instructs the runtime to display the OpenMP thread affinity information, and uses  
36 the format setting defined in the *affinity-format-var* ICV.

1 The runtime does not display the OpenMP thread affinity information when the value of the  
2 **OMP\_DISPLAY\_AFFINITY** environment variable is **FALSE** or undefined. For all values of the  
3 environment variable other than **TRUE** or **FALSE**, the display action is implementation defined.

4 Example:

```
5 setenv OMP_DISPLAY_AFFINITY TRUE
```

6 The above example causes an OpenMP implementation to display OpenMP thread affinity  
7 information during execution of the program, in a format given by the *affinity-format-var* ICV. The  
8 following is a sample output:

```
9 nesting_level= 1, thread_num= 0, thread_affinity= 0,1  
10 nesting_level= 1, thread_num= 1, thread_affinity= 2,3
```

## 11 Cross References

- 12 • Controlling OpenMP thread affinity, see Section 2.8.2 on page 85.
- 13 • **omp\_set\_affinity\_format** routine, see Section 3.2.30 on page 364.
- 14 • **omp\_get\_affinity\_format** routine, see Section 3.2.31 on page 365.
- 15 • **omp\_display\_affinity** routine, see Section 3.2.32 on page 366.
- 16 • **omp\_capture\_affinity** routine, see Section 3.2.33 on page 367.
- 17 • **OMP\_AFFINITY\_FORMAT** environment variable, see Section 6.14 on page 596.

## 18 6.14 OMP\_AFFINITY\_FORMAT

19 The **OMP\_AFFINITY\_FORMAT** environment variable sets the initial value of the  
20 *affinity-format-var* ICV which defines the format when displaying OpenMP thread affinity  
21 information.

22 The value of this environment variable is a character string that may contain as substrings one or  
23 more field specifiers, in addition to other characters. The format of each field specifier is

```
24 %[[[0].] size ] type
```

25 where an individual field specifier must contain the percent symbol (%) and a type. The type can be  
26 a single character short name or its corresponding long name delimited with curly braces, such as  
27 **%n** or **%{thread\_num}**. A literal percent is specified as **%%**. Field specifiers can be provided in  
28 any order.

29 The **0** modifier indicates whether or not to add leading zeros to the output, following any indication  
30 of sign or base. The **.** modifier indicates the output should be right justified when *size* is specified.  
31 By default, output is left justified. The minimum field length is *size*, which is a decimal digit string  
32 with a non-zero first digit. If no *size* is specified, the actual length needed to print the field will be  
33 used. If the **0** modifier is used with *type* of **A**, **{thread\_affinity}**, **H**, **{host}**, or a type that

1 is not printed as a number, the result is unspecified. Any other characters in the format string that  
 2 are not part of a field specifier will be included literally in the output.

**TABLE 6.2:** Available Field Types for Formatting OpenMP Thread Affinity Information

Short Name	Long Name	Meaning
t	team_num	The value returned by <code>omp_get_team_num()</code> .
T	num_teams	The value returned by <code>omp_get_num_teams()</code> .
L	nesting_level	The value returned by <code>omp_get_level()</code> .
n	thread_num	The value returned by <code>omp_get_thread_num()</code> .
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 <code>level</code> is <code>omp_get_level()</code> minus 1.
H	host	The name for the host device on which the OpenMP program is running.
P	process_id	The process identifier used by the implementation.
i	native_thread_id	The native thread identifier used by the implementation.
A	thread_affinity	The list of numerical identifiers, in the format of a comma-separated list of integers or integer ranges, that represent processors on which a thread may execute, subject to OpenMP thread affinity control and/or other external affinity mechanisms.

3 Implementations may define additional field types. If an implementation does not have information  
 4 for a field type, "undefined" is printed for this field when displaying the OpenMP thread affinity  
 5 information.

6 Example:

```
7 setenv OMP_AFFINITY_FORMAT  

8 "Thread Affinity: %0.3L %.8n %.15{thread_affinity} %.12H"
```

9 The above example causes an OpenMP implementation to display OpenMP thread affinity  
 10 information in the following form:

```
11 Thread Affinity: 001 0 0-1,16-17 nid003  

12 Thread Affinity: 001 1 2-3,18-19 nid003
```

## Cross References

- Controlling OpenMP thread affinity, see Section 2.8.2 on page 85.
- `omp_set_affinity_format` routine, see Section 3.2.30 on page 364.
- `omp_get_affinity_format` routine, see Section 3.2.31 on page 365.
- `omp_display_affinity` routine, see Section 3.2.32 on page 366.
- `omp_capture_affinity` routine, see Section 3.2.33 on page 367.
- `OMP_DISPLAY_AFFINITY` environment variable, see Section 6.13 on page 595.

## 6.15 OMP\_DEFAULT\_DEVICE

The `OMP_DEFAULT_DEVICE` environment variable sets the device number to use in device constructs by setting the initial value of the *default-device-var* ICV.

The value of this environment variable must be a non-negative integer value.

## Cross References

- *default-device-var* ICV, see Section 2.6 on page 66.
- device directives, Section 2.15 on page 165.

## 6.16 OMP\_MAX\_TASK\_PRIORITY

The `OMP_MAX_TASK_PRIORITY` environment variable controls the use of task priorities by setting the initial value of the *max-task-priority-var* ICV. The value of this environment variable must be a non-negative integer.

Example:

```
% setenv OMP_MAX_TASK_PRIORITY 20
```

## Cross References

- *max-task-priority-var* ICV, see Section 2.6 on page 66.
- Tasking Constructs, see Section 2.13 on page 142.
- `omp_get_max_task_priority` routine, see Section 3.2.42 on page 374.

## 6.17 OMP\_TARGET\_OFFLOAD

The `OMP_TARGET_OFFLOAD` environment variable sets the initial value of the *target-offload-var* ICV. The value of the `OMP_TARGET_OFFLOAD` environment variable must be one of the following:

**MANDATORY | DISABLED | DEFAULT**

The **MANDATORY** value specifies that program execution is terminated if a device construct or device memory routine is encountered and the device is not available or is not supported by the implementation. Support for the **DISABLED** value is implementation defined. If an implementation supports it, the behavior is as if the only device is the host device.

The **DEFAULT** value specifies the default behavior as described in Section 1.3 on page 21.

Example:

```
% setenv OMP_TARGET_OFFLOAD MANDATORY
```

### Cross References

- *target-offload-var* ICV, see Section 2.6 on page 66.
- Device Directives, see Section 2.15 on page 165.
- Device Memory Routines, see Section 3.6 on page 396.

## 6.18 OMP\_TOOL

The `OMP_TOOL` environment variable sets the *tool-var* ICV, which controls whether an OpenMP runtime will try to register a first party tool.

The value of this environment variable must be one of the following:

**enabled | disabled**

If `OMP_TOOL` is set to any value other than **enabled** or **disabled**, the behavior is unspecified. If `OMP_TOOL` is not defined, the default value for *tool-var* is **enabled**.

Example:

```
% setenv OMP_TOOL enabled
```

### Cross References

- *tool-var* ICV, see Section 2.6 on page 66.
- OMPT Interface, see Chapter 4 on page 425.

## 6.19 OMP\_TOOL\_LIBRARIES

The `OMP_TOOL_LIBRARIES` environment variable sets the *tool-libraries-var* ICV to a list of tool libraries that are considered for use on a device on which an OpenMP implementation is being initialized. The value of this environment variable must be a list of names of dynamically-loadable libraries, separated by an implementation specific, platform typical separator.

If the *tool-var* ICV is not enabled, the value of *tool-libraries-var* is ignored. Otherwise, if `ompt_start_tool` is not visible in the address space on a device where OpenMP is being initialized or if `ompt_start_tool` returns `NULL`, an OpenMP implementation will consider libraries in the *tool-libraries-var* list in a left to right order. The OpenMP implementation will search the list for a library that meets two criteria: it can be dynamically loaded on the current device and it defines the symbol `ompt_start_tool`. If an OpenMP implementation finds a suitable library, no further libraries in the list will be considered.

Example:

```
% setenv OMP_TOOL_LIBRARIES libtoolXY64.so:/usr/local/lib/  
libtoolXY32.so
```

### Cross References

- *tool-libraries-var* ICV, see Section 2.6 on page 66.
- OMPT Interface, see Chapter 4 on page 425.
- `ompt_start_tool` routine, see Section 4.2.1 on page 425.

## 6.20 OMP\_DEBUG

The `OMP_DEBUG` environment variable sets the *debug-var* ICV, which controls whether an OpenMP runtime collects information that an OMPD library may need to support a tool.

The value of this environment variable must be one of the following:

`enabled` | `disabled`

If `OMP_DEBUG` is set to any value other than `enabled` or `disabled` then the behavior is implementation defined.

Example:

```
% setenv OMP_DEBUG enabled
```

### Cross References

- *debug-var* ICV, see Section 2.6 on page 66.
- OMPD Interface, see Chapter 5 on page 529.
- Enabling the Runtime for OMPD, see Section 5.2.1 on page 530.

## 6.21 OMP\_ALLOCATOR

**OMP\_ALLOCATOR** sets the *def-allocator-var* ICV that specifies the default allocator for allocation calls, directives and clauses that do not specify an allocator. The value of this environment variable is a predefined allocator from Table 2.10 on page 161. The value of this environment variable is not case sensitive.

### Cross References

- *def-allocator-var* ICV, see Section 2.6 on page 66.
- Memory allocators, see Section 2.14.2 on page 159.
- **omp\_set\_default\_allocator** routine, see Section 3.8.4 on page 417.
- **omp\_get\_default\_allocator** routine, see Section 3.8.5 on page 418.

## 6.22 OMP\_NUM\_TEAMS

The **OMP\_NUM\_TEAMS** environment variable sets the maximum number of teams created by a **teams** construct by setting the *ntteams-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\_NUM\_TEAMS** is greater than the number of teams that an implementation can support, or if the value is not a positive integer.

### Cross References

- *ntteams-var* ICV, see Section 2.6 on page 66.
- **omp\_get\_max\_teams** routine, see Section 3.2.46 on page 379.

## 6.23 OMP\_TEAMS\_THREAD\_LIMIT

The **OMP\_TEAMS\_THREAD\_LIMIT** environment variable sets the maximum number of OpenMP threads to use in each contention group created by a **teams** construct by setting the *teams-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\_TEAMS\_THREAD\_LIMIT** is greater than the number of threads that an implementation can support, or if the value is not a positive integer.

### Cross References

- *teams-thread-limit-var* ICV, see Section 2.6 on page 66.
- **omp\_get\_teams\_thread\_limit** routine, see Section 3.2.48 on page 381.

*This page intentionally left blank*



# A OpenMP Implementation-Defined Behaviors

This appendix summarizes the behaviors that are described as implementation defined in this API. Each behavior is cross-referenced back to its description in the main specification. An implementation is required to define and to document its behavior in these cases.

- **Processor**: a hardware unit that is implementation defined (see Section 1.2.1 on page 2).
- **Device**: an implementation defined logical execution engine (see Section 1.2.1 on page 2).
- **Device address**: reference to an address in a *device data environment* (see Section 1.2.6 on page 13).
- **Memory model**: the minimum size at which a memory update may also read and write back adjacent variables that are part of another variable (as array or structure elements) is implementation defined but is no larger than required by the base language (see Section 1.4.1 on page 24).
- **requires directive**: support of requirements is implementation defined. All implementation-defined requirements should begin with **ext\_** (see Section 2.4 on page 60).
- **Requires directive**: Support for any feature specified by a requirement clause on a **requires** directive is implementation defined (see Section 2.4 on page 60).
- **Internal control variables**: the initial values of *dyn-var*, *nthreads-var*, *run-sched-var*, *def-sched-var*, *bind-var*, *stacksize-var*, *wait-policy-var*, *thread-limit-var*, *max-active-levels-var*, *place-partition-var*, *affinity-format-var*, *default-device-var* and *def-allocator-var* are implementation defined. The method for initializing a target device's internal control variable is implementation defined (see Section 2.6.2 on page 68).
- **OpenMP context**: the accepted *isa-name* values for the *isa* trait, the accepted *arch-name* values for the *arch* trait, and the accepted *extension-name* values for the *extension* trait are implementation defined (see Section 2.3.1 on page 50).
- **declare variant directive**: whether, for some specific OpenMP context, the prototype of the variant should differ from that of the base function, and if so how it should differ, is implementation defined (see Section 2.3.5 on page 56).
- **Dynamic adjustment of threads**: providing the ability to adjust the number of threads dynamically is implementation defined. Implementations are allowed to deliver fewer threads

1 (but at least one) than indicated in Algorithm 2.1 even if dynamic adjustment is disabled (see  
2 Section 2.8.1 on page 84).

- 3 • **Thread affinity:** For the **close** thread affinity policy, if  $T > P$  and  $P$  does not divide  $T$  evenly,  
4 the exact number of threads in a particular place is implementation defined. For the **spread**  
5 thread affinity, if  $T > P$  and  $P$  does not divide  $T$  evenly, the exact number of threads in a  
6 particular subpartition is implementation defined. The determination of whether the affinity  
7 request can be fulfilled is implementation defined. If not, the mapping of threads in the team to  
8 places is implementation defined (see Section 2.8.2 on page 85).
- 9 • **teams construct:** the number of teams that are created is implementation defined but between  
10 one and the value of the **num\_teams** clause if specified or between one and the value of the  
11 *nteams-var* ICV if its value is greater than zero. The maximum number of threads that  
12 participate in the contention group that each team initiates is implementation defined but  
13 between one and the value of the **thread\_limit** clause if specified or between one and the  
14 value of the *teams-thread-limit-var* ICV if its value is greater than zero. The assignment of the  
15 initial threads to places and the values of the *place-partition-var* and *default-device-var* ICVs for  
16 each initial thread are implementation defined (see Section 2.9 on page 87).
- 17 • **sections construct:** the method of scheduling the structured blocks among threads in the  
18 team is implementation defined (see Section 2.11.1 on page 93).
- 19 • **single construct:** the method of choosing a thread to execute the structured block is  
20 implementation defined (see Section 2.11.2 on page 96)
- 21 • **Worksharing-Loop directive:** the integer type (or kind, for Fortran) used to compute the  
22 iteration count of a collapsed loop is implementation defined. The effect of the  
23 **schedule(runtime)** clause when the *run-sched-var* ICV is set to **auto** is implementation  
24 defined. The value of *simd\_width* for the **simd** schedule modifier is implementation defined. If  
25 any of the associated loop(s) is a generated loop of a **tile** directive, the effect of the **schedule**  
26 clause on the loop schedule is implementation defined (see Section 2.12.2 on page 106).
- 27 • **simd construct:** the integer type (or kind, for Fortran) used to compute the iteration count for  
28 the collapsed loop is implementation defined. The number of iterations that are executed  
29 concurrently at any given time is implementation defined. If the *alignment* parameter is not  
30 specified in the **aligned** clause, the default alignments for the SIMD instructions are  
31 implementation defined (see Section 2.12.3.1 on page 115).
- 32 • **declare simd directive:** if the parameter of the **simdlen** clause is not a constant positive  
33 integer expression, the number of concurrent arguments for the function is implementation  
34 defined. If the *alignment* parameter of the **aligned** clause is not specified, the default  
35 alignments for SIMD instructions are implementation defined (see Section 2.12.3.3 on page 121).
- 36 • **distribute construct:** the integer type (or kind, for Fortran) used to compute the iteration  
37 count for the collapsed loop is implementation defined. If no **dist\_schedule** clause is  
38 specified then the schedule for the **distribute** construct is implementation defined. If any of  
39 the associated loop(s) is a generated loop of a **tile** directive, the effect of the

1       **dist\_schedule** clause on the loop schedule is implementation defined (see Section 2.12.4.1  
2       on page 125).

3       • **taskloop construct**: The number of loop iterations assigned to a task created from a  
4       **taskloop** construct is implementation defined, unless the **grainsize** or **num\_tasks**  
5       clause is specified. The integer type (or kind, for Fortran) used to compute the iteration count for  
6       the collapsed loop is implementation defined. If any of the associated loop(s) is a generated loop  
7       of a **tile** directive, the effect of the **grainsize** and **num\_tasks** clauses on the partitioning  
8       of iterations into tasks is implementation defined (see Section 2.13.2 on page 147).

C++

9       • **taskloop construct**: For **firstprivate** variables of class type, the number of invocations  
10       of copy constructors to perform the initialization is implementation defined (see Section 2.13.2  
11       on page 147).

C++

12       • **Memory spaces**: The actual storage resource that each memory space defined in Table 2.8 on  
13       page 158 represents is implementation defined.

14       • **Memory allocators**: The minimum partitioning size for partitioning of allocated memory over  
15       the storage resources is implementation defined (see Section 2.14.2 on page 159). The default  
16       value for the **pool\_size** allocator trait is implementation defined (see Table 2.9 on page 159).  
17       The associated memory space for each of the predefined **omp\_cgroup\_mem\_alloc**,  
18       **omp\_pteam\_mem\_alloc** and **omp\_thread\_mem\_alloc** allocators is implementation  
19       defined (see Table 2.10 on page 161).

20       • **is\_device\_ptr clause**: Support for pointers created outside of the OpenMP device data  
21       management routines is implementation defined (see Section 2.15.5 on page 174).

22       • **atomic construct**: a compliant implementation may enforce exclusive access between  
23       **atomic** regions that update different storage locations. The circumstances under which this  
24       occurs are implementation defined. If the storage location designated by  $x$  is not size-aligned  
25       (that is, if the byte alignment of  $x$  is not a multiple of the size of  $x$ ), then the behavior of the  
26       atomic region is implementation defined (see Section 2.20.7 on page 237).

Fortran

27       • **Data-sharing attributes**: The data-sharing attributes of dummy arguments without the **VALUE**  
28       attribute are implementation-defined if the associated actual argument is shared, except for the  
29       conditions specified (see Section 2.22.1.2 on page 275).

30       • **threadprivate directive**: if the conditions for values of data in the threadprivate objects of  
31       threads (other than an initial thread) to persist between two consecutive active parallel regions do  
32       not all hold, the allocation status of an allocatable variable in the second region is  
33       implementation defined (see Section 2.22.2 on page 276).

- 1       • **Runtime library definitions:** it is implementation defined whether the include file `omp_lib.h`  
2       or the module `omp_lib` (or both) is provided. It is implementation defined whether any of the  
3       OpenMP runtime library routines that take an argument are extended with a generic interface so  
4       arguments of different **KIND** type can be accommodated (see Section 3.1 on page 335).

## Fortran

- 5       • **omp\_set\_num\_threads routine:** if the argument is not a positive integer the behavior is  
6       implementation defined (see Section 3.2.1 on page 338).
- 7       • **omp\_set\_schedule routine:** for implementation specific schedule kinds, the values and  
8       associated meanings of the second argument are implementation defined (see Section 3.2.12 on  
9       page 347).
- 10       • **omp\_get\_supported\_active\_levels routine:** the number of active levels of  
11       parallelism supported by the implementation is implementation defined, but must be greater than  
12       0 (see Section 3.2.15 on page 351).
- 13       • **omp\_set\_max\_active\_levels routine:** when called from within any explicit `parallel`  
14       region the binding thread set (and binding region, if required) for the  
15       **omp\_set\_max\_active\_levels** region is implementation defined and the behavior is  
16       implementation defined. If the argument is not a non-negative integer then the behavior is  
17       implementation defined (see Section 3.2.16 on page 352).
- 18       • **omp\_get\_max\_active\_levels routine:** when called from within any explicit `parallel`  
19       region the binding thread set (and binding region, if required) for the  
20       **omp\_get\_max\_active\_levels** region is implementation defined (see Section 3.2.17 on  
21       page 353).
- 22       • **omp\_get\_place\_proc\_ids routine:** the meaning of the non-negative numerical identifiers  
23       returned by the `omp_get_place_proc_ids` routine is implementation defined. The order of  
24       the numerical identifiers returned in the array `ids` is implementation defined (see Section 3.2.26  
25       on page 360).
- 26       • **omp\_set\_affinity\_format routine:** when called from within any explicit `parallel`  
27       region, the binding thread set (and binding region, if required) for the  
28       **omp\_set\_affinity\_format** region is implementation defined and the behavior is  
29       implementation defined. If the argument does not conform to the specified format then the result  
30       is implementation defined (see Section 3.2.30 on page 364).
- 31       • **omp\_get\_affinity\_format routine:** when called from within any explicit `parallel`  
32       region the binding thread set (and binding region, if required) for the  
33       **omp\_get\_affinity\_format** region is implementation defined (see Section 3.2.31 on  
34       page 365).
- 35       • **omp\_display\_affinity routine:** if the argument does not conform to the specified format  
36       then the result is implementation defined (see Section 3.2.32 on page 366).

- 1       • **omp\_capture\_affinity routine:** if the *format* argument does not conform to the specified  
2       format then the result is implementation defined (see Section 3.2.33 on page 367).
- 3       • **omp\_get\_initial\_device routine:** the value of the device number of the host device is  
4       implementation defined (see Section 3.2.41 on page 374).
- 5       • **omp\_set\_num\_teams routine:** if the argument is not a positive integer the behavior is  
6       implementation defined (see Section 3.2.45 on page 378).
- 7       • **omp\_set\_teams\_thread\_limit routine:** if the argument is not a positive integer the  
8       behavior is implementation defined (see Section 3.2.47 on page 380).
- 9       • **omp\_target\_memcpy\_rect routine:** the maximum number of dimensions supported is  
10      implementation defined, but must be at least three (see Section 3.6.5 on page 400).
- 11      • **ompt\_callback\_sync\_region\_wait, ompt\_callback\_mutex\_released,**  
12      **ompt\_callback\_dependences, ompt\_callback\_task\_dependence,**  
13      **ompt\_callback\_work, ompt\_callback\_master, ompt\_callback\_target\_map,**  
14      **ompt\_callback\_sync\_region, ompt\_callback\_lock\_init,**  
15      **ompt\_callback\_lock\_destroy, ompt\_callback\_mutex\_acquire,**  
16      **ompt\_callback\_mutex\_acquired, ompt\_callback\_nest\_lock,**  
17      **ompt\_callback\_flush, ompt\_callback\_cancel** and  
18      **ompt\_callback\_dispatch** tool callbacks: if a tool attempts to register a callback with the  
19      string name using the runtime entry point **ompt\_set\_callback**, it is implementation defined  
20      whether the registered callback may never or sometimes invoke this callback for the associated  
21      events (see Table 4.2 on page 433)
- 22      • **Device tracing:** Whether a target device supports tracing or not is implementation defined; if a  
23      target device does not support tracing, a **NULL** may be supplied for the *lookup* function to a  
24      tool's device initializer (see Section 4.2.5 on page 432).
- 25      • **ompt\_set\_trace\_ompt** and **ompt\_buffer\_get\_record\_ompt** runtime entry  
26      **points:** it is implementation defined whether a device-specific tracing interface will define this  
27      runtime entry point, indicating that it can collect traces in OMPT format. The kinds of trace  
28      records available for a device is implementation defined (see Section 4.2.5 on page 432).
- 29      • **ompt\_callback\_target\_data\_op\_t** callback type: it is implementation defined  
30      whether in some operations *src\_addr* or *dest\_addr* might point to an intermediate buffer (see  
31      Section 4.5.2.25 on page 488).
- 32      • **ompt\_set\_callback\_t** entry point type: the subset of the associated event in which the  
33      callback is invoked is implementation defined (see Section 4.6.1.3 on page 499).
- 34      • **ompt\_get\_place\_proc\_ids\_t** entry point type: the meaning of the numerical identifiers  
35      returned is implementation defined. The order of *ids* returned in the array is implementation  
36      defined (see Section 4.6.1.8 on page 503).
- 37      • **ompt\_get\_partition\_place\_nums\_t** entry point type: the order of the identifiers  
38      returned in the array *place\_nums* is implementation defined (see Section 4.6.1.10 on page 505).

- 1       • **ompt\_get\_proc\_id\_t entry point type:** the meaning of the numerical identifier returned is  
2       implementation defined (see Section 4.6.1.11 on page 506).
- 3       • **ompd\_callback\_print\_string\_fn\_t callback function:** the value of *category* is  
4       implementation defined (see Section 5.4.5 on page 548).
- 5       • **ompd\_parallel\_handle\_compare operation:** the means by which parallel region  
6       handles are ordered is implementation defined (see Section 5.5.6.5 on page 564).
- 7       • **ompd\_task\_handle\_compare operation:** the means by which task handles are ordered is  
8       implementation defined (see Section 5.5.7.6 on page 569).
- 9       • **OMPT thread states:** The set of OMPT thread states supported is implementation defined (see  
10       Section 4.4.4.26 on page 454).
- 11       • **OMP\_SCHEDULE environment variable:** if the value does not conform to the specified format  
12       then the result is implementation defined (see Section 6.1 on page 585).
- 13       • **OMP\_NUM\_THREADS environment variable:** if any value of the list specified leads to a number  
14       of threads that is greater than the implementation can support, or if any value is not a positive  
15       integer, then the result is implementation defined (see Section 6.2 on page 586).
- 16       • **OMP\_DYNAMIC environment variable:** if the value is neither **true** nor **false** the behavior is  
17       implementation defined (see Section 6.3 on page 587).
- 18       • **OMP\_PROC\_BIND environment variable:** if the value is not **true**, **false**, or a comma  
19       separated list of **master**, **close**, or **spread**, the behavior is implementation defined. The  
20       behavior is also implementation defined if an initial thread cannot be bound to the first place in  
21       the OpenMP place list (see Section 6.4 on page 588).
- 22       • **OMP\_PLACES environment variable:** the meaning of the numbers specified in the environment  
23       variable and how the numbering is done are implementation defined. The precise definitions of  
24       the abstract names are implementation defined. An implementation may add  
25       implementation-defined abstract names as appropriate for the target platform. When creating a  
26       place list of *n* elements by appending the number *n* to an abstract name, the determination of  
27       which resources to include in the place list is implementation defined. When requesting more  
28       resources than available, the length of the place list is also implementation defined. The behavior  
29       of the program is implementation defined when the execution environment cannot map a  
30       numerical value (either explicitly defined or implicitly derived from an interval) within the  
31       **OMP\_PLACES** list to a processor on the target platform, or if it maps to an unavailable processor.  
32       The behavior is also implementation defined when the **OMP\_PLACES** environment variable is  
33       defined using an abstract name (see Section 6.5 on page 588).
- 34       • **OMP\_STACKSIZE environment variable:** if the value does not conform to the specified format  
35       or the implementation cannot provide a stack of the specified size then the behavior is  
36       implementation defined (see Section 6.6 on page 591).
- 37       • **OMP\_WAIT\_POLICY environment variable:** the details of the **ACTIVE** and **PASSIVE**  
38       behaviors are implementation defined (see Section 6.7 on page 592).

- 1           ● **OMP\_MAX\_ACTIVE\_LEVELS environment variable:** if the value is not a non-negative integer  
2           or is greater than the number of parallel levels an implementation can support then the behavior  
3           is implementation defined (see Section 6.8 on page 592).
- 4           ● **OMP\_NESTED environment variable:** if the value is neither **true** nor **false** the behavior is  
5           implementation defined (see Section 6.9 on page 593).
- 6           ● **Conflicting OMP\_NESTED and OMP\_MAX\_ACTIVE\_LEVELS environment variables:** if  
7           both environment variables are set, the value of **OMP\_NESTED** is **false**, and the value of  
8           **OMP\_MAX\_ACTIVE\_LEVELS** is greater than 1, the behavior is implementation defined (see  
9           Section 6.9 on page 593).
- 10          ● **OMP\_THREAD\_LIMIT environment variable:** if the requested value is greater than the number  
11          of threads an implementation can support, or if the value is not a positive integer, the behavior of  
12          the program is implementation defined (see Section 6.10 on page 593).
- 13          ● **OMP\_DISPLAY\_AFFINITY environment variable:** for all values of the environment variables  
14          other than **TRUE** or **FALSE**, the display action is implementation defined (see Section 6.13 on  
15          page 595).
- 16          ● **OMP\_AFFINITY\_FORMAT environment variable:** if the value does not conform to the  
17          specified format then the result is implementation defined (see Section 6.14 on page 596).
- 18          ● **OMP\_TARGET\_OFFLOAD environment variable:** the support of **disabled** is  
19          implementation defined (see Section 6.17 on page 599).
- 20          ● **OMP\_DEBUG environment variable:** if the value is neither **disabled** nor **enabled** the  
21          behavior is implementation defined (see Section 6.20 on page 600).
- 22          ● **OMP\_NUM\_TEAMS environment variable:** if the value is not a positive integer or is greater than  
23          the number of teams that an implementation can support, the behavior of the program is  
24          implementation defined (see Section 6.22 on page 601).
- 25          ● **OMP\_TEAMS\_THREAD\_LIMIT environment variable:** if the value is not a positive integer or  
26          is greater than the number of threads that an implementation can support, the behavior of the  
27          program is implementation defined (see Section 6.23 on page 601).

*This page intentionally left blank*



# B Features History

This appendix summarizes the major changes between OpenMP API versions since version 2.5.

## B.1 Deprecated Features

The following features have been deprecated in Version 5.0.

- The *nest-var* ICV, the `OMP_NESTED` environment variable, and the `omp_set_nested` and `omp_get_nested` routines were deprecated.
- Lock hints were renamed to synchronization hints. The following lock hint type and constants were deprecated:
  - the C/C++ type `omp_lock_hint_t` and the Fortran kind `omp_lock_hint_kind`;
  - the constants `omp_lock_hint_none`, `omp_lock_hint_uncontended`, `omp_lock_hint_contended`, `omp_lock_hint_nonspeculative`, and `omp_lock_hint_speculative`.

## B.2 Version 5.0 to 5.1 Differences

- C11, C++11, C++14 and C++17 are now fully supported (see Section 1.7 on page 32).
- The OpenMP directive syntax was extended to include C++ attribute specifiers (see Section 2.1 on page 36).
- The `omp_all_memory` reserved locator was added (see Section 2.1 on page 36), and the `depend` clause was extended to allow its use (see Section 2.20.11 on page 257).
- For C/C++, the `declare variant` directive was extended to support elision of preprocessed code and to allow enclosed function definitions to be interpreted as variant functions (see Section 2.3.5 on page 56).
- The `assume` directive was added to allow users to specify invariants (see Section 2.5 on page 63).
- To support device-specific ICV settings the environment variable syntax was extended to support device-specific variables (see Section 2.6.2 on page 68).
- To support clarity in metadirectives, added the `nothing` directive (see Section 2.7.1 on page 77).

- 1       • To allow users to control the compilation process and runtime error actions, added the **error**
- 2       directive (see Section 2.7.2 on page 78).
- 3       • The **scope** directive was added to support reductions without requiring a **parallel** or
- 4       worksharing region (Section 2.10 on page 91).
- 5       • Support was added for specifying a **simd** construct with C++ attribute specifiers (see
- 6       Section 2.12.3.1 on page 115).
- 7       • Support was added for specifying a declare SIMD directive with C++ attribute specifiers (see
- 8       Section 2.12.3.3 on page 121).
- 9       • Loop transformation constructs were added (see Section 2.12.7 on page 140).
- 10      • Support was added so that iterators may be defined and used in a motion clause on a
- 11      **target update** directive (see Section 2.15.6 on page 181) or in a **map** clause (see
- 12      Section 2.22.7.1 on page 318).
- 13      • Support was added for indirect calls to the device version of a procedure or function in **target**
- 14      regions. (see Section 2.15.7 on page 185).
- 15      • The **interop** directive was added to enable portable interoperability with foreign execution
- 16      contexts used to implement OpenMP (see Section 2.15.8 on page 191). Additionally, the
- 17      **omp\_get\_num\_interop\_properties**, **omp\_get\_interop\_property** and
- 18      **omp\_get\_interop\_property\_name** runtime routines were added (see Section 3.7.1 on
- 19      page 410, Section 3.7.2 on page 410 and Section 3.7.3 on page 412) to interpret the value of an
- 20      **omp\_interop\_t** object.
- 21      • To support inout sets, an **inoutset** *dependence-type* was added to the **depend** clause (see
- 22      Section 2.20.11 on page 257).
- 23      • Support for **private** and **firstprivate** as an argument to the **default** clause in C and
- 24      C++ was added (see Section 2.22.4.1 on page 284).
- 25      • New **defaultmap** *implicit-behavior* type of *present* was added (see Section 2.22.7.3 on
- 26      page 328).
- 27      • The **omp\_set\_num\_teams** and **omp\_set\_teams\_thread\_limit** runtime routines were
- 28      added to control the number of teams and the size of those teams on the **teams** construct (see
- 29      Section 3.2.45 on page 378 and Section 3.2.47 on page 380). Additionally, the
- 30      **omp\_get\_max\_teams** and **omp\_get\_teams\_thread\_limit** runtime routines were
- 31      added to retrieve the values that will be used in the next **teams** construct (see Section 3.2.46 on
- 32      page 379 and Section 3.2.48 on page 381).
- 33      • To support asynchronous device memory management, **omp\_target\_memcpy\_async** and
- 34      **omp\_target\_memcpy\_rect\_async** runtime routines were added (see Section 3.6.6 on
- 35      page 402 and Section 3.6.7 on page 404).
- 36      • Support was added to obtain the pointer created by **omp\_target\_associate\_ptr** (see
- 37      Section 3.6.9 on page 408).

- The `OMP_NUM_TEAMS` and `OMP_TEAMS_THREAD_LIMIT` environment variables were added to control the number of teams and size of those teams on the `teams` construct (see Section 6.22 on page 601 and Section 6.23 on page 601).

## B.3 Version 4.5 to 5.0 Differences

- The memory model was extended to distinguish different types of flush operations according to specified flush properties (see Section 1.4.4 on page 26) and to define a happens before order based on synchronizing flush operations (see Section 1.4.5 on page 27).
- Various changes throughout the specification were made to provide initial support of C11, C++11, C++14, C++17 and Fortran 2008 (see Section 1.7 on page 32).
- Fortran 2003 is now fully supported (see Section 1.7 on page 32).
- The `requires` directive (see Section 2.4 on page 60) was added to support applications that require implementation-specific features.
- The `target-offload-var` internal control variable (see Section 2.6 on page 66) and the `OMP_TARGET_OFFLOAD` environment variable (see Section 6.17 on page 599) were added to support runtime control of the execution of device constructs.
- Control over whether nested parallelism is enabled or disabled was integrated into the `max-active-levels-var` internal control variable (see Section 2.6.2 on page 68), the default value of which is now implementation defined, unless determined according to the values of the `OMP_NUM_THREADS` (see Section 6.2 on page 586) or `OMP_PROC_BIND` (see Section 6.4 on page 588) environment variables.
- Support for array shaping (see Section 2.1.4 on page 42) and for array sections with non-unit strides in C and C++ (see Section 2.1.5 on page 42) was added to facilitate specification of discontinuous storage and the `target update` construct (see Section 2.15.6 on page 181) and the `depend` clause (see Section 2.20.11 on page 257) were extended to allow the use of shape-operators (see Section 2.1.4 on page 42).
- Iterators (see Section 2.1.6 on page 46) were added to support expressions in a list that expand to multiple expressions.
- The `metadirective` directive (see Section 2.3.4 on page 54) and `declare variant` directive (see Section 2.3.5 on page 56) were added to support selection of directive variants and declared function variants at a callsite, respectively, based on compile-time traits of the enclosing context.
- The `teams` construct (see Section 2.9 on page 87) was extended to support execution on the host device without an enclosing `target` construct (see Section 2.15.5 on page 174).
- The canonical loop form was defined for Fortran and, for all base languages, extended to permit non-rectangular loop nests (see Section 2.12.1 on page 101).

- 1 • The *relational-op* in the *canonical loop form* for C/C++ was extended to include **!=** (see  
2 Section 2.12.1 on page 101).
- 3 • The default loop schedule modifier for worksharing-loop constructs without the **static**  
4 schedule and the **ordered** clause was changed to **nonmonotonic** (see Section 2.12.2 on  
5 page 106).
- 6 • The collapse of associated loops that are imperfectly nested loops was defined for the  
7 worksharing-loop (see Section 2.12.2 on page 106), **simd** (see Section 2.12.3.1 on page 115),  
8 **taskloop** (see Section 2.13.2 on page 147) and **distribute** (see Section 2.12.4.2 on  
9 page 129) constructs.
- 10 • The **simd** construct (see Section 2.12.3.1 on page 115) was extended to accept the **if**,  
11 **nontemporal** and **order (concurrent)** clauses and to allow the use of **atomic**  
12 constructs within it.
- 13 • The **loop** construct and the **order (concurrent)** clause were added to support compiler  
14 optimization and parallelization of loops for which iterations may execute in any order, including  
15 concurrently (see Section 2.12.5 on page 133).
- 16 • The **scan** directive (see Section 2.12.6 on page 137) and the **inscan** modifier for the  
17 **reduction** clause (see Section 2.22.5.4 on page 302) were added to support inclusive and  
18 exclusive scan computations.
- 19 • To support task reductions, the **task** (see Section 2.13.1 on page 142) and **target** (see  
20 Section 2.15.5 on page 174) constructs were extended to accept the **in\_reduction** clause (see  
21 Section 2.22.5.6 on page 306), the **taskgroup** construct (see Section 2.20.6 on page 235) was  
22 extended to accept the **task\_reduction** clause Section 2.22.5.5 on page 305), and the **task**  
23 modifier was added to the **reduction** clause (see Section 2.22.5.4 on page 302).
- 24 • The **affinity** clause was added to the **task** construct (see Section 2.13.1 on page 142) to  
25 support hints that indicate data affinity of explicit tasks.
- 26 • The **detach** clause for the **task** construct (see Section 2.13.1 on page 142) and the  
27 **omp\_fulfill\_event** runtime routine (see Section 3.5.1 on page 395) were added to support  
28 execution of detachable tasks.
- 29 • To support taskloop reductions, the **taskloop** (see Section 2.13.2 on page 147) and  
30 **taskloop simd** (see Section 2.13.3 on page 152) constructs were extended to accept the  
31 **reduction** (see Section 2.22.5.4 on page 302) and **in\_reduction** (see Section 2.22.5.6 on  
32 page 306) clauses.
- 33 • The **taskloop** construct (see Section 2.13.2 on page 147) was added to the list of constructs  
34 that can be canceled by the **cancel** construct (see Section 2.21.1 on page 265)).
- 35 • To support mutually exclusive inout sets, a **mutexinoutset dependence-type** was added to  
36 the **depend** clause (see Section 2.13.6 on page 156 and Section 2.20.11 on page 257).

- 1       • Predefined memory spaces (see Section 2.14.1 on page 158), predefined memory allocators and  
2       allocator traits (see Section 2.14.2 on page 159) and directives, clauses (see Section 2.14 on  
3       page 158 and API routines (see Section 3.8 on page 413) to use them were added to support  
4       different kinds of memories.
- 5       • The semantics of the **use\_device\_ptr** clause for pointer variables was clarified and the  
6       **use\_device\_addr** clause for using the device address of non-pointer variables inside the  
7       **target data** construct was added (see Section 2.15.2 on page 166).
- 8       • To support reverse offload, the **ancestor** modifier was added to the **device** clause for  
9       **target** constructs (see Section 2.15.5 on page 174).
- 10      • To reduce programmer effort implicit declare target directives for some functions (C, C++,  
11      Fortran) and subroutines (Fortran) were added (see Section 2.15.5 on page 174 and  
12      Section 2.15.7 on page 185).
- 13      • The **target update** construct (see Section 2.15.6 on page 181) was modified to allow array  
14      sections that specify discontinuous storage.
- 15      • The **to** and **from** clauses on the **target update** construct (see Section 2.15.6 on page 181),  
16      the **depend** clause on task generating constructs (see Section 2.20.11 on page 257), and the  
17      **map** clause (see Section 2.22.7.1 on page 318) were extended to allow any lvalue expression as a  
18      list item for C/C++.
- 19      • Support for nested **declare target** directives was added (see Section 2.15.7 on page 185).
- 20      • New combined constructs **master taskloop** (see Section 2.16.7 on page 199),  
21      **parallel master** (see Section 2.16.6 on page 198), **parallel master taskloop** (see  
22      Section 2.16.9 on page 201), **master taskloop simd** (see Section 2.16.8 on page 200),  
23      **parallel master taskloop simd** (see Section 2.16.10 on page 202) were added.
- 24      • The **depend** clause was added to the **taskwait** construct (see Section 2.20.5 on page 233).
- 25      • To support acquire and release semantics with weak memory ordering, the **acq\_rel**,  
26      **acquire**, and **release** clauses were added to the **atomic** construct (see Section 2.20.7 on  
27      page 237) and **flush** construct (see Section 2.20.8 on page 245), and the memory ordering  
28      semantics of implicit flushes on various constructs and runtime routines were clarified (see  
29      Section 2.20.8.1 on page 249).
- 30      • The **atomic** construct was extended with the **hint** clause (see Section 2.20.7 on page 237).
- 31      • The **depend** clause (see Section 2.20.11 on page 257) was extended to support iterators and to  
32      support depend objects that can be created with the new **depobj** construct.
- 33      • Lock hints were renamed to synchronization hints, and the old names were deprecated (see  
34      Section 2.20.12 on page 262).
- 35      • To support conditional assignment to lastprivate variables, the **conditional** modifier was  
36      added to the **lastprivate** clause (see Section 2.22.4.5 on page 290).

- 1 • The description of the **map** clause was modified to clarify the mapping order when multiple  
2 *map-types* are specified for a variable or structure members of a variable on the same construct.  
3 The **close map-type-modifier** was added as a hint for the runtime to allocate memory close to  
4 the target device (see Section 2.22.7.1 on page 318).
- 5 • The capability to map C/C++ pointer variables and to assign the address of device memory that  
6 is mapped by an array section to them was added. Support for mapping of Fortran pointer and  
7 allocatable variables, including pointer and allocatable components of variables, was added (see  
8 Section 2.22.7.1 on page 318).
- 9 • The **defaultmap** clause (see Section 2.22.7.3 on page 328) was extended to allow selecting  
10 the data-mapping or data-sharing attributes for any of the scalar, aggregate, pointer or allocatable  
11 classes on a per-region basis. Additionally it accepts the **none** parameter to support the  
12 requirement that all variables referenced in the construct must be explicitly mapped or privatized.
- 13 • The **declare mapper** directive was added to support mapping of data types with direct and  
14 indirect members (see Section 2.22.7.4 on page 330).
- 15 • The **omp\_set\_nested** (see Section 3.2.10 on page 345) and **omp\_get\_nested** (see  
16 Section 3.2.11 on page 346) routines and the **OMP\_NESTED** environment variable (see  
17 Section 6.9 on page 593) were deprecated.
- 18 • The **omp\_get\_supported\_active\_levels** routine was added to query the number of  
19 active levels of parallelism supported by the implementation (see Section 3.2.15 on page 351).
- 20 • Runtime routines **omp\_set\_affinity\_format** (see Section 3.2.30 on page 364),  
21 **omp\_get\_affinity\_format** (see Section 3.2.31 on page 365), **omp\_set\_affinity**  
22 (see Section 3.2.32 on page 366), and **omp\_capture\_affinity** (see Section 3.2.33 on  
23 page 367) and environment variables **OMP\_DISPLAY\_AFFINITY** (see Section 6.13 on  
24 page 595) and **OMP\_AFFINITY\_FORMAT** (see Section 6.14 on page 596) were added to  
25 provide OpenMP runtime thread affinity information.
- 26 • The **omp\_get\_device\_num** runtime routine (see Section 3.2.37 on page 371) was added to  
27 support determination of the device on which a thread is executing.
- 28 • The **omp\_pause\_resource** and **omp\_pause\_resource\_all** runtime routines were  
29 added to allow the runtime to relinquish resources used by OpenMP (see Section 3.2.43 on  
30 page 375 and Section 3.2.44 on page 377).
- 31 • Support for a first-party tool interface (see Section 4 on page 425) was added.
- 32 • Support for a third-party tool interface (see Section 5 on page 529) was added.
- 33 • Support for controlling offloading behavior with the **OMP\_TARGET\_OFFLOAD** environment  
34 variable was added (see Section 6.17 on page 599).
- 35 • Stubs for Runtime Library Routines (previously Appendix A) were moved to a separate document.
- 36 • Interface Declarations (previously Appendix B) were moved to a separate document.

## B.4 Version 4.0 to 4.5 Differences

- Support for several features of Fortran 2003 was added (see Section 1.7 on page 32 for features that are still not supported).
- A parameter was added to the **ordered** clause of the worksharing-loop construct (see Section 2.12.2 on page 106) and clauses were added to the **ordered** construct (see Section 2.20.9 on page 252) to support doacross loop nests and use of the **simd** construct on loops with loop-carried backward dependences.
- The **linear** clause was added to the worksharing-loop construct (see Section 2.12.2 on page 106).
- The **simdlen** clause was added to the **simd** construct (see Section 2.12.3.1 on page 115) to support specification of the exact number of iterations desired per SIMD chunk.
- The **priority** clause was added to the **task** construct (see Section 2.13.1 on page 142) to support hints that specify the relative execution priority of explicit tasks. The **omp\_get\_max\_task\_priority** routine was added to return the maximum supported priority value (see Section 3.2.42 on page 374) and the **OMP\_MAX\_TASK\_PRIORITY** environment variable was added to control the maximum priority value allowed (see Section 6.16 on page 598).
- Taskloop constructs (see Section 2.13.2 on page 147 and Section 2.13.3 on page 152) were added to support nestable parallel loops that create OpenMP tasks.
- To support interaction with native device implementations, the **use\_device\_ptr** clause was added to the **target data** construct (see Section 2.15.2 on page 166) and the **is\_device\_ptr** clause was added to the **target** construct (see Section 2.15.5 on page 174).
- The **nowait** and **depend** clauses were added to the **target** construct (see Section 2.15.5 on page 174) to improve support for asynchronous execution of **target** regions.
- The **private**, **firstprivate** and **defaultmap** clauses were added to the **target** construct (see Section 2.15.5 on page 174).
- The **declare target** directive was extended to allow mapping of global variables to be deferred to specific device executions and to allow an *extended-list* to be specified in C/C++ (see Section 2.15.7 on page 185).
- To support unstructured data mapping for devices, the **target enter data** (see Section 2.15.3 on page 169) and **target exit data** (see Section 2.15.4 on page 172) constructs were added and the **map** clause (see Section 2.22.7.1 on page 318) was updated.
- To support a more complete set of device construct shortcuts, the **target parallel** (see Section 2.16.16 on page 208), target parallel worksharing-loop (see Section 2.16.17 on page 210), target parallel worksharing-loop SIMD (see Section 2.16.18 on page 211), and **target simd** (see Section 2.16.20 on page 213), combined constructs were added.

- 1 • The **if** clause was extended to take a *directive-name-modifier* that allows it to apply to combined  
2 constructs (see Section 2.18 on page 224).
- 3 • The **hint** clause was added to the **critical** construct (see Section 2.20.1 on page 226).
- 4 • The **source** and **sink** dependence types were added to the **depend** clause (see  
5 Section 2.20.11 on page 257) to support doacross loop nests.
- 6 • The implicit data-sharing attribute for scalar variables in **target** regions was changed to  
7 **firstprivate** (see Section 2.22.1.1 on page 271).
- 8 • Use of some C++ reference types was allowed in some data sharing attribute clauses (see  
9 Section 2.22.4 on page 284).
- 10 • Semantics for reductions on C/C++ array sections were added and restrictions on the use of  
11 arrays and pointers in reductions were removed (see Section 2.22.5.4 on page 302).
- 12 • The **ref**, **val**, and **uval** modifiers were added to the **linear** clause (see Section 2.22.4.6 on  
13 page 292).
- 14 • Support was added to the map clauses to handle structure elements (see Section 2.22.7.1 on  
15 page 318).
- 16 • Query functions for OpenMP thread affinity were added (see Section 3.2.24 on page 359 to  
17 Section 3.2.29 on page 363).
- 18 • The lock API was extended with lock routines that support storing a hint with a lock to select a  
19 desired lock implementation for a lock’s intended usage by the application code (see  
20 Section 3.3.2 on page 385).
- 21 • Device memory routines were added to allow explicit allocation, deallocation, memory transfers  
22 and memory associations (see Section 3.6 on page 396).
- 23 • C/C++ Grammar (previously Appendix B) was moved to a separate document.

## 24 B.5 Version 3.1 to 4.0 Differences

- 25 • Various changes throughout the specification were made to provide initial support of Fortran  
26 2003 (see Section 1.7 on page 32).
- 27 • C/C++ array syntax was extended to support array sections (see Section 2.1.5 on page 42).
- 28 • The **proc\_bind** clause (see Section 2.8.2 on page 85), the **OMP\_PLACES** environment  
29 variable (see Section 6.5 on page 588), and the **omp\_get\_proc\_bind** runtime routine (see  
30 Section 3.2.23 on page 357) were added to support thread affinity policies.
- 31 • SIMD directives were added to support SIMD parallelism (see Section 2.12.3 on page 115).
- 32 • Implementation defined task scheduling points for untied tasks were removed (see Section 2.13.6  
33 on page 156).



- 1 • Device directives (see Section 2.15 on page 165), the `OMP_DEFAULT_DEVICE` environment  
2 variable (see Section 6.15 on page 598), and the `omp_set_default_device`,  
3 `omp_get_default_device`, `omp_get_num_devices`, `omp_get_num_teams`,  
4 `omp_get_team_num`, and `omp_is_initial_device` routines were added to support  
5 execution on devices.
- 6 • The `taskgroup` construct (see Section 2.20.6 on page 235) was added to support more flexible  
7 deep task synchronization.
- 8 • The `atomic` construct (see Section 2.20.7 on page 237) was extended to support atomic swap  
9 with the `capture` clause, to allow new atomic update and capture forms, and to support  
10 sequentially consistent atomic operations with a new `seq_cst` clause.
- 11 • The `depend` clause (see Section 2.20.11 on page 257) was added to support task dependences.
- 12 • The `cancel` construct (see Section 2.21.1 on page 265), the `cancellation point`  
13 construct (see Section 2.21.2 on page 269), the `omp_get_cancellation` runtime routine  
14 (see Section 3.2.9 on page 345) and the `OMP_CANCELLATION` environment variable (see  
15 Section 6.11 on page 594) were added to support the concept of cancellation.
- 16 • The `reduction` clause (see Section 2.22.5.4 on page 302) was extended and the  
17 `declare reduction` construct (see Section 2.22.5.7 on page 307) was added to support user  
18 defined reductions.
- 19 • The `OMP_DISPLAY_ENV` environment variable (see Section 6.12 on page 594) was added to  
20 display the value of ICVs associated with the OpenMP environment variables.
- 21 • Examples (previously Appendix A) were moved to a separate document.

## 22 B.6 Version 3.0 to 3.1 Differences

- 23 • The `bind-var` ICV has been added, which controls whether or not threads are bound to processors  
24 (see Section 2.6.1 on page 66). The value of this ICV can be set with the `OMP_PROC_BIND`  
25 environment variable (see Section 6.4 on page 588).
- 26 • The `nthreads-var` ICV has been modified to be a list of the number of threads to use at each  
27 nested parallel region level and the algorithm for determining the number of threads used in a  
28 parallel region has been modified to handle a list (see Section 2.8.1 on page 84).
- 29 • The `final` and `mergeable` clauses (see Section 2.13.1 on page 142) were added to the `task`  
30 construct to support optimization of task data environments.
- 31 • The `taskyield` construct (see Section 2.13.4 on page 154) was added to allow user-defined  
32 task scheduling points.
- 33 • The `atomic` construct (see Section 2.20.7 on page 237) was extended to include `read`, `write`,  
34 and `capture` forms, and an `update` clause was added to apply the already existing form of the  
35 `atomic` construct.

- Data environment restrictions were changed to allow **intent (in)** and **const**-qualified types for the **firstprivate** clause (see Section 2.22.4.4 on page 288).
- Data environment restrictions were changed to allow Fortran pointers in **firstprivate** (see Section 2.22.4.4 on page 288) and **lastprivate** (see Section 2.22.4.5 on page 290).
- New reduction operators **min** and **max** were added for C and C++ (see Section 2.22.5 on page 295).
- The nesting restrictions in Section 2.23 on page 332 were clarified to disallow closely-nested OpenMP regions within an **atomic** region. This allows an **atomic** region to be consistently defined with other OpenMP regions so that they include all code in the atomic construct.
- The **omp\_in\_final** runtime library routine (see Section 3.2.22 on page 357) was added to support specialization of final task regions.
- Descriptions of examples (previously Appendix A) were expanded and clarified.
- Replaced incorrect use of **omp\_integer\_kind** in Fortran interfaces with **selected\_int\_kind(8)**.

## B.7 Version 2.5 to 3.0 Differences

- The definition of active **parallel** region has been changed: in Version 3.0 a **parallel** region is active if it is executed by a team consisting of more than one thread (see Section 1.2.2 on page 2).
- The concept of tasks has been added to the OpenMP execution model (see Section 1.2.5 on page 11 and Section 1.3 on page 21).
- The OpenMP memory model now covers atomicity of memory accesses (see Section 1.4.1 on page 24). The description of the behavior of **volatile** in terms of **flush** was removed.
- In Version 2.5, there was a single copy of the *nest-var*, *dyn-var*, *nthreads-var* and *run-sched-var* internal control variables (ICVs) for the whole program. In Version 3.0, there is one copy of these ICVs per task (see Section 2.6 on page 66). As a result, the **omp\_set\_num\_threads**, **omp\_set\_nested** and **omp\_set\_dynamic** runtime library routines now have specified effects when called from inside a **parallel** region (see Section 3.2.1 on page 338, Section 3.2.7 on page 343 and Section 3.2.10 on page 345).
- The *thread-limit-var* ICV has been added, which controls the maximum number of threads participating in the OpenMP program. The value of this ICV can be set with the **OMP\_THREAD\_LIMIT** environment variable and retrieved with the **omp\_get\_thread\_limit** runtime library routine (see Section 2.6.1 on page 66, Section 3.2.14 on page 350 and Section 6.10 on page 593).
- The *max-active-levels-var* ICV has been added, which controls the number of nested active **parallel** regions. The value of this ICV can be set with the **OMP\_MAX\_ACTIVE\_LEVELS**

- 1 environment variable and the `omp_set_max_active_levels` runtime library routine, and  
2 it can be retrieved with the `omp_get_max_active_levels` runtime library routine (see  
3 Section 2.6.1 on page 66, Section 3.2.16 on page 352, Section 3.2.17 on page 353 and  
4 Section 6.8 on page 592).
- 5 • The `stacksize-var` ICV has been added, which controls the stack size for threads that the OpenMP  
6 implementation creates. The value of this ICV can be set with the `OMP_STACKSIZE`  
7 environment variable (see Section 2.6.1 on page 66 and Section 6.6 on page 591).
  - 8 • The `wait-policy-var` ICV has been added, which controls the desired behavior of waiting threads.  
9 The value of this ICV can be set with the `OMP_WAIT_POLICY` environment variable (see  
10 Section 2.6.1 on page 66 and Section 6.7 on page 592).
  - 11 • The rules for determining the number of threads used in a `parallel` region have been modified  
12 (see Section 2.8.1 on page 84).
  - 13 • In Version 3.0, the assignment of iterations to threads in a loop construct with a `static`  
14 schedule kind is deterministic (see Section 2.12.2 on page 106).
  - 15 • In Version 3.0, a loop construct may be associated with more than one perfectly nested loop. The  
16 number of associated loops is controlled by the `collapse` clause (see Section 2.12.2 on  
17 page 106).
  - 18 • Random access iterators, and variables of unsigned integer type, may now be used as loop  
19 iterators in loops associated with a loop construct (see Section 2.12.2 on page 106).
  - 20 • The schedule kind `auto` has been added, which gives the implementation the freedom to choose  
21 any possible mapping of iterations in a loop construct to threads in the team (see Section 2.12.2  
22 on page 106).
  - 23 • The `task` construct (see Section 2.13 on page 142) has been added, which provides a  
24 mechanism for creating tasks explicitly.
  - 25 • The `taskwait` construct (see Section 2.20.5 on page 233) has been added, which causes a task  
26 to wait for all its child tasks to complete.
  - 27 • Fortran assumed-size arrays now have predetermined data-sharing attributes (see  
28 Section 2.22.1.1 on page 271).
  - 29 • In Version 3.0, static class members variables may appear in a `threadprivate` directive (see  
30 Section 2.22.2 on page 276).
  - 31 • Version 3.0 makes clear where, and with which arguments, constructors and destructors of  
32 private and threadprivate class type variables are called (see Section 2.22.2 on page 276,  
33 Section 2.22.4.3 on page 287, Section 2.22.4.4 on page 288, Section 2.22.6.1 on page 313 and  
34 Section 2.22.6.2 on page 315).
  - 35 • In Version 3.0, Fortran allocatable arrays may appear in `private`, `firstprivate`,  
36 `lastprivate`, `reduction`, `copyin` and `copyprivate` clauses (see Section 2.22.2 on  
37 page 276, Section 2.22.4.3 on page 287, Section 2.22.4.4 on page 288, Section 2.22.4.5 on

1 page 290, Section 2.22.5.4 on page 302, Section 2.22.6.1 on page 313 and Section 2.22.6.2 on  
2 page 315).

- 3 • In Fortran, **firstprivate** is now permitted as an argument to the **default** clause (see  
4 Section 2.22.4.1 on page 284).
- 5 • For list items in the **private** clause, implementations are no longer permitted to use the storage  
6 of the original list item to hold the new list item on the master thread. If no attempt is made to  
7 reference the original list item inside the **parallel** region, its value is well defined on exit  
8 from the **parallel** region (see Section 2.22.4.3 on page 287).
- 9 • The runtime library routines **omp\_set\_schedule** and **omp\_get\_schedule** have been  
10 added; these routines respectively set and retrieve the value of the *run-sched-var* ICV (see  
11 Section 3.2.12 on page 347 and Section 3.2.13 on page 349).
- 12 • The **omp\_get\_level** runtime library routine has been added, which returns the number of  
13 nested **parallel** regions enclosing the task that contains the call (see Section 3.2.18 on  
14 page 353).
- 15 • The **omp\_get\_ancestor\_thread\_num** runtime library routine has been added, which  
16 returns, for a given nested level of the current thread, the thread number of the ancestor (see  
17 Section 3.2.19 on page 354).
- 18 • The **omp\_get\_team\_size** runtime library routine has been added, which returns, for a given  
19 nested level of the current thread, the size of the thread team to which the ancestor belongs (see  
20 Section 3.2.20 on page 355).
- 21 • The **omp\_get\_active\_level** runtime library routine has been added, which returns the  
22 number of nested active **parallel** regions enclosing the task that contains the call (see  
23 Section 3.2.21 on page 356).
- 24 • In Version 3.0, locks are owned by tasks, not by threads (see Section 3.3 on page 382).

# Index

## Symbols

`_OPENMP` macro, 48, 594–596

## A

acquire flush, 27  
affinity, 85  
**allocate**, 162, 164  
array sections, 42  
array shaping, 42  
**assume**, 63  
**atomic**, 237  
**atomic** construct, 605  
attribute clauses, 284  
attributes, data-mapping, 317  
attributes, data-sharing, 271  
**auto**, 111

## B

**barrier**, 229  
barrier, implicit, 231

## C

**cancel**, 265  
cancellation constructs, 265  
    **cancel**, 265  
    **cancellation point**, 269  
**cancellation point**, 269  
canonical loop form, 101  
**capture**, **atomic**, 237  
clauses  
    **allocate**, 164  
    attribute data-sharing, 284  
    **collapse**, 106, 108  
    **copyin**, 313  
    **copyprivate**, 315

data copying, 313  
data-sharing, 284  
**default**, 284  
**defaultmap**, 328  
**depend**, 257  
**firstprivate**, 288  
hint, 262  
**if** Clause, 224  
**in\_reduction**, 306  
**lastprivate**, 290  
**linear**, 292  
**map**, 318  
**private**, 287  
**reduction**, 302  
**schedule**, 108  
**shared**, 285  
**task\_reduction**, 305  
combined constructs, 193  
    **master taskloop**, 199  
    **master taskloop simd**, 200  
    **parallel loop**, 194  
    **parallel master**, 198  
    **parallel master taskloop**, 201  
    **parallel master taskloop simd**,  
        202  
    **parallel sections**, 195  
    **parallel workshare**, 196  
    parallel worksharing-loop construct,  
        193  
    parallel worksharing-loop SIMD  
        construct, 197  
    **target parallel**, 208  
    **target parallel loop**, 212  
    target parallel worksharing-loop  
        construct, 210

- target parallel worksharing-loop SIMD
  - construct, 211
- target simd**, 213
- target teams**, 215
- target teams distribute**, 215
- target teams distribute parallel
  - worksharing-loop construct, 219
- target teams distribute parallel
  - worksharing-loop SIMD construct, 220
- target teams distribute simd**, 216
- target teams loop construct, 218
- teams distribute**, 203
- teams distribute parallel
  - worksharing-loop construct, 205
- teams distribute parallel
  - worksharing-loop SIMD construct, 206
- teams distribute simd**, 204
- teams loop**, 207
- compilation sentinels, 49
- compliance, 31
- conditional compilation, 48
- constructs
  - atomic**, 237
  - barrier**, 229
  - cancel**, 265
  - cancellation constructs, 265
  - cancellation point**, 269
  - combined constructs, 193
  - critical**, 226
  - declare mapper**, 330
  - depobj**, 256
  - device constructs, 165
  - distribute**, 125
  - distribute parallel do**, 130
  - distribute parallel do simd**, 131
  - distribute parallel for**, 130
  - distribute parallel for simd**, 131
  - distribute parallel worksharing-loop
    - construct, 130
  - distribute parallel worksharing-loop
    - SIMD construct, 131
  - distribute simd**, 129
  - do Fortran**, 106
  - flush**, 245
  - for, C/C++**, 106
  - interop**, 191
  - loop**, 133
  - master**, 225
  - master taskloop**, 199
  - master taskloop simd**, 200
  - ordered**, 252
  - parallel**, 79
  - parallel do Fortran**, 193
  - parallel for C/C++**, 193
  - parallel loop**, 194
  - parallel master**, 198
  - parallel master taskloop**, 201
  - parallel master taskloop simd**, 202
  - parallel sections**, 195
  - parallel workshare**, 196
  - parallel worksharing-loop construct, 193
  - parallel worksharing-loop SIMD
    - construct, 197
  - scope**, 91
  - sections**, 93
  - simd**, 115
  - single**, 96
  - target**, 174
  - target data**, 166
  - target enter data**, 169
  - target exit data**, 172
  - target parallel**, 208
  - target parallel do**, 210
  - target parallel do simd**, 211
  - target parallel for**, 210
  - target parallel for simd**, 211
  - target parallel loop**, 212
  - target parallel worksharing-loop
    - construct, 210

- target parallel worksharing-loop SIMD construct, 211
- target simd**, 213
- target teams**, 215
- target teams distribute**, 215
- target teams distribute parallel worksharing-loop construct, 219
- target teams distribute parallel worksharing-loop SIMD construct, 220
- target teams distribute simd**, 216
- target teams loop**, 218
- target update**, 181
- task**, 142
- taskgroup**, 235
- tasking constructs, 142
- taskloop**, 147
- taskloop simd**, 152
- taskwait**, 233
- taskyield**, 154
- teams**, 87
- teams distribute**, 203
- teams distribute parallel worksharing-loop construct, 205
- teams distribute parallel worksharing-loop SIMD construct, 206
- teams distribute simd**, 204
- teams loop**, 207
- tile**, 140
- workshare**, 98
- worksharing, 93
- worksharing-loop construct, 106
- worksharing-loop SIMD construct, 120
- controlling OpenMP thread affinity, 85
- copyin**, 313
- copyprivate**, 315
- critical**, 226

**D**

- data copying clauses, 313
- data environment, 271
- data terminology, 13

- data-mapping rules and clauses, 317
- data-sharing attribute clauses, 284
- data-sharing attribute rules, 271
- declare mapper**, 330
- declare reduction**, 307
- Declare SIMD, 121
- Declare Target, 185
- declare variant, 56
- default**, 284
- defaultmap**, 328
- depend**, 257
- depend object, 256
- depobj**, 256
- deprecated features, 611
- device constructs
  - declare mapper**, 330
  - device constructs, 165
  - distribute**, 125
  - distribute parallel worksharing-loop construct, 130
  - distribute parallel worksharing-loop SIMD construct, 131
  - distribute simd**, 129
  - target**, 174
  - target update**, 181
  - teams**, 87
- device data environments, 25, 169, 172
- device directives, 165
- device memory routines, 396
- directive format, 36
- directives, 35
  - allocate**, 162
  - assume**, 63
  - declare mapper**, 330
  - declare reduction**, 307
  - Declare SIMD, 121
  - Declare Target, 185
  - declare variant, 56
  - error**, 78
  - memory management directives, 158
  - metadirective**, 54
  - nothing**, 77
  - requires**, 60

- scan** Directive, [137](#)
- threadprivate**, [276](#)
- variant directives, [50](#)
- distribute**, [125](#)
- distribute parallel worksharing-loop
  - construct, [130](#)
- distribute parallel worksharing-loop SIMD
  - construct, [131](#)
- distribute simd**, [129](#)
- do**, *Fortran*, [106](#)
- do simd**, [120](#)
- dynamic**, [110](#)
- dynamic thread adjustment, [603](#)

## E

- environment variables, [585](#)
  - OMP\_AFFINITY\_FORMAT**, [596](#)
  - OMP\_ALLOCATOR**, [601](#)
  - OMP\_CANCELLATION**, [594](#)
  - OMP\_DEBUG**, [600](#)
  - OMP\_DEFAULT\_DEVICE**, [598](#)
  - OMP\_DISPLAY\_AFFINITY**, [595](#)
  - OMP\_DISPLAY\_ENV**, [594](#)
  - OMP\_DYNAMIC**, [587](#)
  - OMP\_MAX\_ACTIVE\_LEVELS**, [592](#)
  - OMP\_MAX\_TASK\_PRIORITY**, [598](#)
  - OMP\_NESTED**, [593](#)
  - OMP\_NUM\_TEAMS**, [601](#)
  - OMP\_NUM\_THREADS**, [586](#)
  - OMP\_PLACES**, [588](#)
  - OMP\_PROC\_BIND**, [588](#)
  - OMP\_SCHEDULE**, [585](#)
  - OMP\_STACKSIZE**, [591](#)
  - OMP\_TARGET\_OFFLOAD**, [599](#)
  - OMP\_TEAMS\_THREAD\_LIMIT**, [601](#)
  - OMP\_THREAD\_LIMIT**, [593](#)
  - OMP\_TOOL**, [599](#)
  - OMP\_TOOL\_LIBRARIES**, [600](#)
  - OMP\_WAIT\_POLICY**, [592](#)
- event, [395](#)
- event callback registration, [430](#)
- event callback signatures, [461](#)
- event routines, [395](#)
- execution environment routines, [338](#)

- execution model, [21](#)

## F

- features history, [611](#)
- firstprivate**, [288](#)
- fixed source form conditional compilation
  - sentinels, [49](#)
- fixed source form directives, [40](#)
- flush**, [245](#)
- flush operation, [26](#)
- flush synchronization, [27](#)
- flush-set, [26](#)
- for**, *C/C++*, [106](#)
- for simd**, [120](#)
- frames, [457](#)
- free source form conditional compilation
  - sentinel, [49](#)
- free source form directives, [40](#)

## G

- glossary, [2](#)
- guided**, [110](#)

## H

- happens before, [27](#)
- header files, [335](#)
- history of features, [611](#)

## I

- ICVs (internal control variables), [66](#)
- if** Clause, [224](#)
- implementation, [603](#)
- implementation terminology, [17](#)
- implicit barrier, [231](#)
- implicit flushes, [249](#)
- in\_reduction**, [306](#)
- include files, [335](#)
- internal control variables, [603](#)
- internal control variables (ICVs), [66](#)
- Interoperability routines, [410](#)
- introduction, [1](#)
- iterators, [46](#)

## L

- lastprivate**, [290](#)



**linear**, 292  
 list item privatization, 281  
 lock routines, 382  
**loop**, 133  
 loop terminology, 9  
 loop transformation constructs, 140

**M**  
**map**, 318  
**master**, 225  
**master taskloop**, 199  
**master taskloop simd**, 200  
 memory allocators, 159  
 memory management, 158  
 memory management directives  
     memory management directives, 158  
 memory management routines, 413  
 memory model, 24  
 memory spaces, 158  
 metadirective, 54  
 modifying and retrieving ICV values, 71  
 modifying ICVs, 68

**N**  
 nesting of regions, 332  
 normative references, 32  
**nothing**, 77

**O**  
**OMP\_AFFINITY\_FORMAT**, 596  
**omp\_alloc**, 419  
**OMP\_ALLOCATOR**, 601  
**OMP\_CANCELLATION**, 594  
**omp\_capture\_affinity**, 367  
**OMP\_DEBUG**, 600  
**OMP\_DEFAULT\_DEVICE**, 598  
**omp\_destroy\_allocator**, 417  
**omp\_destroy\_lock**, 387  
**omp\_destroy\_nest\_lock**, 387  
**OMP\_DISPLAY\_AFFINITY**, 595  
**omp\_display\_affinity**, 366  
**OMP\_DISPLAY\_ENV**, 594  
**OMP\_DYNAMIC**, 587  
**omp\_free**, 420  
**omp\_fulfill\_event**, 395  
**omp\_get\_active\_level**, 356  
**omp\_get\_affinity\_format**, 365  
**omp\_get\_ancestor\_thread\_num**, 354  
**omp\_get\_cancellation**, 345  
**omp\_get\_default\_allocator**, 418  
**omp\_get\_default\_device**, 370  
**omp\_get\_device\_num**, 371  
**omp\_get\_dynamic**, 344  
**omp\_get\_initial\_device**, 374  
**omp\_get\_interop\_property**, 410  
**omp\_get\_interop\_property\_name**,  
     412  
**omp\_get\_level**, 353  
**omp\_get\_mapped\_ptr**, 408  
**omp\_get\_max\_active\_levels**, 353  
**omp\_get\_max\_task\_priority**, 374  
**omp\_get\_max\_teams**, 379  
**omp\_get\_max\_threads**, 340  
**omp\_get\_nested**, 346  
**omp\_get\_num\_devices**, 370  
**omp\_get\_num\_interop\_properties**,  
     410  
**omp\_get\_num\_places**, 359  
**omp\_get\_num\_procs**, 342  
**omp\_get\_num\_teams**, 372  
**omp\_get\_num\_threads**, 339  
**omp\_get\_partition\_num\_places**,  
     362  
**omp\_get\_partition\_place\_nums**,  
     363  
**omp\_get\_place\_num**, 361  
**omp\_get\_place\_num\_procs**, 360  
**omp\_get\_place\_proc\_ids**, 360  
**omp\_get\_proc\_bind**, 357  
**omp\_get\_schedule**, 349  
**omp\_get\_supported\_active**  
     **\_levels**, 351  
**omp\_get\_team\_num**, 372  
**omp\_get\_team\_size**, 355  
**omp\_get\_teams\_thread\_limit**, 381  
**omp\_get\_thread\_limit**, 350  
**omp\_get\_thread\_num**, 341

omp\_get\_wtick, 394  
 omp\_get\_wtime, 393  
 omp\_in\_final, 357  
 omp\_in\_parallel, 342  
 omp\_init\_allocator, 415  
 omp\_init\_lock, 384, 385  
 omp\_init\_nest\_lock, 384, 385  
 omp\_is\_initial\_device, 373  
 OMP\_MAX\_ACTIVE\_LEVELS, 592  
 OMP\_MAX\_TASK\_PRIORITY, 598  
 OMP\_NESTED, 593  
 OMP\_NUM\_TEAMS, 601  
 OMP\_NUM\_THREADS, 586  
 omp\_pause\_resource, 375  
 omp\_pause\_resource\_all, 377  
 OMP\_PLACES, 588  
 OMP\_PROC\_BIND, 588  
 OMP\_SCHEDULE, 585  
 omp\_set\_affinity\_format, 364  
 omp\_set\_default\_allocator, 417  
 omp\_set\_default\_device, 369  
 omp\_set\_dynamic, 343  
 omp\_set\_lock, 388  
 omp\_set\_max\_active\_levels, 352  
 omp\_set\_nest\_lock, 388  
 omp\_set\_nested, 345  
 omp\_set\_num\_teams, 378  
 omp\_set\_num\_threads, 338  
 omp\_set\_schedule, 347  
 omp\_set\_teams\_thread\_limit, 380  
 OMP\_STACKSIZE, 591  
 omp\_target\_alloc, 396  
 omp\_target\_associate\_ptr, 406  
 omp\_target\_disassociate\_ptr, 409  
 omp\_target\_free, 397  
 omp\_target\_is\_present, 398  
 omp\_target\_memcpy, 399  
 omp\_target\_memcpy\_async, 402  
 omp\_target\_memcpy\_rect, 400  
 omp\_target\_memcpy\_rect\_async, 404  
 OMP\_TARGET\_OFFLOAD, 599  
 OMP\_TEAMS\_THREAD\_LIMIT, 601  
 omp\_test\_lock, 391  
 omp\_test\_nest\_lock, 391  
 OMP\_THREAD\_LIMIT, 593  
 OMP\_TOOL, 599  
 OMP\_TOOL\_LIBRARIES, 600  
 omp\_unset\_lock, 389  
 omp\_unset\_nest\_lock, 389  
 OMP\_WAIT\_POLICY, 592  
 ompd\_bp\_device\_begin, 583  
 ompd\_bp\_device\_end, 584  
 ompd\_bp\_parallel\_begin, 580  
 ompd\_bp\_parallel\_end, 580  
 ompd\_bp\_task\_begin, 581  
 ompd\_bp\_task\_end, 582  
 ompd\_bp\_thread\_begin, 582  
 ompd\_bp\_thread\_end, 583  
 ompd\_callback\_device\_host  
     \_fn\_t, 547  
 ompd\_callback\_get\_thread  
     \_context\_for\_thread\_id  
     \_fn\_t, 541  
 ompd\_callback\_memory\_alloc  
     \_fn\_t, 539  
 ompd\_callback\_memory\_free  
     \_fn\_t, 540  
 ompd\_callback\_memory\_read  
     \_fn\_t, 544  
 ompd\_callback\_memory\_write  
     \_fn\_t, 546  
 ompd\_callback\_print\_string  
     \_fn\_t, 548  
 ompd\_callback\_sizeof\_fn\_t, 542  
 ompd\_callback\_symbol\_addr  
     \_fn\_t, 543  
 ompd\_callbacks\_t, 548  
 ompd\_dll\_locations\_valid, 531  
 ompd\_dll\_locations, 530  
 ompt\_callback\_buffer  
     \_complete\_t, 487  
 ompt\_callback\_buffer  
     \_request\_t, 486  
 ompt\_callback\_cancel\_t, 481  
 ompt\_callback\_control

- `_tool_t`, 495
- `ompt_callback_dependences_t`, 470
- `ompt_callback_dispatch_t`, 467
- `ompt_callback_device`
  - `_finalize_t`, 484
- `ompt_callback_device`
  - `_initialize_t`, 482
- `ompt_callback_flush_t`, 481
- `ompt_callback_implicit`
  - `_task_t`, 473
- `ompt_callback_master_t`, 474
- `ompt_callback_mutex`
  - `_acquire_t`, 477
- `ompt_callback_mutex_t`, 478
- `ompt_callback_nest_lock_t`, 479
- `ompt_callback_parallel`
  - `_begin_t`, 463
- `ompt_callback_parallel`
  - `_end_t`, 465
- `ompt_callback_sync_region_t`, 475
- `ompt_callback_device_load_t`, 485
- `ompt_callback_device`
  - `_unload_t`, 486
- `ompt_callback_target_data`
  - `_op_t`, 488
- `ompt_callback_target_map_t`, 491
- `ompt_callback_target`
  - `_submit_t`, 493
- `ompt_callback_target_t`, 490
- `ompt_callback_task_create_t`, 469
- `ompt_callback_task`
  - `_dependence_t`, 471
- `ompt_callback_task`
  - `_schedule_t`, 472
- `ompt_callback_thread`
  - `_begin_t`, 462
- `ompt_callback_thread_end_t`, 462
- `ompt_callback_work_t`, 466
- OpenMP compliance, 31
- `ordered`, 252

**P**

- `parallel`, 79
- `parallel loop`, 194

- `parallel master construct`, 198
- `parallel master taskloop`, 201
- `parallel master taskloop simd`, 202
- `parallel sections`, 195
- `parallel workshare`, 196
- `parallel worksharing-loop construct`, 193
- `parallel worksharing-loop SIMD`
  - `construct`, 197
- `private`, 287

## R

- `read, atomic`, 237
- `reduction`, 302
- reduction clauses, 295
- release flush, 27
- `requires`, 60
- `runtime`, 111
- runtime library definitions, 335
- runtime library routines, 335

## S

- `scan Directive`, 137
- scheduling, 156
- `scope`, 91
- `sections`, 93
- `shared`, 285
- `simd`, 115
- SIMD Directives, 115
- Simple Lock Routines, 383
- `single`, 96
- stand-alone directives, 41
- `static`, 110
- strong flush, 26
- synchronization constructs, 226
- synchronization constructs and clauses, 226
- synchronization hints, 262
- synchronization terminology, 9

## T

- `target`, 174
- `target data`, 166
- target memory routines, 396
- `target parallel`, 208
- `target parallel loop`, 212

- target parallel worksharing-loop construct
  - construct, [210](#)
- target parallel worksharing-loop SIMD
  - construct, [211](#)
- target simd**, [213](#)
- target teams**, [215](#)
- target teams distribute**, [215](#)
- target teams distribute parallel
  - worksharing-loop construct, [219](#)
- target teams distribute parallel
  - worksharing-loop SIMD
    - construct, [220](#)
- target teams distribute simd**, [216](#)
- target teams loop**, [218](#)
- target update**, [181](#)
- task**, [142](#)
- task scheduling, [156](#)
- task\_reduction**, [305](#)
- taskgroup**, [235](#)
- tasking constructs, [142](#)
- tasking terminology, [11](#)
- taskloop**, [147](#)
- taskloop simd**, [152](#)
- taskwait**, [233](#)
- taskyield**, [154](#)
- teams**, [87](#)
- teams distribute**, [203](#)
- teams distribute parallel worksharing-loop
  - construct, [205](#)
- teams distribute parallel worksharing-loop
  - SIMD construct, [206](#)
- teams distribute simd**, [204](#)
- teams loop**, [207](#)
- thread affinity, [85](#)
- threadprivate**, [276](#)
- tile**, [140](#)
- timer, [393](#)
- timing routines, [393](#)
- tool control, [421](#)
- tool initialization, [428](#)
- tool interfaces definitions, [425](#), [530](#)
- tools header files, [425](#), [530](#)
- tracing device activity, [432](#)

## U

- update, atomic**, [237](#)
- utility directives, [77](#)

## V

- variables, environment, [585](#)
- variant directives, [50](#)

## W

- wait identifier, [459](#)
- wall clock timer, [393](#)
- erorr**, [78](#)
- workshare**, [98](#)
- worksharing
  - constructs, [93](#)
  - parallel, [193](#)
  - scheduling, [114](#)
- worksharing constructs, [93](#)
- worksharing-loop construct, [106](#)
- worksharing-loop SIMD construct, [120](#)
- write, atomic**, [237](#)