



OpenMP Technical Report 11: Version 6.0 Preview 1

This Technical Report is the first preview for the OpenMP Application Programming Specification Version 6.0. This version removes features that have been deprecated in versions 5.0, 5.1, and 5.2. This preview extends the loop transforming directives with new loop transformations, improves device support with changes and extensions to the device directives, adds new memory traits for greater control of memory allocation and provides new API routines to define and to query memory spaces. See Appendix B.2 for the complete list of changes relative to version 5.2.

EDITORS

Bronis R. de Supinski

Michael Klemm

November 10, 2022

Expires November 9, 2023

We actively solicit comments. Please provide feedback on this document either to the editors directly or by emailing to info@openmp.org

OpenMP Architecture Review Board – www.openmp.org – info@openmp.org
OpenMP ARB, 9450 SW Gemini Dr., PMB 63140, Beaverton, OR 77008, 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, supports 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 6.0 Preview 1, November 2022

Copyright ©1997-2022 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 page intentionally left blank in published version.

This is a preview of the OpenMP API Specification Version 6.0 and includes the following internal GitHub issues (corresponding Trac ticket numbers in parentheses when they exist) applied to the 5.2 LaTeX sources: 1495 (563), 1584 (652), 1843 (911), 1935, 2019, 2186-2187, 2691, 2721, 2734, 2736, 2740, 2757, 2902, 3027, 3058, 3161, 3168, 3171, 3180, 3184-3185, 3189-3190, 3193, 3199, 3201, 3206-3210, 3216-3217, 3220 3222, 3229-3230, 3232, 3238, 3240-3241, 3245, 3255, 3259, 3267-3268, 3278-3279, 3281, 3285, 3290, 3293, 3296, 3303, 3327, 3335-3337, 3341, 3345, 3347, 3351, 3353, 3367, 3396

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

Contents

1	Overview of the OpenMP API	1
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	10
1.2.5	Tasking Terminology	12
1.2.6	Data Terminology	14
1.2.7	Implementation Terminology	21
1.2.8	Tool Terminology	22
1.3	Execution Model	24
1.4	Memory Model	28
1.4.1	Structure of the OpenMP Memory Model	28
1.4.2	Device Data Environments	29
1.4.3	Memory Management	30
1.4.4	The Flush Operation	30
1.4.5	Flush Synchronization and <i>Happens Before</i>	32
1.4.6	OpenMP Memory Consistency	33
1.5	Tool Interfaces	34
1.5.1	OMPT	34
1.5.2	OMPD	35
1.6	OpenMP Compliance	36
1.7	Normative References	36
1.8	Organization of this Document	38

2	Internal Control Variables	40
2.1	ICV Descriptions	40
2.2	ICV Initialization	42
2.3	Modifying and Retrieving ICV Values	45
2.4	How the Per-Data Environment ICVs Work	48
2.5	ICV Override Relationships	49
3	Directive and Construct Syntax	50
3.1	Directive Format	51
3.1.1	Fixed Source Form Directives	57
3.1.2	Free Source Form Directives	58
3.2	Clause Format	59
3.2.1	OpenMP Argument Lists	63
3.2.2	Reserved Locators	65
3.2.3	OpenMP Operations	65
3.2.4	Array Shaping	66
3.2.5	Array Sections	67
3.2.6	iterator Modifier	70
3.3	Conditional Compilation	73
3.3.1	Fixed Source Form Conditional Compilation Sentinels	73
3.3.2	Free Source Form Conditional Compilation Sentinel	74
3.4	if Clause	75
3.5	destroy Clause	76
4	Base Language Formats and Restrictions	77
4.1	OpenMP Types and Identifiers	77
4.2	OpenMP Stylized Expressions	79
4.3	Structured Blocks	79
4.3.1	OpenMP Context-Specific Structured Blocks	80
4.4	Loop Concepts	88
4.4.1	Canonical Loop Nest Form	88
4.4.2	OpenMP Loop-Iteration Spaces and Vectors	94
4.4.3	collapse Clause	96
4.4.4	ordered Clause	97

4.4.5	Consistent Loop Schedules	98
5	Data Environment	99
5.1	Data-Sharing Attribute Rules	99
5.1.1	Variables Referenced in a Construct	99
5.1.2	Variables Referenced in a Region but not in a Construct	103
5.2	threadprivate Directive	104
5.3	List Item Privatization	109
5.4	Data-Sharing Attribute Clauses	111
5.4.1	default Clause	112
5.4.2	shared Clause	113
5.4.3	private Clause	113
5.4.4	firstprivate Clause	114
5.4.5	lastprivate Clause	117
5.4.6	linear Clause	120
5.4.7	is_device_ptr Clause	123
5.4.8	use_device_ptr Clause	123
5.4.9	has_device_addr Clause	124
5.4.10	use_device_addr Clause	125
5.5	Reduction Clauses and Directives	126
5.5.1	OpenMP Reduction Identifiers	126
5.5.2	OpenMP Reduction Expressions	126
5.5.3	Implicitly Declared OpenMP Reduction Identifiers	130
5.5.4	initializer Clause	132
5.5.5	Properties Common to All Reduction Clauses	132
5.5.6	Reduction Scoping Clauses	135
5.5.7	Reduction Participating Clauses	136
5.5.8	reduction Clause	136
5.5.9	task_reduction Clause	139
5.5.10	in_reduction Clause	140
5.5.11	declare reduction Directive	141
5.6	scan Directive	143
5.6.1	inclusive Clause	145
5.6.2	exclusive Clause	146

5.7	Data Copying Clauses	146
5.7.1	copyin Clause	146
5.7.2	copyprivate Clause	148
5.8	Data-Mapping Control	150
5.8.1	Implicit Data-Mapping Attribute Rules	150
5.8.2	Mapper Identifiers and mapper Modifiers	152
5.8.3	map Clause	152
5.8.4	enter Clause	161
5.8.5	link Clause	162
5.8.6	defaultmap Clause	162
5.8.7	declare mapper Directive	164
5.9	Data-Motion Clauses	167
5.9.1	to Clause	169
5.9.2	from Clause	169
5.10	uniform Clause	170
5.11	aligned Clause	171
5.12	groupprivate Directive	172
5.13	local Clause	174
6	Memory Management	175
6.1	Memory Spaces	175
6.2	Memory Allocators	176
6.3	align Clause	179
6.4	allocator Clause	180
6.5	allocate Directive	180
6.6	allocate Clause	182
6.7	allocators Construct	184
6.8	uses_allocators Clause	185
7	Variant Directives	187
7.1	OpenMP Contexts	187
7.2	Context Selectors	189
7.3	Matching and Scoring Context Selectors	192

7.4	Metadirectives	193
7.4.1	when Clause	194
7.4.2	otherwise Clause	195
7.4.3	metadirective	195
7.4.4	begin metadirective	196
7.5	Declare Variant Directives	197
7.5.1	match Clause	198
7.5.2	adjust_args Clause	199
7.5.3	append_args Clause	200
7.5.4	declare variant Directive	201
7.5.5	begin declare variant Directive	202
7.6	dispatch Construct	204
7.6.1	interop Clause	205
7.6.2	novariants Clause	206
7.6.3	nocontext Clause	206
7.7	declare simd Directive	207
7.7.1	<i>branch</i> Clauses	209
7.8	Declare Target Directives	209
7.8.1	declare target Directive	211
7.8.2	begin declare target Directive	213
7.8.3	indirect Clause	214
8	Informational and Utility Directives	216
8.1	at Clause	216
8.2	requires Directive	216
8.2.1	<i>requirement</i> Clauses	218
8.3	Assumption Directives	219
8.3.1	<i>assumption</i> Clauses	220
8.3.2	assumes Directive	221
8.3.3	assume Directive	221
8.3.4	begin assumes Directive	222
8.4	nothing Directive	222
8.5	error Directive	222
8.5.1	severity Clause	223

8.5.2	message Clause	224
9	Loop-Transforming Constructs	225
9.1	tile Construct	225
9.1.1	sizes Clause	226
9.2	unroll Construct	227
9.2.1	full Clause	227
9.2.2	partial Clause	228
9.3	reverse Construct	228
9.4	interchange Construct	229
9.4.1	permutation Clause	230
9.5	apply Clause	230
10	Parallelism Generation and Control	232
10.1	omp_curr_progress_width Identifier	232
10.2	parallel Construct	232
10.2.1	Determining the Number of Threads for a parallel Region	236
10.2.2	num_threads Clause	236
10.2.3	Controlling OpenMP Thread Affinity	237
10.2.4	proc_bind Clause	239
10.2.5	safesync Clause	239
10.3	teams Construct	240
10.3.1	num_teams Clause	243
10.4	order Clause	243
10.5	simd Construct	245
10.5.1	nontemporal Clause	246
10.5.2	safelen Clause	247
10.5.3	simdlen Clause	247
10.6	masked Construct	248
10.6.1	filter Clause	249
11	Work-Distribution Constructs	250
11.1	single Construct	250
11.2	scope Construct	252

11.3	sections Construct	253
11.3.1	section Directive	254
11.4	workshare Construct	255
11.5	Worksharing-Loop Constructs	258
11.5.1	for Construct	260
11.5.2	do Construct	261
11.5.3	schedule Clause	262
11.6	distribute Construct	264
11.6.1	dist_schedule Clause	266
11.7	loop Construct	267
11.7.1	bind Clause	268
12	Tasking Constructs	270
12.1	untied Clause	270
12.2	mergeable Clause	270
12.3	final Clause	271
12.4	priority Clause	271
12.5	task Construct	272
12.5.1	affinity Clause	274
12.5.2	detach Clause	275
12.6	taskloop Construct	276
12.6.1	grainsize Clause	279
12.6.2	num_tasks Clause	279
12.7	taskyield Construct	280
12.8	Initial Task	281
12.9	Task Scheduling	282
13	Device Directives and Clauses	285
13.1	device_type Clause	285
13.2	device Clause	286
13.3	thread_limit Clause	287
13.4	Device Initialization	288
13.5	target data Construct	289
13.6	target enter data Construct	290

13.7	target exit data Construct	292
13.8	target Construct	294
13.9	target update Construct	299
14	Interoperability	301
14.1	interop Construct	301
14.1.1	OpenMP Foreign Runtime Identifiers	303
14.1.2	init Clause	303
14.1.3	use Clause	304
14.2	Interoperability Requirement Set	304
15	Synchronization Constructs and Clauses	306
15.1	Synchronization Hints	306
15.1.1	Synchronization Hint Type	306
15.1.2	hint Clause	308
15.2	critical Construct	309
15.3	Barriers	310
15.3.1	barrier Construct	310
15.3.2	Implicit Barriers	312
15.3.3	Implementation-Specific Barriers	313
15.4	taskgroup Construct	314
15.5	taskwait Construct	315
15.6	nowait Clause	317
15.7	nogroup Clause	318
15.8	OpenMP Memory Ordering	319
15.8.1	<i>memory-order</i> Clauses	319
15.8.2	<i>atomic</i> Clauses	319
15.8.3	<i>extended-atomic</i> Clauses	319
15.8.4	memscope Clause	320
15.8.5	atomic Construct	321
15.8.6	flush Construct	325
15.8.7	Implicit Flushes	327
15.9	OpenMP Dependences	331
15.9.1	<i>task-dependence-type</i> Modifier	331

15.9.2	Depend Objects	332
15.9.3	update Clause	332
15.9.4	depobj Construct	333
15.9.5	depend Clause	334
15.9.6	doacross Clause	337
15.10	ordered Construct	339
15.10.1	Stand-alone ordered Construct	340
15.10.2	Block-associated ordered Construct	341
15.10.3	<i>parallelization-level</i> Clauses	342
16	Cancellation Constructs	343
16.1	cancel Construct	343
16.2	cancellation point Construct	347
17	Composition of Constructs	349
17.1	Nesting of Regions	349
17.2	Clauses on Combined and Composite Constructs	350
17.3	Combined and Composite Directive Names	353
17.4	Combined Construct Semantics	354
17.5	Composite Construct Semantics	355
18	Runtime Library Routines	356
18.1	Runtime Library Definitions	356
18.2	Thread Team Routines	359
18.2.1	omp_set_num_threads	359
18.2.2	omp_get_num_threads	360
18.2.3	omp_get_max_threads	361
18.2.4	omp_get_thread_num	362
18.2.5	omp_in_parallel	362
18.2.6	omp_set_dynamic	363
18.2.7	omp_get_dynamic	364
18.2.8	omp_get_cancellation	364
18.2.9	omp_set_schedule	365
18.2.10	omp_get_schedule	367
18.2.11	omp_get_thread_limit	367

18.2.12	<code>omp_get_supported_active_levels</code>	368
18.2.13	<code>omp_set_max_active_levels</code>	368
18.2.14	<code>omp_get_max_active_levels</code>	369
18.2.15	<code>omp_get_level</code>	370
18.2.16	<code>omp_get_ancestor_thread_num</code>	370
18.2.17	<code>omp_get_team_size</code>	371
18.2.18	<code>omp_get_active_level</code>	372
18.3	Thread Affinity Routines	373
18.3.1	<code>omp_get_proc_bind</code>	373
18.3.2	<code>omp_get_num_places</code>	374
18.3.3	<code>omp_get_place_num_procs</code>	375
18.3.4	<code>omp_get_place_proc_ids</code>	376
18.3.5	<code>omp_get_place_num</code>	376
18.3.6	<code>omp_get_partition_num_places</code>	377
18.3.7	<code>omp_get_partition_place_nums</code>	378
18.3.8	<code>omp_set_affinity_format</code>	378
18.3.9	<code>omp_get_affinity_format</code>	379
18.3.10	<code>omp_display_affinity</code>	381
18.3.11	<code>omp_capture_affinity</code>	381
18.4	Teams Region Routines	382
18.4.1	<code>omp_get_num_teams</code>	383
18.4.2	<code>omp_get_team_num</code>	383
18.4.3	<code>omp_set_num_teams</code>	384
18.4.4	<code>omp_get_max_teams</code>	385
18.4.5	<code>omp_set_teams_thread_limit</code>	385
18.4.6	<code>omp_get_teams_thread_limit</code>	386
18.5	Tasking Routines	387
18.5.1	<code>omp_get_max_task_priority</code>	387
18.5.2	<code>omp_in_explicit_task</code>	388
18.5.3	<code>omp_in_final</code>	388
18.6	Resource Relinquishing Routines	389
18.6.1	<code>omp_pause_resource</code>	389
18.6.2	<code>omp_pause_resource_all</code>	391

18.7	Device Information Routines	391
18.7.1	<code>omp_get_num_procs</code>	392
18.7.2	<code>omp_get_max_progress_width</code>	392
18.7.3	<code>omp_set_default_device</code>	393
18.7.4	<code>omp_get_default_device</code>	393
18.7.5	<code>omp_get_num_devices</code>	394
18.7.6	<code>omp_get_device_num</code>	395
18.7.7	<code>omp_is_initial_device</code>	395
18.7.8	<code>omp_get_initial_device</code>	396
18.8	Device Memory Routines	396
18.8.1	<code>omp_target_alloc</code>	396
18.8.2	<code>omp_target_free</code>	398
18.8.3	<code>omp_target_is_present</code>	400
18.8.4	<code>omp_target_is_accessible</code>	401
18.8.5	<code>omp_target_memcpy</code>	402
18.8.6	<code>omp_target_memcpy_rect</code>	403
18.8.7	<code>omp_target_memcpy_async</code>	405
18.8.8	<code>omp_target_memcpy_rect_async</code>	407
18.8.9	<code>omp_target_associate_ptr</code>	410
18.8.10	<code>omp_target_disassociate_ptr</code>	412
18.8.11	<code>omp_get_mapped_ptr</code>	414
18.9	Lock Routines	415
18.9.1	<code>omp_init_lock</code> and <code>omp_init_nest_lock</code>	417
18.9.2	<code>omp_init_lock_with_hint</code> and <code>omp_init_nest_lock_with_hint</code>	418
18.9.3	<code>omp_destroy_lock</code> and <code>omp_destroy_nest_lock</code>	419
18.9.4	<code>omp_set_lock</code> and <code>omp_set_nest_lock</code>	420
18.9.5	<code>omp_unset_lock</code> and <code>omp_unset_nest_lock</code>	422
18.9.6	<code>omp_test_lock</code> and <code>omp_test_nest_lock</code>	423
18.10	Timing Routines	425
18.10.1	<code>omp_get_wtime</code>	425
18.10.2	<code>omp_get_wtick</code>	425

18.11	Event Routine	426
18.11.1	<code>omp_fulfill_event</code>	426
18.12	Interoperability Routines	427
18.12.1	<code>omp_get_num_interop_properties</code>	429
18.12.2	<code>omp_get_interop_int</code>	429
18.12.3	<code>omp_get_interop_ptr</code>	430
18.12.4	<code>omp_get_interop_str</code>	431
18.12.5	<code>omp_get_interop_name</code>	431
18.12.6	<code>omp_get_interop_type_desc</code>	432
18.12.7	<code>omp_get_interop_rc_desc</code>	433
18.13	Memory Management Routines	433
18.13.1	Memory Management Types	434
18.13.2	Routines to get memory spaces	437
18.13.3	<code>omp_init_allocator</code>	439
18.13.4	Routines to Get Allocators	441
18.13.5	<code>omp_destroy_allocator</code>	443
18.13.6	<code>omp_set_default_allocator</code>	444
18.13.7	<code>omp_get_default_allocator</code>	445
18.13.8	<code>omp_alloc</code> and <code>omp_aligned_alloc</code>	446
18.13.9	<code>omp_free</code>	448
18.13.10	<code>omp_calloc</code> and <code>omp_aligned_calloc</code>	449
18.13.11	<code>omp_realloc</code>	451
18.13.12	<code>omp_get_memspace_num_resources</code>	453
18.13.13	<code>omp_get_submemspace</code>	454
18.14	Tool Control Routine	455
18.15	Environment Display Routine	458
19	OMPT Interface	460
19.1	OMPT Interfaces Definitions	460
19.2	Activating a First-Party Tool	460
19.2.1	<code>ompt_start_tool</code>	460
19.2.2	Determining Whether a First-Party Tool Should be Initialized	462
19.2.3	Initializing a First-Party Tool	463
19.2.4	Monitoring Activity on the Host with OMPT	466

19.2.5	Tracing Activity on Target Devices with OMPT	467
19.3	Finalizing a First-Party Tool	471
19.4	OMPT Data Types	471
19.4.1	Tool Initialization and Finalization	472
19.4.2	Callbacks	472
19.4.3	Tracing	474
19.4.4	Miscellaneous Type Definitions	476
19.5	OMPT Tool Callback Signatures and Trace Records	493
19.5.1	Initialization and Finalization Callback Signature	493
19.5.2	Event Callback Signatures and Trace Records	495
19.6	OMPT Runtime Entry Points for Tools	531
19.6.1	Entry Points in the OMPT Callback Interface	532
19.6.2	Entry Points in the OMPT Device Tracing Interface	548
19.6.3	Lookup Entry Points: ompt_function_lookup_t	560
20	OMPD Interface	562
20.1	OMPD Interfaces Definitions	563
20.2	Activating a Third-Party Tool	563
20.2.1	Enabling Runtime Support for OMPD	563
20.2.2	ompd_dll_locations	563
20.2.3	ompd_dll_locations_valid	564
20.3	OMPD Data Types	565
20.3.1	Size Type	565
20.3.2	Wait ID Type	565
20.3.3	Basic Value Types	566
20.3.4	Address Type	566
20.3.5	Frame Information Type	566
20.3.6	System Device Identifiers	567
20.3.7	Native Thread Identifiers	568
20.3.8	OMPD Handle Types	568
20.3.9	OMPD Scope Types	569
20.3.10	Team Generator Types	570
20.3.11	ICV ID Type	570
20.3.12	Tool Context Types	571

20.3.13	Return Code Types	572
20.3.14	Primitive Type Sizes	573
20.4	OMPD Third-Party Tool Callback Interface	574
20.4.1	Memory Management of OMPD Library	574
20.4.2	Context Management and Navigation	576
20.4.3	Accessing Memory in the OpenMP Program or Runtime	579
20.4.4	Data Format Conversion: <code>ompd_callback_device_host_fn_t</code>	583
20.4.5	<code>ompd_callback_print_string_fn_t</code>	584
20.4.6	The Callback Interface	585
20.5	OMPD Tool Interface Routines	586
20.5.1	Per OMPD Library Initialization and Finalization	587
20.5.2	Per OpenMP Process Initialization and Finalization	590
20.5.3	Thread and Signal Safety	594
20.5.4	Address Space Information	594
20.5.5	Thread Handles	596
20.5.6	Parallel Region Handles	601
20.5.7	Task Handles	605
20.5.8	Querying Thread States	612
20.5.9	Display Control Variables	614
20.5.10	Accessing Scope-Specific Information	616
20.6	Breakpoint Symbol Names for OMPD	620
20.6.1	Beginning Parallel Regions	620
20.6.2	Ending Parallel Regions	621
20.6.3	Beginning Teams Regions	622
20.6.4	Ending Teams Regions	622
20.6.5	Beginning Task Regions	623
20.6.6	Ending Task Regions	623
20.6.7	Beginning OpenMP Threads	624
20.6.8	Ending OpenMP Threads	624
20.6.9	Beginning Target Regions	625
20.6.10	Ending Target Regions	625
20.6.11	Initializing OpenMP Devices	626
20.6.12	Finalizing OpenMP Devices	626

21	Environment Variables	628
21.1	Parallel Region Environment Variables	629
21.1.1	<code>OMP_DYNAMIC</code>	629
21.1.2	<code>OMP_NUM_THREADS</code>	629
21.1.3	<code>OMP_THREAD_LIMIT</code>	630
21.1.4	<code>OMP_MAX_ACTIVE_LEVELS</code>	630
21.1.5	<code>OMP_PLACES</code>	631
21.1.6	<code>OMP_PROC_BIND</code>	633
21.2	Program Execution Environment Variables	634
21.2.1	<code>OMP_SCHEDULE</code>	634
21.2.2	<code>OMP_STACKSIZE</code>	635
21.2.3	<code>OMP_WAIT_POLICY</code>	635
21.2.4	<code>OMP_DISPLAY_AFFINITY</code>	636
21.2.5	<code>OMP_AFFINITY_FORMAT</code>	637
21.2.6	<code>OMP_CANCELLATION</code>	639
21.2.7	<code>OMP_AVAILABLE_DEVICES</code>	639
21.2.8	<code>OMP_DEFAULT_DEVICE</code>	640
21.2.9	<code>OMP_TARGET_OFFLOAD</code>	641
21.2.10	<code>OMP_MAX_TASK_PRIORITY</code>	641
21.3	OMPT Environment Variables	642
21.3.1	<code>OMP_TOOL</code>	642
21.3.2	<code>OMP_TOOL_LIBRARIES</code>	642
21.3.3	<code>OMP_TOOL_VERBOSE_INIT</code>	643
21.4	OMPD Environment Variables	644
21.4.1	<code>OMP_DEBUG</code>	644
21.5	Memory Allocation Environment Variables	644
21.5.1	<code>OMP_ALLOCATOR</code>	644
21.6	Teams Environment Variables	645
21.6.1	<code>OMP_NUM_TEAMS</code>	645
21.6.2	<code>OMP_TEAMS_THREAD_LIMIT</code>	646
21.7	<code>OMP_DISPLAY_ENV</code>	646
A	OpenMP Implementation-Defined Behaviors	647

B	Features History	656
B.1	Deprecated Features	656
B.2	Version 5.2 to 6.0 Differences	656
B.3	Version 5.1 to 5.2 Differences	658
B.4	Version 5.0 to 5.1 Differences	660
B.5	Version 4.5 to 5.0 Differences	663
B.6	Version 4.0 to 4.5 Differences	667
B.7	Version 3.1 to 4.0 Differences	668
B.8	Version 3.0 to 3.1 Differences	669
B.9	Version 2.5 to 3.0 Differences	670
	Index	673

List of Figures

19.1 First-Party Tool Activation Flow Chart	462
---	-----

List of Tables

2.1	ICV Scopes and Descriptions	40
2.2	ICV Initial Values	42
2.3	Ways to Modify and to Retrieve ICV Values	45
2.4	ICV Override Relationships	49
3.1	Syntactic Properties for Clauses, Arguments and Modifiers	61
5.1	Implicitly Declared C/C++ Reduction Identifiers	131
5.2	Implicitly Declared Fortran Reduction Identifiers	131
5.3	Map-Type Decay of Map Type Combinations	165
6.1	Predefined Memory Spaces	175
6.2	Allocator Traits	176
6.3	Predefined Allocators	179
11.1	ompt_callback_work Callback Work Types for Worksharing-Loop	259
12.1	ompt_callback_task_create Callback Flags Evaluation	273
18.1	Required Values of the omp_interop_property_t enum Type	428
18.2	Required Values for the omp_interop_rc_t enum Type	428
18.3	Standard Tool Control Commands	456
19.1	OMPT Callback Interface Runtime Entry Point Names and Their Type Signatures	465
19.2	Callbacks for which ompt_set_callback Must Return ompt_set_always	467
19.3	Callbacks for which ompt_set_callback May Return Any Non-Error Code	468
19.4	OMPT Tracing Interface Runtime Entry Point Names and Their Type Signatures	469
20.1	Mapping of Scope Type and OMPD Handles	570
21.1	Predefined Abstract Names for OMP_PLACES	632
21.2	Available Field Types for Formatting OpenMP Thread Affinity Information	638

1 Overview of the OpenMP API

The collection of compiler directives, library routines, and environment variables that this document describes 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:

`https://www.openmp.org`

The directives, library routines, environment variables, and tool support that this document defines 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, work-distribution 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 some or all 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 dependences, data conflicts, race conditions, or deadlocks. Compliant implementations also are not required to check for any 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 *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 *primary 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*).

divergent threads Two *OpenMP threads* that have reached different points in user code or otherwise have reached a common point via calls from different points in user code.

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

hardware thread An indivisible hardware execution unit on which only one *OpenMP thread* can execute at a time.

progress unit An implementation-defined set of consecutive *hardware threads* on which *OpenMP threads* may execute a common stream of instructions and serially execute diverging user code when any two threads become *divergent threads*.

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* with respect to which the current *device* performs an operation, as specified by a *device construct* or an OpenMP device memory routine.

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](#) for a listing of current *base languages* for the OpenMP API.

1	base program	A program written in a <i>base language</i> .
2	preprocessed code	For C/C++, a sequence of preprocessing tokens that result from the first six phases of
3		translation, as defined by the <i>base language</i> .
4	program order	An ordering of operations performed by the same <i>thread</i> as determined by the
5		execution sequence of operations specified by the <i>base language</i> .
6		COMMENT: For versions of C and C++ that include base language
7		support for threading, <i>program order</i> corresponds to the <i>sequenced before</i>
8		relation between operations performed by the same <i>thread</i> .
9	structured block	For C/C++, an executable statement, possibly compound, with a single entry at the
10		top and a single exit at the bottom, or an OpenMP <i>construct</i> .
11		For Fortran, a <i>strictly structured block</i> , a <i>loosely structured block</i> , or an OpenMP
12		context-specific structured block (see Section 4.3.1).
13	structured block	For C/C++, a sequence of zero or more executable statements (including OpenMP
14	sequence	<i>constructs</i>) that together have a single entry at the top and a single exit at the bottom.
15		For Fortran, a block of zero or more executable constructs (including OpenMP
16		<i>constructs</i>) with a single entry at the top and a single exit at the bottom.
17	strictly structured	A single Fortran BLOCK construct, with a single entry at the top and a single exit at
18	block	the bottom.
19	loosely structured	A block of zero or more executable constructs (including OpenMP <i>constructs</i>),
20	block	where the first executable construct (if any) is not a Fortran BLOCK construct, with a
21		single entry at the top and a single exit at the bottom.
22	compilation unit	For C/C++, a translation unit.
23		For Fortran, a program unit.
24	enclosing context	For C/C++, the innermost scope enclosing an OpenMP <i>directive</i> .
25		For Fortran, the innermost scoping unit enclosing an OpenMP <i>directive</i> .
26	directive	A <i>base language</i> mechanism to specify <i>OpenMP program</i> behavior.
27		COMMENT: See Section 3.1 for a description of OpenMP <i>directive</i>
28		syntax in each <i>base language</i> .
29	white space	A non-empty sequence of space and/or horizontal tab characters.
30	OpenMP program	A program that consists of a <i>base program</i> that is annotated with OpenMP <i>directives</i>
31		or that calls OpenMP API runtime library routines.
32	conforming program	An <i>OpenMP program</i> that follows all rules and restrictions of the OpenMP
33		specification.

1	implementation code	Implicit code that is introduced by the OpenMP implementation.
2	metadirective	A <i>directive</i> that conditionally resolves to another <i>directive</i> .
3	declarative directive	An OpenMP <i>directive</i> that may only be placed in a declarative context and results in
4		one or more declarations only; it is not associated with the immediate execution of
5		any user code or <i>implementation code</i> .
6	executable directive	An OpenMP <i>directive</i> that appears in an executable context and results in
7		<i>implementation code</i> and/or prescribes the manner in which associated user code
8		must execute.
9	informational directive	An OpenMP <i>directive</i> that is neither declarative nor executable, but otherwise
10		conveys user code properties to the compiler.
11	utility directive	An OpenMP <i>directive</i> that facilitates interactions with the compiler and/or supports
12		code readability; it may be either informational or executable.
13	stand-alone directive	An OpenMP <i>construct</i> in which no user code is associated, but may produce
14		implementation code.
15	construct	An OpenMP <i>executable directive</i> and its paired end directive (if any) and the
16		associated <i>structured block</i> (if any) not including the code in any called routines.
17		That is, the lexical extent of an <i>executable directive</i> .
18	subsidiary directive	An OpenMP <i>directive</i> that is not an <i>executable directive</i> and that appears only as part
19		of an OpenMP <i>construct</i> .
20	combined construct	A <i>construct</i> that is a shortcut for specifying one <i>construct</i> immediately nested inside
21		another <i>construct</i> . A <i>combined construct</i> is semantically identical to that of
22		explicitly specifying the first <i>construct</i> containing one instance of the second
23		<i>construct</i> and no other statements.
24	composite construct	A <i>construct</i> that is composed of two <i>constructs</i> but does not have identical semantics
25		to specifying one of the <i>constructs</i> immediately nested inside the other. A <i>composite</i>
26		<i>construct</i> either adds semantics not included in the <i>constructs</i> from which it is
27		composed or provides an effective nesting of the one <i>construct</i> inside the other that
28		would otherwise be non-conforming.
29	constituent construct	For a given combined or composite <i>construct</i> , a <i>construct</i> from which it, or any one
30		of its <i>constituent constructs</i> , is composed.

31 COMMENT: The *constituent constructs* of a
32 **target teams distribute parallel for simd** construct are the
33 following constructs: **target**,
34 **teams distribute parallel for simd**, **teams**,
35 **distribute parallel for simd**, **distribute**,
36 **parallel for simd**, **parallel**, **for simd**, **for**, and **simd**.

1	leaf construct	For a given combined or composite <i>construct</i> , a <i>constituent construct</i> that is not itself
2		a combined or composite <i>construct</i> .
3		COMMENT: The <i>leaf constructs</i> of a
4		target teams distribute parallel for simd construct are the
5		following constructs: target , teams , distribute , parallel ,
6		for , and simd .
7	combined target	A <i>combined construct</i> that is composed of a target construct along with another
8	construct	construct.
9	region	All code encountered during a specific instance of the execution of a given <i>construct</i> ,
10		structured block sequence or OpenMP library routine. A <i>region</i> includes any code in
11		called routines as well as any <i>implementation code</i> . The generation of a <i>task</i> at the
12		point where a <i>task generating construct</i> is encountered is a part of the <i>region</i> of the
13		<i>encountering thread</i> . However, an <i>explicit task region</i> that corresponds to a <i>task</i>
14		<i>generating construct</i> is not part of the <i>region</i> of the <i>encountering thread</i> unless it is
15		an <i>included task region</i> . The point where a target or teams directive is
16		encountered is a part of the <i>region</i> of the <i>encountering thread</i> , but the <i>region</i> that
17		corresponds to the target or teams directive is not.
18		COMMENTS:
19		A <i>region</i> may also be thought of as the dynamic or runtime extent of a
20		<i>construct</i> or of an OpenMP library routine.
21		During the execution of an <i>OpenMP program</i> , a <i>construct</i> may give rise
22		to many <i>regions</i> .
23	active parallel region	A parallel <i>region</i> that is executed by a <i>team</i> consisting of more than one <i>thread</i> .
24	inactive parallel region	A parallel <i>region</i> that is executed by a <i>team</i> of only one <i>thread</i> .
25	active target region	A target <i>region</i> that is executed on a <i>device</i> other than the <i>device</i> that encountered
26		the target <i>construct</i> .
27	inactive target region	A target <i>region</i> that is executed on the same <i>device</i> that encountered the target
28		<i>construct</i> .
29	sequential part	All code encountered during the execution of an <i>initial task region</i> that is not part of
30		a parallel <i>region</i> corresponding to a parallel <i>construct</i> or a task <i>region</i>
31		corresponding to a task <i>construct</i> .
32		COMMENTS:
33		A <i>sequential part</i> is enclosed by an <i>implicit parallel region</i> .

1 Executable statements in called routines may be in both a *sequential part*
2 and any number of explicit **parallel regions** at different points in the
3 program execution.

4 **primary thread** An *OpenMP thread* that has *thread number 0*. A *primary thread* may be an *initial*
5 *thread* or the *thread* that encounters a **parallel construct**, creates a *team*,
6 generates a set of *implicit tasks*, and then executes one of those *tasks* as *thread*
7 number 0.

8 **worker thread** An *OpenMP thread* that is not the *primary thread* of a *team* and that executes one of
9 the *implicit tasks* of a *parallel region*.

10 **parent thread** The *thread* that encountered the **parallel construct** and generated a **parallel**
11 *region* is the *parent thread* of each of the *threads* in the *team* of that **parallel**
12 *region*. The *primary thread* of a **parallel region** is the same *thread* as its *parent*
13 *thread* with respect to any resources associated with an *OpenMP thread*.

14 **child thread** When a *thread* encounters a **parallel construct**, each of the *threads* in the
15 generated **parallel region's team** are *child threads* of the encountering *thread*.
16 The **target** or **teams region's initial thread** is not a *child thread* of the *thread* that
17 encountered the **target** or **teams construct**.

18 **ancestor thread** For a given *thread*, its *parent thread* or one of its *parent thread's ancestor threads*.

19 **team** A set of one or more *threads* participating in the execution of a **parallel region**.

20 COMMENTS:

21 For an *active parallel region*, the *team* comprises the *primary thread* and
22 at least one additional *thread*.

23 For an *inactive parallel region*, the *team* comprises only the *primary*
24 *thread*.

25 **league** The set of *teams* created by a **teams construct**.

26 **implicit parallel region** An *inactive parallel region* that is not generated from a **parallel construct**.
27 *Implicit parallel regions* surround the whole *OpenMP program*, all **target regions**,
28 and all **teams regions**.

29 **initial thread** The *thread* that executes an *implicit parallel region*.

30 **initial team** The *team* that comprises an *initial thread* executing an *implicit parallel region*.

31 **nested construct** A *construct* (lexically) enclosed by another *construct*.

32 **closely nested construct** A *construct* nested inside another *construct* with no other *construct* nested between
33 them.

34 **explicit region** A *region* that corresponds to either a *construct* of the same name or a library routine
35 call that explicitly appears in the program.

1	nested region	A <i>region</i> (dynamically) enclosed by another <i>region</i> . That is, a <i>region</i> generated from
2		the execution of another <i>region</i> or one of its <i>nested regions</i> .
3		COMMENT: Some nestings are <i>conforming</i> and some are not. See
4		Section 17.1 for the restrictions on nesting.
5	closely nested region	A <i>region nested</i> inside another <i>region</i> with no parallel <i>region nested</i> between
6		them.
7	strictly nested region	A <i>region nested</i> inside another <i>region</i> with no other <i>explicit region nested</i> between
8		them.
9	all threads	All OpenMP <i>threads</i> participating in the <i>OpenMP program</i> .
10	current team	All <i>threads</i> in the <i>team</i> executing the innermost enclosing parallel <i>region</i> .
11	encountering thread	For a given <i>region</i> , the <i>thread</i> that encounters the corresponding <i>construct</i> .
12	all tasks	All <i>tasks</i> participating in the <i>OpenMP program</i> .
13	current team tasks	All <i>tasks</i> encountered by the corresponding <i>team</i> . The <i>implicit tasks</i> constituting the
14		parallel <i>region</i> and any <i>descendent tasks</i> encountered during the execution of
15		these <i>implicit tasks</i> are included in this set of tasks.
16	generating task	For a given <i>region</i> , the task for which execution by a <i>thread</i> generated the <i>region</i> .
17	binding thread set	The set of <i>threads</i> that are affected by, or provide the context for, the execution of a
18		<i>region</i> .
19		The <i>binding thread set</i> for a given <i>region</i> can be <i>all threads</i> on a specified set of
20		devices, <i>all threads</i> that are executing tasks in a <i>contention group</i> , all <i>primary</i>
21		<i>threads</i> that are executing the initial tasks of an enclosing teams <i>region</i> , the <i>current</i>
22		<i>team</i> , or the <i>encountering thread</i> .
23		COMMENT: The <i>binding thread set</i> for a particular <i>region</i> is described in
24		its corresponding subsection of this specification.
25	binding task set	The set of <i>tasks</i> that are affected by, or provide the context for, the execution of a
26		<i>region</i> .
27		The <i>binding task set</i> for a given <i>region</i> can be <i>all tasks</i> , the <i>current team tasks</i> , <i>all</i>
28		<i>tasks of the current team that are generated in the region</i> , the <i>binding implicit task</i> , or
29		the <i>generating task</i> .
30		COMMENT: The <i>binding task set</i> for a particular <i>region</i> (if applicable) is
31		described in its corresponding subsection of this specification.
32	binding region	The enclosing <i>region</i> that determines the execution context and limits the scope of
33		the effects of the bound <i>region</i> is called the <i>binding region</i> .

1		<i>Binding region</i> is not defined for <i>regions</i> for which the <i>binding thread</i> set is <i>all</i>
2		<i>threads</i> or the <i>encountering thread</i> , nor is it defined for <i>regions</i> for which the <i>binding</i>
3		<i>task set</i> is <i>all tasks</i> .
4	orphaned construct	A <i>construct</i> that gives rise to a <i>region</i> for which the <i>binding thread set</i> is the <i>current</i>
5		<i>team</i> , but is not nested within another <i>construct</i> that gives rise to the <i>binding region</i> .
6	work-distribution	A <i>construct</i> that is cooperatively executed by <i>threads</i> in the <i>binding thread set</i> of the
7	construct	corresponding <i>region</i> .
8	worksharing construct	A <i>work-distribution construct</i> that is executed by the thread team of the innermost
9		enclosing parallel <i>region</i> and includes, by default, an implicit barrier.
10	device construct	An OpenMP <i>construct</i> that accepts the device clause.
11	cancellable construct	An OpenMP <i>construct</i> that can be cancelled.
12	device routine	A function (for C/C++ and Fortran) or subroutine (for Fortran) that can be executed
13		on a <i>target device</i> , as part of a target <i>region</i> .
14	target variant	A version of a <i>device routine</i> that can only be executed as part of a target <i>region</i> .
15	foreign runtime	A runtime environment that exists outside the OpenMP runtime with which the
16	environment	OpenMP implementation may interoperate.
17	foreign execution	A context that is instantiated from a <i>foreign runtime environment</i> in order to facilitate
18	context	execution on a given <i>device</i> .
19	foreign task	A unit of work executed in a <i>foreign execution context</i> .
20	indirect device	An indirect call to the <i>device</i> version of a procedure on a <i>device</i> other than the host
21	invocation	<i>device</i> , through a function pointer (C/C++), a pointer to a member function (C++) or
22		a procedure pointer (Fortran) that refers to the host version of the procedure.
23	place	An unordered set of <i>processors</i> on a <i>device</i> .
24	place list	The ordered list that describes all OpenMP <i>places</i> available to the execution
25		environment.
26	place partition	An ordered list that corresponds to a contiguous interval in the OpenMP <i>place list</i> . It
27		describes the <i>places</i> currently available to the execution environment for a given
28		parallel <i>region</i> .
29	place number	A number that uniquely identifies a <i>place</i> in the <i>place list</i> , with zero identifying the
30		first <i>place</i> in the <i>place list</i> , and each consecutive whole number identifying the next
31		<i>place</i> in the <i>place list</i> .
32	accessible devices	The host <i>device</i> and all non-host <i>devices</i> accessible for execution.
33	supported devices	The host <i>device</i> and all non-host <i>devices</i> supported by the implementation for
34		execution of target code for which the device-related requirements of the requires
35		directive are fulfilled.

1	thread affinity	A binding of <i>threads</i> to <i>places</i> within the current <i>place partition</i> .
2	SIMD instruction	A single machine instruction that can operate on multiple data elements.
3	SIMD lane	A software or hardware mechanism capable of processing one data element from a
4		<i>SIMD instruction</i> .
5	SIMD chunk	A set of iterations executed concurrently, each by a <i>SIMD lane</i> , by a single <i>thread</i> by
6		means of <i>SIMD instructions</i> .
7	memory	A storage resource to store and to retrieve variables accessible by OpenMP <i>threads</i> .
8	memory space	A representation of storage resources from which <i>memory</i> can be allocated or
9		deallocated. More than one memory space may exist.
10	target memory space	A <i>memory space</i> that is associated with at least one device that is not the current
11		device when it is created.
12	memory allocator	An OpenMP object that fulfills requests to allocate and to deallocate <i>memory</i> for
13		program variables from the storage resources of its associated <i>memory space</i> .
14	associated memory space (of an allocator)	The <i>memory space</i> that is specified when an allocator is created.
15	associated device (of an allocator)	If the <i>associated memory space</i> is a predefined memory space, the current device.
16		For any other <i>associated memory space</i> , any device that is specified when the
17		memory space is created.
18	handle	An opaque reference that uniquely identifies an abstraction.

19 1.2.3 Loop Terminology

20	canonical loop nest	A loop nest that complies with the rules and restrictions defined in Section 4.4.1 .
21	loop-associated directive	An OpenMP <i>executable directive</i> for which the associated user code must be a
22		<i>canonical loop nest</i> .
23	associated loop	A loop from a <i>canonical loop nest</i> that is controlled by a given <i>loop-associated</i>
24		<i>directive</i> .
25	loop nest depth	For a <i>canonical loop nest</i> , the maximal number of loops, including the outermost
26		loop, that can be associated with a <i>loop-associated directive</i> .
27	logical iteration space	For a <i>loop-associated directive</i> , the sequence $0, \dots, N - 1$ where N is the number of
28		iterations of the loops associated with the directive. The logical numbering denotes
29		the sequence in which the iterations would be executed if the set of associated loops
30		were executed sequentially.

1	logical iteration	An iteration from the associated loops of a <i>loop-associated directive</i> , designated by a logical number from the <i>logical iteration space</i> of the associated loops.
2		
3	logical iteration vector space	For a <i>loop-associated directive</i> with n associated nested loops, the set of n -tuples (i_1, \dots, i_n) . For the k^{th} associated loop, from outermost to innermost, i_k is its <i>logical iteration number</i> as if it was the only associated loop.
4		
5		
6	logical iteration vector	An iteration from the associated nested loops of a <i>loop-associated directive</i> , where n is the number of associated loops, designated by an n -tuple from the <i>logical iteration vector space</i> of the associated loops.
7		
8		
9	lexicographic order	The total order of two <i>logical iteration vectors</i> $\omega_a = (i_1, \dots, i_n)$ and $\omega_b = (j_1, \dots, j_n)$, denoted by $\omega_a \leq_{\text{lex}} \omega_b$, where either $\omega_a = \omega_b$ or $\exists m \in \{1, \dots, n\}$ such that $i_m < j_m$ and $i_k = j_k$ for all $k \in \{1, \dots, m - 1\}$.
10		
11		
12	product order	The partial order of two <i>logical iteration vectors</i> $\omega_a = (i_1, \dots, i_n)$ and $\omega_b = (j_1, \dots, j_n)$, denoted by $\omega_a \leq_{\text{product}} \omega_b$, where $i_k \leq j_k$ for all $k \in \{1, \dots, n\}$.
13		
14	loop-transforming construct	A construct that is replaced by the loops that result from applying the transformation as defined by its directive to its associated loops.
15		
16	generated loop	A loop that is generated by a <i>loop-transforming construct</i> and is one of the resulting loops that replace the construct.
17		
18	SIMD loop	A loop that includes at least one <i>SIMD chunk</i> .
19	non-rectangular loop	For a loop nest, a loop for which a loop bound references the iteration variable of a surrounding loop in the loop nest.
20		
21	perfectly nested loop	A loop that has no intervening code between it and the body of its surrounding loop. The outermost loop of a loop nest is always perfectly nested.
22		
23	doacross loop nest	A loop nest, consisting of loops that may be associated with the same <i>loop-associated directive</i> , that has cross-iteration dependences. An iteration is dependent on one or more lexicographically earlier iterations.
24		
25		
26		COMMENT: The ordered clause parameter on a worksharing-loop
27		directive identifies the loops associated with the <i>doacross loop nest</i> .

28 1.2.4 Synchronization Terminology

29	barrier	A point in the execution of a program encountered by a <i>team of threads</i> , beyond which no <i>thread</i> in the team may execute until all <i>threads</i> in the <i>team</i> have reached the barrier and all <i>explicit tasks</i> generated by the <i>team</i> have executed to completion. If <i>cancellation</i> has been requested, <i>threads</i> may proceed to the end of the canceled <i>region</i> even if some <i>threads</i> in the team have not reached the <i>barrier</i> .
30		
31		
32		
33		

1	cancellation	An action that cancels (that is, aborts) an OpenMP <i>region</i> and causes executing
2		<i>implicit</i> or <i>explicit</i> tasks to proceed to the end of the canceled <i>region</i> .
3	cancellation point	A point at which implicit and explicit tasks check if cancellation has been requested.
4		If cancellation has been observed, they perform the <i>cancellation</i> .
5	flush	An operation that a <i>thread</i> performs to enforce consistency between its view and
6		other <i>threads</i> ' view of memory.
7	thread-set	The set of threads for which a flush operation may enforce memory consistency.
8	flush property	A property that determines the manner in which a <i>flush</i> operation enforces memory
9		consistency. The defined flush properties are:
10		<ul style="list-style-type: none"> • <i>strong</i>: flushes a set of variables from the current thread's temporary view of
11		the memory to the memory;
12		<ul style="list-style-type: none"> • <i>release</i>: orders memory operations that precede the flush before memory
13		operations performed by a different thread with which it synchronizes;
14		<ul style="list-style-type: none"> • <i>acquire</i>: orders memory operations that follow the flush after memory
15		operations performed by a different thread that synchronizes with it.
16		COMMENT: Any <i>flush</i> operation has one or more <i>flush properties</i> .
17	strong flush	A <i>flush</i> operation that has the <i>strong flush property</i> .
18	release flush	A <i>flush</i> operation that has the <i>release flush property</i> .
19	acquire flush	A <i>flush</i> operation that has the <i>acquire flush property</i> .
20	atomic operation	An operation that is specified by an atomic construct or is implicitly performed by
21		the OpenMP implementation and that atomically accesses and/or modifies a specific
22		storage location.
23	atomic scope	The set of the threads that may concurrently access or modify a given storage location
24		with <i>atomic operations</i> , where at least one of the operations modifies the location.
25	atomic read	An <i>atomic operation</i> that is specified by an atomic construct on which the read
26		clause is present.
27	atomic write	An <i>atomic operation</i> that is specified by an atomic construct on which the write
28		clause is present.
29	atomic update	An <i>atomic operation</i> that is specified by an atomic construct on which the
30		update clause is present.
31	atomic captured	An <i>atomic update</i> operation that is specified by an atomic construct on which the
32	update	capture clause is present.

1	atomic conditional update	An <i>atomic update</i> operation that is specified by an atomic construct on which the compare clause is present.
2		
3	read-modify-write	An <i>atomic operation</i> that reads and writes to a given storage location.
4		COMMENT: Any <i>atomic update</i> is a <i>read-modify-write</i> operation.
5	sequentially consistent atomic construct	An atomic construct for which the seq_cst clause is specified.
6	non-sequentially consistent atomic construct	An atomic construct for which the seq_cst clause is not specified
7	sequentially consistent atomic operation	An <i>atomic operation</i> that is specified by a <i>sequentially consistent atomic construct</i> .

8 1.2.5 Tasking Terminology

9	task	A specific instance of executable code and its <i>data environment</i> that the OpenMP implementation can schedule for execution by <i>threads</i> .
10		
11	task region	A <i>region</i> consisting of all code encountered during the execution of a <i>task</i> .
12	implicit task	A <i>task</i> generated by an <i>implicit parallel region</i> or generated when a parallel construct is encountered during execution.
13		
14	binding implicit task	The <i>implicit task</i> of the current thread team assigned to the encountering thread.
15	explicit task	A <i>task</i> that is not an <i>implicit task</i> .
16	initial task	An <i>implicit task</i> associated with an <i>implicit parallel region</i> .
17	current task	For a given <i>thread</i> , the <i>task</i> corresponding to the <i>task region</i> in which it is executing.
18	encountering task	For a given <i>region</i> , the <i>current task</i> of the <i>encountering thread</i> .
19	child task	A <i>task</i> is a <i>child task</i> of its generating <i>task region</i> . A <i>child task region</i> is not part of its generating <i>task region</i> .
20		
21	sibling tasks	<i>Tasks</i> that are <i>child tasks</i> of the same <i>task region</i> .
22	descendent task	A <i>task</i> that is the <i>child task</i> of a <i>task region</i> or of one of its <i>descendent task regions</i> .
23	contention group	All <i>implicit tasks</i> and their <i>descendent tasks</i> that are generated in an <i>implicit parallel region</i> , <i>R</i> , and in all <i>nested regions</i> for which <i>R</i> the innermost enclosing <i>implicit parallel region</i> .
24		
25		

1	task completion	A condition that is satisfied when a thread reaches the end of the executable code that is associated with the <i>task</i> and any <i>allow-completion event</i> that is created for the <i>task</i> has been fulfilled.
2		
3		
4		COMMENT: Completion of the <i>initial task</i> that is generated when the
5		program begins occurs at program exit.
6	task scheduling point	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> .
7		
8		
9	task switching	The act of a <i>thread</i> switching from the execution of one <i>task</i> to another <i>task</i> .
10	tied task	A <i>task</i> that, when its <i>task region</i> is suspended, can be resumed only by the same <i>thread</i> that was executing it before suspension. That is, the <i>task</i> is tied to that <i>thread</i> .
11		
12	untied task	A <i>task</i> that, when its <i>task region</i> is suspended, can be resumed by any <i>thread</i> in the team. That is, the <i>task</i> is not tied to any <i>thread</i> .
13		
14	undelayed task	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.
15		
16		
17	included task	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> .
18		
19	merged task	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> .
20		
21	mergeable task	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> .
22	final task	A <i>task</i> that forces all of its <i>child tasks</i> to become <i>final</i> and <i>included tasks</i> .
23	detachable task	An <i>explicit task</i> that only completes after an associated <i>event</i> variable that represents an <i>allow-completion event</i> is fulfilled and execution of the associated structured block has completed.
24		
25		
26	task dependence	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.
27		
28		
29	dependent task	A <i>task</i> that because of a <i>task dependence</i> cannot be executed until its <i>predecessor tasks</i> have completed.
30		
31	mutually exclusive tasks	<i>Tasks</i> that may be executed in any order, but not at the same time.
32	predecessor task	A <i>task</i> that must complete before its <i>dependent tasks</i> can be executed.
33	task synchronization construct	A taskwait , taskgroup , or a barrier <i>construct</i> .

1	task generating construct	A <i>construct</i> that generates one or more <i>explicit tasks</i> that are <i>child tasks</i> of the <i>encountering task</i> .
2		
3	target task	A <i>mergeable</i> and <i>untied task</i> that is generated by a <i>device construct</i> or a call to a device memory routine and that coordinates activity between the current <i>device</i> and the <i>target device</i> .
4		
5		
6	taskgroup set	A set of tasks that are logically grouped by a taskgroup <i>region</i> , such that a task is a member of the <i>taskgroup set</i> if and only if its task <i>region</i> is nested in the taskgroup <i>region</i> and it binds to the same parallel <i>region</i> as the taskgroup <i>region</i> .
7		
8		
9		

1.2.6 Data Terminology

11	variable	A named data storage block, for which the value can be defined and redefined during the execution of a program; for C/C++, this includes const -qualified types when explicitly permitted.
12		
13		
14		COMMENT: An array element or structure element is a variable that is
15		part of another variable.
16	scalar variable	For C/C++, a scalar variable, as defined by the base language.
17		For Fortran, a scalar variable with intrinsic type, as defined by the base language,
18		excluding character type.
19	aggregate variable	A variable, such as an array or structure, composed of other variables. For Fortran, a variable of character type is considered an aggregate variable.
20		
21	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.
22		
23	array item	An array, an array section, or an array element.
24	shape-operator	For C/C++, an array shaping operator that reinterprets a pointer expression as an array with one or more specified dimensions.
25		
26	C pointer	For C/C++, a base language pointer variable.
27		For Fortran, a variable of type C_PTR .
28	implicit array	For C/C++, the set of array elements of non-array type <i>T</i> that may be accessed by applying a sequence of [] operators to a given pointer that is either a pointer to type <i>T</i> or a pointer to a multidimensional array of elements of type <i>T</i> .
29		
30		
31		For Fortran, the set of array elements for a given array pointer.

1		COMMENT: For C/C++, the implicit array for pointer p with type T
2		$(*)[10]$ consists of all accessible elements $p[i][j]$, for all i and $j=0,1,\dots,9$.
3	base pointer	For C/C++, an lvalue pointer expression that is used by a given lvalue expression or
4		array section to refer indirectly to its storage, where the lvalue expression or array
5		section is part of the implicit array for that lvalue pointer expression.
6		For Fortran, a data pointer that appears last in the designator for a given variable or
7		array section, where the variable or array section is part of the pointer target for that
8		data pointer.
9		COMMENT: For the array section
10		$(*p0).x0[k1].p1 \rightarrow p2[k2].x1[k3].x2[4][0:n]$, where identifiers pi have a
11		pointer type declaration and identifiers xi have an array type declaration,
12		the <i>base pointer</i> is: $(*p0).x0[k1].p1 \rightarrow p2$.
13	corresponding base	For a given data entity that has a base pointer, an assignment to the base pointer such
14	pointer initialization	that any lexical reference to the data entity or a subobject of the data entity in a
15		target region refers to its corresponding data entity or subobject in the device data
16		environment.
17	assumed-size array	For C/C++, an array section for which the number of array elements is assumed.
18		For Fortran, an assumed-size array in the base language.
19	zero-offset	An <i>assumed-size array</i> for which the lower bound is zero.
	assumed-size array	
20	named pointer	For C/C++, the <i>base pointer</i> of a given lvalue expression or array section, or the <i>base</i>
21		<i>pointer</i> of one of its <i>named pointers</i> .
22		For Fortran, the <i>base pointer</i> of a given variable or array section, or the <i>base pointer</i>
23		of one of its <i>named pointers</i> .
24		COMMENT: For the array section
25		$(*p0).x0[k1].p1 \rightarrow p2[k2].x1[k3].x2[4][0:n]$, where identifiers pi have a
26		pointer type declaration and identifiers xi have an array type declaration,
27		the <i>named pointers</i> are: $p0$, $(*p0).x0[k1].p1$, and $(*p0).x0[k1].p1 \rightarrow p2$.
28	containing array	For C/C++, a non-subscripted array (a <i>containing array</i>) to which a series of zero or
29		more array subscript operators and/or $.$ (dot) operators are applied to yield a given
30		lvalue expression or array section for which storage is contained by the array.
31		For Fortran, an array (a <i>containing array</i>) without the POINTER attribute and
32		without a subscript list to which a series of zero or more array subscript operators
33		and/or component selectors are applied to yield a given variable or array section for
34		which storage is contained by the array.

1 COMMENT: An array is a containing array of itself. For the array section
2 $(*p0).x0[k1].p1 \rightarrow p2[k2].x1[k3].x2[4][0:n]$, where identifiers p_i have a
3 pointer type declaration and identifiers x_i have an array type declaration,
4 the *containing arrays* are: $(*p0).x0[k1].p1 \rightarrow p2[k2].x1$ and
5 $(*p0).x0[k1].p1 \rightarrow p2[k2].x1[k3].x2$.

6 **containing structure** For C/C++, a structure to which a series of zero or more . (dot) operators and/or
7 array subscript operators are applied to yield a given lvalue expression or array
8 section for which storage is contained by the structure.

9 For Fortran, a structure to which a series of zero or more component selectors and/or
10 array subscript selectors are applied to yield a given variable or array section for
11 which storage is contained by the structure.

12 COMMENT: A structure is a containing structure of itself. For C/C++, a
13 structure pointer p to which the \rightarrow operator applies is equivalent to the
14 application of a . (dot) operator to $(*p)$ for the purposes of determining
15 containing structures.

16 For the array section $(*p0).x0[k1].p1 \rightarrow p2[k2].x1[k3].x2[4][0:n]$, where
17 identifiers p_i have a pointer type declaration and identifiers x_i have an
18 array type declaration, the *containing structures* are: $(*p0).x0[k1].p1$,
19 $(*p0).x0[k1].p1.p2[k2]$ and $(*p0).x0[k1].p1.p2[k2].x1[k3]$

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

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

24 COMMENT: For the array section
25 $(*p0).x0[k1].p1 \rightarrow p2[k2].x1[k3].x2[4][0:n]$, where identifiers p_i have a
26 pointer type declaration and identifiers x_i have an array type declaration,
27 the *base array* is: $(*p0).x0[k1].p1 \rightarrow p2[k2].x1[k3].x2$.

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

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

32 COMMENT: For the array section
33 $(*p0).x0[k1].p1 \rightarrow p2[k2].x1[k3].x2[4][0:n]$, where identifiers p_i have a
34 pointer type declaration and identifiers x_i have an array type declaration,
35 the *named arrays* are: $(*p0).x0$, $(*p0).x0[k1].p1 \rightarrow p2[k2].x1$, and
36 $(*p0).x0[k1].p1 \rightarrow p2[k2].x1[k3].x2$.

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

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 *base expression* is: $(*p0).x0[k1].p1 \rightarrow p2[k2].x1[k3].x2$.

5 More examples for C/C++:

- 6 • The *base expression* for $x[i]$ and for $x[i:n]$ is x , if x is an array or
7 pointer.
- 8 • The *base expression* for $x[5][i]$ and for $x[5][i:n]$ is x , if x is a pointer
9 to an array or x is 2-dimensional array.
- 10 • The *base expression* for $y[5][i]$ and for $y[5][i:n]$ is $y[5]$, if y is an
11 array of pointers or y is a pointer to a pointer.

12 Examples for Fortran:

- 13 • The *base expression* for $x(i)$ and for $x(i:j)$ is x .

14 **base variable** For a given data entity that is a variable or array section, a variable denoted by a base
15 language identifier that is either the data entity or is a *containing array* or *containing*
16 *structure* of the data entity.

17 COMMENT:

18 Examples for C/C++:

- 19 • The data entities x , $x[i]$, $x[:n]$, $x[i].y[j]$ and $x[i].y[:n]$, where x and y
20 have array type declarations, all have the *base variable* x .
- 21 • The lvalue expressions and array sections $p[i]$, $p[:n]$, $p[i].y[j]$ and
22 $p[i].y[:n]$, where p has a pointer type and $p[i].y$ has an array type, has
23 a *base pointer* p but does not have a *base variable*.

24 Examples for Fortran:

- 25 • The data objects x , $x(i)$, $x(:n)$, $x(i)\%y(j)$ and $x(i)\%y(:n)$, where x and
26 y have array type declarations, all have the *base variable* x .
- 27 • The data objects $p(i)$, $p(:n)$, $p(i)\%y(j)$ and $p(i)\%y(:n)$, where p has a
28 pointer type and $p(i)\%y$ has an array type, has a *base pointer* p but
29 does not have a *base variable*.
- 30 • For the associated pointer p , p is both its *base variable* and *base*
31 *pointer*.

32 **base storage location** If a data entity has a *base pointer*, the first storage location of the implicit array of its
33 *base pointer*; otherwise, if the data entity has a *base variable*, the first storage
34 location of its *base variable*; otherwise, the first storage location of the data entity.

35 **base address** The address of the *base storage location* of a data entity.

1	attached pointer	A pointer variable in a <i>device data environment</i> that, as a result of a mapping operation, becomes the base pointer of a given data entity that also exists in the <i>device data environment</i> .
2		
3		
4	simply contiguous array section	An array section that statically can be determined to have contiguous storage or that, in Fortran, has the CONTIGUOUS attribute.
5		
6	structure	A structure is a variable that contains one or more variables.
7		For C/C++: Implemented using struct types.
8		For C++: Implemented using class types.
9		For Fortran: Implemented using derived types.
10	string literal	For C/C++, a string literal.
11		For Fortran, a character literal constant.
12	private variable	With respect to a given set of <i>task regions</i> or <i>SIMD lanes</i> that bind to the same parallel region , a <i>variable</i> for which the name provides access to a different block of storage for each <i>task region</i> or <i>SIMD lane</i> .
13		
14		
15		A <i>variable</i> that is part of another variable (as an array element or a structure element) cannot be made private independently of other components. If a <i>variable</i> is privatized, its components are also private.
16		
17		
18	shared variable	With respect to a given set of <i>task regions</i> that bind to the same parallel region , a <i>variable</i> for which the name provides access to the same block of storage for each <i>task region</i> .
19		
20		
21		A <i>variable</i> that is part of another variable (as an array element or a structure element) cannot be <i>shared</i> independently of the other components, except for static data members of C++ classes.
22		
23		
24	threadprivate variable	A <i>variable</i> that is replicated, one instance per <i>thread</i> , by the OpenMP implementation. Its name then provides access to a different block of storage for each <i>thread</i> .
25		
26		
27		A <i>variable</i> that is part of another variable (as an array element or a structure element) cannot be made <i>threadprivate</i> independently of the other components, except for static data members of C++ classes. If a <i>variable</i> is made <i>threadprivate</i> , its components are also <i>threadprivate</i> .
28		
29		
30		
31	threadprivate memory	The set of <i>threadprivate variables</i> associated with each <i>thread</i> .

1	groupprivate variable	A <i>variable</i> that is replicated, one instance per a specified group of tasks, by the
2		OpenMP implementation. Its name provides access to a different block of storage for
3		each specified group.
4		A <i>variable</i> that is part of another variable (as an array element or a structure element)
5		cannot be made <i>groupprivate</i> independently of the other components, except for
6		static data members of C++ classes. If a <i>variable</i> is made <i>groupprivate</i> , its
7		components are also <i>groupprivate</i> with respect to the same group.
8	device local variable	A <i>variable</i> with <i>static storage duration</i> that is replicated for each <i>device</i> by the
9		OpenMP implementation. Its name provides access to a different block of storage for
10		each <i>device</i> .
11		A <i>variable</i> that is part of another variable (as an array element or a structure element)
12		cannot be made <i>device local</i> independently of the other components, except for static
13		data members of C++ classes. If a <i>variable</i> is made <i>device local</i> , its components are
14		also <i>device local</i> .
15	data environment	The <i>variables</i> associated with the execution of a given <i>region</i> .
16	device data environment	The initial <i>data environment</i> associated with a device.
17	device address	An address of an object that may be referenced on a <i>target device</i> .
18	device pointer	An <i>implementation-defined handle</i> that refers to a <i>device address</i> and is represented
19		by a <i>C pointer</i> .
20	mapped variable	An original <i>variable</i> in a <i>data environment</i> with a corresponding <i>variable</i> in a <i>device</i>
21		<i>data environment</i> .
22		COMMENT: The original and corresponding <i>variables</i> may share storage.
23	mapping operation	An operation that establishes or removes a correspondence between a <i>variable</i> in one
24		data environment and another variable in a <i>device data environment</i> .
25	mapper	An operation that defines how variables of given type are to be mapped or updated
26		with respect to a <i>device data environment</i> .
27	user-defined mapper	A <i>mapper</i> that is defined by a declare mapper directive.
28	map-type decay	The process that determines the final map types of the map operations that result
29		from mapping a variable with a <i>user-defined mapper</i> .
30	mappable type	A type that is valid for a <i>mapped variable</i> . If a type is composed from other types
31		(such as the type of an array element or a structure element) and any of the other
32		types are not mappable then the type is not mappable.

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

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

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

5 In addition, for class types:

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

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

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

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

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

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

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

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

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

21 **class type** For C++, *variables* declared with one of the **class**, **struct**, or **union** keywords.

22 **storage location** A data storage block in memory.

23 **static storage duration** For C/C++, the lifetime of an object with static storage duration, as defined by the
24 base language.

25 For Fortran, the lifetime of a variable with a **SAVE** attribute, implicit or explicit, a
26 common block object or a variable declared in a module.

27 **NULL** A null pointer. For C, the value **NULL**. For C++, the value **NULL** or the value
28 **nullptr**. For Fortran, the value **C_NULL_PTR**.

29 **non-null value** A value that is not *NULL*.

30 **non-null pointer** A pointer that is not *NULL*.

1.2.7 Implementation Terminology

2	supported active levels of parallelism	An implementation-defined maximum number of <i>active parallel regions</i> that may
3		enclose any region of code in the program.
4	OpenMP API support	Support of at least one active level of parallelism.
5	nested parallelism support	Support of more than one active level of parallelism.
6	internal control variable	A conceptual variable that specifies runtime behavior of a set of <i>threads</i> or <i>tasks</i> in an <i>OpenMP program</i> .
7		
8		COMMENT: The acronym ICV is used interchangeably with the term
9		<i>internal control variable</i> throughout this specification.
10	OpenMP Additional Definitions document	A document that exists outside of the OpenMP specification and defines additional values that may be used in a <i>conforming program</i> . The <i>OpenMP Additional Definitions document</i> is available via
11		https://www.openmp.org/specifications/ .
12		
13		
14	compliant implementation	An implementation of the OpenMP specification that compiles and executes any <i>conforming program</i> as defined by the specification.
15		
16		COMMENT: A <i>compliant implementation</i> may exhibit <i>unspecified behavior</i> when compiling or executing a <i>non-conforming program</i> .
17		
18	unspecified behavior	A behavior or result that is not specified by the OpenMP specification or not known prior to the compilation or execution of an <i>OpenMP program</i> .
19		
20		Such <i>unspecified behavior</i> may result from:
21		• Issues documented by the OpenMP specification as having <i>unspecified behavior</i> .
22		• A <i>non-conforming program</i> .
23		• A <i>conforming program</i> exhibiting an <i>implementation-defined</i> behavior.
24		
25	implementation defined	Behavior that must be documented by the implementation, and is allowed to vary among different <i>compliant implementations</i> . An implementation is allowed to define this behavior as <i>unspecified</i> .
26		
27		
28		COMMENT: All features that have <i>implementation-defined</i> behavior are documented in Appendix A .
29		
30	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. Deprecated features may not be fully specified. In general, a deprecated feature was
31		
32		

1 fully specified in the version of the specification immediately prior to the one in
2 which it is deprecated. In most cases, a new feature replaces the deprecated feature.
3 Unless otherwise specified, whether any modifications provided by the replacement
4 feature apply to the deprecated feature is implementation defined.

5 1.2.8 Tool Terminology

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

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

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

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

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

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

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

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

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

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

20 **wait identifier** A unique opaque *handle* associated with each data object (for example, a lock) that
21 the OpenMP runtime uses to enforce mutual exclusion and potentially to cause a
22 *thread* to wait actively or passively.

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

26 **canonical frame** An address associated with a procedure *frame* on a call stack that was the value of the
27 **address** stack pointer immediately prior to calling the procedure for which the *frame*
28 represents the invocation.

29 **runtime entry point** A function interface provided by an OpenMP runtime for use by a *tool*. A *runtime*
30 *entry point* is typically not associated with a global function symbol.

31 **trace record** A data structure in which to store information associated with an occurrence of an
32 *event*.

33 **native trace record** A *trace record* for an OpenMP *device* that is in a device-specific format.

1	signal	A software interrupt delivered to a <i>thread</i> .
2	signal handler	A function called asynchronously when a <i>signal</i> is delivered to a <i>thread</i> .
3	async signal safe	The guarantee that interruption by <i>signal</i> delivery will not interfere with a set of
4		operations. An async signal safe <i>runtime entry point</i> is safe to call from a <i>signal</i>
5		<i>handler</i> .
6	code block	A contiguous region of memory that contains code of an OpenMP program to be
7		executed on a <i>device</i> .
8	OMPT	An interface that helps a <i>first-party tool</i> monitor the execution of an OpenMP
9		program.
10	OMPT interface state	A state that indicates the permitted interactions between a <i>first-party tool</i> and the
11		OpenMP implementation.
12	OMPT active	An <i>OMPT interface state</i> in which the OpenMP implementation is prepared to accept
13		runtime calls from a <i>first-party tool</i> and will dispatch any registered callbacks and in
14		which a <i>first-party tool</i> can invoke <i>runtime entry points</i> if not otherwise restricted.
15	OMPT pending	An <i>OMPT interface state</i> in which the OpenMP implementation can only call
16		functions to initialize a <i>first-party tool</i> and in which a <i>first-party tool</i> cannot invoke
17		<i>runtime entry points</i> .
18	OMPT inactive	An <i>OMPT interface state</i> in which the OpenMP implementation will not make any
19		callbacks and in which a <i>first-party tool</i> cannot invoke <i>runtime entry points</i> .
20	OMPD	An interface that helps a <i>third-party tool</i> inspect the OpenMP state of a program that
21		has begun execution.
22	OMPD library	A dynamically loadable library that implements the <i>OMPD</i> interface.
23	image file	An executable or shared library.
24	address space	A collection of logical, virtual, or physical memory address ranges that contain code,
25		stack, and/or data. Address ranges within an <i>address space</i> need not be contiguous.
26		An <i>address space</i> consists of one or more <i>segments</i> .
27	segment	A portion of an <i>address space</i> associated with a set of address ranges.
28	OpenMP architecture	The architecture on which an OpenMP <i>region</i> executes.
29	tool architecture	The architecture on which an <i>OMPD tool</i> executes.
30	OpenMP process	A collection of one or more <i>threads</i> and <i>address spaces</i> . A process may contain
31		<i>threads</i> and <i>address spaces</i> for multiple <i>OpenMP architectures</i> . At least one <i>thread</i>
32		in an OpenMP process is an OpenMP <i>thread</i> . A process may be live or a core file.
33	address space handle	A <i>handle</i> that refers to an <i>address space</i> within an OpenMP process.

1	thread handle	A <i>handle</i> that refers to an OpenMP <i>thread</i> .
2	parallel handle	A <i>handle</i> that refers to an OpenMP <i>parallel region</i> .
3	task handle	A <i>handle</i> that refers to an OpenMP <i>task region</i> .
4	descendent handle	An output <i>handle</i> that is returned from the <i>OMPD</i> library in a function that accepts an input <i>handle</i> : the output <i>handle</i> is a descendent of the input <i>handle</i> .
5		
6	ancestor handle	An input <i>handle</i> that is passed to the <i>OMPD</i> library in a function that returns an output <i>handle</i> : the input <i>handle</i> is an ancestor of the output <i>handle</i> . For a given
7		<i>handle</i> , the ancestors of the <i>handle</i> are also the ancestors of the <i>handle</i> 's descendent.
8		
9		COMMENT: A <i>tool</i> cannot use a <i>handle</i> in an <i>OMPD</i> call if any ancestor
10		of the <i>handle</i> has been released, except for <i>OMPD</i> calls that release it.
11	tool context	An opaque reference provided by a <i>tool</i> to an <i>OMPD</i> library. A <i>tool context</i> uniquely
12		identifies an abstraction.
13	address space context	A <i>tool context</i> that refers to an <i>address space</i> within a process.
14	thread context	A <i>tool context</i> that refers to a <i>native thread</i> .
15	native thread identifier	An identifier for a <i>native thread</i> defined by a <i>thread</i> implementation.

16 1.3 Execution Model

17 A compliant implementation must follow the abstract execution model defined by the supported
18 base language and OpenMP specification, as observable from the results of user code in a
19 conforming program. These results do not include output from external monitoring tools or tools
20 that use the OpenMP tool interfaces, which may reflect deviations from the execution model such as
21 the unprescribed use of additional threads, SIMD instructions, alternate loop transformations, or
22 other target devices to facilitate parallel execution of the program.

23 The OpenMP API uses the fork-join model of parallel execution. Multiple threads of execution
24 perform tasks defined implicitly or explicitly by OpenMP directives. The OpenMP API is intended
25 to support programs that will execute correctly both as parallel programs (multiple threads of
26 execution and a full OpenMP support library) and as sequential programs (directives ignored and a
27 simple OpenMP stubs library). However, a conforming OpenMP program may execute correctly as
28 a parallel program but not as a sequential program, or may produce different results when executed
29 as a parallel program compared to when it is executed as a sequential program. Further, using
30 different numbers of threads may result in different numeric results because of changes in the
31 association of numeric operations. For example, a serial addition reduction may have a different
32 pattern of addition associations than a parallel reduction. These different associations may change
33 the results of floating-point addition.

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

4 The thread that executes the implicit parallel region that surrounds the whole program executes on
5 the *host device*. An implementation may support other devices besides the host device. If
6 supported, these devices are available to the host device for *offloading* code and data. Each device
7 has its own threads that are distinct from threads that execute on another device. Threads cannot
8 migrate from one device to another device. Each device is identified by a device number. The
9 device number for the host device is the value of the total number of non-host devices, while each
10 non-host device has a unique device number that is greater than or equal to zero and less than the
11 device number for the host device. Additionally, the constant `omp_initial_device` can be
12 used as an alias for the host device and the constant `omp_invalid_device` can be used to
13 specify an invalid device number. A *conforming device number* is either a non-negative integer that
14 is less than or equal to `omp_get_num_devices()` or equal to `omp_initial_device` or
15 `omp_invalid_device`.

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

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

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

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

33 If a construct creates a data environment, the data environment is created at the time the construct is
34 encountered. The description of a construct defines whether it creates a data environment.

35 When any thread encounters a **parallel** construct, the thread creates a team of itself and zero or
36 more additional threads and becomes the primary thread of the new team. A set of implicit tasks,
37 one per thread, is generated. The code for each task is defined by the code inside the **parallel**
38 construct. Each task is assigned to a different thread in the team and becomes tied; that is, it is
39 always executed by the thread to which it is initially assigned. The task region of the task being
40 executed by the encountering thread is suspended, and each member of the new team executes its

1 implicit task. An implicit barrier occurs at the end of the **parallel** region. Only the primary
2 thread resumes execution beyond the end of the **parallel** construct, resuming the task region
3 that was suspended upon encountering the **parallel** construct. Any number of **parallel**
4 constructs can be specified in a single program.

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

11 When any team encounters a worksharing construct, the work inside the construct is divided among
12 the members of the team, and executed cooperatively instead of being executed by every thread. An
13 implicit barrier occurs at the end of any region that corresponds to a worksharing construct for
14 which the **nowait** clause is not specified. Redundant execution of code by every thread in the
15 team resumes after the end of the worksharing construct.

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

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

30 When a **loop** construct is encountered, the iterations of the loop associated with the construct are
31 executed in the context of its encountering threads, as determined according to its binding region. If
32 the **loop** region binds to a **teams** region, the region is encountered by the set of primary threads
33 that execute the **teams** region. If the **loop** region binds to a **parallel** region, the region is
34 encountered by the team of threads that execute the **parallel** region. Otherwise, the region is
35 encountered by a single thread.

36 If the **loop** region binds to a **teams** region, the encountering threads may continue execution
37 after the **loop** region without waiting for all iterations to complete; the iterations are guaranteed to
38 complete before the end of the **teams** region. Otherwise, all iterations must complete before the
39 encountering threads continue execution after the **loop** region. All threads that encounter the
40 **loop** construct may participate in the execution of the iterations. Only one thread may execute any

1 given iteration.

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

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

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

18 When *compile-time error termination* is performed, the effect is as if an **error** directive for which
19 *sev-level* is **fatal** and *action-time* is **compilation** is encountered. When *runtime error*
20 *termination* is performed, the effect is as if an **error** directive for which *sev-level* is **fatal** and
21 *action-time* is **execution** is encountered.

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

25 The OpenMP specification makes no guarantee that input or output to the same file is synchronous
26 when executed in parallel. In this case, the programmer is responsible for synchronizing input and
27 output processing with the assistance of OpenMP synchronization constructs or library routines.
28 For the case where each thread accesses a different file, the programmer does not need to
29 synchronize access.

30 All concurrency semantics defined by the base language with respect to threads of execution apply
31 to OpenMP threads, unless otherwise specified. An OpenMP thread *makes progress* when it
32 performs a flush operation, performs input or output processing, terminates, or makes progress as
33 defined by the base language. A set of threads in the same progress unit are not guaranteed to make
34 progress if one thread from the set is waiting for another thread in the set to synchronize with it, and
35 the threads are divergent threads. Otherwise, OpenMP threads will eventually make progress. The
36 generation and execution of explicit tasks by threads in the current team does not prevent any of the
37 threads from making progress if executing the explicit tasks as included tasks would ensure they
38 make progress.

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*. A given storage location in the memory may be associated with one or more devices, such that only threads on associated devices have access to it. 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 from within the region corresponding to the directive is unspecified; see [Section 5.4.3](#) 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 [Chapter 5](#).

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

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.

Two memory operations are considered unordered if the order in which they must complete, as seen by their affected threads, is not specified by the memory consistency guarantees listed in [Section 1.4.6](#). If multiple threads write to the same memory unit (defined consistently with the above access considerations) then a data race occurs if the writes are unordered. Similarly, if at least one thread reads from a memory unit and at least one thread writes to that same memory unit then a data race occurs if the read and write are unordered. If a data race occurs then the result of

1 the program is unspecified.

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

8 A storage location in memory that is associated with a given device has a device address that may
9 be dereferenced by a thread executing on that device, but it may not be generally accessible from
10 other devices. A different device may obtain a device pointer that refers to this device address. The
11 manner in which a program can obtain the referenced device address from a device pointer, outside
12 of mechanisms specified by OpenMP, is implementation defined. Unless otherwise specified, the
13 atomic scope of a storage location is all threads on the current device.

14 1.4.2 Device Data Environments

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

20 When an original variable is mapped to a device data environment and a corresponding variable is
21 not present in the device data environment, a new corresponding variable (of the same type and size
22 as the original variable) is created in the device data environment. Conversely, the original variable
23 becomes the new variable's corresponding variable in the device data environment of the device
24 that performs a mapping operation.

25 The corresponding variable in the device data environment may share storage with the original
26 variable. Writes to the corresponding variable may alter the value of the original variable. The
27 impact of this possibility on memory consistency is discussed in [Section 1.4.6](#). When a task
28 executes in the context of a device data environment, references to the original variable refer to the
29 corresponding variable in the device data environment. If an original variable is not currently
30 mapped and a corresponding variable does not exist in the device data environment then accesses to
31 the original variable result in unspecified behavior unless the **unified_shared_memory**
32 clause is specified on a **requires** directive for the compilation unit.

33 The relationship between the value of the original variable and the initial or final value of the
34 corresponding variable depends on the *map-type*. Details of this issue, as well as other issues with
35 mapping a variable, are provided in [Section 5.8.3](#).

36 The original variable in a data environment and a corresponding variable in a device data
37 environment may share storage. Without intervening synchronization data races can occur.

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

4 **1.4.3 Memory Management**

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

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

17 **1.4.4 The Flush Operation**

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

24 A flush operation has an associated *thread-set* that constrains the threads for which it enforces
25 memory consistency. Consistency is only guaranteed to be enforced between the view of memory
26 of these threads. Unless otherwise stated, the thread-set of a flush operation only includes all
27 threads on the current device.

28 If a flush operation is a strong flush, it enforces consistency between a thread's temporary view and
29 memory. A strong flush operation is applied to a set of variables called the *flush-set*. A strong flush
30 restricts how an implementation may reorder memory operations. Implementations must not
31 reorder the code for a memory operation for a given variable, or the code for a flush operation for
32 the variable, with respect to a strong flush operation that refers to the same variable.

33 If a thread has performed a write to its temporary view of a shared variable since its last strong flush
34 of that variable then, when it executes another strong flush of the variable, the strong flush does not
35 complete until the value of the variable has been written to the variable in memory. If a thread
36 performs multiple writes to the same variable between two strong flushes of that variable, the strong
37 flush ensures that the value of the last write is written to the variable in memory. A strong flush of a

1 variable executed by a thread also causes its temporary view of the variable to be discarded, so that
2 if its next memory operation for that variable is a read, then the thread will read from memory and
3 capture the value in its temporary view. When a thread executes a strong flush, no later memory
4 operation by that thread for a variable involved in that strong flush is allowed to start until the strong
5 flush completes. The completion of a strong flush executed by a thread is defined as the point at
6 which all writes to the flush-set performed by the thread before the strong flush are visible in
7 memory to all other threads, and at which that thread's temporary view of the flush-set is discarded.

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

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

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

- 31 1. The value is written to the variable by the first thread;
- 32 2. The first thread performs a release flush;
- 33 3. The second thread performs an acquire flush; and
- 34 4. The value is read from the variable by the second thread.

35
36 **Note** – OpenMP synchronization operations, described in [Chapter 15](#) and in [Section 18.9](#), are
37 recommended for enforcing this order. Synchronization through variables is possible but is not

1 recommended because the proper timing of flushes is difficult.
2

3 The flush properties that define whether a flush operation is a strong flush, a release flush, or an
4 acquire flush are not mutually disjoint. A flush operation may be a strong flush and a release flush;
5 it may be a strong flush and an acquire flush; it may be a release flush and an acquire flush; or it
6 may be all three.

7 1.4.5 Flush Synchronization and *Happens Before*

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

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

- 17 • If a release flush is performed on entry to an atomic operation, that atomic operation starts its
18 release sequence.
- 19 • If a release flush is performed in an implicit **flush** region, an atomic operation that is
20 provided by the implementation and that modifies an internal synchronization variable starts
21 its release sequence.
- 22 • If a release flush is performed by an explicit **flush** region, any atomic operation that
23 modifies a shared variable and follows the **flush** region in its thread's program order starts
24 an associated release sequence.

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

- 28 • If an acquire flush is performed on exit from an atomic operation, that atomic operation is its
29 associated atomic operation.
- 30 • If an acquire flush is performed in an implicit **flush** region, an atomic operation that is
31 provided by the implementation and that reads an internal synchronization variable is its
32 associated atomic operation.
- 33 • If an acquire flush is performed by an explicit **flush** region, any atomic operation that reads
34 a shared variable and precedes the **flush** region in its thread's program order is an
35 associated atomic operation.

1 The atomic scope of the internal synchronization variable that is used in implicit **flush** regions is
2 the intersection of the thread-sets of the synchronizing flush operations.

3 A release flush synchronizes with an acquire flush if the following conditions are satisfied:

- 4 • An atomic operation associated with the acquire flush reads a value written by a modification
5 from a release sequence associated with the release flush; and
- 6 • The thread that performs each flush is in both of their respective thread-sets.

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

- 9 1. *X* and *Y* are performed by the same thread, and *X* precedes *Y* in the thread's program order;
- 10 2. *X* synchronizes with *Y* according to the flush synchronization conditions explained above or
11 according to the base language's definition of *synchronizes with*, if such a definition exists; or
- 12 3. Another operation, *Z*, exists such that *X* simply happens before *Z* and *Z* simply happens
13 before *Y*.

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

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

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



20 **1.4.6 OpenMP Memory Consistency**

21 The following rules guarantee an observable completion order for a given pair of memory
22 operations in race-free programs, as seen by all affected threads. If both memory operations are
23 strong flushes, the affected threads are all threads in both of their respective thread-sets. If exactly
24 one of the memory operations is a strong flush, the affected threads are all threads in its thread-set.
25 Otherwise, the affected threads are all threads.

- 26 • If two operations performed by different threads are sequentially consistent atomic operations
27 or they are strong flushes that flush the same variable, then they must be completed as if in
28 some sequential order, seen by all affected threads.
- 29 • If two operations performed by the same thread are sequentially consistent atomic operations
30 or they access, modify, or, with a strong flush, flush the same variable, then they must be
31 completed as if in that thread's program order, as seen by all affected threads.
- 32 • If two operations are performed by different threads and one happens before the other, then
33 they must be completed as if in that *happens before* order, as seen by all affected threads, if:

- 1 – both operations access or modify the same variable;
- 2 – both operations are strong flushes that flush the same variable; or
- 3 – both operations are sequentially consistent atomic operations.
- 4 • Any two atomic memory operations from different **atomic** regions must be completed as if
- 5 in the same order as the strong flushes implied in their respective regions, as seen by all
- 6 affected threads.

7 The flush operation can be specified using the **flush** directive, and is also implied at various
8 locations in an OpenMP program: see [Section 15.8.6](#) for details.

9 
10 **Note** – Since flush operations by themselves cannot prevent data races, explicit flush operations are
11 only useful in combination with non-sequentially consistent atomic directives.
12 

13 OpenMP programs that:

- 14 • Do not use non-sequentially consistent atomic directives;
- 15 • Do not rely on the accuracy of a *false* result from **omp_test_lock** and
- 16 **omp_test_nest_lock**; and
- 17 • Correctly avoid data races as required in [Section 1.4.1](#),

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

21 1.5 Tool Interfaces

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

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

28 1.5.1 OMPT

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

- 30 • A mechanism to initialize a *first-party tool*;
- 31 • Routines that enable a *tool* to determine the capabilities of an OpenMP implementation;

- Routines that enable a *tool* to examine OpenMP state information associated with a *thread*;
- Mechanisms that enable a *tool* to map implementation-level calling contexts back to their source-level representations;
- A callback interface that enables a *tool* to receive notification of OpenMP *events*;
- A tracing interface that enables a *tool* to trace activity on OpenMP *target devices*; and
- A runtime library routine that an application can use to control a *tool*.

OpenMP implementations may differ with respect to the *thread states* that they support, the mutual exclusion implementations that they employ, and the OpenMP *events* for which *tool callbacks* are invoked. For some OpenMP *events*, OpenMP implementations must guarantee that a registered callback will be invoked for each occurrence of the *event*. For other OpenMP *events*, OpenMP implementations are permitted to invoke a registered callback for some or no occurrences of the *event*; for such OpenMP *events*, however, OpenMP implementations are encouraged to invoke *tool callbacks* on as many occurrences of the *event* as is practical. [Section 19.2.4](#) specifies the subset of OMPT callbacks that an OpenMP implementation must support for a minimal implementation of the OMPT interface.

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

[Chapter 20](#) 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 interfaces, according to the syntax and semantics laid out in Chapters 1 through 20. Appendices A and B as well as sections designated as Notes (see [Section 1.8](#)) 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.

- 1 • ISO/IEC 14882:2011, *Information Technology - Programming Languages - C++*.
2 This OpenMP API specification refers to ISO/IEC 14882:2011 as C++11.
- 3 • ISO/IEC 14882:2014, *Information Technology - Programming Languages - C++*.
4 This OpenMP API specification refers to ISO/IEC 14882:2014 as C++14.
- 5 • ISO/IEC 14882:2017, *Information Technology - Programming Languages - C++*.
6 This OpenMP API specification refers to ISO/IEC 14882:2017 as C++17.
- 7 • ISO/IEC 14882:2020, *Information Technology - Programming Languages - C++*.
8 This OpenMP API specification refers to ISO/IEC 14882:2020 as C++20.
- 9 • ISO/IEC 1539:1980, *Information Technology - Programming Languages - Fortran*.
10 This OpenMP API specification refers to ISO/IEC 1539:1980 as Fortran 77.
- 11 • ISO/IEC 1539:1991, *Information Technology - Programming Languages - Fortran*.
12 This OpenMP API specification refers to ISO/IEC 1539:1991 as Fortran 90.
- 13 • ISO/IEC 1539-1:1997, *Information Technology - Programming Languages - Fortran*.
14 This OpenMP API specification refers to ISO/IEC 1539-1:1997 as Fortran 95.
- 15 • ISO/IEC 1539-1:2004, *Information Technology - Programming Languages - Fortran*.
16 This OpenMP API specification refers to ISO/IEC 1539-1:2004 as Fortran 2003.
- 17 • ISO/IEC 1539-1:2010, *Information Technology - Programming Languages - Fortran*.
18 This OpenMP API specification refers to ISO/IEC 1539-1:2010 as Fortran 2008.
- 19 • ISO/IEC 1539-1:2018, *Information Technology - Programming Languages - Fortran*.
20 This OpenMP API specification refers to ISO/IEC 1539-1:2018 as Fortran 2018. While
21 future versions of the OpenMP specification are expected to address the following features,
22 currently their use may result in unspecified behavior.
 - 23 – Declared type of a polymorphic allocatable component in structure constructor
 - 24 – **SELECT RANK** construct
 - 25 – Assumed-rank dummy argument
 - 26 – Assumed-type dummy argument
 - 27 – Interoperable procedure enhancements

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

1.8 Organization of this Document

The remainder of this document is structured as normative chapters that define the directives, including their syntax and semantics, the runtime routines and the tool interfaces that comprise the OpenMP API. The document also includes appendices that facilitate maintaining a compliant implementation of the API.

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

▼————— C / C++ —————▼
C/C++ specific text...
▲————— C / C++ —————▲

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

▼————— C —————▼
C specific text...
▲————— C —————▲

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

▼————— C++ —————▼
C++ specific text...
▲————— C++ —————▲

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

▼————— Fortran —————▼
Fortran specific text...
▲————— Fortran —————▲

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

▼----- Fortran (cont.) -----▼

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

▼----- C/C++ (cont.) -----▼

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

1
2
3
4

▼
Note – Non-normative text...
▲

COMMENT: Non-normative text...

2 Internal Control Variables

An OpenMP implementation must act as if internal control variables (ICVs) control the behavior of an OpenMP program. These ICVs store information such as the number of threads to use for future **parallel** regions. One copy exists of each ICV per instance of its scope. Possible ICV scopes are: global; device; implicit task; and data environment. If an ICV has global scope then one copy exists for the whole program. 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 2.2](#).

2.1 ICV Descriptions

Table [2.1](#) shows the scope and description of each ICV.

TABLE 2.1: ICV Scopes and Descriptions

ICV	Scope	Description
<i>active-levels-var</i>	data environment	Number of nested active parallel regions such that all parallel regions are enclosed by the outermost initial task region on the device
<i>affinity-format-var</i>	device	Controls the thread affinity format when displaying thread affinity
<i>available-devices-var</i>	global	Controls target-device availability and the device-number assignment
<i>bind-var</i>	data environment	Controls the binding of OpenMP threads to places; when binding is requested, indicates that the execution environment is advised not to move threads between places; can also provide default thread affinity policies

ICV	Scope	Description
<i>cancel-var</i>	global	Controls the desired behavior of the cancel construct and cancellation points
<i>debug-var</i>	global	Controls whether an OpenMP implementation will collect information that an OMPD library can access to satisfy requests from a tool
<i>def-allocator-var</i>	implicit task	Controls the memory allocator used by memory allocation routines, directives and clauses that do not specify one explicitly
<i>default-device-var</i>	global	Controls the default target device
<i>device-num-var</i>	device	Device number of a given device
<i>display-affinity-var</i>	global	Controls the display of thread affinity
<i>dyn-var</i>	data environment	Enables dynamic adjustment of the number of threads used for encountered parallel regions
<i>explicit-task-var</i>	data environment	Whether a given task is an explicit task
<i>final-task-var</i>	data environment	Whether a given task is a final task
<i>league-size-var</i>	data environment	Number of initial teams in a league
<i>levels-var</i>	data environment	Number of nested parallel regions such that all parallel regions are enclosed by the outermost initial task region on the device
<i>max-active-levels-var</i>	data environment	Controls the maximum number of nested active parallel regions when the innermost parallel region is generated by a given task
<i>max-task-priority-var</i>	global	Controls the maximum value that can be specified in the priority clause
<i>nteam-var</i>	device	Controls the number of teams requested for encountered teams regions
<i>nthreads-var</i>	data environment	Controls the number of threads requested for encountered parallel regions
<i>num-devices-var</i>	global	Number of available non-host devices
<i>num-procs-var</i>	device	The number of processors available on the device
<i>place-partition-var</i>	implicit task	Controls the place partition available for encountered parallel regions
<i>run-sched-var</i>	data environment	Controls the schedule used for worksharing-loop regions that specify the runtime schedule kind
<i>stacksize-var</i>	device	Controls the stack size for threads that the OpenMP implementation creates
<i>target-offload-var</i>	global	Controls the offloading behavior

ICV	Scope	Description
<i>team-generator-var</i>	data environment	Generator type of current team that refers to a construct name or the OpenMP program
<i>team-num-var</i>	data environment	Initial team number of a given thread
<i>team-size-var</i>	data environment	Size of the current team
<i>teams-thread-limit-var</i>	device	Controls the maximum number of threads that may execute tasks in parallel in each contention group that a teams construct creates
<i>thread-limit-var</i>	data environment	Controls the maximum number of threads that may execute tasks in the contention group in parallel
<i>thread-num-var</i>	data environment	Thread number of an implicit task within its binding team
<i>tool-libraries-var</i>	global	List of absolute paths to tool libraries
<i>tool-var</i>	global	Indicates that a tool will be registered
<i>tool-verbose-init-var</i>	global	Controls whether an OpenMP implementation will verbosely log the registration of a tool
<i>wait-policy-var</i>	device	Controls the desired behavior of waiting threads

Cross References

- Team Generator Types, see [Section 20.3.10](#)

2.2 ICV Initialization

Table 2.2 shows the ICVs, associated environment variables, and initial values.

TABLE 2.2: ICV Initial Values

ICV	Environment Variable	Initial Value
<i>active-levels-var</i>	(none)	<i>Zero</i>
<i>affinity-format-var</i>	OMP_AFFINITY_FORMAT	Implementation defined
<i>available-devices-var</i>	OMP_AVAILABLE_DEVICES	*
<i>bind-var</i>	OMP_PROC_BIND	Implementation defined
<i>cancel-var</i>	OMP_CANCELLATION	<i>False</i>
<i>debug-var</i>	OMP_DEBUG	disabled
<i>def-allocator-var</i>	OMP_ALLOCATOR	Implementation defined
<i>default-device-var</i>	OMP_DEFAULT_DEVICE	See below
<i>device-num-var</i>	(none)	<i>Zero</i>

ICV	Environment Variable	Initial Value
<i>display-affinity-var</i>	OMP_DISPLAY_AFFINITY	<i>False</i>
<i>dyn-var</i>	OMP_DYNAMIC	Implementation defined
<i>explicit-task-var</i>	(none)	<i>False</i>
<i>final-task-var</i>	(none)	<i>False</i>
<i>league-size-var</i>	(none)	<i>One</i>
<i>levels-var</i>	(none)	<i>Zero</i>
<i>max-active-levels-var</i>	OMP_MAX_ACTIVE_LEVELS, OMP_NUM_THREADS, OMP_PROC_BIND	Implementation defined
<i>max-task-priority-var</i>	OMP_MAX_TASK_PRIORITY	<i>Zero</i>
<i>nteam-var</i>	OMP_NUM_TEAMS	<i>Zero</i>
<i>nthreads-var</i>	OMP_NUM_THREADS	Implementation defined
<i>num-devices-var</i>	(none)	Implementation defined
<i>num-procs-var</i>	(none)	Implementation defined
<i>place-partition-var</i>	OMP_PLACES	Implementation defined
<i>run-sched-var</i>	OMP_SCHEDULE	Implementation defined
<i>stacksize-var</i>	OMP_STACKSIZE	Implementation defined
<i>target-offload-var</i>	OMP_TARGET_OFFLOAD	default
<i>team-generator-var</i>	(none)	<i>Zero</i>
<i>team-num-var</i>	(none)	<i>Zero</i>
<i>team-size-var</i>	(none)	<i>One</i>
<i>teams-thread-limit-var</i>	OMP_TEAMS_THREAD_LIMIT	<i>Zero</i>
<i>thread-limit-var</i>	OMP_THREAD_LIMIT	Implementation defined
<i>thread-num-var</i>	(none)	<i>Zero</i>
<i>tool-libraries-var</i>	OMP_TOOL_LIBRARIES	empty string
<i>tool-var</i>	OMP_TOOL	enabled
<i>tool-verbose-init-var</i>	OMP_TOOL_VERBOSE_INIT	disabled
<i>wait-policy-var</i>	OMP_WAIT_POLICY	Implementation defined

1 If an ICV has an associated environment variable and that ICV does not have global scope then the
2 ICV has a set of associated device-specific environment variables that extend the associated
3 environment variable with the following syntax:

4 `<ENVIRONMENT VARIABLE> _ALL`

5 or

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

1 where `<ENVIRONMENT VARIABLE>` is the associated environment variable and `<device>` is the
2 device number as specified in the `device` clause (see [Section 13.2](#)); the semantic and precedence
3 is described in [Chapter 21](#).

4 **Semantics**

- 5 • The initial value of `dyn-var` is implementation defined if the implementation supports
6 dynamic adjustment of the number of threads; otherwise, the initial value is *false*.
- 7 • If `target-offload-var` is **mandatory** and the number of available non-host devices is zero
8 then the `default-device-var` is initialized to `omp_invalid_device`. Otherwise, the initial
9 value is an implementation-defined non-negative integer that is less than or, if
10 `target-offload-var` is not **mandatory**, equal to `omp_get_initial_device()`.
- 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 non-host device ICVs are initialized before any OpenMP API construct or OpenMP
14 API 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_AFFINITY_FORMAT`, see [Section 21.2.5](#)
- 20 • `OMP_ALLOCATOR`, see [Section 21.5.1](#)
- 21 • `OMP_CANCELLATION`, see [Section 21.2.6](#)
- 22 • `OMP_DEBUG`, see [Section 21.4.1](#)
- 23 • `OMP_DEFAULT_DEVICE`, see [Section 21.2.8](#)
- 24 • `OMP_DISPLAY_AFFINITY`, see [Section 21.2.4](#)
- 25 • `OMP_DYNAMIC`, see [Section 21.1.1](#)
- 26 • `OMP_MAX_ACTIVE_LEVELS`, see [Section 21.1.4](#)
- 27 • `OMP_MAX_TASK_PRIORITY`, see [Section 21.2.10](#)
- 28 • `OMP_NUM_TEAMS`, see [Section 21.6.1](#)
- 29 • `OMP_NUM_THREADS`, see [Section 21.1.2](#)
- 30 • `OMP_PLACES`, see [Section 21.1.5](#)
- 31 • `OMP_PROC_BIND`, see [Section 21.1.6](#)
- 32 • `OMP_SCHEDULE`, see [Section 21.2.1](#)
- 33 • `OMP_STACKSIZE`, see [Section 21.2.2](#)

- `OMP_TARGET_OFFLOAD`, see [Section 21.2.9](#)
- `OMP_TEAMS_THREAD_LIMIT`, see [Section 21.6.2](#)
- `OMP_THREAD_LIMIT`, see [Section 21.1.3](#)
- `OMP_TOOL`, see [Section 21.3.1](#)
- `OMP_TOOL_LIBRARIES`, see [Section 21.3.2](#)
- `OMP_WAIT_POLICY`, see [Section 21.2.3](#)

2.3 Modifying and Retrieving ICV Values

Table 2.3 shows methods for modifying and retrieving the ICV values. If *(none)* is listed for an ICV, the OpenMP API does not support its modification or retrieval. Calls to OpenMP API routines retrieve or modify data environment scoped ICVs in the data environment of their binding tasks.

TABLE 2.3: Ways to Modify and to Retrieve ICV Values

ICV	Ways to Modify Value	Ways to Retrieve Value
<i>active-levels-var</i>	(none)	<code>omp_get_active_level</code>
<i>affinity-format-var</i>	<code>omp_set_affinity_format</code>	<code>omp_get_affinity_format</code>
<i>available-devices-var</i>	(none)	(none)
<i>bind-var</i>	(none)	<code>omp_get_proc_bind</code>
<i>cancel-var</i>	(none)	<code>omp_get_cancellation</code>
<i>debug-var</i>	(none)	(none)
<i>def-allocator-var</i>	<code>omp_set_default_allocator</code>	<code>omp_get_default_allocator</code>
<i>default-device-var</i>	<code>omp_set_default_device</code>	<code>omp_get_default_device</code>
<i>device-num-var</i>	(none)	<code>omp_get_device_num</code>
<i>display-affinity-var</i>	(none)	(none)
<i>dyn-var</i>	<code>omp_set_dynamic</code>	<code>omp_get_dynamic</code>
<i>explicit-task-var</i>	(none)	<code>omp_in_explicit_task</code>
<i>final-task-var</i>	(none)	<code>omp_in_final</code>
<i>league-size-var</i>	(none)	<code>omp_get_num_teams</code>
<i>levels-var</i>	(none)	<code>omp_get_level</code>
<i>max-active-levels-var</i>	<code>omp_set_max_active_levels</code>	<code>omp_get_max_active_levels</code>
<i>max-task-priority-var</i>	(none)	<code>omp_get_max_task_priority</code>
<i>ntteams-var</i>	<code>omp_set_num_teams</code>	<code>omp_get_max_teams</code>
<i>nthreads-var</i>	<code>omp_set_num_threads</code>	<code>omp_get_max_threads</code>
<i>num-devices-var</i>	(none)	<code>omp_get_num_devices</code>
<i>num-procs-var</i>	(none)	<code>omp_get_num_procs</code>

ICV	Ways to Modify Value	Ways to Retrieve Value
<i>place-partition-var</i>	(none)	<code>omp_get_partition_num_places</code> , <code>omp_get_partition_place_nums</code> , <code>omp_get_place_num_procs</code> , <code>omp_get_place_proc_ids</code>
<i>run-sched-var</i>	<code>omp_set_schedule</code>	<code>omp_get_schedule</code>
<i>stacksize-var</i>	(none)	(none)
<i>target-offload-var</i>	(none)	(none)
<i>team-generator-var</i>	(none)	(none)
<i>team-num-var</i>	(none)	<code>omp_get_team_num</code>
<i>team-size-var</i>	(none)	<code>omp_get_num_threads</code>
<i>teams-thread-limit-var</i>	<code>omp_set_teams_thread_limit</code>	<code>omp_get_teams_thread_limit</code>
<i>thread-limit-var</i>	<code>thread_limit</code>	<code>omp_get_thread_limit</code>
<i>thread-num-var</i>	(none)	<code>omp_get_thread_num</code>
<i>tool-libraries-var</i>	(none)	(none)
<i>tool-var</i>	(none)	(none)
<i>tool-verbose-init-var</i>	(none)	(none)
<i>wait-policy-var</i>	(none)	(none)

Semantics

- 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.
- 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.
- Detailed values in the *place-partition-var* ICV are retrieved using the listed runtime calls.
- The `thread_limit` clause sets the *thread-limit-var* ICV for the region of the construct on which it appears.

Cross References

- `thread_limit` clause, see [Section 13.3](#)
- `omp_get_active_level`, see [Section 18.2.18](#)
- `omp_get_affinity_format`, see [Section 18.3.9](#)
- `omp_get_cancellation`, see [Section 18.2.8](#)
- `omp_get_default_allocator`, see [Section 18.13.7](#)
- `omp_get_default_device`, see [Section 18.7.4](#)
- `omp_get_dynamic`, see [Section 18.2.7](#)

- 1 • `omp_get_level`, see [Section 18.2.15](#)
- 2 • `omp_get_max_active_levels`, see [Section 18.2.14](#)
- 3 • `omp_get_max_task_priority`, see [Section 18.5.1](#)
- 4 • `omp_get_max_teams`, see [Section 18.4.4](#)
- 5 • `omp_get_max_threads`, see [Section 18.2.3](#)
- 6 • `omp_get_num_procs`, see [Section 18.7.1](#)
- 7 • `omp_get_num_threads`, see [Section 18.2.2](#)
- 8 • `omp_get_partition_num_places`, see [Section 18.3.6](#)
- 9 • `omp_get_partition_place_nums`, see [Section 18.3.7](#)
- 10 • `omp_get_place_num_procs`, see [Section 18.3.3](#)
- 11 • `omp_get_place_proc_ids`, see [Section 18.3.4](#)
- 12 • `omp_get_proc_bind`, see [Section 18.3.1](#)
- 13 • `omp_get_schedule`, see [Section 18.2.10](#)
- 14 • `omp_get_supported_active_levels`, see [Section 18.2.12](#)
- 15 • `omp_get_teams_thread_limit`, see [Section 18.4.6](#)
- 16 • `omp_get_thread_limit`, see [Section 18.2.11](#)
- 17 • `omp_get_thread_num`, see [Section 18.2.4](#)
- 18 • `omp_in_final`, see [Section 18.5.3](#)
- 19 • `omp_set_affinity_format`, see [Section 18.3.8](#)
- 20 • `omp_set_default_allocator`, see [Section 18.13.6](#)
- 21 • `omp_set_default_device`, see [Section 18.7.3](#)
- 22 • `omp_set_dynamic`, see [Section 18.2.6](#)
- 23 • `omp_set_max_active_levels`, see [Section 18.2.13](#)
- 24 • `omp_set_num_teams`, see [Section 18.4.3](#)
- 25 • `omp_set_num_threads`, see [Section 18.2.1](#)
- 26 • `omp_set_schedule`, see [Section 18.2.9](#)
- 27 • `omp_set_teams_thread_limit`, see [Section 18.4.5](#)

2.4 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, unless otherwise specified.

When a **parallel** construct is encountered, the value of each ICV with implicit task scope is inherited 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 implicit tasks 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 implicit tasks 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.

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

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

If a **target** construct with a **thread_limit** clause is encountered, the *thread-limit-var* ICV from the data environment of the generated initial task is instead set to an implementation defined value between one and the value specified in the clause.

If a **target** construct with no **thread_limit** clause is encountered, the *thread-limit-var* ICV from the data environment of the generated initial task is set to an implementation defined value that is greater than zero.

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

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

If a **target** construct, **teams** construct, or **parallel** construct is encountered, the *team-generator-var* ICV for the data environments of the generated implicit tasks is instead set to the value of the appropriate team generator type as specified in [Section 20.3.10](#).

When encountering a worksharing-loop region for which the **runtime** schedule kind is specified, all implicit task regions that constitute the binding parallel region must have the same value for

1 *run-sched-var* in their data environments. Otherwise, the behavior is unspecified.

2 **Cross References**

- 3 • Team Generator Types, see [Section 20.3.10](#)

4 **2.5 ICV Override Relationships**

5 Table 2.4 shows the override relationships among construct clauses and ICVs. The table only lists
6 ICVs that can be overridden by a clause.

TABLE 2.4: ICV Override Relationships

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

7 **Semantics**

- 8 • The **num_threads** clause overrides the value of the first element of the *nthreads-var* ICV.
- 9 • If a **schedule** clause specifies a modifier then that modifier overrides any modifier that is
10 specified in the *run-sched-var* ICV.
- 11 • If *bind-var* is not set to *false* then the **proc_bind** clause overrides the value of the first
12 element of the *bind-var* ICV; otherwise, the **proc_bind** clause has no effect.

13 **Cross References**

- 14 • **allocate** clause, see [Section 6.6](#)
- 15 • **allocator** clause, see [Section 6.4](#)
- 16 • **num_teams** clause, see [Section 10.3.1](#)
- 17 • **num_threads** clause, see [Section 10.2.2](#)
- 18 • **proc_bind** clause, see [Section 10.2.4](#)
- 19 • **schedule** clause, see [Section 11.5.3](#)
- 20 • **thread_limit** clause, see [Section 13.3](#)

3 Directive and Construct Syntax

This chapter describes the syntax of OpenMP directives, clauses and any related base language code. OpenMP directives are specified with various base-language mechanisms that allow compilers to 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

The following restrictions apply to OpenMP directives:

- Unless otherwise specified, a program must not depend on any ordering of the evaluations of the expressions that appear in the clauses specified on a directive.
- Unless otherwise specified, a program must not depend on any side effects of the evaluations of the expressions that appear in the clauses specified on a directive.

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

C++

- A **throw** executed inside a region that arises from a thread-limiting directive must cause execution to resume within the same region, and the same thread that threw the exception must catch it. If the directive is also exception-aborting then whether the exception is caught or the **throw** results in runtime error termination is implementation defined.

C++

Fortran

- A directive may not appear in a pure procedure unless it is pure.
- A directive may not appear in a **WHERE**, **FORALL** or **DO CONCURRENT** construct.
- If more than one image is executing the program, any image control statement, **ERROR STOP** statement, **FAIL IMAGE** statement, collective subroutine call or access to a coindexed object that appears in an explicit OpenMP region will result in unspecified behavior.

Fortran

3.1 Directive Format

This section defines several categories of directives and constructs. OpenMP directives are specified with a *directive-specification*. A *directive-specification* consists of the *directive-specifier* and any clauses that may optionally be associated with the OpenMP directive:

```
| directive-specifier [[ , ] clause [ , ] clause ... ]
```

The *directive-specifier* is:

```
| directive-name
```

or for argument-modified directives:

```
| directive-name[ (directive-arguments) ]
```

White space in a *directive-name* is not optional.

Some OpenMP directives specify a paired **end** directive, where the *directive-name* of the paired **end** directives is:

- If *directive-name* starts with **begin**, the *end-directive-name* replaces **begin** with **end**
- otherwise it is **end** *directive-name* unless otherwise specified.

The *directive-specification* of a paired **end** directive may include one or more optional *end-clause*:

```
| directive-specifier [[ , ] end-clause [ , ] end-clause...]
```

where *end-clause* has the *end-clause* property, which explicitly allows it on a paired **end** directive.

An OpenMP *directive* may be specified as a pragma directive:

```
| #pragma omp directive-specification new-line
```

or a pragma operator:

```
| _Pragma ("omp directive-specification")
```

The use of **omp** as the first preprocessing token of a pragma directive is reserved for OpenMP directives that are defined in this specification. The use of **omp_x** as the first preprocessing token of a pragma directive is reserved for implementation-defined extensions to the OpenMP directives.

Note – In this *directive*, *directive-name* is **depobj**, *directive-arguments* is **o**. *directive-specifier* is **depobj(o)** and *directive-specification* is **depobj(o) depend(inout: d)**.

```
#pragma omp depobj(o) depend(inout: d)
```

White space can be used before and after the #. Preprocessing tokens in *directive-specification* of **#pragma** and **_Pragma** pragmas are subject to macro expansion.

C / C++

C++

In C++11 and higher, an OpenMP *directive* may be specified as a C++ attribute specifier:

```
[[ omp :: directive-attr ]]
```

or

```
[[ using omp : directive-attr ]]
```

where *directive-attr* is

```
directive( directive-specification )
```

or

```
sequence( [omp::]directive-attr [[, [omp::]directive-attr] ... ] )
```

Multiple attributes on the same statement are allowed. Attribute directives that apply to the same statement are unordered unless the **sequence** attribute is specified, in which case the right-to-left ordering applies. The **omp::** namespace qualifier within a **sequence** attribute is optional. The application of multiple attributes in a **sequence** attribute is ordered as if each directive had been specified as a pragma directive on subsequent lines.


Note – This example shows the expected transformation:

```
[[ omp::sequence(directive(parallel), directive(for)) ]]  
for(...) {}  
// becomes  
#pragma omp parallel  
#pragma omp for  
for(...) {}
```


The use of **omp** as the attribute namespace of an attribute specifier, or as the optional namespace qualifier within a **sequence** attribute, is reserved for OpenMP directives that are defined in this specification. The use of **omp_x** as the attribute namespace of an attribute specifier, or as the

1 optional namespace qualifier within a **sequence** attribute, is reserved for implementation-defined
2 extensions to the OpenMP directives.

3 The pragma and attribute forms are interchangeable for any OpenMP directive. Some OpenMP
4 directives may be composed of consecutive attribute specifiers if specified in their syntax. Any two
5 consecutive attribute specifiers may be reordered or expressed as a single attribute specifier, as
6 permitted by the base language, without changing the behavior of the OpenMP directive.



7 Directives are case-sensitive. Each expression used in the OpenMP syntax inside of a clause must
8 be a valid *assignment-expression* of the base language unless otherwise specified.



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



10 An OpenMP *directive* for Fortran is specified with a stylized comment as follows:

```
11 | sentinel directive-specification
```

12 All OpenMP compiler directives must begin with a directive *sentinel*. The format of a sentinel
13 differs between fixed form and free form source files, as described in [Section 3.1.1](#) and
14 [Section 3.1.2](#). In order to simplify the presentation, free form is used for the syntax of OpenMP
15 directives for Fortran throughout this document, except as noted.

16 Directives are case insensitive. Directives cannot be embedded within continued statements, and
17 statements cannot be embedded within directives. Each expression used in the OpenMP syntax
18 inside of a clause must be a valid *expression* of the base language unless otherwise specified.



19 A directive may be categorized as one of the following:

- 20 ● meta
- 21 ● declarative
- 22 ● executable
- 23 ● informational
- 24 ● utility
- 25 ● subsidiary

26 Base language code can be associated with directives. The directive’s association can be
27 categorized as:

- 28 ● none

- 1 • block-associated
- 2 • loop-associated
- 3 • declaration-associated
- 4 • delimited
- 5 • separating

C++

A declarative directive that is declaration-associated may alternatively be expressed as an attribute specifier where *directive-attr* is:

```
decl ( directive-specification )
```

A declarative directive with an association of none that accepts a variable list or extended list as a directive argument or clause argument may alternatively be expressed with an attribute specifier that also uses the **decl** attribute, applies to variable and/or function declarations, and omits the variable list or extended list argument. The effect is as if the omitted list argument is the list of declared variables and/or functions to which the attribute specifier applies.

C++

A *directive* and its associated base language code constitute a syntactic formation that follows the syntax given below unless otherwise specified. The *end-directive* in a specified formation refers to the paired **end** directive for the *directive*. An OpenMP construct is a formation for which the *directive* is executable.

Directives with an association of none are not associated with any base language code. The resulting formation therefore has the following syntax:

```
directive
```

Formations that result from a block-associated directive have the following syntax:

C / C++

```
directive  
structured-block
```

C / C++

Fortran

```
directive  
structured-block  
[end-directive]
```

If *structured-block* is a loosely structured block, *end-directive* is required. If *structured-block* is a strictly structured block, *end-directive* is optional. If *structured-block* is an OpenMP context-specific structured block, the *end-directive* is required, unless otherwise specified. An *end-directive* that immediately follows a *directive* and its associated strictly structured block or OpenMP context-specific structured block is always paired with that *directive*.

Fortran

1 Loop-associated directives are block-associated directives for which the associated *structured-block*
2 is *loop-nest*, a canonical loop nest.

3  **Fortran**

For a loop-associated directive, the paired **end** directive is optional.

4  **Fortran**

5  **C / C++**

Formations that result from a declaration-associated directive have the following syntax:

6 ***declaration-associated-specification***

where *declaration-associated-specification* is either:

7 ***directive***

8 ***function-definition-or-declaration***

or:

9 ***directive***

10 ***declaration-associated-specification***

11 In all cases the *directive* is associated with the *function-definition-or-declaration*.

12  **C / C++**

13  **Fortran**

The formation that results from a declaration-associated directive in Fortran has the same syntax as the formation for a directive with an association of none.

15 If a directive appears in the specification part of a module then the behavior is as if that directive
16 appears in the specification part of any program unit that references the module with a **USE**
17 statement.

18  **Fortran**

The formation that results from a delimited directive has the following syntax:

19 ***directive***

20 ***base-language-code***

21 ***end-directive***

22 Separating directives are used to split a *structured-block* that is associated with a construct into
23 multiple *structured-block-sequences*.

24 Separating directives and the containing structured block have the following syntax:

25 ***structured-block-sequence***

26 ***directive***

27 ***structured-block-sequence***

28 **[*directive***

29 ***structured-block-sequence ...]***

1 wrapped in a single compound statement for C/C++ or optionally wrapped in a single **BLOCK**
2 construct for Fortran.

C++

3 Formations that result from directives that are specified as attribute specifiers that use the
4 **directive** attribute are specified as follows. If the directive has an association of none, the
5 resulting formation consists of the attribute specifier and a null statement (i.e., “;”) to which the
6 specifier applies. If the directive is block-associated or loop-associated, the resulting formation
7 consists of the attribute specifier and a structured block to which the specifier applies. If the
8 directives are separating or delimited then the resulting formation is as previously specified except
9 the attribute specifier for each directive, including the **end** directive, applies to a null statement.

10 Formations that result from directives that are specified as attribute specifiers and are
11 declaration-associated or use the **decl** attribute are specified as follows. If the directives are
12 declaration-associated then the resulting formation consists of the attribute specifiers and the
13 *function-definition-or-declaration* to which the specifiers apply. If the directive uses the **decl**
14 attribute then the resulting formation consists of the attribute specifier and the variable and/or
15 function declarations to which the specifier applies.

C++

16 Restrictions

17 Restrictions to directive format are as follows:

- 18 • Orphaned separating directives are prohibited. That is, the separating directives must appear
19 within the structured block associated with the same construct with which it is associated and
20 must not be encountered elsewhere in the region of that associated construct.
- 21 • A stand-alone directive may be placed only at a point where a base language executable
22 statement is allowed.

Fortran

- 23 • OpenMP directives, except **simd** and declarative directives, may not appear in pure
24 procedures.
- 25 • OpenMP directives may not appear in the **WHERE**, **FORALL** or **DO CONCURRENT** constructs.
- 26 • A declarative directive must be specified in the specification part after all **USE**, **IMPORT** and
27 **IMPLICIT** statements.

Fortran

C++

- A directive that uses the attribute syntax cannot be applied to the same statement or associated declaration as a directive that uses the pragma syntax.
- For any directive that has a paired **end** directive, both directives must use either the attribute syntax or the pragma syntax.
- Neither a stand-alone directive nor a declarative directive may be used in place of a substatement in a selection statement or iteration statement, or in place of the statement that follows a label.
- 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 the function as permitted by the base language.

C++

C

- Neither a stand-alone directive nor a declarative directive may be used in place of a substatement in a selection statement, in place of the loop body in an iteration statement, or in place of the statement that follows a label.

C

Fortran

3.1.1 Fixed Source Form Directives

The following sentinels are recognized in fixed form source files:

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

The sentinels that end with **omp** are reserved for OpenMP directives that are defined in this specification. The sentinels that end with **omx** are reserved for implementation-defined extensions to the OpenMP directives.

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.

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

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

11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29


3.1.2 Free Source Form Directives

12 The following sentinels are recognized in free form source files:

```
13 | !$omp | !$ompx
```

15 The **!\$omp** sentinel is reserved for OpenMP directives that are defined in this specification. The
16 **!\$omp~~x~~** sentinel is reserved for implementation-defined extensions to the OpenMP directives.

17 The sentinel can appear in any column as long as it is preceded only by white space. It must appear
18 as a single word with no intervening white space. Fortran free form line length and white space
19 rules apply to the directive line. Initial directive lines must have a space after the sentinel. The
20 initial line of a directive must not be a continuation line for a base language statement. Fortran free
21 form continuation rules apply. Thus, continued directive lines must have an ampersand (&) as the
22 last non-blank character on the line, prior to any comment placed inside the directive; continuation
23 directive lines can have an ampersand after the directive sentinel with optional white space before
24 and after the ampersand.

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

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

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

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

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

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

Fortran

3.2 Clause Format

This section defines the format and categories of OpenMP clauses. OpenMP clauses are specified as part of a *directive-specification*. Clauses are optional and, thus, may be omitted from a *directive-specification* unless otherwise specified. The order in which clauses appear on directives is not significant unless otherwise specified. A *clause-specification* specifies each OpenMP clause in a *directive-specification* where *clause-specification* for inarguable clauses is simply:

clause-name

Inarguable clauses often form natural groupings that have similar semantic effect and so are frequently specified as a clause grouping. For argument-modified clauses, *clause-specification* is:

clause-name [(*clause-argument-specification* [; *clause-argument-specification* [; ...]])]

C / C++

White space in a *clause-name* is prohibited. White space within a *clause-argument-specification* and between another *clause-argument-specification* is optional.

C / C++

An implementation may allow clauses with clause names that start with the `ompx_` prefix for use on any OpenMP directive, and the format and semantics of any such clause is implementation defined. All other clause names are reserved.

For argument-modified clauses, the first *clause-argument-specification* is required unless otherwise explicitly stated while additional ones are only permitted on clauses that explicitly allow them. When the first one is omitted, the syntax is identical to an inarguable clause. Clause arguments may be unmodified or modified. For an unmodified argument, *clause-argument-specification* is:

clause-argument-list

1 Unless otherwise specified, modified arguments are pre-modified, for which the format is:

2 `[modifier-specification-list :]clause-argument-list`

3 A few modified arguments are explicitly specified as post-modified, for which the format is:

4 `clause-argument-list[: modifier-specification-list]`

5 For many OpenMP clauses, *clause-argument-list* is an OpenMP argument list, which is a
6 comma-separated list of a specific kind of list items (see [Section 3.2.1](#)), in which case the format of
7 *clause-argument-list* is:

8 `argument-name`

9 For all other OpenMP clauses, *clause-argument-list* is a comma-separated list of arguments so the
10 format is:

11 `argument-name [, argument-name [, ...]]`

12 In most of these cases, the list only has a single item so the format of *clause-argument-list* is again:

13 `argument-name`

14 In all cases, white space in *clause-argument-list* is optional.

15 A *modifier-specification-list* is a comma-separated list of clause argument modifiers for which the
16 format is:

17 `modifier-specification [[, modifier-specification] , ...]`

18 Clause argument modifiers may be simple or complex. Almost all clause argument modifiers are
19 simple, for which the format of *modifier-specification* is:

20 `modifier-name`

21 The format of a complex modifier is:

22 `modifier-name (modifier-parameter-specification)`

23 where *modifier-parameter-specification* is a comma-separated list of arguments as defined above for
24 *clause-argument-list*. The position of each *modifier-argument-name* in the list is significant.

25 Each *argument-name* and *modifier-name* is an OpenMP term that may be used in the definitions of
26 the clause and any directives on which the clause may appear. Syntactically, each of these terms is
27 one of the following:

- 28 • *keyword*: An OpenMP keyword
- 29 • *OpenMP identifier*: An OpenMP identifier
- 30 • *OpenMP argument list*: An OpenMP argument list
- 31 • *expression*: An expression of some OpenMP type
- 32 • *OpenMP stylized expression*: An OpenMP stylized expression

1 A particular lexical instantiation of an argument specifies a parameter of the clause, while a lexical
 2 instantiation of a modifier and its parameters affects how or when the argument is applied.

3 The order of arguments must match the order in the *clause-specification*. The order of modifiers in
 4 a *clause-argument-specification* is not significant unless otherwise specified.

5 General syntactic properties govern the use of clauses, clause and directive arguments, and
 6 modifiers in an OpenMP directive. These properties are summarized in Table 3.1, along with the
 7 respective default properties for clauses, arguments and modifiers.

TABLE 3.1: Syntactic Properties for Clauses, Arguments and Modifiers

Property	Property Description	Inverse Property	Clause defaults	Argument defaults	Modifier defaults
required	must be present	optional	optional	required	optional
unique	may appear at most once	repeatable	repeatable	unique	unique
exclusive	must appear alone	compatible	compatible	compatible	compatible
ultimate	must lexically appear last (or first for a modifier in a post-modified clause)	free	free	free	free

8 A clause, argument or modifier with a given property implies that it does not have the
 9 corresponding inverse property, and vice versa. The ultimate property implies the unique property.
 10 If all arguments and modifiers of an argument-modified clause or directive are optional and omitted
 11 then the parentheses of the syntax for the clause or directive is also omitted.

12 Arguments and modifiers that are expressions may additionally have any of the following value
 13 properties: constant, positive, non-negative, and region-invariant.

14
 15 **Note** – In this example, *clause-specification* is **depend (inout : d)**, *clause-name* is **depend**
 16 and *clause-argument-specification* is **inout : d**. The **depend** clause has an argument for which
 17 *argument-name* is *locator-list*, which syntactically is the OpenMP locator list **d** in the example.
 18 Similarly, the **depend** clause accepts a simple clause modifier with the name
 19 *task-dependence-type*. Syntactically, *task-dependence-type* is the keyword **inout** in the example.

20 **#pragma omp depobj(o) depend(inout : d)**

21
 22 The clauses that a directive accepts may form sets. These sets may imply restrictions on their use
 23 on that directive or may otherwise capture properties for the clauses on the directive. While specific
 24 properties may be defined for a clause set on a particular directive, the following clause-set
 25 properties have general meanings and implications as indicated by the restrictions below: required,
 26 unique, and exclusive.

1 All clauses that are specified as a clause grouping form a clause set for which properties are
2 specified with the specification of the grouping. Some directives accept a clause grouping for which
3 each member is a *directive-name* of a directive that has a specific property. These groupings are
4 required, unique and exclusive unless otherwise specified.

5 The restrictions for a directive apply to the union of the clauses on the directive and its paired **end**
6 directive.

7 **Restrictions**

8 Restrictions to clauses and clause sets are as follows:

- 9 • A required clause for a directive must appear on the directive.
- 10 • A unique clause for a directive may appear at most once on the directive.
- 11 • An exclusive clause for a directive must not appear if a clause with a different *clause-name*
12 also appears on the directive.
- 13 • An ultimate clause for a directive must be the lexically last clause to appear on the directive.
- 14 • If a clause set has the required property, at least one clause in the set must be present on the
15 directive for which the clause set is specified.
- 16 • If a clause is a member of a set that has the unique property for a directive then the clause has
17 the unique property for that directive regardless of whether it has the unique property when it
18 is not part of such a set.
- 19 • If one clause of a clause set with the exclusive property appears on a directive, no other
20 clauses with a different *clause-name* in that set may appear on the directive.
- 21 • A required argument must appear in the *clause-specification*, unless otherwise specified.
- 22 • A unique argument may appear at most once in a *clause-argument-specification*.
- 23 • An exclusive argument must not appear if an argument with a different *argument-name*
24 appears in the *clause-argument-specification*.
- 25 • A required modifier must appear in the *clause-argument-specification*.
- 26 • A unique modifier may appear at most once in a *clause-argument-specification*.
- 27 • An exclusive modifier must not appear if a modifier with a different *modifier-name* also
28 appears in the *clause-argument-specification*.
- 29 • If a clause is pre-modified, an ultimate modifier must be the last modifier in a
30 *clause-argument-specification* in which any modifier appears.
- 31 • If a clause is post-modified, an ultimate modifier must be the first modifier in a
32 *clause-argument-specification* in which any modifier appears.
- 33 • A modifier that is an expression must neither lexically match the name of a simple modifier
34 defined for the clause that is an OpenMP keyword nor *modifier-name parenthesized-tokens*,

1 where *modifier-name* is the *modifier-name* of a complex modifier defined for the clause and
2 *parenthesized-tokens* is a token sequence that starts with (and ends with).

- 3 • A constant argument or parameter must be a compile-time constant.
- 4 • A positive argument or parameter must be greater than zero; a non-negative argument or
5 parameter must be greater than or equal to zero.
- 6 • A region-invariant argument or parameter must have the same value throughout any given
7 execution of the construct or, for declarative directives, execution of the function or
8 subroutine with which the declaration is associated.

9 Cross References

- 10 • Directive Format, see [Section 3.1](#)
- 11 • OpenMP Argument Lists, see [Section 3.2.1](#)
- 12 • OpenMP Stylized Expressions, see [Section 4.2](#)
- 13 • OpenMP Types and Identifiers, see [Section 4.1](#)

14 3.2.1 OpenMP Argument Lists

15 The OpenMP API defines several kinds of lists, each of which can be used as syntactic instances of
16 clause arguments. A list of any OpenMP type consists of a comma-separated collection of one or
17 more expressions of that OpenMP type. A variable list consists of a comma-separated collection of
18 one or more *variable list items*. An extended list consists of a comma-separated collection of one or
19 more *extended list items*. A locator list consists of a comma-separated collection of one or more
20 *locator list items*. A parameter list consists of a comma-separated collection of one or more
21 *parameter list items*. A type-name list consists of a comma-separated collection of one or more
22 *type-name list items*. A directive-name list consists of a comma-separated collection of one or more
23 *directive-name list items*, each of which is the *directive-name* of some OpenMP directive. A
24 directive specification list consists of a comma-separated collection of one or more
25 *directive-specification* items. A foreign runtime preference list consists of a comma-separated
26 collection of one or more *foreign-runtime list items*, each of which is an OpenMP *foreign-runtime*
27 identifier. An OpenMP operation list consists of a comma-separated collection of one or more
28 *OpenMP operation list items*, each of which is an OpenMP operation defined in [Section 3.2.3](#).

▼ C / C++ ▼

29 A *variable list item* is a variable or an array section. An *extended list item* is a *variable list item* or a
30 function name. A *locator list item* is any lvalue expression including variables, array sections, and
31 reserved locators. A *parameter list item* is the name of a function parameter. A *type-name list item*
32 is a type name.

▲ C / C++ ▲

1 A *variable list item* is one of the following:

- 2 • a variable that is not coindexed and that is not a substring;
- 3 • an array section that is not coindexed and that does not contain an element that is a substring;
- 4 • a *named constant*;
- 5 • an associate name that may appear in a variable definition context; or
- 6 • a common block name (enclosed in slashes).

7 An *extended list item* is a *variable list item* or a procedure name. A *locator list item* is a *variable list*
 8 *item*, a function reference with data pointer result, or a reserved locator. A *parameter list item* is a
 9 dummy argument of a subroutine or function. A *type-name list item* is a type specifier that must not
 10 be **CLASS (*)** or an abstract type.

11 A *named constant* as a *list item* can appear only in clauses where it is explicitly allowed.

12 When a named common block appears in an OpenMP argument list, it has the same meaning and
 13 restrictions as if every explicit member of the common block appeared in the list. An explicit
 14 member of a common block is a variable that is named in a **COMMON** statement that specifies the
 15 common block name and is declared in the same scoping unit in which the clause appears. Named
 16 common blocks do not include the blank common block.

17 Although variables in common blocks can be accessed by use association or host association,
 18 common block names cannot. As a result, a common block name specified in a clause must be
 19 declared to be a common block in the same scoping unit in which the clause appears.

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

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

26 Restrictions

27 The restrictions to OpenMP lists are as follows:

- 28 • Unless otherwise specified, OpenMP list items must be directive-wide unique, i.e., a list item
 29 can only appear once in one OpenMP list of all arguments, clauses, and modifiers of the
 30 directive.
- 31 • All list items must be visible, according to the scoping rules of the base language.
- 32 • The *directive-specifier* and the clauses in a *directive-specification* item must not be
 33 comma-separated.

C

- Unless otherwise specified, a variable that is part of another variable (as an array element or a structure element) must not be a variable list item, an extended list item or a locator list item.

C

C++

- Unless otherwise specified, a variable that is part of another variable (as an array element or a structure element) must not be a variable list item, an extended list item or locator list item except if the list appears on a clause that 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

- Unless otherwise specified, a variable that is part of another variable (as an array element or a structure element) must not be a variable list item, an extended list item or locator list item.

Fortran

3.2.2 Reserved Locators

On some directives, some clauses accept the use of reserved locators as special identifiers that represent system storage not necessarily bound to any base language storage item. Reserved locators may only appear in clauses and directives where they are explicitly allowed and may not otherwise be referenced in the program. The list of reserved locators is:

`omp_all_memory`

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

3.2.3 OpenMP Operations

On some directives, some clauses accept the use of OpenMP operations. An OpenMP operation named `<generic_name>` is a special expression that may be specified in an OpenMP operation list and that is used to construct an object of the `<generic_name>` OpenMP type (see [Section 4.1](#)). In general, the format of an OpenMP operation is the following:

`<generic_name> (operation-parameter-specification)`

3.2.4 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 := ( [s1] [s2] . . . [sn] ) 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 for which 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.

- 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 / C++

3.2.5 Array Sections

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

C / C++

To specify an array section in an OpenMP directive, array subscript expressions are extended with one of the following syntaxes:

```
[ lower-bound : length : stride ]
[ lower-bound : length : ]
[ lower-bound : length ]
[ lower-bound : : stride ]
[ lower-bound : : ]
[ lower-bound : ]
[ : length : stride ]
[ : length : ]
[ : length ]
[ : : stride ]
[ : : ]
[ : ]
```

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 *stride* is absent it defaults to 1.

When the *length* is absent and the size of the dimension is known, it defaults to $\lceil (size - lower-bound) / stride \rceil$, where *size* is the size of the array dimension. When the *length* is absent and the size of the dimension is not known, the array section is an assumed-size array.

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]
a[1:10]
a[1:]
a[:10:2]
b[10][:][:]
b[10][:][:0]
c[42][0:6][:]
c[42][0:6:2][:]
c[1:10][42][0:6]
S.c[:100]
p->y[:10]
this->a[:N]
(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

1 section. The only operator that may be applied to an array section is a subscript operator for which
2 the array section appears as the postfix expression.
3
4

▲ C / C++ ▲

▼ Fortran ▼

5 Fortran has built-in support for array sections although some restrictions apply to their use in
6 OpenMP directives, as enumerated in the following section.

▲ Fortran ▲

7 **Restrictions**

8 Restrictions to array sections are as follows:

- 9
- An array section can appear only in clauses for which it is explicitly allowed.
 - A *stride* expression may not be specified unless otherwise stated.
- 10

▼ C / C++ ▼

- 11
- An assumed-size array can appear only in clauses for which it is explicitly allowed.
 - An element of an array section with a non-zero size must have a complete type.
 - The base expression of an array section must have an array or pointer type.
 - If a consecutive sequence of array subscript expressions appears in an array section, and the first subscript expression in the sequence uses the extended array section syntax defined in this section, then only the last subscript expression in the sequence may select array elements that have a pointer type.
- 12
13
14
15
16
17

▲ C / C++ ▲

▼ C++ ▼

- 18
- If the type of the base expression of an array section is a reference to a type *T*, then the type will be considered to be *T* for all purposes of the array section.
 - An array section cannot be used in an overloaded `[]` operator.
- 19
20

▲ C++ ▲

▼ Fortran ▼

- 21
- If a stride expression is specified, it must be positive.
 - The upper bound for the last dimension of an assumed-size dummy array must be specified.
 - If a list item is an array section with vector subscripts, the first array element must be the lowest in the array element order of the array section.
 - If a list item is an array section, the last *part-ref* of the list item must have a section subscript list.
- 22
23
24
25
26

▲ Fortran ▲

3.2.6 iterator Modifier

Modifiers

Name	Modifies	Type	Properties
<i>iterator</i>	<i>locator-list</i>	Complex, name: iterator Arguments: iterator-specifier OpenMP expression (repeatable)	unique

Clauses

affinity, depend, from, map, to

An *iterator* modifier is a unique, complex modifier that defines a set of iterators, each of which is an *iterator-identifier* and an associated set of values. An *iterator-identifier* expands to those values in the clause argument for which it is specified. Each member of the *modifier-parameter-specification* list of an *iterator* modifier is an *iterator-specifier* with this format:

	C / C++		
[<i>iterator-type</i>] <i>iterator-identifier</i> = <i>range-specification</i>			
	C / C++		
	Fortran		
[<i>iterator-type</i> ::] <i>iterator-identifier</i> = <i>range-specification</i>			
	Fortran		

where:

- *iterator-identifier* is a base-language identifier.
- *iterator-type* is a type that is permitted in a type-name list.
- *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 <i>iterator-specifier</i> , if the <i>iterator-type</i> is not specified then that iterator is of int type.			
	C / C++		
	Fortran		
In an <i>iterator-specifier</i> , if the <i>iterator-type</i> is not specified then that iterator has default integer type.			
	Fortran		

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

2 An iterator only exists in the context of the clause argument that it modifies. An iterator also hides
3 all accessible symbols with the same name in the context of that clause argument.

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

C / C++

6 The values of the iterator are the set of values i_0, \dots, i_{N-1} where:

- 7 • $i_0 = (\text{iterator-type}) \text{ begin};$
- 8 • $i_j = (\text{iterator-type}) (i_{j-1} + \text{step}),$ where $j \geq 1;$ and
- 9 • if $\text{step} > 0,$
 - 10 – $i_0 < (\text{iterator-type}) \text{ end};$
 - 11 – $i_{N-1} < (\text{iterator-type}) \text{ end};$ and
 - 12 – $(\text{iterator-type}) (i_{N-1} + \text{step}) \geq (\text{iterator-type}) \text{ end};$
- 13 • if $\text{step} < 0,$
 - 14 – $i_0 > (\text{iterator-type}) \text{ end};$
 - 15 – $i_{N-1} > (\text{iterator-type}) \text{ end};$ and
 - 16 – $(\text{iterator-type}) (i_{N-1} + \text{step}) \leq (\text{iterator-type}) \text{ end}.$

C / C++

Fortran

17 The values of the iterator are the set of values i_1, \dots, i_N where:

- 18 • $i_1 = \text{begin};$
- 19 • $i_j = i_{j-1} + \text{step},$ where $j \geq 2;$ and
- 20 • if $\text{step} > 0,$
 - 21 – $i_1 \leq \text{end};$
 - 22 – $i_N \leq \text{end};$ and
 - 23 – $i_N + \text{step} > \text{end};$
- 24 • if $\text{step} < 0,$
 - 25 – $i_1 \geq \text{end};$
 - 26 – $i_N \geq \text{end};$ and
 - 27 – $i_N + \text{step} < \text{end}.$

Fortran

1 The set of values will be empty if no possible value complies with the conditions above.
2 If an *iterator-identifier* appears in a list-item expression of the modified argument, the effect is as if
3 the list item is instantiated within the clause for each member of the iterator value set, substituting
4 each occurrence of *iterator-identifier* in the list-item expression with the iterator value. If the
5 iterator value set is empty then the effect is as if the list item was not specified.

6 Restrictions

7 Restrictions to *iterator* modifiers are as follows:

- 8 • The *iterator-type* must not declare a new type.
- 9 • For each value i in an iterator value set, the mathematical result of $i + step$ must be
10 representable in *iterator-type*.

▼ C / C++ ▼

- 11 • The *iterator-type* must be an integral or pointer type.
- 12 • The *iterator-type* must not be **const** qualified.

▲ C / C++ ▲

▼ Fortran ▼

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

▲ Fortran ▲

- 14 • If the *step* expression of a *range-specification* equals zero, the behavior is unspecified.
- 15 • Each *iterator-identifier* can only be defined once in the *modifier-parameter-specification*.
- 16 • Iterators cannot appear in the *range-specification*.

17 Cross References

- 18 • **affinity** clause, see [Section 12.5.1](#)
- 19 • **depend** clause, see [Section 15.9.5](#)
- 20 • **from** clause, see [Section 5.9.2](#)
- 21 • **map** clause, see [Section 5.8.3](#)
- 22 • **to** clause, see [Section 5.9.1](#)

3.3 Conditional Compilation

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

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

Fortran

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

Fortran

Fortran

3.3.1 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; and
- 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
```

```

1  #ifdef _OPENMP
2      10 iam = omp_get_thread_num() +
3          &          index
4  #endif

```



3.3.2 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 blank character after the sentinel; and
- 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):

```

22  c23456789
23  !$ iam = omp_get_thread_num() +      &
24  !$&  index
25
26  #ifdef _OPENMP
27      iam = omp_get_thread_num() +      &
28      index
29  #endif

```



3.4 `if` Clause

Name: <code>if</code>	Properties: <i>default</i>
-----------------------	----------------------------

Arguments

Name	Type	Properties
<i>if-expression</i>	expression of logical type	<i>default</i>

Modifiers

Name	Modifies	Type	Properties
<i>directive-name-modifier</i>	<i>if-expression</i>	Keyword: directive-name	unique

Directives

`cancel`, `parallel`, `simd`, `target`, `target data`, `target enter data`, `target exit data`, `target update`, `task`, `taskloop`, `teams`

Semantics

If no *directive-name-modifier* is specified then the effect is as if a *directive-name-modifier* was specified with the *directive-name* of the directive on which the clause appears.

The effect of the `if` clause depends on the construct to which it is applied. If the construct is not a combined or composite construct then the effect is described in the section that describes that construct. For combined or composite constructs, the `if` clause only applies to the semantics of the construct named in the *directive-name-modifier*. For a combined or composite construct, if no *directive-name-modifier* is specified then the `if` clause applies to all constituent constructs to which an `if` clause can apply.

Restrictions

Restrictions to the `if` clause are as follows:

- At most one `if` clause can be specified that applies to the semantics of any construct or constituent construct of a *directive-specification*.
- The *directive-name-modifier* must specify the *directive-name* of the construct or of a constituent construct of the *directive-specification* on which the `if` clause appears.

Cross References

- `cancel` directive, see [Section 16.1](#)
- `parallel` directive, see [Section 10.2](#)
- `simd` directive, see [Section 10.5](#)
- `target` directive, see [Section 13.8](#)
- `target data` directive, see [Section 13.5](#)

- 1 • **target enter data** directive, see [Section 13.6](#)
- 2 • **target exit data** directive, see [Section 13.7](#)
- 3 • **target update** directive, see [Section 13.9](#)
- 4 • **task** directive, see [Section 12.5](#)
- 5 • **taskloop** directive, see [Section 12.6](#)
- 6 • **teams** directive, see [Section 10.3](#)

3.5 destroy Clause

Name: <code>destroy</code>	Properties: <i>default</i>
-----------------------------------	-----------------------------------

Arguments

Name	Type	Properties
<i>destroy-var</i>	variable of OpenMP variable type	<i>default</i>

Directives

[depobj](#), [interop](#)

Semantics

When the **destroy** clause appears on a **depobj** construct, the state of *destroy-var* is set to uninitialized.

When the **destroy** clause appears on an **interop** construct, the *interop-type* is inferred based on the *interop-type* used to initialize *destroy-var*, and *destroy-var* is set to the value of **omp_interop_none** after resources associated with *destroy-var* are released. The object referred to by *destroy-var* is unusable after destruction and the effect of using values associated with it is unspecified until it is initialized again by another **interop** construct.

Restrictions

- *destroy-var* must be non-const.
- If the **destroy** clause appears on a **depobj** construct, *destroy-var* must refer to the same depend object as the *depobj* argument of the construct.
- If the **destroy** clause appears on an **interop** construct *destroy-var* must refer to a variable of OpenMP **interop** type.

Cross References

- **depobj** directive, see [Section 15.9.4](#)
- **interop** directive, see [Section 14.1](#)

4 Base Language Formats and Restrictions

This section defines concepts and restrictions on base language code used in OpenMP. The concepts help support base language neutrality for OpenMP directives and their associated semantics.

Restrictions

The following restrictions apply generally for base language code in an OpenMP program:

- Programs must not declare names that begin with the `omp_` or `omp_x_` prefix, as these are reserved for the OpenMP implementation.

C++

- Programs must not declare a namespace with the `omp` or `omp_x` names, as these are reserved for the OpenMP implementation.

C++

4.1 OpenMP Types and Identifiers

An OpenMP identifier is a special identifier for use within OpenMP directives and clauses for some specific purpose. For example, OpenMP reduction identifiers specify the combiner operation to use in a reduction, OpenMP mapper identifiers specify the name of a user-defined mapper, and OpenMP foreign runtime identifiers specify the name of a foreign runtime.

An OpenMP context-specific constant is a special identifier for use within user code that the implementation implicitly declares and evaluates to a compile-time constant value when referenced in a given context.

Generic OpenMP types specify the type of expression or variable that is used in OpenMP contexts regardless of the base language. These types support the definition of many important OpenMP concepts independently of the base language in which they are used.

The assignable OpenMP type instance is defined to facilitate base language neutrality. An assignable OpenMP type instance can be used as an argument of an OpenMP construct in order for the implementation to modify the value of that instance.

C / C++

An assignable OpenMP type instance is an lvalue expression of that OpenMP type.

C / C++

▼ Fortran ▼
1 An assignable OpenMP type instance is a variable or a function reference with data pointer result of
2 that OpenMP type.
▲

▲ Fortran ▲
3 The OpenMP logical type supports logical variables and expressions in any base language.

▼ C / C++ ▼
4 Any OpenMP logical expression is a scalar expression. This document uses *true* as a generic term
5 for a non-zero integer value and *false* as a generic term for an integer value of zero.
▲

▲ C / C++ ▲
▼ Fortran ▼
6 Any OpenMP logical expression is a scalar logical expression. This document uses *true* as a generic
7 term for a logical value of **.TRUE.** and *false* as a generic term for a logical value of **.FALSE.**
▲

▲ Fortran ▲
8 The OpenMP integer type supports integer variables and expressions in any base language.

▼ C / C++ ▼
9 Any OpenMP integer expression is an integer expression.
▲

▲ C / C++ ▲
▼ Fortran ▼
10 Any OpenMP integer expression is a scalar integer expression.
▲

▲ Fortran ▲
11 The OpenMP string type supports character string variables and expressions in any base language.

▼ C / C++ ▼
12 Any OpenMP string expression is an expression of type qualified or unqualified **const char ***
13 or **char *** pointing to a null-terminated character string.
▲

▲ C / C++ ▲
▼ Fortran ▼
14 Any OpenMP string expression is a character string of default kind.
▲

▲ Fortran ▲

15 OpenMP function identifiers support procedure names in any base language. Regardless of the base
16 language, any OpenMP function identifier is the name of a procedure as a base language identifier.

17 Each OpenMP type other than those specifically defined in this section has a generic name,
18 *<generic_name>*, by which it is referred throughout this document and that is used to construct the
19 base language construct that corresponds to that OpenMP type.

▼ C / C++ ▼
20 A variable of *<generic_name>* OpenMP type is a variable of type **omp_<generic_name>_t**.
▲

▲ C / C++ ▲

Fortran

A variable of `<generic_name>` OpenMP type is a scalar integer variable of kind `omp_<generic_name>_kind`.

Fortran

Cross References

- OpenMP Foreign Runtime Identifiers, see [Section 14.1.1](#)
- OpenMP Reduction Identifiers, see [Section 5.5.1](#)
- `mapper` modifier, see [Section 5.8.2](#)

4.2 OpenMP Stylized Expressions

An OpenMP stylized expression is a base language expression that is subject to restrictions that enable its use within an OpenMP implementation. These expressions often make use of special variable identifiers that the implementation binds to well-defined internal state.

Cross References

- OpenMP Combiner Expressions, see [Section 5.5.2.1](#)
- OpenMP Initializer Expressions, see [Section 5.5.2.2](#)

4.3 Structured Blocks

This section specifies the concept of a structured block. A structured block:

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

C / C++

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

C / C++

Fortran

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

Fortran

C / C++

1 A structured block sequence that consists of no statements or more than one statement may appear
2 only for executable directives that explicitly allow it. The corresponding compound statement
3 obtained by enclosing the sequence in { and } must be a structured block and the structured block
4 sequence then should be considered to be a structured block with all of its restrictions.

C / C++

Restrictions

5 Restrictions to structured blocks are as follows:

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

C / C++

- 8 • The point of entry to a structured block must not be a call to `set jmp`.
- 9 • `long jmp` must not violate the entry/exit criteria of structured blocks.

C / C++

C++

- 10 • `throw`, `co_await`, `co_yield` and `co_return` must not violate the entry/exit criteria of
11 structured blocks.

C++

Fortran

- 12 • If a **BLOCK** construct appears in a structured block, that **BLOCK** construct must not contain
13 any **ASYNCHRONOUS** or **VOLATILE** statements, nor any specification statements that
14 include the **ASYNCHRONOUS** or **VOLATILE** attributes.

Fortran

4.3.1 OpenMP Context-Specific Structured Blocks

15 An OpenMP context-specific structured block consists of statements that conform to specific
16 restrictions so that OpenMP can treat them as a structured block or a structured block sequence.
17 The restrictions depend on the context in which the context-specific structured block can be used.

4.3.1.1 OpenMP Allocator Structured Blocks

Fortran

18 An OpenMP *allocator structured block* consists of *allocate-stmt*, where *allocate-stmt* is a Fortran
19 **ALLOCATE** statement. For an **allocators** directive, the paired **end** directive is optional.

Fortran

Cross References

- `allocators` directive, see [Section 6.7](#)

4.3.1.2 OpenMP Function Dispatch Structured Blocks

An OpenMP *function dispatch structured block* is a context-specific structured block that identifies the location of a function dispatch.

C / C++

A function dispatch structured block is an expression statement with one of the following forms:

```
lvalue-expression = target-call ( [expression-list] );
```

or

```
target-call ( [expression-list] );
```

C / C++

Fortran

A function dispatch structured block is an expression statement with one of the following forms, where *expression* can be a variable or a function reference with data pointer result:

```
expression = target-call ( [arguments] )
```

or

```
CALL target-call [ ( [arguments] ) ]
```

For a `dispatch` directive, the paired `end` directive is optional.

Fortran

Restrictions

Restrictions to the function dispatch structured blocks are as follows:

C++

- The *target-call* expression can only be a direct call.

C++

Fortran

- *target-call* must be a procedure name.
- *target-call* must not be a procedure pointer.

Fortran

Cross References

- `dispatch` directive, see [Section 7.6](#)

4.3.1.3 OpenMP Atomic Structured Blocks

An OpenMP *atomic structured block* is a context-specific structured block that can appear in an **atomic** construct. The form of an atomic structured block depends on the atomic semantics that the directive enforces.

In the following definitions:

C / C++

- x , r (result), and v (as applicable) are lvalue expressions with scalar type.
- e (expected) is an expression with scalar type.
- d (desired) is an expression with scalar type.
- e and v may refer to, or access, the same storage location.
- $expr$ is an expression with scalar type.
- The order operation, $ordop$, is one of $<$, or $>$.
- $binop$ is one of $+$, $*$, $-$, $/$, $\&$, \wedge , $|$, \ll , or \gg .
- $==$ comparisons are performed by comparing the value representation of operand values for equality after the usual arithmetic conversions; if the object representation does not have any padding bits, the comparison is performed as if with **memcmp**.
- For forms that allow multiple occurrences of x , the number of times that x is evaluated is unspecified but will be at least one.
- For forms that allow multiple occurrences of $expr$, the number of times that $expr$ is evaluated is unspecified but will be at least one.
- The number of times that r is evaluated is unspecified but will be at least one.
- Whether d is evaluated if $x == e$ evaluates to *false* is unspecified.

C / C++

Fortran

- x and v (as applicable) are either scalar variables or function references with scalar data pointer result of non-character intrinsic type.
- e (expected) and d (desired) are scalar expressions.
- $expr$ is a scalar expression.
- r (result) is a scalar logical variable.
- $expr-list$ is a comma-separated, non-empty list of scalar expressions.
- *intrinsic-procedure-name* is one of **MAX**, **MIN**, **IAND**, **IOR**, or **IEOR**.
- *operator* is one of $+$, $*$, $-$, $/$, **.AND.**, **.OR.**, **.EQV.**, or **.NEQV.**

- 1 • *equalop* is `==`, `.EQ.`, or `.EQV.`.
- 2 • `==` or `.EQ.` comparisons are performed by comparing the physical representation of operand
- 3 values for equality after the usual conversions as described in the base language, while
- 4 ignoring padding bits, if any.
- 5 • `.EQV.` comparisons are performed as described in the base language.
- 6 • For forms that allow multiple occurrences of *x*, the number of times that *x* is evaluated is
- 7 unspecified but will be at least one.
- 8 • For forms that allow multiple occurrences of *expr*, the number of times that *expr* is evaluated
- 9 is unspecified but will be at least one.
- 10 • The number of times that *r* is evaluated is unspecified but will be at least one.
- 11 • Whether *d* is evaluated if *x equalop e* evaluates to *false* is unspecified.

Fortran

A *read-atomic* structured block can be specified for **atomic** directives that enforce atomic read semantics but not capture semantics.

C / C++

A *read-atomic* structured block is *read-expr-stmt*, a read expression statement that has the following form:

```
v = x;
```

C / C++

Fortran

A *read-atomic* structured block is *read-statement*, a read statement that has the following form:

```
v = x
```

Fortran

A *write-atomic* structured block can be specified for **atomic** directives that enforce atomic write semantics but not capture semantics.

C / C++

A *write-atomic* structured block is *write-expr-stmt*, a write expression statement that has the following form:

```
x = expr;
```

C / C++

Fortran

A *write-atomic* structured block is *write-statement*, a write statement that has the following form:

```
x = expr
```

Fortran

1 An *update-atomic* structured block can be specified for **atomic** directives that enforce atomic
2 update semantics but not capture semantics.

C / C++

3 An *update-atomic* structured block is *update-expr-stmt*, an update expression statement that has one
4 of the following forms:

```
5 x++;  
6 x--;  
7 ++x;  
8 --x;  
9 x binop= expr;  
10 x = x binop expr;  
11 x = expr binop x;
```

C / C++

Fortran

12 An *update-atomic* structured block is *update-statement*, an update statement that has one of the
13 following forms:

```
14 x = x operator expr  
15 x = expr operator x  
16 x = intrinsic-procedure-name (x, expr-list)  
17 x = intrinsic-procedure-name (expr-list, x)
```

Fortran

18 A *conditional-update-atomic* structured block can be specified for **atomic** directives that enforce
19 atomic conditional update semantics but not capture semantics.

C / C++

20 A *conditional-update-atomic* structured block is either *cond-expr-stmt*, a conditional expression
21 statement that has one of the following forms:

```
22 x = expr ordop x ? expr : x;  
23 x = x ordop expr ? expr : x;  
24 x = x == e ? d : x;
```

25 or *cond-update-stmt*, a conditional update statement that has one of the following forms:

```
26 if (expr ordop x) { x = expr; }  
27 if (x ordop expr) { x = expr; }  
28 if (x == e) { x = d; }
```

C / C++

Fortran

A *conditional-update-atomic* structured block is *conditional-update-statement*, a conditional update statement that has one of the following forms:

```
if ( x equalop e ) then  
    x = d  
end if
```

or

```
if ( x equalop e ) x = d
```

For an **atomic** construct with *read-atomic*, *write-atomic*, *update-atomic*, or *conditional-update-atomic* structured block, the paired **end** directive is optional.

Fortran

A *capture-atomic* structured block can be specified for **atomic** directives that enforce capture semantics. It is further categorized as a *write-capture-atomic*, *update-capture-atomic*, or *conditional-update-capture-atomic* structured block, which can be specified for **atomic** directives that enforce write, update or conditional update atomic semantics in addition to capture semantics.

C / C++

A *capture-atomic* structured block is *capture-stmt*, a capture statement that has one of the following forms:

```
v = expr-stmt  
{ v = x; expr-stmt }  
{ expr-stmt v = x; 
```

If *expr-stmt* is *write-expr-stmt* or *expr-stmt* is *update-expr-stmt* as specified above then it is an *update-capture-atomic* structured block. If *expr-stmt* is *cond-expr-stmt* as specified above then it is a *conditional-update-capture-atomic* structured block. In addition, a *conditional-update-capture-atomic* structured block can have one of the following forms:

```
{ v = x; cond-update-stmt }  
{ cond-update-stmt v = x;   
if(x == e) { x = d;  } else { v = x;  }  
{ r = x == e; if(r) { x = d; }  }  
{ r = x == e; if(r) { x = d; } else { v = x; }  }
```

C / C++

Fortran

A *capture-atomic* structured block has one of the following forms:

```
statement  
capture-statement
```

or

```
1 | capture-statement  
2 | statement
```

3 where *capture-statement* has the following form:

```
4 | v = x
```

5 If *statement* is *write-statement* as specified above then it is a *write-capture-atomic* structured block.
6 If *statement* is *update-statement* as specified above then it is an *update-capture-atomic* structured
7 block. If *statement* is *conditional-update-statement* as specified above then it is a
8 *conditional-update-capture-atomic* structured block. In addition, for a
9 *conditional-update-capture-atomic* structured block, *statement* can have the following form:

```
10 | x = expr
```

11 In addition, a *conditional-update-capture-atomic* structured block can have one of the following
12 forms:

```
13 | if (x equalop e) then  
14 |     x = d  
15 | else  
16 |     v = x  
17 | end if
```

18 or

```
19 | r = x equalop e  
20 | if (r) x = d
```

21 or

```
22 | r = x equalop e  
23 | if (r) then  
24 |     x = d  
25 | else  
26 |     v = x  
27 | endif
```

Fortran

28 Restrictions

29 Restrictions to OpenMP atomic structured blocks are as follows:

C / C++

- 30 • In forms where *e* is assigned it must be an lvalue.
- 31 • *r* must be of integral type.
- 32 • During the execution of an **atomic** region, multiple syntactic occurrences of *x* must
33 designate the same storage location.

- 1 • During the execution of an **atomic** region, multiple syntactic occurrences of *r* must
2 designate the same storage location.
- 3 • During the execution of an **atomic** region, multiple syntactic occurrences of *expr* must
4 evaluate to the same value.
- 5 • None of *v*, *x*, *r*, *d* and *expr* (as applicable) may access the storage location designated by any
6 other symbol in the list.
- 7 • In forms that capture the original value of *x* in *v*, *v* and *e* may not refer to, or access, the same
8 storage location.
- 9 • *binop*, *binop=*, *ordop*, *==*, *++*, and *--* are not overloaded operators.
- 10 • The expression *x binop expr* must be numerically equivalent to *x binop (expr)*. This
11 requirement is satisfied if the operators in *expr* have precedence greater than *binop*, or by
12 using parentheses around *expr* or subexpressions of *expr*.
- 13 • The expression *expr binop x* must be numerically equivalent to *(expr) binop x*. This
14 requirement is satisfied if the operators in *expr* have precedence equal to or greater than
15 *binop*, or by using parentheses around *expr* or subexpressions of *expr*.
- 16 • The expression *x ordop expr* must be numerically equivalent to *x ordop (expr)*. This
17 requirement is satisfied if the operators in *expr* have precedence greater than *ordop*, or by
18 using parentheses around *expr* or subexpressions of *expr*.
- 19 • The expression *expr ordop x* must be numerically equivalent to *(expr) ordop x*. This
20 requirement is satisfied if the operators in *expr* have precedence equal to or greater than
21 *ordop*, or by using parentheses around *expr* or subexpressions of *expr*.
- 22 • The expression *x == e* must be numerically equivalent to *x == (e)*. This requirement is
23 satisfied if the operators in *e* have precedence equal to or greater than *==*, or by using
24 parentheses around *e* or subexpressions of *e*.



- 25 • *x* must not have the **ALLOCATABLE** attribute.
- 26 • During the execution of an **atomic** region, multiple syntactic occurrences of *x* must
27 designate the same storage location.
- 28 • During the execution of an **atomic** region, multiple syntactic occurrences of *r* must
29 designate the same storage location.
- 30 • During the execution of an **atomic** region, multiple syntactic occurrences of *expr* must
31 evaluate to the same value.
- 32 • None of *v*, *x*, *d*, *r*, *expr*, and *expr-list* (as applicable) may access the same storage location as
33 any other symbol in the list.

- In forms that capture the original value of x in v , v may not access the same storage location as e .
- If *intrinsic-procedure-name* refers to **IAND**, **IOR**, or **IEOR**, exactly one expression must appear in *expr-list*.
- The expression x *operator* $expr$ must be, depending on its type, either mathematically or logically equivalent to x *operator* ($expr$). This requirement is satisfied if the operators in $expr$ have precedence greater than *operator*, or by using parentheses around $expr$ or subexpressions of $expr$.
- The expression $expr$ *operator* x must be, depending on its type, either mathematically or logically equivalent to ($expr$) *operator* x . This requirement is satisfied if the operators in $expr$ have precedence equal to or greater than *operator*, or by using parentheses around $expr$ or subexpressions of $expr$.
- The expression x *equalop* e must be, depending on its type, either mathematically or logically equivalent to x *equalop* (e). This requirement is satisfied if the operators in e have precedence equal to or greater than *equalop*, or by using parentheses around e or subexpressions of e .
- *intrinsic-procedure-name* must refer to the intrinsic procedure name and not to other program entities.
- *operator* must refer to the intrinsic operator and not to a user-defined operator.
- All assignments must be intrinsic assignments.

Fortran

Cross References

- **atomic** directive, see [Section 15.8.5](#)

4.4 Loop Concepts

OpenMP semantics frequently involve loops that occur in the base language code. As detailed in this section, OpenMP defines several concepts that facilitate the specification of those semantics and their associated syntax.

4.4.1 Canonical Loop Nest Form

A loop nest has *canonical loop nest form* if it conforms to *loop-nest* in the following grammar:

loop-nest One of the following:

C / C++

```

for (init-expr; test-expr; incr-expr)
  loop-body
```

C / C++

1 or

C++

```
2 for (range-decl: range-expr)
3   loop-body
```

4 A range-based **for** loop is equivalent to a regular **for** loop using iterators, as
5 defined in the base language. A range-based **for** loop has no iteration variable.

C++

6 or

Fortran

```
7 DO [ label ] var = lb , ub [ , incr ]
8   [intervening-code]
9   loop-body
10  [intervening-code]
11  [ label ] END DO
```

12 If the *loop-nest* is a *nonblock-do-construct*, it is treated as a *block-do-construct*
13 for each **DO** construct.

14 The value of *incr* is the increment of the loop. If not specified, its value is
15 assumed to be 1.

Fortran

16 or

```
17 loop-transforming-construct
```

18 or

```
19 generated-canonical-loop
```

20 *loop-body* One of the following:

```
21 loop-nest
```



22 or

C / C++

```
23 {
24   [intervening-code]
25   loop-body
26   [intervening-code]
27 }
```



C / C++

28 or



9 *loop-transformation-construct*
10 A loop-transforming construct.

11 *generated-canonical-loop*
12 A generated loop from a loop-transforming construct that has canonical loop nest
13 form and for which the loop body matches *loop-body*.



21 Additionally, intervening code must not contain OpenMP directives or calls to the
22 OpenMP runtime API in its corresponding region. If intervening code is present,
23 then a loop at the same depth within the loop nest is not a perfectly nested loop.

24 *final-loop-body* A structured block that terminates the scope of loops in the loop nest. If the loop
25 nest is associated with a loop-associated directive, loops in this structured block
26 cannot be associated with that directive.

1	<i>init-expr</i>	One of the following:	
2		<i>var = lb</i>	
3		<i>integer-type var = lb</i>	
4		<i>pointer-type var = lb</i>	C
			C
5		<i>random-access-iterator-type var = lb</i>	C++
			C++
6	<i>test-expr</i>	One of the following:	
7		<i>var relational-op ub</i>	
8		<i>ub relational-op var</i>	
9	<i>relational-op</i>	One of the following:	
10		<	
11		<=	
12		>	
13		>=	
14		!=	
15	<i>incr-expr</i>	One of the following:	
16		<i>++var</i>	
17		<i>var++</i>	
18		<i>--var</i>	
19		<i>var --</i>	
20		<i>var += incr</i>	
21		<i>var -= incr</i>	
22		<i>var = var + incr</i>	
23		<i>var = incr + var</i>	
24		<i>var = var - incr</i>	
25		The value of <i>incr</i> , respectively 1 and -1 for the increment and decrement	
26		operators, is the increment of the loop.	
			C / C++
27	<i>var</i>	One of the following:	
28		A variable of a signed or unsigned integer type.	C / C++
			C / C++

1	A variable of a pointer type.	C
2	A variable of a random access iterator type.	C++
3	A scalar variable of integer type.	Fortran

4 *var* is the iteration variable of the loop. It must not be modified during the
5 execution of *intervening-code* or *loop-body* in the loop.

6 *lb, ub* One of the following:
7 Expressions of a type compatible with the type of *var* that are loop invariant with
8 respect to the outermost loop.

9 or

- 10 One of the following:
11 *var-outer*
12 *var-outer + a2*
13 *a2 + var-outer*
14 *var-outer - a2*

15 where *var-outer* is of a type compatible with the type of *var*.

16 or

- 17 If *var* is of an integer type, one of the following:
18 *a2 - var-outer*
19 *a1 * var-outer*
20 *a1 * var-outer + a2*
21 *a2 + a1 * var-outer*
22 *a1 * var-outer - a2*
23 *a2 - a1 * var-outer*
24 *var-outer * a1*
25 *var-outer * a1 + a2*
26 *a2 + var-outer * a1*
27 *var-outer * a1 - a2*
28 *a2 - var-outer * a1*

1 where *var-outer* is of an integer type.

2 *lb* and *ub* are loop bounds. A loop for which *lb* or *ub* refers to *var-outer* is a
3 non-rectangular loop. If *var* is of an integer type, *var-outer* must be of an integer
4 type with the same signedness and bit precision as the type of *var*.

5 The coefficient in a loop bound is 0 if the bound does not refer to *var-outer*. If a
6 loop bound matches a form in which *a1* appears, the coefficient is *-a1* if the
7 product of *var-outer* and *a1* is subtracted from *a2*, and otherwise the coefficient
8 is *a1*. For other matched forms where *a1* does not appear, the coefficient is -1 if
9 *var-outer* is subtracted from *a2*, and otherwise the coefficient is 1.

10 *a1, a2, incr* Integer expressions that are loop invariant with respect to the outermost loop of
11 the loop nest.

12 If the loop is associated with a loop-associated directive, the expressions are
13 evaluated before the construct formed from that directive.

14 *var-outer* The loop iteration variable of a surrounding loop in the loop nest.



15 *range-decl* A declaration of a variable as defined by the base language for range-based **for**
16 loops.

17 *range-expr* An expression that is valid as defined by the base language for range-based **for**
18 loops. It must be invariant with respect to the outermost loop of the loop nest and
19 the iterator derived from it must be a random access iterator.



20 Restrictions

21 Restrictions to canonical loop nests are as follows:



- 22
- 23 ● If *test-expr* is of the form *var relational-op b* and *relational-op* is $<$ or $<=$ then *incr-expr* must
24 cause *var* to increase on each iteration of the loop. If *test-expr* is of the form *var*
25 *relational-op b* and *relational-op* is $>$ or $>=$ then *incr-expr* must cause *var* to decrease on
each iteration of the loop. Increase and decrease are using the order induced by *relational-op*.
 - 26 ● If *test-expr* is of the form *ub relational-op var* and *relational-op* is $<$ or $<=$ then *incr-expr*
27 must cause *var* to decrease on each iteration of the loop. If *test-expr* is of the form *ub*
28 *relational-op var* and *relational-op* is $>$ or $>=$ then *incr-expr* must cause *var* to increase on
29 each iteration of the loop. Increase and decrease are using the order induced by *relational-op*.
 - 30 ● If *relational-op* is $!=$ then *incr-expr* must cause *var* to always increase by 1 or always
31 decrease by 1 and the increment must be a constant expression.

- 1 • *final-loop-body* must not contain any **break** statement that would cause the termination of
2 the innermost loop.

← C / C++ →

← Fortran →

- 3 • *final-loop-body* must not contain any **EXIT** statement that would cause the termination of the
4 innermost loop.

← Fortran →

- 5 • A *loop-nest* must also be a structured block.
- 6 • For a non-rectangular loop, if *var-outer* is referenced in *lb* and *ub* then they must both refer to
7 the same iteration variable.
- 8 • For a non-rectangular loop, let a_{lb} and a_{ub} be the respective coefficients in *lb* and *ub*,
9 $incr_{inner}$ the increment of the non-rectangular loop and $incr_{outer}$ the increment of the loop
10 referenced by *var-outer*. $incr_{inner}(a_{ub} - a_{lb})$ must be a multiple of $incr_{outer}$.
- 11 • The loop iteration variable may not appear in a **threadprivate** directive.

12 **Cross References**

- 13 • **threadprivate** directive, see [Section 5.2](#)
- 14 • Loop-Transforming Constructs, see [Chapter 9](#)

15 **4.4.2 OpenMP Loop-Iteration Spaces and Vectors**

16 A loop-associated directive controls some number of the outermost loops of an associated loop
17 nest, called the associated loops, in accordance with its specified clauses. These associated loops
18 and their loop iteration variables form an OpenMP *loop-iteration space*. OpenMP *loop-iteration*
19 *vectors* allow other directives to refer to points in that loop-iteration space.

20 A loop-transforming construct that appears inside a loop nest is replaced according to its semantics
21 before any loop can be associated with a loop-associated directive that is applied to the loop nest.
22 The depth of the loop nest is determined according to the loops in the loop nest, after any such
23 replacements have taken place. A loop counts towards the depth of the loop nest if it is a base
24 language loop statement or generated loop and it matches *loop-nest* while applying the production
25 rules for canonical loop nest form to the loop nest.

26 The canonical loop nest form allows the iteration count of all associated loops to be computed
27 before executing the outermost loop.

28 For any associated loop, the iteration count is computed as follows:

C / C++

- If *var* has a signed integer type and the *var* operand of *test-expr* after usual arithmetic conversions has an unsigned integer type then the loop iteration count is computed from *lb*, *test-expr* and *incr* using an unsigned integer type corresponding to the type of *var*.
- Otherwise, if *var* has an integer type then the loop iteration count is computed from *lb*, *test-expr* and *incr* using the type of *var*.

C / C++

C

- If *var* has a pointer type then the loop iteration count is computed from *lb*, *test-expr* and *incr* using the type `ptrdiff_t`.

C

C++

- If *var* has a random access iterator type then the loop iteration count is computed from *lb*, *test-expr* and *incr* using the type `std::iterator_traits<random-access-iterator-type>::difference_type`.
- For range-based **for** loops, the loop iteration count is computed from *range-expr* using the type `std::iterator_traits<random-access-iterator-type>::difference_type` where *random-access-iterator-type* is the iterator type derived from *range-expr*.

C++

Fortran

- The loop iteration count is computed from *lb*, *ub* and *incr* using the type of *var*.

Fortran

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

No synchronization is implied during the evaluation of the *lb*, *ub*, *incr* or *range-expr* expressions. Whether, in what order, or how many times any side effects within the *lb*, *ub*, *incr*, or *range-expr* expressions occur is unspecified.

Let the number of loops associated with a construct be *n*. The OpenMP loop-iteration space is the *n*-dimensional space defined by the values of *var_i*, $1 \leq i \leq n$, the iteration variables of the associated loops, with *i* = 1 referring to the outermost loop of the loop nest. An OpenMP loop-iteration vector, which may be used as an argument of OpenMP directives and clauses, then has the form:

$$var_1 [\pm offset_1], var_2 [\pm offset_2], \dots, var_n [\pm offset_n]$$

where *offset_i* is a compile-time constant non-negative OpenMP integer expression that facilitates identification of relative points in the loop-iteration space.

The iterations of some number of associated loops can be collapsed into one larger iteration space that is called the logical iteration space. The particular integer type used to compute the iteration count for the collapsed loop is implementation defined, but its bit precision must be at least that of the widest type that the implementation would use for the iteration count of each loop if it was the only associated loop. OpenMP defines a special loop-iteration vector, `omp_cur_iteration`, for which $offset_i = 0 \forall i$. This loop-iteration vector enables identification of relative points in the logical iteration space as:

$$omp_cur_iteration [\pm logical_offset]$$

where *logical_offset* is a compile-time constant non-negative OpenMP integer expression.

For directives that result in the execution of a collapsed logical iteration space, the number of times that any intervening code between any two loops of the same logical iteration space will be executed is unspecified but will be the same for all intervening code at the same depth, at least once per iteration of the loop that encloses the intervening code and at most once per logical iteration. If the iteration count of any loop is zero and that loop does not enclose the intervening code, the behavior is unspecified.

4.4.3 collapse Clause

Name: <code>collapse</code>	Properties: unique
------------------------------------	---------------------------

Arguments

Name	Type	Properties
<i>n</i>	expression of integer type	<i>default</i>

Directives

`distribute`, `do`, `for`, `loop`, `simd`, `taskloop`

Semantics

The `collapse` clause associates one or more loops with the directive on which it appears for the purpose of identifying the portion of the depth of the canonical loop nest to which to apply the semantics of the directive. The argument *n* specifies the number of loops of the associated loop nest to which to apply those semantics. On all directives on which the `collapse` clause may appear, the effect is as if a value of one was specified for *n* if the `collapse` clause is not specified.

Restrictions

- *n* must not evaluate to a value greater than the depth of the associated loop nest.

Cross References

- **ordered** clause, see [Section 4.4.4](#)
- **distribute** directive, see [Section 11.6](#)
- **do** directive, see [Section 11.5.2](#)
- **for** directive, see [Section 11.5.1](#)
- **loop** directive, see [Section 11.7](#)
- **simd** directive, see [Section 10.5](#)
- **taskloop** directive, see [Section 12.6](#)

4.4.4 ordered Clause

Name: <code>ordered</code>	Properties: unique
-----------------------------------	---------------------------

Arguments

Name	Type	Properties
<i>n</i>	expression of integer type	optional, constant, positive

Directives

[do](#), [for](#), [simd](#)

Semantics

The **ordered** clause associates one or more loops with the directive on which it appears for the purpose of identifying cross-iteration dependences. The argument *n* specifies the number of loops of the associated loop to use for that purpose. If *n* is not specified then the behavior is as if *n* is specified with the same value as is specified for the **collapse** clause on the construct.

Restrictions

- None of the associated loops may be non-rectangular loops.
- The **ordered** clause must not appear on a worksharing-loop directive if the associated loops include the generated loops of a **tile** directive.
- *n* must not evaluate to a value greater than the depth of the associated loop nest.
- If *n* is explicitly specified, the associated loops must be perfectly nested.
- If *n* is explicitly specified and the **collapse** clause is also specified for the **ordered** clause on the same construct, *n* must be greater than or equal to the *n* specified for the **collapse** clause.
- If *n* is explicitly specified, a **linear** clause must not be specified on the same directive.

- 1
- If n is explicitly specified, none of the associated loops may be a range-based **for** loop.

2

Cross References

- 3
- **collapse** clause, see [Section 4.4.3](#)

4

 - **linear** clause, see [Section 5.4.6](#)

5

 - **do** directive, see [Section 11.5.2](#)

6

 - **for** directive, see [Section 11.5.1](#)

7

 - **simd** directive, see [Section 10.5](#)

8

 - **tile** directive, see [Section 9.1](#)

9

4.4.5 Consistent Loop Schedules

10 For constructs formed from loop-associated directives that have consistent schedules, the
11 implementation will guarantee that memory effects of a logical iteration in the first loop nest
12 happen before the execution of the same logical iteration in the second loop nest.

13 Two constructs formed from loop-associated directives have consistent schedules if all of the
14 following conditions hold:

- 15
- The constructs have the same *directive-name*;

16

 - The regions that correspond to the two constructs have the same binding region;

17

 - The constructs have the same reproducible schedule;

18

 - The associated loop nests have identical logical iteration vector spaces; and

19

 - The associated loop nests are either both rectangular or both non-rectangular.

5 Data Environment

This chapter presents directives and clauses for controlling data environments. These clauses and directives include the *data-environment attribute clauses*, which explicitly determine the attributes of list items specified in a *list* parameter. The data-environment attribute clauses form a general clause set for which certain restrictions apply to their use on directives that accept any members of the set. In addition, these clauses are divided into two subsets that also form general clause sets: *data-sharing attribute clauses* and *data-mapping attribute clauses*. Data-sharing attribute clauses control the data-sharing attributes of variables in a construct, indicating whether a variable is shared or private in the outermost scope of the construct. Data-mapping attribute clauses control the data-mapping attributes of variables in a data environment, indicating whether a variable is mapped from the data environment to another device data environment. Additional restrictions apply to the use of these sets on directives that accept any members of them.

5.1 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 5.1.1](#) describes the data-sharing attribute rules for variables referenced in a construct.
- [Section 5.1.2](#) describes the data-sharing attribute rules for variables referenced in a region, but outside any construct.

5.1.1 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 **copyprivate** clause or a data-sharing attribute clause other than the **private** clause on 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.

Fortran

A type parameter inquiry or complex part designator that is referenced in a construct is treated as if its designator is referenced.

Fortran

1 Certain variables and objects have *predetermined* data-sharing attributes for the construct in which
2 they are referenced. The first matching rule from the following list of predetermined data-sharing
3 attribute rules applies for variables and objects that are referenced in a construct.

Fortran

- Variables declared within a **BLOCK** construct inside a construct that do not have the **SAVE** attribute are private.

Fortran

- Variables and common blocks (in Fortran) that appear as arguments in **threadprivate** directives or variables with the **_Thread_local** (in C) or **thread_local** (in C++) storage-class specifier are threadprivate.
- Variables and common blocks (in Fortran) that appear as arguments in **groupprivate** directives are groupprivate.
- Variables and common blocks (in Fortran) that appear as arguments in **local** clauses on **declare target** directives are device local.

C

- Variables with automatic storage duration that are declared in a scope inside the construct are private.

C

C++

- Variables of non-reference type with automatic storage duration that are declared in a scope inside the construct are private.

C++

C / C++

- Objects with dynamic storage duration are shared.

C / C++

- The loop iteration variable in the associated loop of a **simd** construct with just one associated loop is linear with a *linear-step* that is the increment of the associated loop.
- The loop iteration variables in the associated loops of a **simd** construct with multiple associated loops are lastprivate.
- The loop iteration variable in any associated loop of a **loop** construct is lastprivate.
- The loop iteration variable in any associated loop of a loop-associated construct is otherwise private.

C++

- The implicitly declared variables of a range-based **for** loop are private.

C++

Fortran

- Loop iteration variables inside **parallel**, **teams**, or task generating constructs are private in the innermost such construct that encloses the loop.
- Implied-do, **FORALL** and **DO CONCURRENT** indices are private.

Fortran

C / C++

- Variables with static storage duration that are declared in a scope inside the construct are shared.
- If a list item in a **has_device_addr** clause or 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 the construct has a base pointer then the base pointer is private.
- Static data members are shared.
- The **__func__** variable and similar function-local predefined variables are shared.

C / C++

Fortran

- Assumed-size arrays and *named constants* are shared in constructs that are not data-mapping.
- *Named constants* are firstprivate in **target** constructs.
- An associate name that may appear in a variable definition context is shared if its association occurs outside of the construct and otherwise it has the same data-sharing attribute as the selector with which it is associated.

Fortran

Variables with predetermined data-sharing attributes may not be listed in data-sharing attribute clauses, except for the cases listed below. For these exceptions only, listing a predetermined variable in a data-sharing attribute clause is allowed and overrides the variable's predetermined data-sharing attributes.

- The loop iteration variable in any associated loop of a loop-associated construct may be listed in a **private** or **lastprivate** clause.
- If a **simd** construct has just one associated loop then its loop iteration variable may be listed in a **linear** clause with a *linear-step* that is the increment of the associated loop.

C / C++

- Variables with **const**-qualified type with no mutable members may be listed in a **firstprivate** clause, even if they are static data members.
- The `__func__` variable and similar function-local predefined variables may be listed in a **shared** or **firstprivate** clause.

C / C++

Fortran

- Loop iteration variables of loops that are not associated with any OpenMP directive may be listed in data-sharing attribute clauses on the surrounding **teams**, **parallel** or task generating construct, and on enclosed constructs, subject to other restrictions.
- Assumed-size arrays may be listed in a **shared** clause.
- *Named constants* may be listed in a **shared** or **firstprivate** clause.

Fortran

Additional restrictions on the variables that may appear in individual clauses are described with each clause in [Section 5.4](#).

Variables with *explicitly determined* data-sharing attributes are those that are referenced in a given construct and are listed in a data-sharing attribute clause on the construct.

Variables with *implicitly determined* data-sharing attributes are those that are referenced in a given construct and do not have predetermined or explicitly determined data-sharing attributes in that construct.

Rules for variables with implicitly determined data-sharing attributes are as follows:

- In a **parallel**, **teams**, or task generating construct, the data-sharing attributes of these variables are determined by the **default** clause, if present (see [Section 5.4.1](#)).
- In a **parallel** construct, if no **default** clause is present, these variables are shared.
- For constructs other than task generating constructs, if no **default** clause is present, these variables reference the variables with the same names that exist in the enclosing context.
- In a **target** construct, variables that are not mapped after applying data-mapping attribute rules (see [Section 5.8](#)) are **firstprivate**.

C++

- In an orphaned task generating construct, if no **default** clause is present, formal arguments passed by reference are **firstprivate**.

C++

Fortran

- In an orphaned task generating construct, if no **default** clause is present, dummy arguments are firstprivate.

Fortran

- In a task generating construct, if no **default** clause is present, a variable for which the data-sharing attribute is not determined by the rules above and that in the enclosing context is determined to be shared by all implicit tasks bound to the current team is shared.
- In a task generating construct, if no **default** clause is present, a variable for which the data-sharing attribute is not determined by the rules above is firstprivate.

A program is non-conforming if a variable in a task generating construct is implicitly determined to be firstprivate according to the above rules but is not permitted to appear in a **firstprivate** clause according to the restrictions specified in [Section 5.4.4](#).

5.1.2 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 the corresponding 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 as arguments in a **threadprivate** or **groupprivate** directive.
- Objects with dynamic storage duration are shared.
- Static data members are shared unless they appear as arguments in a **threadprivate** or **groupprivate** 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 as arguments in a **threadprivate** or **groupprivate** 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 as arguments in a **threadprivate** or **groupprivate** directive.
- Dummy arguments of called routines in the region that have the **VALUE** attribute are private.
- A dummy argument of a called routine in the region that does not have the **VALUE** attribute is private if the associated actual argument is not shared.
- A dummy argument of a called routine in the region that does not have the **VALUE** attribute is 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.
- Implied-do indices, **DO CONCURRENT** indices, **FORALL** indices, and other local variables declared in called routines in the region are private.

Fortran

5.2 threadprivate Directive

Name: `threadprivate`
Category: declarative

Association: none
Properties: *default*

Arguments

threadprivate (*list*)

Name	Type	Properties
<i>list</i>	list of variable list item type	<i>default</i>

Semantics

The **threadprivate** directive specifies that variables are replicated, with each thread having its own copy. 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.

1 Each copy of a block-scope threadprivate variable that has a dynamic initializer is initialized the
 2 first time its thread encounters its definition; if its thread does not encounter its definition, its
 3 initialization is unspecified.

4 The content of a threadprivate variable can change across a task scheduling point if the executing
 5 thread switches to another task that modifies the variable. For more details on task scheduling, see
 6 [Section 1.3](#) and [Chapter 12](#).

7 In **parallel** regions, references by the primary thread are to the copy of the variable in the thread
 8 that encountered the **parallel** region.

9 During a sequential part, references are to the initial thread's copy of the variable. The values of
 10 data in the initial thread's copy of a threadprivate variable are guaranteed to persist between any
 11 two consecutive references to the variable in the program, provided that no **teams** construct that is
 12 not nested inside of a **target** construct is encountered between the references and that the initial
 13 thread is not executing code inside of a **teams** region. For initial threads that are executing code
 14 inside of a **teams** region, the values of data in the copies of a threadprivate variable of those initial
 15 threads are guaranteed to persist between any two consecutive references to the variable inside that
 16 **teams** region.

17 The values of data in the threadprivate variables of threads that are not initial threads are
 18 guaranteed to persist between two consecutive active **parallel** regions only if all of the
 19 following conditions hold:

- 20 • Neither **parallel** region is nested inside another explicit **parallel** region;
- 21 • The sizes of the thread teams used to execute both **parallel** regions are the same;
- 22 • The thread affinity policies used to execute both **parallel** regions are the same;
- 23 • The value of the *dyn-var* internal control variable in the enclosing task region is *false* at entry
 24 to both **parallel** regions;
- 25 • No **teams** construct that is not nested inside of a **target** construct is encountered between
 26 the **parallel** regions;
- 27 • No construct with an **order** clause that specifies **concurrent** is encountered between the
 28 **parallel** regions; and
- 29 • Neither the **omp_pause_resource** nor **omp_pause_resource_all** routine is
 30 called.

31 If these conditions all hold, and if a threadprivate variable is referenced in both regions, then threads
 32 with the same thread number in their respective regions reference the same copy of that variable.

C / C++

If the above conditions hold, the storage duration, lifetime, and value of a thread's copy of a threadprivate variable that does not appear in any **copyin** clause on the corresponding construct of the second region spans the two consecutive active **parallel** regions. Otherwise, the storage duration, lifetime, and value of a thread's copy of the variable in the second region is unspecified.

C / C++

Fortran

If the above conditions hold, the definition, association, or allocation status of a thread's copy of a threadprivate variable or a variable in a threadprivate common block that is not affected by any **copyin** clause that appears on the corresponding construct of the second region (a variable is affected by a **copyin** clause if the variable appears in the **copyin** clause or it is in a common block that appears in the **copyin** clause) spans the two consecutive active **parallel** regions. Otherwise, the definition and association status of a thread's copy of the variable in the second region are undefined, and the allocation status of an allocatable variable are implementation defined.

If a threadprivate variable or a variable in a threadprivate common block is not affected by any **copyin** clause that appears on the corresponding construct of the first **parallel** region in which it is referenced, the thread's copy of the variable inherits the declared type parameter and the default parameter values from the original variable. The variable or any subobject of the variable is initially defined or undefined according to the following rules:

- If it has the **ALLOCATABLE** attribute, each copy created has an initial allocation status of unallocated;
- If it has the **POINTER** attribute, each copy has the same association status as the initial association status.
- If it does not have either the **POINTER** or the **ALLOCATABLE** attribute:
 - If it is initially defined, either through explicit initialization or default initialization, each copy created is so defined;
 - Otherwise, each copy created is undefined.

Fortran

C++

The order in which any constructors for different threadprivate variables of class type are called is unspecified. The order in which any destructors for different threadprivate variables of class type are called is unspecified. A variable that is part of another variable (as an array element or a structure element) may appear in a **threadprivate** directive only if it is a static data member of a C++ class.

C++

Restrictions

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

- A thread must not reference another thread's copy of a **threadprivate** variable.
- A **threadprivate** variable must not appear as the base variable of a list item in any clause except for the **copyin** and **copyprivate** clauses.
- A program in which an untied task accesses **threadprivate** storage is non-conforming.

C / C++


- Each list item must be a file-scope, namespace-scope, or static block-scope variable.
- No list item may have an incomplete type.
- The address of a **threadprivate** variable must not be an address constant.
- 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, the behavior is unspecified.
- 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 must have static storage duration.
- 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.



- Each list item must be a named variable or a named common block; a named common block must appear between slashes.
- The *list* argument must not include any coarrays or associate names.
- 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 variable may only appear as an argument 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 as an argument in a **threadprivate** directive must be declared in the scope of a module or have the **SAVE** attribute, either explicitly or implicitly.
- The effect of an access to a `threadprivate` variable in a **DO CONCURRENT** construct is unspecified.



Cross References

- **copyin** clause, see [Section 5.7.1](#)
- **order** clause, see [Section 10.4](#)
- *dyn-var* ICV, see [Table 2.1](#)
- Determining the Number of Threads for a **parallel** Region, see [Section 10.2.1](#)

5.3 List Item Privatization

Some data-sharing attribute clauses, including reduction clauses, specify that list items that appear in their *list* argument may be privatized for the construct on which they appear. 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 then, within the same logical iteration of the loops, 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, whether references are to a new list item or the original list item is unspecified. 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, whether the task or SIMD lane receives a new list item is unspecified.

The value and/or allocation status of the original list item will change only:

- If accessed and modified via a 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, whether accesses anywhere in the region through the implicit **this** pointer refer to the new list item or the original list item is unspecified.

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 the 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
13 will 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
17 same as 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 base language constructs
20 such as **EQUIVALENCE** or **COMMON**. If A is a variable that is privatized by a construct and B is a
21 variable 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,
24 and association status of B to become undefined; and
- 25 • Any definition of B , or of its allocation or association status, causes the contents, allocation,
26 and 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++**

- 7 • A variable of class type (or array thereof) that is privatized requires an accessible,
unambiguous default constructor for the class type.

8 **C++**

9 **C / C++**

- 10 • A variable that is privatized must not have a **const**-qualified type unless it is of class type
11 with 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 type.

12 **C / C++**

13 **Fortran**

- 14 • Variables that appear in namelist statements, in variable format expressions, and in
15 expressions for statement function definitions, must not be privatized.
- 16 • Pointers with the **INTENT (IN)** attribute must not be privatized. This restriction does not
17 apply to the **firstprivate** clause.
- 18 • A private variable must not be coindexed or appear as an actual argument to a procedure
where the corresponding dummy argument is a coarray.
- Assumed-size arrays must not be privatized.

19 **Fortran**

20 **5.4 Data-Sharing Attribute Clauses**

21 Several constructs accept clauses that allow a user to control the data-sharing attributes of variables
22 referenced in the construct. Not all of the clauses listed in this section are valid on all directives.

23 The set of clauses that is valid on a particular directive is described with the directive. The
reduction data-sharing attribute clauses are explained in [Section 5.5](#).

24 A list item may be specified in both **firstprivate** and **lastprivate** clauses.

25 **C++**

26 If a variable referenced in a data-sharing attribute clause has a type derived from a template and the
27 program does not otherwise reference that variable, any behavior related to that variable is
unspecified.

C++

Fortran

If individual members of a common block appear in a data-sharing attribute clause other than the **shared** clause, the variables no longer have a Fortran storage association with the common block.

Fortran

5.4.1 default Clause

Name: default	Properties: unique
----------------------	---------------------------

Arguments

Name	Type	Properties
<i>data-sharing-attribute</i>	Keyword: firstprivate , none , private , shared	<i>default</i>

Directives

parallel, **task**, **taskloop**, **teams**

Semantics

The **default** clause determines the implicit data-sharing attribute of certain variables that are referenced in the construct, in accordance with the rules given in [Section 5.1.1](#).

If *data-sharing-attribute* is not **none**, the data-sharing attribute of all variables referenced in the construct that have implicitly determined data-sharing attributes will be *data-sharing-attribute*. If *data-sharing-attribute* is **none**, the data-sharing attribute is not implicitly determined.

Restrictions

Restrictions to the **default** clause are as follows:

- If *data-sharing-attribute* is **none**, each variable that is referenced in the construct and 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.

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, is referenced in the construct, and 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.

C / C++

Cross References

- **parallel** directive, see [Section 10.2](#)
- **task** directive, see [Section 12.5](#)
- **taskloop** directive, see [Section 12.6](#)
- **teams** directive, see [Section 10.3](#)

5.4.2 shared Clause

Name: <code>shared</code>	Properties: data-environment attribute, data-sharing attribute
----------------------------------	---

Arguments

Name	Type	Properties
<i>list</i>	list of variable list item type	<i>default</i>

Directives

`parallel`, `task`, `taskloop`, `teams`

Semantics

The **shared** clause declares one or more list items to be shared by tasks generated by the construct on which it appears. 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 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. A reference to the shared storage that is associated with the dummy argument by any other task must be synchronized with the reference to the procedure to avoid possible data races.

Fortran

Cross References

- `parallel` directive, see [Section 10.2](#)
- `task` directive, see [Section 12.5](#)
- `taskloop` directive, see [Section 12.6](#)
- `teams` directive, see [Section 10.3](#)

5.4.3 private Clause

Name: <code>private</code>	Properties: data-environment attribute, data-sharing attribute, privatization
-----------------------------------	--

Arguments

Name	Type	Properties
<i>list</i>	list of variable list item type	<i>default</i>

Directives

[distribute](#), [do](#), [for](#), [loop](#), [parallel](#), [scope](#), [sections](#), [simd](#), [single](#), [target](#),
[task](#), [taskloop](#), [teams](#)

Semantics

The **private** clause specifies that its list items are to be privatized according to [Section 5.3](#). 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 an **order** clause that specifies **concurrent** is also present.

Restrictions

Restrictions to the **private** clause are as specified in [Section 5.3](#).

Cross References

- **distribute** directive, see [Section 11.6](#)
- **do** directive, see [Section 11.5.2](#)
- **for** directive, see [Section 11.5.1](#)
- **loop** directive, see [Section 11.7](#)
- **parallel** directive, see [Section 10.2](#)
- **scope** directive, see [Section 11.2](#)
- **sections** directive, see [Section 11.3](#)
- **simd** directive, see [Section 10.5](#)
- **single** directive, see [Section 11.1](#)
- **target** directive, see [Section 13.8](#)
- **task** directive, see [Section 12.5](#)
- **taskloop** directive, see [Section 12.6](#)
- **teams** directive, see [Section 10.3](#)
- List Item Privatization, see [Section 5.3](#)

5.4.4 firstprivate Clause

Name: <code>firstprivate</code>	Properties: data-environment attribute, data-sharing attribute, privatization
--	--

Arguments

Name	Type	Properties
<i>list</i>	list of variable list item type	<i>default</i>

Directives

`distribute`, `do`, `for`, `parallel`, `scope`, `sections`, `single`, `target`, `task`,
`taskloop`, `teams`

Semantics

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 5.4.3](#), except as noted. In addition, the new list item is initialized from the original list item that exists 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 construct that is not a work-distribution 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 work-distribution construct, the initial value of the new list item for each implicit task of the threads that execute the construct is the value of the original list item that exists in the implicit task immediately prior to the point in time that the 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++ ▲

▼ C++ ▼

For each variable of class type:

- If the **firstprivate** clause is not on a **target** construct then a copy constructor is invoked to perform the initialization; and
- If the **firstprivate** clause is on a **target** construct then how many copy constructors, if any, are invoked is unspecified.

If copy constructors are called, the order in which copy constructors for different variables of class type are called is unspecified.

▲ C++ ▲

Fortran

1 If the original list item does not have the **POINTER** attribute, initialization of the new list items
2 occurs as if by intrinsic assignment unless the original list item has a compatible type-bound
3 defined assignment, in which case initialization of the new list items occurs as if by the defined
4 assignment. If the original list item that does not have the **POINTER** attribute has the allocation
5 status of unallocated, the new list items will have the same status.

6 If the original list item has the **POINTER** attribute, the new list items receive the same association
7 status as the original list item, as if by pointer assignment.

8 The list items that appear in a **firstprivate** clause may include *named constants*.

Fortran

Restrictions

9 Restrictions to the **firstprivate** clause are as follows:
10

- 11 • A list item that is private within a **parallel** region must not appear in a **firstprivate**
12 clause on a worksharing construct if any of the worksharing regions that arise from the
13 worksharing construct ever bind to any of the **parallel** regions that arise from the
14 **parallel** construct.
- 15 • A list item that is private within a **teams** region must not appear in a **firstprivate**
16 clause on a **distribute** construct if any of the **distribute** regions that arise from the
17 **distribute** construct ever bind to any of the **teams** regions that arise from the **teams**
18 construct.
- 19 • A list item that appears in a **reduction** clause of a **parallel** construct must not appear
20 in a **firstprivate** clause on a worksharing, **task**, or **taskloop** construct if any of the
21 worksharing or **task** regions that arise from the worksharing, **task**, or **taskloop**
22 construct ever bind to any of the **parallel** regions that arise from the **parallel**
23 construct.
- 24 • A list item that appears in a **reduction** clause of a **teams** construct must not appear in a
25 **firstprivate** clause on a **distribute** construct if any of the **distribute** regions
26 that arise from the **distribute** construct ever bind to any of the **teams** regions that arise
27 from the **teams** construct.
- 28 • A list item that appears in a **reduction** clause of a worksharing construct must not appear
29 in a **firstprivate** clause in a **task** construct encountered during execution of any of the
30 worksharing regions that arise from the worksharing construct.

C++

- 31 • A variable of class type (or array thereof) that appears in a **firstprivate** clause requires
32 an accessible, unambiguous copy constructor for the class type.

- If the original list item in a **firstprivate** clause on a work-distribution construct has a reference type then it must bind to the same object for all threads in the binding thread set of the work-distribution region.

C++

Fortran

- If the list item is a polymorphic variable with the **ALLOCATABLE** attribute, the behavior is unspecified.

Fortran

Cross References

- **private** clause, see [Section 5.4.3](#)
- **distribute** directive, see [Section 11.6](#)
- **do** directive, see [Section 11.5.2](#)
- **for** directive, see [Section 11.5.1](#)
- **parallel** directive, see [Section 10.2](#)
- **scope** directive, see [Section 11.2](#)
- **sections** directive, see [Section 11.3](#)
- **single** directive, see [Section 11.1](#)
- **target** directive, see [Section 13.8](#)
- **task** directive, see [Section 12.5](#)
- **taskloop** directive, see [Section 12.6](#)
- **teams** directive, see [Section 10.3](#)

5.4.5 lastprivate Clause

Name: <code>lastprivate</code>	Properties: data-environment attribute, data-sharing attribute, privatization
---------------------------------------	--

Arguments

Name	Type	Properties
<i>list</i>	list of variable list item type	<i>default</i>

Modifiers

Name	Modifies	Type	Properties
<i>lastprivate-modifier</i>	<i>list</i>	Keyword: conditional	default

Directives

`distribute`, `do`, `for`, `loop`, `sections`, `simd`, `taskloop`

Semantics

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 5.4.3. In addition, when a **lastprivate** clause without the **conditional** modifier appears on a directive and the list item is not an iteration variable of any associated loop, the value of each new list item from the sequentially last iteration of the associated loops, or the lexically last structured block sequence associated with a **sections** construct, is assigned to the original list item. When the **conditional** modifier appears on the clause or the list item is an iteration variable of one of the associated loops, if sequential execution of the loop nest would assign a value to the list item then the original list item is assigned the value that the list item would have after sequential execution of the loop nest.

C++

For class types, the copy assignment operator is invoked. The order in which copy assignment operators for different variables of the same class type are invoked is unspecified.

C++

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, any list item that is not an iteration variable of the associated loops and that is not assigned a value by the sequentially last iteration of the loops, or by the lexically last structured block sequence associated with a **sections** construct, has an unspecified value after the construct. When the **conditional** modifier does not appear on the **lastprivate** clause, a list item that is the iteration variable of an associated loop and that would not be assigned a value during sequential execution of the loop nest has an unspecified value 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.

1 Otherwise, if the **lastprivate** clause is used on a construct to which the **nowait** or the
2 **nogroup** clauses are applied, accesses to the original list item may create a data race. To avoid
3 this data race, if an assignment to the original list item occurs then synchronization must be inserted
4 to ensure that the assignment completes and the original list item is flushed to memory.

5 If a list item that appears in a **lastprivate** clause with the **conditional** modifier is
6 modified in the region by an assignment outside the construct or not to the list item then the value
7 assigned to the original list item is unspecified.

8 Restrictions

9 Restrictions to the **lastprivate** clause are as follows:

- 10 ● A list item must not appear in a **lastprivate** clause on a work-distribution construct if
11 the corresponding region binds to the region of a parallelism-generating construct in which
12 the list item is private.
- 13 ● A list item that appears in a **lastprivate** clause with the **conditional** modifier must
14 be a scalar variable.

C++

- 15 ● A variable of class type (or array thereof) that appears in a **lastprivate** clause requires
16 an accessible, unambiguous default constructor for the class type, unless the list item is also
17 specified in a **firstprivate** clause.
- 18 ● A variable of class type (or array thereof) that appears in a **lastprivate** clause requires
19 an accessible, unambiguous copy assignment operator for the class type.
- 20 ● If an original list item in a **lastprivate** clause on a work-distribution construct has a
21 reference type then it must bind to the same object for all threads in the binding thread set of
22 the work-distribution region.

C++

Fortran

- 23 ● A variable that appears in a **lastprivate** clause must be definable.
- 24 ● If the original list item has the **ALLOCATABLE** attribute, the corresponding list item of
25 which the value is assigned to the original item must have an allocation status of allocated
26 upon exit from the sequentially last iteration or lexically last structured block sequence
27 associated with a **sections** construct.
- 28 ● If the list item is a polymorphic variable with the **ALLOCATABLE** attribute, the behavior is
29 unspecified.

Fortran

Cross References

- **private** clause, see [Section 5.4.3](#)
- **distribute** directive, see [Section 11.6](#)
- **do** directive, see [Section 11.5.2](#)
- **for** directive, see [Section 11.5.1](#)
- **loop** directive, see [Section 11.7](#)
- **sections** directive, see [Section 11.3](#)
- **simd** directive, see [Section 10.5](#)
- **taskloop** directive, see [Section 12.6](#)

5.4.6 linear Clause

Name: linear	Properties: data-environment attribute, data-sharing attribute, privatization, post-modified
----------------------------	---

Arguments

Name	Type	Properties
<i>list</i>	list of variable list item type	<i>default</i>

Modifiers

Name	Modifies	Type	Properties
<i>step-simple-modifier</i>	<i>list</i>	OpenMP integer expression	exclusive, region-invariant, unique
<i>step-complex-modifier</i>	<i>list</i>	Complex, name: step Arguments: <i>linear-step</i> expression of integer type (region-invariant)	unique
<i>linear-modifier</i>	<i>list</i>	Keyword: ref , uval , val	unique

Directives

[declare](#) [simd](#), [do](#), [for](#), [simd](#)

Semantics

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 5.4.3](#), except as noted. If the *step-simple-modifier* is specified, the behavior is as if the *step-complex-modifier* is instead specified with *step-simple-modifier* as its *linear-step* argument. If *linear-step* is not specified, it is assumed to be 1.

1 When a **linear** clause is specified on a construct, the value of the new list item on each logical
2 iteration of the associated loops corresponds to the value of the original list item before entering the
3 construct plus the logical number of the iteration times *linear-step*. The value corresponding to the
4 sequentially last logical iteration of the associated loops is assigned to the original list item.

5 When a **linear** clause is specified on a **declare simd** directive, the list items refer to
6 parameters of the procedure to which the directive applies. For a given call to the procedure, the
7 clause determines whether the SIMD version generated by the directive may be called. If the clause
8 does not specify the **ref linear-modifier**, the SIMD version requires that the value of the
9 corresponding argument at the callsite is equal to the value of the argument from the first lane plus
10 the logical number of the lane times the *linear-step*. If the clause specifies the **ref linear-modifier**,
11 the SIMD version requires that the storage locations of the corresponding arguments at the callsite
12 from each SIMD lane correspond to locations within a hypothetical array of elements of the same
13 type, indexed by the logical number of the lane times the *linear-step*.

14 Restrictions

15 Restrictions to the **linear** clause are as follows:

- 16 ● Only a loop iteration variable of a loop that is associated with the construct may appear as a
17 list item in a **linear** clause if a **reduction** clause with the **inscan** modifier also
18 appears on the construct.
- 19 ● A *linear-modifier* may be specified as **ref** or **uval** only on a **declare simd** directive.
- 20 ● For a **linear** clause that appears on a loop-associated construct, the difference between the
21 value of a list item at the end of a logical iteration and its value at the beginning of the logical
22 iteration must be equal to *linear-step*.
- 23 ● If *linear-modifier* is **uval** for a list item in a **linear** clause that is specified on a
24 **declare simd** directive and the list item is modified during a call to the SIMD version of
25 the procedure, the program must not depend on the value of the list item upon return from the
26 procedure.
- 27 ● If *linear-modifier* is **uval** for a list item in a **linear** clause that is specified on a
28 **declare simd** directive, the program must not depend on the storage of the argument in
29 the procedure being the same as the storage of the corresponding argument at the callsite.

- 30
- ▼ C ▲
- All list items must be of integral or pointer type.
 - If specified, *linear-modifier* must be **val**.

- 31
- ▲ C ▼
- ▼ C++ ▲
- 32 ● If *linear-modifier* is not **ref**, all list items must be of integral or pointer type, or must be a
33 reference to an integral or pointer type.
 - 34 ● If *linear-modifier* is **ref** or **uval**, all list items must be 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 a list item in a **linear** clause that is specified on a **declare simd** directive is of a reference type and *linear-modifier* is not **ref**, the difference between the value of the argument on exit from the function and its value on entry to the function must be the same for all SIMD lanes.

C++

Fortran

- If *linear-modifier* is not **ref**, all list items must be of type **integer**.
- If *linear-modifier* is **ref** or **uval**, all list items must be dummy arguments without the **VALUE** attribute.
- List items must not be variables that have the **POINTER** attribute.
- If *linear-modifier* is not **ref** and a list item has the **ALLOCATABLE** attribute, the allocation status of the list item in the sequentially last iteration must be allocated upon exit from that iteration.
- If *linear-modifier* is **ref**, list items must be polymorphic variables, assumed-shape arrays, or variables with the **ALLOCATABLE** attribute.
- If a list item in a **linear** clause that is specified on a **declare simd** directive is a dummy argument without the **VALUE** attribute and *linear-modifier* is not **ref**, the difference between the value of the argument on exit from the procedure and its value on entry to the procedure must be the same for all SIMD lanes.
- A common block name must not appear in a **linear** clause.

Fortran

Cross References

- **private** clause, see [Section 5.4.3](#)
- **declare simd** directive, see [Section 7.7](#)
- **do** directive, see [Section 11.5.2](#)
- **for** directive, see [Section 11.5.1](#)
- **simd** directive, see [Section 10.5](#)
- **taskloop** directive, see [Section 12.6](#)

5.4.7 `is_device_ptr` Clause

Name: <code>is_device_ptr</code>	Properties: data-environment attribute, data-sharing attribute
---	---

Arguments

Name	Type	Properties
<i>list</i>	list of variable list item type	<i>default</i>

Directives

[dispatch](#), [target](#)

Semantics

The `is_device_ptr` clause indicates that its list items are device pointers. Support for device pointers created outside of OpenMP, specifically outside of any OpenMP mechanism that returns a device pointer, is implementation defined.

If the `is_device_ptr` clause is specified on a `target` construct, each list item privatized inside the construct and the new list item is initialized to the device address to which the original list item refers.

Restrictions

Restrictions to the `is_device_ptr` clause are as follows:

- Each list item must be a valid device pointer for the device data environment.

Cross References

- `has_device_addr` clause, see [Section 5.4.9](#)
- `dispatch` directive, see [Section 7.6](#)
- `target` directive, see [Section 13.8](#)

5.4.8 `use_device_ptr` Clause

Name: <code>use_device_ptr</code>	Properties: data-environment attribute, data-sharing attribute
--	---

Arguments

Name	Type	Properties
<i>list</i>	list of variable list item type	<i>default</i>

Directives

[target data](#)

Semantics

Each list item in the `use_device_ptr` clause results in a new list item that is a device pointer that refers to a device address, determined as follows. A list item is treated as if a zero-offset assumed-size array at the storage location to which the list item points is mapped by a `map` clause on the construct with a *map-type* of `alloc`. If a matched candidate is found for the assumed-size array (see [Section 5.8.3](#)), the new list item refers to the device address that is the base address of the array section that corresponds to the assumed-size array in the device data environment. Otherwise, the new list item refers to the address stored in the original list item. All references to the list item inside the structured block associated with the construct are replaced with the new list item.

Restrictions

Restrictions to the `use_device_ptr` clause are as follows:

- Each list item must be a C pointer for which the value is the address of an object that has corresponding storage in the device data environment or is accessible on the target device.

Cross References

- `target data` directive, see [Section 13.5](#)

5.4.9 has_device_addr Clause

Name: <code>has_device_addr</code>	Properties: data-environment attribute, data-sharing attribute
---	---

Arguments

Name	Type	Properties
<i>list</i>	list of variable list item type	<i>default</i>

Directives

[target](#)

Semantics

The `has_device_addr` clause indicates that its list items already have device addresses and therefore they may be directly accessed from a target device. If the device address of a list item is not for the device on which the region that is associated with the construct on which the clause appears executes, accessing the list item inside the region results in unspecified behavior. The list items may include array sections.

▼ [Fortran](#) ▼

For a list item in a `has_device_addr` clause, the `CONTIGUOUS` attribute, storage location, storage size, array bounds, character length, association status and allocation status (as applicable) are the same inside the construct on which the clause appears as for the original list item. The result of inquiring about other list item properties inside the structured block is implementation defined. For a list item that is an array section, the array bounds and result when invoking `C_LOC` inside the structured block is the same as if the base expression had been specified in the clause instead.

▲ [Fortran](#) ▲

Restrictions

Restrictions to the `has_device_addr` clause are as follows:

C / C++

- Each list item must have a valid device address for the device data environment.

C / C++

Fortran

- A list item must either have a valid device address for the device data environment, be an unallocated allocatable variable, or be a disassociated data pointer.
- The association status of a list item that is a pointer must not be undefined unless it is a structure component and it results from a predefined default mapper.

Fortran

Cross References

- `target` directive, see [Section 13.8](#)

5.4.10 use_device_addr Clause

Name: <code>use_device_addr</code>	Properties: data-environment attribute, data-sharing attribute
------------------------------------	--

Arguments

Name	Type	Properties
<i>list</i>	list of variable list item type	<i>default</i>

Directives

`target data`

Semantics

Each list item in a `use_device_addr` clause that is present in the device data environment is treated as if it is implicitly mapped by a `map` clause on the construct with a *map-type* of `alloc`. If a corresponding list item or part of a corresponding list item has storage in the device data environment and the list item has a base variable, all references to the list item inside the structured block associated with the construct are replaced with references to the corresponding list item. Otherwise, all references are to the original list item. The list items in a `use_device_addr` clause may include array sections and assumed-size arrays.

C / C++

If a list item is an array section that has a base pointer, all references to the base pointer inside the structured block are replaced with a new pointer that contains the base address of the corresponding list item. This conversion may be elided if a corresponding list item is not present.

C / C++

Restrictions

Restrictions to the `use_device_addr` clause are as follows:

- Each list item must have a corresponding list item in the device data environment or be accessible on the target device.
- If a list item is an array section, the base expression must be a base language identifier.

Cross References

- `target data` directive, see [Section 13.5](#)

5.5 Reduction Clauses and Directives

The reduction clauses are data-sharing attribute clauses that can be used to perform some forms of recurrence calculations in parallel. Reduction clauses include reduction scoping clauses and reduction participating clauses. Reduction scoping clauses define the region in which a reduction is computed. Reduction participating clauses define the participants in the reduction.

5.5.1 OpenMP Reduction Identifiers

The syntax of an OpenMP reduction identifier is defined as follows:

C
A reduction identifier is either an *identifier* or one of the following operators: `+`, `*`, `&`, `|`, `^`, `&&` and `||`.

C++
A reduction identifier is either an *id-expression* or one of the following operators: `+`, `*`, `&`, `|`, `^`, `&&` and `||`.

Fortran
A 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`.

5.5.2 OpenMP Reduction Expressions

A *reduction expression* is an OpenMP stylized expression that is relevant to reduction clauses. It is either a *combiner expression* or an *initializer expression*.

Restrictions

Restrictions to reduction expressions are as follows:

- If execution of a reduction expression results in the execution of an OpenMP construct or an OpenMP API call, the behavior is unspecified.

C / C++

- If a reduction expression corresponds to a reduction identifier that is used in a **target** region, a **declare target** directive must be specified for any function that can be accessed through the expression.

C / C++

Fortran

- Any subroutine or function used in a reduction expression must be an intrinsic function, or must have an accessible interface.
- Any user-defined operator, defined assignment or extended operator used in a reduction expression must have an accessible interface.
- If any subroutine, function, user-defined operator, defined assignment or extended operator is used in a reduction expression, it must be accessible to the subprogram in which the corresponding **reduction** clause is specified.
- Any subroutine used in a reduction expression must not have any alternate returns appear in the argument list.
- If the list item in the corresponding **reduction** clause is an array or array section, any procedure used in a reduction expression must either be elemental or have dummy arguments that are scalar.
- Any procedure called in the region of a reduction expression must be pure and may not reference any host-associated variables.
- If a reduction expression corresponds to a reduction identifier that is used in a **target** region, a **declare target** directive must be specified for any function or subroutine that can be accessed through the expression.

Fortran

5.5.2.1 OpenMP Combiner Expressions

A *combiner expression* specifies how a reduction combines partial results into a single value.

Fortran

A combiner expression is an assignment statement or a subroutine name followed by an argument list.

Fortran

1 In the definition of a combiner expression, `omp_in` and `omp_out` correspond to two special
2 variable identifiers that refer to storage of the type of the reduction list item to which the reduction
3 applies. If the list item is an array or array section, the identifiers to which `omp_in` and `omp_out`
4 correspond each refer to an array element. Each of the two special variable identifiers denotes one
5 of the values to be combined before executing the combiner expression. The special `omp_out`
6 identifier refers to the storage that holds the resulting combined value after executing the combiner
7 expression. The number of times that the combiner expression is executed and the order of these
8 executions for any reduction clause are unspecified.

Fortran

9 If the combiner expression is a subroutine name with an argument list, the combiner expression is
10 evaluated by calling the subroutine with the specified argument list. If the combiner expression is an
11 assignment statement, the combiner expression is evaluated by executing the assignment statement.

12 If a generic name is used in a combiner expression and the list item in the corresponding reduction
13 clause is an array or array section, it is resolved to the specific procedure that is elemental or only
14 has scalar dummy arguments.

Fortran

Restrictions

15 Restrictions to combiner expressions are as follows:

- 16 • The only variables allowed in a combiner expression are `omp_in` and `omp_out`.

Fortran

- 17 • Any selectors in the designator of `omp_in` and `omp_out` must be *component selectors*.

Fortran

5.5.2.2 OpenMP Initializer Expressions

19 An *initializer expression* determines the initializer for the private copies of reduction list items. If
20 the initialization of the copies is not determined *a priori*, the syntax of an initializer expression is as
21 follows:
22

C

```
23 | omp_priv = initializer
```

C

24 or

C++

```
25 | omp_priv initializer
```

C++

26 or

1 `function-name (argument-list)` C / C++

2 or

3 `omp_priv = expression` Fortran

4 or

5 `subroutine-name (argument-list)` Fortran

6 In the definition of an initializer expression, the `omp_priv` special variable identifier refers to the
7 storage to be initialized. The special variable identifier `omp_orig` can be used in an initializer
8 expression to refer to the storage of the original variable to be reduced. The number of times that an
9 initializer expression is evaluated and the order of these evaluations are unspecified.

10 If an initializer expression is a function name with an argument list, it is evaluated by calling the
11 function with the specified argument list. Otherwise, an initializer expression specifies how
12 `omp_priv` is declared and initialized. C / C++

Fortran

13 If an initializer expression is a subroutine name with an argument list, the *initializer-expr* is
14 evaluated by calling the subroutine with the specified argument list. If an initializer expression is an
15 assignment statement, the initializer expression is evaluated by executing the assignment statement.

Fortran

C

16 The *a priori* initialization of private copies that are created for reductions follows the rules for
17 initialization of objects with static storage duration.

C

C++

18 The *a priori* initialization of private copies that are created for reductions follows the rules for
19 *default-initialization*.

C++

Fortran

The rules for *a priori* initialization of private copies that are created for reductions are as follows:

- For **complex**, **real**, or **integer** types, the value 0 will be used.
- For **logical** types, the value **.false.** will be used.
- For derived types for which default initialization is specified, default initialization will be used.
- Otherwise, the behavior is unspecified.

Fortran

Restrictions

Restrictions to initializer expressions are as follows:

- The only variables allowed in an initializer expression are **omp_priv** and **omp_orig**.
- If an initializer expression modifies the variable **omp_orig**, the behavior is unspecified.

C

- If an initializer expression is a function name with an argument list, one of the arguments must be the address of **omp_priv**.

C

C++

- If an initializer expression is a function name with an argument list, one of the arguments must be **omp_priv** or the address of **omp_priv**.

C++

Fortran

- If an initializer expression is a subroutine name with an argument list, one of the arguments must be **omp_priv**.

Fortran

5.5.3 Implicitly Declared OpenMP Reduction Identifiers

C / C++

Table 5.1 lists each reduction identifier that is implicitly declared at every scope and semantic initializer expression. The actual initializer value is that value as expressed in the data type of the reduction list item if that list item is an arithmetic type. In C++, list items of class type are assigned or constructed with an integral value that matches the initializer value as specified in Section 5.5.5.

TABLE 5.1: Implicitly Declared C/C++ Reduction Identifiers

Identifier	Initializer	Combiner
+	<code>omp_priv = 0</code>	<code>omp_out += omp_in</code>
*	<code>omp_priv = 1</code>	<code>omp_out *= omp_in</code>
&	<code>omp_priv = ~ 0</code>	<code>omp_out &= omp_in</code>
	<code>omp_priv = 0</code>	<code>omp_out = omp_in</code>
^	<code>omp_priv = 0</code>	<code>omp_out ^= omp_in</code>
&&	<code>omp_priv = 1</code>	<code>omp_out = omp_in && omp_out</code>
	<code>omp_priv = 0</code>	<code>omp_out = omp_in omp_out</code>
max	<code>omp_priv = Minimal representable number in the reduction list item type</code>	<code>omp_out = omp_in > omp_out ? omp_in : omp_out</code>
min	<code>omp_priv = Maximal representable number in the reduction list item type</code>	<code>omp_out = omp_in < omp_out ? omp_in : omp_out</code>



1 Table 5.2 lists each reduction identifier that is implicitly declared for numeric and logical types and
2 its semantic initializer value. The actual initializer value is that value as expressed in the data type
3 of the reduction list item.

TABLE 5.2: Implicitly Declared Fortran Reduction Identifiers

Identifier	Initializer	Combiner
+	<code>omp_priv = 0</code>	<code>omp_out = omp_in + omp_out</code>
*	<code>omp_priv = 1</code>	<code>omp_out = omp_in * omp_out</code>
<code>.and.</code>	<code>omp_priv = .true.</code>	<code>omp_out = omp_in .and. omp_out</code>
<code>.or.</code>	<code>omp_priv = .false.</code>	<code>omp_out = omp_in .or. omp_out</code>
<code>.eqv.</code>	<code>omp_priv = .true.</code>	<code>omp_out = omp_in .eqv. omp_out</code>
<code>.neqv.</code>	<code>omp_priv = .false.</code>	<code>omp_out = omp_in .neqv. omp_out</code>

table continued on next page

table continued from previous page

Identifier	Initializer	Combiner
<code>max</code>	<code>omp_priv = Minimal representable number in the reduction list item type</code>	<code>omp_out = max(omp_in, omp_out)</code>
<code>min</code>	<code>omp_priv = Maximal 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

5.5.4 initializer Clause

Name: <code>initializer</code>	Properties: unique
--------------------------------	--------------------

Arguments

Name	Type	Properties
<i>initializer-expr</i>	expression of initializer type	<i>default</i>

Directives

`declare reduction`

Semantics

The `initializer` clause can be used to specify *initializer-expr* as the initializer expression for a user-defined reduction.

Cross References

- `declare reduction` directive, see [Section 5.5.11](#)

5.5.5 Properties Common to All Reduction Clauses

The *clause-specification* of a reduction clause has a *clause-argument-specification* that specifies an OpenMP variable list argument and has a required *reduction-identifier* modifier that specifies the reduction identifier to use for the reduction. The reduction identifier must match a previously declared reduction identifier of the same name and type for each of the list items. This match is done by means of a name lookup in the base language.

The list items that appear in a reduction clause may include array sections.

C++

1 If the type is a derived class then any reduction identifier that matches its base classes is also a
2 match if no specific match for the type has been specified.

3 If the reduction identifier is an implicitly declared reduction identifier or otherwise not an
4 *id-expression* then it is implicitly converted to one by prepending the keyword operator (for
5 example, `+` becomes *operator+*). This conversion is valid for the `+`, `*`, `/`, `&&` and `||` operators.

6 If the reduction identifier is qualified then a qualified name lookup is used to find the declaration.

7 If the reduction identifier is unqualified then an *argument-dependent name lookup* must be
8 performed using the type of each list item.

9 If the type is of class type and the reduction identifier is implicitly declared, then it must provide the
10 operator as described above as well as one of:

- 11 • A default constructor and an assignment operator that accepts a type that can be implicitly
12 constructed from an integer expression.

```
13     template<typename T>  
14     requires (T&& t) {  
15         T();  
16         t = 0;  
17     };
```

- 18 • A single-argument constructor that accepts a type that can be implicitly constructed from an
19 integer expression.

```
20     template<typename T>  
21     requires () {  
22         T(0);  
23     };
```

24 The first of these that matches will be used, with the initializer value being passed to the assignment
25 operator or constructor.

C++

26 If a list item is an array or array section, it will be treated as if a reduction clause would be applied
27 to each separate element of the array section.

28 If a list item is an array section, the elements of any copy of the array section will be stored
29 contiguously.

Fortran

30 If the original list item has the **POINTER** attribute, any copies of the list item are associated with
31 private targets.

Fortran

1 Any copies of a list item associated with the reduction are initialized with the initializer value of the
2 reduction identifier. Any copies are combined using the combiner associated with the reduction
3 identifier.

4 Execution Model Events

5 The *reduction-begin* event occurs before a task begins to perform loads and stores that belong to the
6 implementation of a reduction and the *reduction-end* event occurs after the task has completed
7 loads and stores associated with the reduction. If a task participates in multiple reductions, each
8 reduction may be bracketed by its own pair of *reduction-begin/reduction-end* events or multiple
9 reductions may be bracketed by a single pair of events. The interval defined by a pair of
10 *reduction-begin/reduction-end* events may not contain a task scheduling point.

11 Tool Callbacks

12 A thread dispatches a registered `ompt_callback_reduction` with
13 `ompt_sync_region_reduction` in its *kind* argument and `ompt_scope_begin` as its
14 *endpoint* argument for each occurrence of a *reduction-begin* event in that thread. Similarly, a thread
15 dispatches a registered `ompt_callback_reduction` with
16 `ompt_sync_region_reduction` in its *kind* argument and `ompt_scope_end` as its
17 *endpoint* argument for each occurrence of a *reduction-end* event in that thread. These callbacks
18 occur in the context of the task that performs the reduction and has the type signature
19 `ompt_callback_sync_region_t`.

20 Restrictions

21 Restrictions common to reduction clauses are as follows:

- 22 • Any array element must be specified at most once in all list items on a directive.
- 23 • For a reduction identifier declared in a `declare reduction` directive, the directive must
24 appear before its use in a reduction clause.
- 25 • If a list item is an array section, it must specify contiguous storage, it cannot be a zero-length
26 array section and its base expression must be a base language identifier.
- 27 • If a list item is an array section or an array element, accesses to the elements of the array
28 outside the specified array section or array element result in unspecified behavior.

▼ C / C++ ▼

- 29 • The type of a list item that appears in a reduction clause must be valid for the reduction
30 identifier. For a `max` or `min` reduction in C, the type of the list item must be an allowed
31 arithmetic data type: `char`, `int`, `float`, `double`, or `_Bool`, possibly modified with
32 `long`, `short`, `signed`, or `unsigned`. For a `max` or `min` reduction in C++, the type of
33 the list item must be an allowed arithmetic data type: `char`, `wchar_t`, `int`, `float`,
34 `double`, or `bool`, possibly modified with `long`, `short`, `signed`, or `unsigned`.
- 35 • A list item that appears in a reduction clause must not be `const`-qualified.
- 36 • The reduction identifier for any list item must be unambiguous and accessible.

▲ C / C++ ▲

Fortran

- The type, type parameters and rank of a list item that appears in a reduction clause must be valid for the combiner expression and the initializer expression.
- A list item that appears in a reduction clause must be definable.
- A procedure pointer must not appear in a reduction clause.
- A pointer with the **INTENT (IN)** attribute must not appear in a reduction clause.
- An original list item with the **POINTER** attribute or any pointer component of an original list item that is referenced in a combiner expression must be associated at entry to the construct that contains the reduction clause. Additionally, the list item or the pointer component of the list item must not be deallocated, allocated, or pointer assigned within the region.
- An original list item with the **ALLOCATABLE** attribute or any allocatable component of an original list item that corresponds to a special variable identifier in the combiner expression or the initializer expression must be in the allocated state at entry to the construct that contains the reduction clause. Additionally, the list item or the allocatable component of the list item must be neither deallocated nor allocated, explicitly or implicitly, within the region.
- If the reduction identifier is defined in a **declare reduction** directive, the **declare reduction** directive must be in the same subprogram, or accessible by host or use association.
- If the reduction identifier is a user-defined operator, the same explicit interface for that operator must be accessible at the location of the **declare reduction** directive that defines the reduction identifier.
- If the reduction identifier is defined in a **declare reduction** directive, any procedure referenced in the **initializer** clause or the combiner expression must be an intrinsic function, or must have an explicit interface where the same explicit interface is accessible as at the **declare reduction** directive.

Fortran

Cross References

- `ompt_callback_sync_region_t`, see [Section 19.5.2.13](#)
- `ompt_scope_endpoint_t`, see [Section 19.4.4.11](#)
- `ompt_sync_region_t`, see [Section 19.4.4.14](#)

5.5.6 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 5.5.5](#), 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. To avoid data races, concurrent reads or updates of the original list item must be synchronized with that update of the original list item, which may occur after the construct on which the reduction scoping clause appears, for example, due to the use of the **nowait** clause.

The location in the OpenMP program at which values are combined and the order in which values are combined are unspecified. Thus, when comparing sequential and parallel executions, or when comparing one parallel execution to another (even if the number of threads used is the same), bitwise-identical results are not guaranteed. Similarly, side effects (such as floating-point exceptions) may not be identical and may not occur at the same location in the OpenMP program.

5.5.7 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 5.5.5](#), 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.

5.5.8 reduction Clause

Name: reduction	Properties: data-environment attribute, data-sharing attribute, privatization, reduction scoping, reduction participating
-------------------------------	--

Arguments

Name	Type	Properties
<i>list</i>	list of variable list item type	<i>default</i>

Modifiers

Name	Modifies	Type	Properties
<i>reduction-identifier</i>	<i>list</i>	An OpenMP reduction identifier	required, ultimate
<i>reduction-modifier</i>	<i>list</i>	Keyword: default , inscan , task	default

Directives

[do](#), [for](#), [loop](#), [parallel](#), [scope](#), [sections](#), [simd](#), [taskloop](#), [teams](#)

Semantics

The **reduction** clause is a reduction scoping clause and a reduction participating clause, as described in [Section 5.5.6](#) and [Section 5.5.7](#). For each list item, a private copy is created for 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*.

If *reduction-modifier* is not present or the **default** *reduction-modifier* is present, the behavior is as follows. For **parallel** 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 5.3](#). 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 5.6](#)). The list items are privatized in the construct according to the description and restrictions in [Section 5.3](#). At the end of the region, each original list item is assigned the value described in [Section 5.6](#).

If the **task** *reduction-modifier* is present for a **parallel** or worksharing construct, then each list item is privatized according to the description and restrictions in [Section 5.3](#), and an unspecified number of additional private copies may be 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 5.5.10](#)). After the end of the region, the original list item contains the result of the reduction.

Restrictions

Restrictions to the **reduction** clause are as follows:

- All restrictions common to all reduction clauses, as listed in [Section 5.5.5](#), apply to this clause.
- A list item that appears in a **reduction** clause on a worksharing construct must be shared in the **parallel** region to which a corresponding worksharing region binds.
- If an array section or array element appears as a list item in a **reduction** clause on a worksharing construct, all threads of the team must specify the same storage location.

- Each list item specified with the **inscan** *reduction-modifier* must appear as a list item in an **inclusive** or **exclusive** clause on a **scan** directive enclosed by the construct.
- If the **inscan** *reduction-modifier* is specified, a **reduction** clause without the **inscan** *reduction-modifier* must not appear on the same construct.
- A **reduction** clause with the **task** *reduction-modifier* may only appear on a **parallel** construct, a worksharing construct or a combined or composite construct for which any of the aforementioned constructs is a constituent construct and neither **simd** nor **loop** are constituent constructs.
- A **reduction** clause with the **inscan** *reduction-modifier* may only appear on a worksharing-loop construct, a **simd** construct or a combined or composite construct for which any of the aforementioned constructs is a constituent construct and **distribute** is not a constituent construct.
- The **inscan** *reduction-modifier* must not be specified on a construct for which the **ordered** or **schedule** clause is specified.
- A list item that appears in a **reduction** clause of the innermost enclosing worksharing or **parallel** construct must not be accessed in an explicit task generated by a construct for which an **in_reduction** clause over the same list item does not appear.
- The **task** *reduction-modifier* must not appear in a **reduction** clause if the **nowait** clause is specified on the same construct.

C / C++

- If a list item in a **reduction** clause on a worksharing construct has a reference type then it must bind to the same object for all threads of the team.
- If a list item in a **reduction** clause on a worksharing construct is an array section or an array element then the base pointer must point to the same variable for all threads of the team.
- A variable of class type (or array thereof) that appears in a **reduction** clause with the **inscan** *reduction-modifier* requires an accessible, unambiguous default constructor for the class type; the number of calls to it while performing the scan computation is unspecified.
- A variable of class type (or array thereof) that appears in a **reduction** clause with the **inscan** *reduction-modifier* requires an accessible, unambiguous copy assignment operator for the class type; the number of calls to it while performing the scan computation is unspecified.

C / C++

Cross References

- **ordered** clause, see [Section 4.4.4](#)
- **private** clause, see [Section 5.4.3](#)
- **schedule** clause, see [Section 11.5.3](#)
- **do** directive, see [Section 11.5.2](#)
- **for** directive, see [Section 11.5.1](#)
- **loop** directive, see [Section 11.7](#)
- **parallel** directive, see [Section 10.2](#)
- **scan** directive, see [Section 5.6](#)
- **scope** directive, see [Section 11.2](#)
- **sections** directive, see [Section 11.3](#)
- **simd** directive, see [Section 10.5](#)
- **taskloop** directive, see [Section 12.6](#)
- **teams** directive, see [Section 10.3](#)
- List Item Privatization, see [Section 5.3](#)

5.5.9 task_reduction Clause

Name: <code>task_reduction</code>	Properties: data-environment attribute, data-sharing attribute, privatization, reduction scoping
--	---

Arguments

Name	Type	Properties
<i>list</i>	list of variable list item type	<i>default</i>

Modifiers

Name	Modifies	Type	Properties
<i>reduction-identifier</i>	<i>list</i>	An OpenMP reduction identifier	required, ultimate

Directives

[taskgroup](#)

Semantics

The **task_reduction** clause is a reduction scoping clause, as described in [Section 5.5.6](#), that specifies a reduction among tasks. 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 that participate in the reduction. After the end of the region, the original list item contains the result of the reduction.

Restrictions

Restrictions to the **task_reduction** clause are as follows:

- All restrictions common to all reduction clauses, as listed in [Section 5.5.5](#), apply to this clause.

Cross References

- **taskgroup** directive, see [Section 15.4](#)

5.5.10 in_reduction Clause

Name: <code>in_reduction</code>	Properties: data-environment attribute, data-sharing attribute, privatization, reduction participating
--	---

Arguments

Name	Type	Properties
<i>list</i>	list of variable list item type	<i>default</i>

Modifiers

Name	Modifies	Type	Properties
<i>reduction-identifier</i>	<i>list</i>	An OpenMP reduction identifier	required, ultimate

Directives

[target](#), [task](#), [taskloop](#)

Semantics

The **in_reduction** clause is a reduction participating clause, as described in [Section 5.5.7](#), that specifies that a task participates in a reduction. For a given 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 **task** as the *reduction-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.

1 For the **task** construct, the generated task becomes the participating task. For each list item, a
 2 private copy may be created as if the **private** clause had been used.

3 For the **target** construct, the target task becomes the participating task. For each list item, a
 4 private copy may be created in the data environment of the target task as if the **private** clause
 5 had been used. This private copy will be implicitly mapped into the device data environment of the
 6 target device, if the target device is not the parent device.

7 At the end of the task region, if a private copy was created its value is combined with a copy created
 8 by a reduction scoping clause or with the original list item.

9 **Restrictions**

10 Restrictions to the **in_reduction** clause are as follows:

- 11 • All restrictions common to all reduction clauses, as listed in [Section 5.5.5](#), apply to this
 12 clause.
- 13 • A list item that appears in a **task_reduction** clause or a **reduction** clause with **task**
 14 as the *reduction-modifier* that is specified on a construct that corresponds to a region in which
 15 the region of the participating task is closely nested must match each list item. The construct
 16 that corresponds to the innermost enclosing region that meets this condition must specify the
 17 same *reduction-identifier* for the matching list item as the **in_reduction** clause.

18 **Cross References**

- 19 • **target** directive, see [Section 13.8](#)
- 20 • **task** directive, see [Section 12.5](#)
- 21 • **taskloop** directive, see [Section 12.6](#)

22 **5.5.11 declare reduction Directive**

23 Name: <code>declare reduction</code>	Association: none
Category: declarative	Properties: pure

24 **Arguments**

25 **declare reduction** (*reduction-specifier*)

26 Name	Type	Properties
<i>reduction-specifier</i>	OpenMP reduction specifier	<i>default</i>

27 **Clauses**

28 [initializer](#)

Semantics

The **declare reduction** directive declares a *reduction-identifier* that can be used in a reduction clause as a user-defined reduction. The directive argument *reduction-specifier* uses the following syntax:

```
reduction-identifier : typename-list : combiner
```

where *reduction-identifier* is a reduction identifier, *typename-list* is a type-name list, and *combiner* is an OpenMP combiner expression.

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 types specified in the **declare reduction** directive. If the directive specifies several types then the behavior is as if a **declare reduction** directive was specified for each type. The visibility and accessibility of a user-defined reduction are the same as those of a variable declared at the same location in the program.

C++

The **declare reduction** directive can also appear at the locations in a program where 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 location in the program.

C++

The enclosing context of the *combiner* and of the *initializer-expr* that is specified by the **initializer** clause 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 location in the program.

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 parameters with which the variable is declared.

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 one of the allowed intrinsic procedures, and if the operator or procedure name appears in an accessibility statement in the same module, the accessibility of the corresponding **declare reduction** directive is determined by the accessibility attribute of the statement.

If the *reduction-identifier* is the same as a generic name that is one of the allowed intrinsic procedures and is accessible, and if it has the same name as a derived type in the same module, the accessibility of the corresponding **declare reduction** directive is determined by the accessibility of the generic name according to the base language.

Fortran

Restrictions

Restrictions to the **declare reduction** directive are as follows:

- 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.
- 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++

Fortran

- If the length type parameter is specified for a type, it must be a constant, a colon (:) or an asterisk (*).
- 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.

Fortran

Cross References

- **initializer** clause, see [Section 5.5.4](#)
- OpenMP Combiner Expressions, see [Section 5.5.2.1](#)
- OpenMP Initializer Expressions, see [Section 5.5.2.2](#)
- OpenMP Reduction Identifiers, see [Section 5.5.1](#)

5.6 scan Directive

Name: <code>scan</code> Category: subsidiary	Association: separating Properties: <i>default</i>
---	---

Separated directives

`do`, `for`, `simd`

Clauses

`exclusive`, `inclusive`

Clause set

Properties: unique, required, exclusive	Members: <code>exclusive</code> , <code>inclusive</code>
--	---

Semantics

The **scan** directive separates the *final-loop-body* of an enclosing **simd** construct or worksharing-loop construct (or a composite construct that combines them) into a structured block sequence that serves as an *input phase* and a structured block sequence that serves as a *scan 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. Thus, it specifies that a scan computation updates each list item on each logical iteration of the enclosing loop nest that is associated with the separated directive.

If the **inclusive** clause is specified, the input phase includes the preceding structured block sequence and the scan phase includes the following structured block sequence and, thus, the directive specifies that an inclusive scan computation is performed for each list item of *list*. If the **exclusive** clause is specified, the input phase excludes the preceding structured block sequence and instead includes the following structured block sequence, while the scan phase includes the preceding structured block sequence and, thus, the directive specifies that an exclusive scan computation is performed for each list item of *list*.

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

- if $N = 1$, a_1
- if $N > 1$, $op(\text{PRESUM}_{\text{last}}(op, a_1, \dots, a_j), \text{PRESUM}_{\text{last}}(op, a_k, \dots, a_N))$,
 $1 \leq j + 1 = k \leq N$.

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

Let *orig-val* be the value of the original list item on entry to the separated construct. Let *combiner* be the combiner expression for the *reduction-identifier* specified by the **reduction** clause on the construct. Let u_i be the update value of a list item for iteration *i*. For list items that appear in an **inclusive** clause on the **scan** directive, at the beginning of the scan phase for iteration *i* the list item is assigned the result of the operation $\text{PRESUM}_{\text{last}}(\text{combiner}, \text{orig-val}, u_0, \dots, u_i)$. For list items that appear in an **exclusive** clause on the **scan** directive, at the beginning of the scan phase for iteration $i = 0$ the list item is assigned the value *orig-val*, and at the beginning of the scan phase for iteration $i > 0$ the list item is assigned the result of the operation $\text{PRESUM}_{\text{last}}(\text{combiner}, \text{orig-val}, u_0, \dots, u_{i-1})$.

For list items that appear in an **inclusive** clause, at the end of the separated construct, the original list item is assigned the private copy from the last logical iteration of the loops associated with the separated construct. For list items that appear in an **exclusive** clause, let *k* be the last logical iteration of the loops associated with the separated construct. At the end of the separated

construct, the original list item is assigned the result of the operation $\text{PRESUM}_{\text{last}}(\text{combiner}, \text{orig-val}, u_0, \dots, u_k)$.

Restrictions

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

- A separated construct must have at most one **scan** directive as a separating directive.
- The loops that are associated with the directive to which the **scan** directive is associated must all be perfectly nested.
- Each list item that appears in the **inclusive** or **exclusive** clause must appear in a **reduction** clause with the **inscan** modifier on the separated construct.
- Each list item that appears in a **reduction** clause with the **inscan** modifier on the separated construct must appear in a clause on the separating **scan** directive.
- Cross-iteration dependences across different logical iterations must not exist, except for dependences for the list items specified in an **inclusive** or **exclusive** clause.
- Intra-iteration dependences from a statement in the structured block sequence that precede a **scan** directive to a statement in the structured block sequence that follows a **scan** directive must not exist, except for dependences for the list items specified in an **inclusive** or **exclusive** clause.
- The private copy of list items that appear in the **inclusive** or **exclusive** clause must not be modified in the *scan phase*.

Cross References

- **exclusive** clause, see [Section 5.6.2](#)
- **inclusive** clause, see [Section 5.6.1](#)
- **reduction** clause, see [Section 5.5.8](#)
- **do** directive, see [Section 11.5.2](#)
- **for** directive, see [Section 11.5.1](#)
- **simd** directive, see [Section 10.5](#)

5.6.1 inclusive Clause

Name: inclusive	Properties: unique
-------------------------------	---------------------------

Arguments

Name	Type	Properties
<i>list</i>	list of variable list item type	<i>default</i>

Directives

[scan](#)

Semantics

The **inclusive** clause is used on a separating directive that separates a structured block into two structured block sequences. The clause determines the association of the structured block sequence that precedes the directive on which the clause appears to a phase of that directive.

The list items that appear in an **inclusive** clause may include array sections.

Cross References

- [scan](#) directive, see [Section 5.6](#)

5.6.2 exclusive Clause

Name: <code>exclusive</code>	Properties: unique
-------------------------------------	---------------------------

Arguments

Name	Type	Properties
<i>list</i>	list of variable list item type	<i>default</i>

Directives

[scan](#)

Semantics

The **exclusive** clause is used on a separating directive that separates a structured block into two structured block sequences. The clause determines the association of the structured block sequence that precedes the directive on which the clause appears to a phase of that directive.

The list items that appear in an **exclusive** clause may include array sections.

Cross References

- [scan](#) directive, see [Section 5.6](#)

5.7 Data Copying Clauses

This section describes the **copyin** clause and the **copyprivate** clause. These two clauses support copying data values from private or threadprivate variables of an implicit task or thread to the corresponding variables of other implicit tasks or threads in the team.

5.7.1 copyin Clause

Name: <code>copyin</code>	Properties: data copying
----------------------------------	---------------------------------

Arguments

Name	Type	Properties
<i>list</i>	list of variable list item type	<i>default</i>

Directives

`parallel`

Semantics

The **copyin** clause provides a mechanism to copy the value of a threadprivate variable of the primary thread to the threadprivate variable of each other member of the team that is executing the **parallel** region.

C / C++

The copy is performed after the team is formed and prior to the execution of the associated structured block. For variables of non-array type, the copy is 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 primary thread to the corresponding element of the array of all other threads.

C / C++

C++

For class types, the copy assignment operator is invoked. The order in which copy assignment operators for different variables of the same class type are invoked is unspecified.

C++

Fortran

The copy is performed, as if by assignment, after the team is formed and prior to the execution of the associated structured block.

Named variables that appear in a threadprivate common block may be specified. The whole common block does not need to be specified.

On entry to any **parallel** region, each thread's copy of a variable that is affected by a **copyin** clause for the **parallel** region will acquire the type parameters, allocation, association, and definition status of the copy of the primary thread, according to the following rules:

- If the original list item has the **POINTER** attribute, each copy receives the same association status as that of the copy of the primary thread as if by pointer assignment.
- If the original list item does not have the **POINTER** attribute, each copy becomes defined with the value of the copy of the primary thread as if by intrinsic assignment unless the list item has a type bound procedure as a defined assignment. If the original list item that does not have the **POINTER** attribute has the allocation status of unallocated, each copy will have the same status.
- If the original list item is unallocated or unassociated, each copy inherits the declared type parameters and the default type parameter values from the original list item.

Fortran

Restrictions

Restrictions to the **copyin** clause are as follows:

- A list item that appears in a **copyin** clause must be threadprivate.

C++

- A variable of class type (or array thereof) that appears in a **copyin** clause requires an accessible, unambiguous copy assignment operator for the class type.

C++

Fortran

- A common block name that appears in a **copyin** clause must be declared to be a common block in the same scoping unit in which the **copyin** clause appears.

- A polymorphic variable with the **ALLOCATABLE** attribute must not be a list item.

Fortran

Cross References

- **parallel** directive, see [Section 10.2](#)
- **threadprivate** directive, see [Section 5.2](#)

5.7.2 copyprivate Clause

Name: copyprivate	Properties: end-clause, data copying
---------------------------------	---

Arguments

Name	Type	Properties
<i>list</i>	list of variable list item type	<i>default</i>

Directives

single

Semantics

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. The effect of the **copyprivate** clause on the specified list items occurs after the execution of the structured block associated with the associated construct, and before any of the threads in the team have left the barrier at the end of the construct. 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 if, for example, the **nowait** clause is used to remove the barrier.

C / C++

1 In all other implicit tasks that belong to the parallel region, each specified list item becomes defined
2 with the value of the corresponding list item in the implicit task associated with the thread that
3 executed the structured block. For variables of non-array type, the definition occurs by copy
4 assignment. For an array of elements of non-array type, each element is copied by copy assignment
5 from an element of the array in the data environment of the implicit task that is associated with the
6 thread that executed the structured block to the corresponding element of the array in the data
7 environment of the other implicit tasks.

C / C++

C++

8 For class types, a copy assignment operator is invoked. The order in which copy assignment
9 operators for different variables of class type are called is unspecified.

C++

Fortran

10 If a list item does not have the **POINTER** attribute, then in all other implicit tasks that belong to the
11 parallel region, the list item becomes defined as if by intrinsic assignment with the value of the
12 corresponding list item in the implicit task that is associated with the thread that executed the
13 structured block. If the list item has a type bound procedure as a defined assignment, the
14 assignment is performed by the defined assignment.

15 If the list item has the **POINTER** attribute then in all other implicit tasks that belong to the parallel
16 region the list item receives, as if by pointer assignment, the same association status as the
17 corresponding list item in the implicit task that is associated with the thread that executed the
18 structured block.

19 The order in which any final subroutines for different variables of a finalizable type are called is
20 unspecified.

Fortran

Restrictions

21 Restrictions to the **copyprivate** clause are as follows:

- 22 • All list items that appear in a **copyprivate** clause must be either threadprivate or private
23 in the enclosing context.

C++

- 24 • A variable of class type (or array thereof) that appears in a **copyprivate** clause requires
25 an accessible unambiguous copy assignment operator for the class type.

C++

Fortran

- A common block that appears in a **copyprivate** clause must be **threadprivate**.
- Pointers with the **INTENT (IN)** attribute must not appear in a **copyprivate** clause.
- Any list item with the **ALLOCATABLE** attribute must have the allocation status of **allocated** when the intrinsic assignment is performed.
- If a list item is a polymorphic variable with the **ALLOCATABLE** attribute, the behavior is unspecified.

Fortran

Cross References

- **firstprivate** clause, see [Section 5.4.4](#)
- **private** clause, see [Section 5.4.3](#)
- **single** directive, see [Section 11.1](#)

5.8 Data-Mapping Control

This section describes the available mechanisms for controlling how data are mapped to device data environments. It covers implicit data-mapping attribute rules for variables referenced in **target** constructs, explicit clauses for specifying how data should be mapped, and clauses for making available variables with static lifetimes and procedures on other devices. It also describes how mappers may be defined and referenced to control the mapping of data with user-defined types. When storage is mapped, the programmer must ensure, by adding proper synchronization or by explicit unmapping, that the storage does not reach the end of its lifetime before it is unmapped.

5.8.1 Implicit Data-Mapping Attribute Rules

When specified, explicit data-environment attribute clauses on **target** directives determine the attributes for variables referenced in a **target** construct. Otherwise, the first matching rule from the following list determines the implicit data-mapping (or data-sharing) attribute for variables referenced in a **target** construct that do not have a predetermined data-sharing attribute according to [Section 5.1.1](#). References to structure elements or array elements are treated as references to the structure or array, respectively, for the purposes of determining implicit data-mapping or data-sharing attributes of variables in a **target** construct.

- If a variable appears in an **enter** or **link** clause on a declare target directive that does not have a **device_type** clause with the **nohost** *device-type-description* then it is treated as if it had appeared in a **map** clause with a *map-type* of **tofrom**.

- 1 • If a variable is the base variable of a list item in a **reduction**, **lastprivate** or **linear**
2 clause on a combined target construct then the list item is treated as if it had appeared in a
3 **map** clause with a *map-type* of **tofrom** if [Section 17.2](#) specifies this behavior.
- 4 • If a variable is the base variable of a list item in an **in_reduction** clause on a **target**
5 construct then it is treated as if the list item had appeared in a **map** clause with a *map-type* of
6 **tofrom** and a *map-type-modifier* of **always**.
- 7 • If a **defaultmap** clause is present for the category of the variable and specifies an implicit
8 behavior other than **default**, the data-mapping or data-sharing attribute is determined by
9 that clause.

C++

- 10 • If the **target** construct is within a class non-static member function, and a variable is an
11 accessible data member of the object for which the non-static data member function is
12 invoked, the variable is treated as if the **this[:1]** expression had appeared in a **map** clause
13 with a *map-type* of **tofrom**. Additionally, if the variable is of type pointer or reference to
14 pointer, it is also treated as if it is the base expression of a zero-offset assumed-size array that
15 appears in a **map** clause with the **alloc** *map-type*.
- 16 • If the **this** keyword is referenced inside a **target** construct within a class non-static
17 member function, it is treated as if the **this[:1]** expression had appeared in a **map** clause
18 with a *map-type* of **tofrom**.

C++

- 19 • A variable that is of type pointer, but is neither a pointer to function nor (for C++) a pointer
20 to a member function, is treated as if it is the base expression of an zero-offset assumed-size
21 array that appears in a **map** clause with the **alloc** *map-type*.

C / C++

- 22 • A variable that is of type reference to pointer, but is neither a reference to point to function
23 nor a reference to a pointer to a member function is treated as if it is the base expression of an
24 zero-offset assumed-size array that appears in a **map** clause with the **alloc** *map-type*.

C++

- 25 • If a variable is not a scalar then it is treated as if it had appeared in a **map** clause with a
26 *map-type* of **tofrom**.

Fortran

- 27 • If a scalar variable has the **TARGET**, **ALLOCATABLE** or **POINTER** attribute then it is treated
28 as if it had appeared in a **map** clause with a *map-type* of **tofrom**.

Fortran

- 29 • If the above rules do not apply then a scalar variable is not mapped but instead has an implicit
30 data-sharing attribute of **firstprivate** (see [Section 5.1.1](#)).

5.8.2 Mapper Identifiers and `mapper` Modifiers

Modifiers

Name	Modifies	Type	Properties
<i>mapper</i>	<i>locator-list</i>	Complex, name: mapper Arguments: <i>mapper-identifier</i> OpenMP identifier (<i>default</i>)	unique

Clauses

`from`, `map`, `to`

Mapper identifiers can be used to uniquely identify the mapper used in a `map` or data-motion clause through a *mapper* modifier, which is a unique, complex modifier. A **declare mapper** directive defines a mapper identifier that can later be specified in a *mapper* modifier as its *modifier-parameter-specification*. Each mapper identifier is a base-language identifier or **default** where **default** is the default mapper for all types.

A non-structure type *T* has a predefined default mapper that is defined as if by the following **declare mapper** directive:

```

13  #pragma omp declare mapper(T v) map(tofrom: v)
14  !$omp declare mapper(T :: v) map(tofrom: v)

```

C / C++ Fortran
C / C++ Fortran

A structure type *T* has a predefined default mapper that is defined as if by a **declare mapper** directive that specifies *v* in a `map` clause with the **alloc map-type** and each structure element of *v* in a `map` clause with the **tofrom map-type**.

A **declare mapper** directive that uses the **default mapper** identifier overrides the predefined default mapper for the given type, making it the default mapper for variables of that type.

Cross References

- `from` clause, see [Section 5.9.2](#)
- `map` clause, see [Section 5.8.3](#)
- `to` clause, see [Section 5.9.1](#)

5.8.3 `map` Clause

Name: <code>map</code>	Properties: data-environment attribute, data-mapping attribute
-------------------------------	---

Arguments

Name	Type	Properties
<i>locator-list</i>	list of locator list item type	<i>default</i>

Modifiers

Name	Modifies	Type	Properties
<i>map-type-modifier</i>	<i>locator-list</i>	Keyword: always , close , present	default
<i>mapper</i>	<i>locator-list</i>	Complex, name: mapper Arguments: <i>mapper-identifier</i> OpenMP identifier (<i>default</i>)	unique
<i>iterator</i>	<i>locator-list</i>	Complex, name: iterator Arguments: <i>iterator-specifier</i> OpenMP expression (repeatable)	unique
<i>map-type</i>	<i>locator-list</i>	Keyword: alloc , delete , from , release , to , tofrom	default

Directives

declare mapper, **target**, **target data**, **target enter data**, **target exit data**

Semantics

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. If a *map-type* is not specified, the *map-type* defaults to **tofrom** unless the list item is an assumed-size array, in which case the *map-type* defaults to **alloc**. The **map** clause is *map-entering* if the *map-type* is **to**, **tofrom** or **alloc**. The **map** clause is *map-exiting* if the *map-type* is **from**, **tofrom**, **release** or **delete**.

The list items that appear in a **map** clause may include array sections, assumed-size arrays, and structure elements. A list item in a **map** clause may reference any *iterator-identifier* defined in its *iterator* modifier. A list item may appear more than once in the **map** clauses that are specified on the same directive.

▼ C / C++ ▼

If a list item is a zero-length array section that has a single array subscript, the behavior is as if the list item is an assumed-size array that is instead mapped with the **alloc** *map-type*.

▲ C / C++ ▲

1 When a list item in a **map** clause that is not an assumed-size array is mapped on a map-entering
2 construct and corresponding storage is created in the device data environment on entry to the
3 region, the mapped list item becomes a *matchable candidate* with an associated *starting address*,
4 *ending address*, and *base address*. The starting address is the address of the first storage location of
5 the list item, the ending address is the address of the last storage location of the list item, and the
6 base address is the address of the base storage location of the list item. The *mapped address range*
7 of the list item is the range of addresses that starts from the starting address and ends with the
8 ending address. The *extended address range* of the list item is the range of addresses that starts
9 from the minimum of the starting address and the base address and that ends with the maximum of
10 the ending address and the base address. The current set of matchable candidates consists of any
11 **map** clause list item on the construct that is a matchable candidate and all previously mapped
12 matchable candidates that are still mapped.

13 A list item in a **map** clause that is an assumed-size array is treated as if an array section, with a base
14 expression, lower bound and length determined as follows, is substituted in its place if a *matched*
15 *candidate* is found. If the assumed-size array is an array section, the base expression of the
16 substitute array section is the same as for the assumed-size array; otherwise, the base expression is
17 the assumed-size array. If the mapped address range of a matchable candidate includes the first
18 storage location of the assumed-size array, it is a matched candidate. Otherwise, if the extended
19 address range of a matchable candidate includes the first storage location of the assumed-size array,
20 it is a matched candidate. If there are multiple matched candidates, an arbitrary one of them is the
21 found matched candidate. The lower bound and length of the substitute array section are set such
22 that its storage is identical to the storage of the found matched candidate. If a matched candidate is
23 not found then a substitute array section is not formed and no further actions that are described in
24 this section are performed for the list item.

25 If a *mapper* modifier is not present, the behavior is as if a *mapper* modifier was specified with the
26 **default** parameter. The map behavior of a list item in a **map** clause is modified by a visible
27 user-defined mapper (see [Section 5.8.7](#)) if the *mapper-identifier* of the *mapper* modifier is defined
28 for a base-language type that matches the type of the list item. Otherwise, the predefined default
29 mapper for the type of the list item applies. The effect of the *mapper* is to remove the list item from
30 the **map** clause, if the **present** modifier does not also appear, and to apply the clauses specified in
31 the declared mapper to the construct on which the **map** clause appears. In the clauses applied by the
32 *mapper*, references to *var* are replaced with references to the list item and the *map-type* is replaced
33 with a final map type that is determined according to the rules of map-type decay (see
34 [Section 5.8.7](#)).

35 A list item that is an array or array section of a type for which a user-defined mapper exists is
36 mapped as if the map type decays to **alloc**, **release**, or **delete**, and then each array element
37 is mapped with the original map type, as if by a separate construct, according to the mapper.

Fortran

38 If a component of a derived type list item is a **map** clause list item that results from the predefined
39 default mapper for that derived type, and if the derived type component is not an explicit list item or
40 the base expression of an explicit list item in a **map** clause on the construct, then:

- If it has the **POINTER** attribute, it is *attach-ineligible*; and
- If it has the **ALLOCATABLE** attribute and an allocated allocation status, and it is present in the device data environment when the construct is encountered, the **map** clause may treat its allocation status as if it is unallocated if the corresponding component does not have allocated storage.

If a list item in a **map** clause is an associated pointer that is not attach-ineligible and the pointer is not the base pointer of another list item in a **map** clause on the same construct, then it is treated as if its pointer target is implicitly mapped in the same clause. For the purposes of the **map** clause, the mapped pointer target is treated as if its base pointer is the associated pointer.



If a list item has a closure type that is associated with a lambda expression, it is mapped as if it has a structure type. For each variable that is captured by reference by the lambda expression, references to the variable in the function call operator for the new list item refer to its corresponding storage in the device data environment, if it exists prior to a task encountering the construct associated with the **map** clause, and otherwise refer to its original storage. For each pointer that is not a function pointer that is captured by the lambda expression, the behavior is as if the pointer or, for capture by copy, the corresponding pointer member of the closure object is the base expression of an zero-offset assumed-size array that appears in a **map** clause with the **alloc map-type**.

If the **this** pointer is captured by a lambda expression in class scope, and a variable of the associated closure type is later mapped explicitly or implicitly with its full static type, the behavior is as if the object to which **this** points is also mapped as an array section, of length one, for which the base pointer is the non-static data member that corresponds to the **this** pointer in the closure object.



If a **map** clause with a **present map-type-modifier** appears on a construct and on entry to the region the corresponding list item is not present in the device data environment, runtime error termination is performed.

The **map** clauses on a construct collectively determine the set of *mappable storage blocks* for that construct. All **map** clause list items that have the same containing structure or share storage result in a single mappable storage block that contains the storage of the list items. The storage for each other **map** clause list item becomes a distinct mappable storage block.

For each mappable storage block that is determined by the **map** clauses on a map-entering construct, on entry to the region the following sequence of steps occurs as if performed as a single atomic operation:

1. If a corresponding storage block is not present in the device data environment then:
 - a) A corresponding storage block, which may share storage with the original storage block, is created in the device data environment of the device;

- b) The corresponding storage block receives a reference count that is initialized to zero. This reference count also applies to any part of the corresponding storage block.
2. The reference count of the corresponding storage block is incremented by one.
3. For each **map** clause list item on the construct that is contained by the mappable storage block:
 - a) If the reference count of the corresponding storage block is one, a new list item with language-specific attributes derived from the original list item is created in the corresponding storage block. The reference count of the new list item is always equal to the reference count of its storage.
 - b) If the reference count of the corresponding list item is one or if the **always** *map-type-modifier* is specified, and if the *map-type* is **to** or **tofrom**, the corresponding list item is updated as if the list item appeared in a **to** clause on a **target update** directive.

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.

For **map** clauses on map-entering constructs, if any list item has a base pointer for which a corresponding pointer exists in the data environment after all mappable storage blocks are mapped, and either a new list item or the corresponding pointer is created in the device data environment on entry to the region, then the corresponding pointer becomes an attached pointer to the corresponding list item via corresponding base pointer initialization.

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, if it appears on a mapping-only construct that applies to the device data environment of the host device, or if the corresponding list item has an attached pointer that shares storage with the original pointer to which it corresponds.

For each mappable storage block that is determined by the **map** clauses on a map-exiting construct, and for which corresponding storage is present in the device data environment, on exit from the region the following sequence of steps occurs as if performed as a single atomic operation:

1. For each **map** clause list item that is contained by the mappable storage block:

- 1 a) If the reference count of the corresponding list item is one or if the **always**
- 2 *map-type-modifier* is specified, and if the *map-type* is **from** or **tofrom**, the original
- 3 list item is updated as if the list item appeared in a **from** clause on a **target update**
- 4 directive.
- 5 2. If the *map-type* is not **delete** and the reference count of the corresponding storage block is
- 6 finite then the reference count is decremented by one.
- 7 3. If the *map-type* is **delete** and the reference count of the corresponding storage block is
- 8 finite then the reference count is set to zero.
- 9 4. If the reference count of the corresponding storage block is zero, all storage to which that
- 10 reference count applies is removed from the device data environment.

11 If the effect of the **map** clauses on a construct would assign the value of a corresponding list item to

12 an original list item more than once, then an implementation is allowed to ignore additional

13 assignments of the same value to the original list item.

14 In all cases on exit from the region, concurrent reads or updates of any part of the original list item

15 must be synchronized with any update of the original list item that occurs as a result of the **map**

16 clause to avoid data races.

17 If a single contiguous part of the original storage of a list item that results from an implicit

18 data-mapping attribute has corresponding storage in the device data environment prior to a task

19 encountering the construct on which the **map** clause appears, only that part of the original storage

20 will have corresponding storage in the device data environment as a result of the **map** clause.

21 If a list item with an implicit data-mapping attribute does not have any corresponding storage in the

22 device data environment prior to a task encountering the construct associated with the **map** clause,

23 and one or more contiguous parts of the original storage are either list items or base pointers to list

24 items that are explicitly mapped on the construct, only those parts of the original storage will have

25 corresponding storage in the device data environment as a result of the **map** clauses on the

26 construct.

▼ C / C++ ▼

27 If a new list item is created then the new list item will have the same static type as the original list

28 item, and language-specific attributes of the new list item, including size and alignment, are

29 determined by that type.

▲ C / C++ ▲

▼ C++ ▼

30 If corresponding storage that differs from the original mappable storage block is created in a device


31 data environment, all new list items that are created in that corresponding storage are default

32 initialized. Default initialization for new list items of class type, including their data members, is

33 performed as if with an implicitly-declared default constructor and as if non-static data member

34 initializers are ignored.

1 If the type of a new list item is a reference to a type T then it is initialized to refer to the object in
2 the device data environment that corresponds to the object referenced by the original list item. The
3 effect is as if the object were mapped through a pointer with an array section of length one and
4 elements of type T .



5 If a new list item is created then the new list item will have the same type, type parameter, and rank
6 as the original list item. The new list item inherits all default values for the type parameters from
7 the original list item.

8 If the allocation status of an original list item that has the **ALLOCATABLE** attribute is changed
9 while a corresponding list item is present in the device data environment, the allocation status of the
10 corresponding list item is unspecified until the list item is again mapped with an **always** modifier
11 on entry to a map-entering region.



12 The **close map-type-modifier** is a hint to the runtime to allocate memory close to the target device.

13 Execution Model Events

14 The *target-map* event occurs in a thread that executes the outermost region that corresponds to an
15 encountered device construct with a **map** clause, after the *target-task-begin* event for the device
16 construct and before any mapping operations are performed.

17 The *target-data-op-begin* event occurs before a thread initiates a data operation on the target device
18 that is associated with a **map** clause, in the outermost region that corresponds to the encountered
19 construct.

20 The *target-data-op-end* event occurs after a thread initiates a data operation on the target device
21 that is associated with a **map** clause, in the outermost region that corresponds to the encountered
22 construct.

23 Tool Callbacks

24 A thread dispatches one or more registered **ompt_callback_target_map** or
25 **ompt_callback_target_map_emi** callbacks for each occurrence of a *target-map* event in
26 that thread. The callback occurs in the context of the target task and has type signature
27 **ompt_callback_target_map_t** or **ompt_callback_target_map_emi_t**,
28 respectively.

29 A thread dispatches a registered **ompt_callback_target_data_op_emi** callback with
30 **ompt_scope_begin** as its endpoint argument for each occurrence of a *target-data-op-begin*
31 event in that thread. Similarly, a thread dispatches a registered
32 **ompt_callback_target_data_op_emi** callback with **ompt_scope_end** as its endpoint
33 argument for each occurrence of a *target-data-op-end* event in that thread. These callbacks have
34 type signature **ompt_callback_target_data_op_emi_t**.

1 A thread dispatches a registered `ompt_callback_target_data_op` callback for each
2 occurrence of a `target-data-op-end` event in that thread. The callback occurs in the context of the
3 target task and has type signature `ompt_callback_target_data_op_t`.

4 **Restrictions**

5 Restrictions to the `map` clause are as follows:

- 6 • Two list items of the `map` clauses on the same construct must not share original storage
7 unless one of the following is true: they are the same list item, one is the containing structure
8 of the other, at least one is an assumed-size array, or at least one is implicitly mapped due to
9 the list item also appearing in a `use_device_addr` clause.
- 10 • If the same list item appears more than once in `map` clauses on the same construct, the `map`
11 clauses must specify the same *mapper* modifier.
- 12 • A variable that is groupprivate or device local must not appear as a list item in a `map` clause.
- 13 • If a list item is an array section, it must specify contiguous storage.
- 14 • If an expression that is used to form a list item in a `map` clause contains an iterator identifier,
15 the list item instances that would result from different values of the iterator must not have the
16 same containing array and must not have base pointers that share original storage.
- 17 • If multiple list items are explicitly mapped on the same construct and have the same
18 containing array or have base pointers that share original storage, and if any of the list items
19 do not have corresponding list items that are present in the device data environment prior to a
20 task encountering the construct, then the list items must refer to the same array elements of
21 either the containing array or the implicit array of the base pointers.
- 22 • If any part of the original storage of a list item that is explicitly mapped by a `map` clause has
23 corresponding storage in the device data environment prior to a task encountering the
24 construct associated with the `map` clause, all of the original storage must have corresponding
25 storage in the device data environment prior to the task encountering the construct.
- 26 • If an array appears as a list item in a `map` clause and it has corresponding storage in the
27 device data environment, the corresponding storage must correspond to a single mappable
28 storage block that was previously mapped.
- 29 • If a list item is an element of a structure, and a different element of the structure has a
30 corresponding list item in the device data environment prior to a task encountering the
31 construct associated with the `map` clause, then the list item must also have a corresponding
32 list item in the device data environment prior to the task encountering the construct.
- 33 • A list item must have a mappable type.
- 34 • If a mapper modifier appears in a `map` clause, the type on which the specified mapper
35 operates must match the type of the list items in the clause.
- 36 • Memory spaces and memory allocators must not appear as a list item in a `map` clause.

1 • If a list item is an assumed-size array, multiple matched candidates must not exist unless they
2 are subobjects of the same containing structure.

3 • If a list item is an assumed-size array, the *map-type* must be **alloc**.

C++

4 • If a list item has a polymorphic class type and its static type does not match its dynamic type,
5 the behavior is unspecified if the **map** clause is specified on a map-entering construct and a
6 corresponding list item is not present in the device data environment prior to a task
7 encountering the construct.

8 • No type mapped through a reference may contain a reference to its own type, or any
9 references to types that could produce a cycle of references.

C++

C / C++

10 • A list item cannot be a variable that is a member of a structure of 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
13 the lifetime of the list item to which the corresponding pointer is attached in the device data
14 environment.

C / C++

Fortran

15 • The association status of a list item that is a pointer must not be undefined unless it is a
16 structure component and it results from a predefined default mapper.

17 • If a list item of a **map** clause is an allocatable variable or is the subobject of an allocatable
18 variable, the original allocatable variable may not be allocated, deallocated or reshaped while
19 the corresponding allocatable variable has allocated storage.

20 • A pointer that has a corresponding attached pointer and is associated with a given pointer
21 target must not become associated with a different pointer target for the duration of the
22 lifetime of the list item to which the corresponding pointer is attached in the device data
23 environment.

24 • If an array section is mapped and the size of the section is smaller than that of the whole
25 array, the behavior of referencing the whole array in the **target** region is unspecified.

26 • A list item must not be a complex part designator.

Fortran

Cross References

- `declare mapper` directive, see [Section 5.8.7](#)
- `target` directive, see [Section 13.8](#)
- `target data` directive, see [Section 13.5](#)
- `target enter data` directive, see [Section 13.6](#)
- `target exit data` directive, see [Section 13.7](#)
- `target update` directive, see [Section 13.9](#)
- Array Sections, see [Section 3.2.5](#)
- `iterator` modifier, see [Section 3.2.6](#)
- `mapper` modifier, see [Section 5.8.2](#)
- `ompt_callback_target_data_op_emi_t` and `ompt_callback_target_data_op_t`, see [Section 19.5.2.25](#)
- `ompt_callback_target_map_emi_t` and `ompt_callback_target_map_t`, see [Section 19.5.2.27](#)

5.8.4 enter Clause

Name: <code>enter</code>	Properties: data-environment attribute, data-mapping attribute
---------------------------------	---

Arguments

Name	Type	Properties
<i>list</i>	list of extended list item type	<i>default</i>

Directives

`declare target`

Semantics

The `enter` clause is a data-mapping clause.

▼ [C / C++](#) ▼

If a function appears in an `enter` clause in the same compilation unit in which the definition of the function occurs then a device-specific version of the function is created for all devices to which the directive of the clause applies.

If a variable appears in an `enter` clause in the same compilation unit in which the definition of the variable occurs then the original list item is allocated a corresponding list item in the device data environment of all devices to which the directive of the clause applies.

▲ [C / C++](#) ▲

Fortran

1 If a procedure appears in an **enter** clause in the same compilation unit in which the definition of
2 the procedure occurs then a device-specific version of the procedure is created for all devices to
3 which the directive of the clause applies.

4 If a variable that is host associated appears in an **enter** clause then the original list item is
5 allocated a corresponding list item in the device data environment of all devices to which the
6 directive of the clause applies.

Fortran

7 If a variable appears in an **enter** clause then the corresponding list item in the device data
8 environment of each device to which the directive of the clause applies is initialized once, in the
9 manner specified by the program, but at an unspecified point in the program prior to the first
10 reference to that list item. The list item is never removed from those device data environments as if
11 its reference count was initialized to positive infinity.

Cross References

- **declare target** directive, see [Section 7.8.1](#)

5.8.5 link Clause

Name: <code>link</code>	Properties: data-environment attribute
--------------------------------	---

Arguments

Name	Type	Properties
<i>list</i>	list of variable list item type	<i>default</i>

Directives

[declare target](#)

Semantics

21 The **link** clause supports compilation of device routines that refer to variables with static storage
22 duration that appear as list items in the clause. The **declare target** directive on which the
23 clause appears does not map the list items. Instead, they are mapped according to the data-mapping
24 rules described in [Section 5.8](#).

Cross References

- **declare target** directive, see [Section 7.8.1](#)
- Data-Mapping Control, see [Section 5.8](#)

5.8.6 defaultmap Clause

Name: <code>defaultmap</code>	Properties: unique, post-modified
--------------------------------------	--

Arguments

Name	Type	Properties
<i>implicit-behavior</i>	Keyword: alloc , default , firstprivate , from , none , present , to , tofrom	<i>default</i>

Modifiers

Name	Modifies	Type	Properties
<i>variable-category</i>	<i>implicit-behavior</i>	Keyword: aggregate , all , allocatable , pointer , scalar	default

Directives

target

Semantics

The **defaultmap** clause determines the implicit data-mapping or data-sharing attribute of certain variables that are referenced in a **target** construct, in accordance with the rules given in [Section 5.8.1](#). The *variable-category* specifies the variables for which the attribute may be set, and the attribute is specified by *implicit-behavior*. If no *variable-category* is specified in the clause then the effect is as if **all** was specified for the *variable-category*.

▼ C / C++ ▲

The **scalar** *variable-category* specifies non-pointer variables of scalar type.

▲ C / C++ ▲

▼ Fortran ▼

The **scalar** *variable-category* specifies non-pointer and non-allocatable variables of scalar type. The **allocatable** *variable-category* specifies variables with the **ALLOCATABLE** attribute.

▲ Fortran ▲

The **pointer** *variable-category* specifies variables of pointer type. The **aggregate** *variable-category* specifies variables of aggregate type (arrays or structures). Finally, the **all** *variable-category* specifies all variables.

If *implicit-behavior* is the name of a map type, the attribute is a data-mapping attribute determined by an implicit **map** clause with the specified map type. If *implicit-behavior* is **firstprivate**, the attribute is a data-sharing attribute of firstprivate. If *implicit-behavior* is **present**, the attribute is a data-mapping attribute determined by an implicit **map** clause with the *map-type* of **alloc** and *map-type-modifier* of **present**. If *implicit-behavior* is **none** then no implicit data-mapping or data-sharing attributes are defined for variables in *variable-category*, except for variables that appear in the **enter** or **link** clause of a **declare target** directive. If *implicit-behavior* is **default** then the clause has no effect.

Restrictions

Restrictions to the **defaultmap** clause are as follows:

- A given *variable-category* may be specified in at most one **defaultmap** clause on a construct.
- If a **defaultmap** clause specifies the **all** *variable-category*, no other **defaultmap** clause may appear on the construct.
- If *implicit-behavior* is **none**, each variable that is specified by *variable-category* and is referenced in the construct but does not have a predetermined data-sharing and does not appear in an **enter** or **link** clause on a **declare target** directive must be explicitly listed in a data-environment attribute clause on the construct.

C / C++

- The specified *variable-category* must not be **allocatable**.

C / C++

Cross References

- **target** directive, see [Section 13.8](#)
- Implicit Data-Mapping Attribute Rules, see [Section 5.8.1](#)

5.8.7 declare mapper Directive

Name: <code>declare mapper</code> Category: declarative	Association: none Properties: <i>default</i>
--	---

Arguments

declare mapper (*mapper-specifier*)

Name	Type	Properties
<i>mapper-specifier</i>	OpenMP mapper specifier	<i>default</i>

Clauses

map

Semantics

User-defined mappers can be defined using the **declare mapper** directive. The *mapper-specifier* directive argument declares the mapper using the following syntax:

C / C++

[*mapper-identifier* :] type var

C / C++

```
[ mapper-identifier : ] type :: var
```

where *mapper-identifier* is a mapper identifier, *type* is a type that is permitted in a type-name list, and *var* is a base-language identifier.

The *type* and an optional *mapper-identifier* uniquely identify the mapper for use in a **map** clause or motion clause later in the program. The visibility and accessibility of this declaration are the same as those of a variable declared at the same location in the program.

If *mapper-identifier* is not specified, the behavior is as if *mapper-identifier* is **default**.

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 effect that a user-defined mapper has on either a **map** clause that maps a list item of the given base language type or a motion clause that invokes the mapper and updates a list item of the given base language type is to replace the map or update with a set of **map** clauses or updates derived from the **map** clauses specified by the mapper, as described in [Section 5.8.3](#) and [Section 5.9](#).

The final map types that a mapper applies for a **map** clause that maps a list item of the given type are determined according to the rules of map-type decay, defined according to [Table 5.3](#). [Table 5.3](#) shows the final map type that is determined by the combination of two map types, where the rows represent the map type specified by the mapper and the columns represent the map type specified by a **map** clause that invokes the mapper. For a **target exit data** construct that invokes a mapper with a **map** clause that has the **from** map type, if a **map** clause in the mapper specifies an **alloc** or **to** map type then the result is a **release** map type.

A list item in a **map** clause that appears on a **declare mapper** directive may include array sections.

All **map** clauses that are introduced by a mapper are further subject to mappers that are in scope, except a **map** clause with list item *var* maps *var* without invoking a mapper.

TABLE 5.3: Map-Type Decay of Map Type Combinations

	alloc	to	from	tofrom	release	delete
alloc	alloc	alloc	alloc (release)	alloc	release	delete
to	alloc	to	alloc (release)	to	release	delete
from	alloc	alloc	from	from	release	delete
tofrom	alloc	to	from	tofrom	release	delete

C++

1 The **declare mapper** directive can also appear at locations in the program at which a static data
2 member could be declared. In this case, the visibility and accessibility of the declaration are the
3 same as those of a static data member declared at the same location in the program.

C++

Restrictions

4 Restrictions to the **declare mapper** directive are as follows:

- 6 • No instance of *type* can be mapped as part of the mapper, either directly or indirectly through
7 another base language type, except the instance *var* that is passed as the list item. If a set of
8 **declare mapper** directives results in a cyclic definition then the behavior is unspecified.
- 9 • The *type* must not declare a new base language type.
- 10 • At least one **map** clause that maps *var* or at least one element of *var* is required.
- 11 • List items in **map** clauses on the **declare mapper** directive may only refer to the declared
12 variable *var* and entities that could be referenced by a procedure defined at the same location.
- 13 • Neither the **release** or **delete map-type** may be specified on any **map** clause.
- 14 • If a *mapper-modifier* is specified for a **map** clause, its parameter must be **default**.
- 15 • Multiple **declare mapper** directives that specify the same *mapper-identifier* for the same
16 base language type or for compatible base language types, according to the base language
17 rules, may not appear in the same scope.

C

- 18 • *type* must be a **struct** or **union** type.

C

C++

- 19 • *type* must be a **struct**, **union**, or **class** type.
- 20 • If *type* is **struct** or **class**, it must not be derived from any virtual base class.

C++

Fortran

- 21 • *type* must not be an intrinsic type, an abstract type, or a parameterized derived type.

Fortran

Cross References

- 22 • **map** clause, see [Section 5.8.3](#)

23

5.9 Data-Motion Clauses

Data-motion clauses specify data movement between a device set that is specified by the construct on which they appear. One member of that device set is always the *encountering device*, which is the device on which the encountering task for that construct executes. How the other devices, which are the *targeted devices*, are determined is defined by the construct specification. Each data-motion clause specifies the direction of the data movement relative to the targeted devices.

A data-motion clause specifies an OpenMP locator list as its argument. A corresponding list item and an original list item exist for each list item. If the corresponding list item is not present in the device data environment then no assignment occurs between the corresponding and original list items. Otherwise, each corresponding list item in the device data environment has an original list item in the data environment of the encountering task. Assignment is performed to either the original or corresponding list item as specified with the specific data-motion clauses. List items may reference any *iterator-identifier* defined in its *iterator* modifier. The list items may include array sections with *stride* expressions.

C / C++

The list items may use shape-operators.

C / C++

If a list item is an array or array section then it is treated as if it is replaced by each of its array elements in the clause.

If the *mapper* modifier is not specified, the behavior is as if the modifier was specified with the **default** *mapper-identifier*. The effect of a data-motion clause on a list item is modified by a visible user-defined mapper if *mapper-identifier* is specified for a type that matches the type of the list item. Otherwise, the predefined default mapper for the type of the list item applies. Each list item is replaced with the list items that the given mapper specifies are to be mapped with a map type that is compatible with the data movement direction associated with the clause.

If a **present** *expectation* is specified and the corresponding list item is not present in the device data environment then runtime error termination is performed. For a list item that is replaced with a set of list items as a result of a user-defined mapper, the *expectation* only applies to those mapper list items that share storage with the original list item.

Fortran

If a list item or a subobject of a list item has the **ALLOCATABLE** attribute, its assignment is performed only if its allocation status is allocated and only with respect to the allocated storage. If a list item has the **POINTER** attribute and its association status is associated, the effect is as if the assignment is performed with respect to the pointer target.

On exit from the associated region, if the corresponding list item is an attached pointer, the original list item, if associated, will be associated with the same pointer target with which it was associated on entry to the region and 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

C / C++

1 On exit from the associated region, if the corresponding list item is an attached pointer, the original
2 list item will have the value it had on entry to the region and the corresponding list item will have
3 the value it had on entry to the region.

C / C++

4 For each list item that is not an attached pointer, the value of the assigned list item is assigned the
5 value of the other list item. To avoid data races, concurrent reads or updates of the assigned list
6 item must be synchronized with the update of an assigned list item that occurs as a result of a
7 data-motion clause.

Restrictions

8 Restrictions to data-motion clauses are as follows:

- 9 • Each list item clause must have a mappable type.
- 10 • If an array appears as a list item in a motion clause and it has corresponding storage in the
11 device data environment, the corresponding storage must correspond to a single mappable
12 storage block that was previously mapped.

Fortran

- 14 • The association status of a list item that is a pointer must not be undefined unless it is a
15 structure component and it results from a predefined default mapper.

Fortran

Cross References

- 16 • **device** clause, see [Section 13.2](#)
- 17 • **from** clause, see [Section 5.9.2](#)
- 18 • **to** clause, see [Section 5.9.1](#)
- 19 • **declare mapper** directive, see [Section 5.8.7](#)
- 20 • **target update** directive, see [Section 13.9](#)
- 21 • Array Sections, see [Section 3.2.5](#)
- 22 • Array Shaping, see [Section 3.2.4](#)
- 23 • **iterator** modifier, see [Section 3.2.6](#)
- 24

5.9.1 to Clause

Name: <code>to</code>	Properties: data-motion attribute
------------------------------	--

Arguments

Name	Type	Properties
<i>locator-list</i>	list of locator list item type	<i>default</i>

Modifiers

Name	Modifies	Type	Properties
<i>expectation</i>	<i>Generic</i>	Keyword: present	default
<i>mapper</i>	<i>locator-list</i>	Complex, name: mapper Arguments: <i>mapper-identifier</i> OpenMP identifier (<i>default</i>)	unique
<i>iterator</i>	<i>locator-list</i>	Complex, name: iterator Arguments: <i>iterator-specifier</i> OpenMP expression (repeatable)	unique

Directives

`target update`

Semantics

The `to` clause is a data motion clause that specifies movement to the targeted devices from the encountering device so the corresponding list items are the assigned list items and the compatible map types are `to` and `tofrom`.

Cross References

- `target update` directive, see [Section 13.9](#)
- `iterator` modifier, see [Section 3.2.6](#)

5.9.2 from Clause

Name: <code>from</code>	Properties: data-motion attribute
--------------------------------	--

Arguments

Name	Type	Properties
<i>locator-list</i>	list of locator list item type	<i>default</i>

Modifiers

Name	Modifies	Type	Properties
<i>expectation</i>	<i>Generic</i>	Keyword: present	default
<i>mapper</i>	<i>locator-list</i>	Complex, name: mapper Arguments: mapper-identifier OpenMP identifier (<i>default</i>)	unique
<i>iterator</i>	<i>locator-list</i>	Complex, name: iterator Arguments: iterator-specifier OpenMP expression (repeatable)	unique

Directives

[target update](#)

Semantics

The **from** clause is a data motion clause that specifies movement from the targeted devices to the encountering device so the original list items are the assigned list items and the compatible map types are **from** and **tofrom**.

Cross References

- **target update** directive, see [Section 13.9](#)
- **iterator** modifier, see [Section 3.2.6](#)

5.10 uniform Clause

Name: uniform	Properties: data-environment attribute
-----------------------------	---

Arguments

Name	Type	Properties
<i>parameter-list</i>	list of parameter list item type	<i>default</i>

Directives

[declare simd](#)

Semantics

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

Cross References

- **declare simd** directive, see [Section 7.7](#)

5.11 aligned Clause

Name: <code>aligned</code>	Properties: data-environment attribute, post-modified
-----------------------------------	--

Arguments

Name	Type	Properties
<i>list</i>	list of variable list item type	<i>default</i>

Modifiers

Name	Modifies	Type	Properties
<i>alignment</i>	<i>list</i>	OpenMP integer expression	positive, region invariant, ultimate, unique

Directives

`declare simd, simd`

Semantics

C / C++

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

C / C++

Fortran

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

Fortran

The *alignment* modifier specifies the alignment that the program ensures related to the list items. If the *alignment* modifier is not specified, implementation-defined default alignments for SIMD instructions on the target platforms are assumed.

Restrictions

Restrictions to the **aligned** clause are as follows:

C

- The type of list items must be array or pointer.

C

C++

- The type of list items must be array, pointer, reference to array, or reference to pointer.

C++

Fortran

- Each list item must be an array.

Fortran

Cross References

- `declare simd` directive, see [Section 7.7](#)
- `simd` directive, see [Section 10.5](#)

5.12 groupprivate Directive

Name: <code>groupprivate</code> Category: declarative	Association: none Properties: <i>default</i>
--	---

Arguments

`groupprivate` (*list*)

Name	Type	Properties
<i>list</i>	list of variable list item type	<i>default</i>

Clauses

`device_type`

Semantics

The `groupprivate` directive specifies that list items are replicated such that each contention group receives its own copy. Each copy of the list item is uninitialized upon creation. The lifetime of a `groupprivate` variable is limited to the lifetime of all tasks in the contention group.

For a `device_type` clause that is specified implicitly or explicitly on the directive, the behavior is as if the list items appear in a `local` clause on a `declare` target directive on which the same `device_type` clause is specified and at the same program point.

All references to a variable in *list* in any task will refer to the `groupprivate` copy of that variable that is created for the contention group of the innermost enclosing implicit parallel region.

Restrictions

Restrictions to the `groupprivate` directive are as follows:

- A task that executes in a particular contention group must not access the storage of a `groupprivate` copy of the list item that is created for a different contention group.
- A variable that is declared with an initializer must not appear in a `groupprivate` directive.

C / C++

- Each list item must be a file-scope, namespace-scope, or static block-scope variable.
- No list item may have an incomplete type.
- The address of a `groupprivate` variable must not be an address constant.
- If any list item is a file-scope variable, the directive must appear outside any definition or declaration, and must lexically precede all references to any of the variables in the list.

- 1 • If any list is a namespace-scope variable, the directive must appear outside any definition or
2 declaration other than the namespace definition itself and must lexically precede all
3 references to any of the variables in the list.
- 4 • Each variable in the list of a **groupprivate** directive at file, namespace, or class scope
5 must refer to a variable declaration at file, namespace, or class scope that lexically precedes
6 the directive.
- 7 • If any list item is a static block-scope variable, the directive must appear in the scope of the
8 variable and not in a nested scope and must lexically precede all references to any of the
9 variables in the list.
- 10 • Each variable in the list of a **groupprivate** directive in block scope must have static
11 storage duration and must refer to a variable declaration in the same scope that lexically
12 precedes the directive.
- 13 • If a variable is specified in a **groupprivate** directive in one translation unit, it must be
14 specified in a **groupprivate** directive in every translation unit in which it is declared.

C / C++

C++

- 15 • If any list item is a static class member variable, the directive must appear in the class
16 definition, in the same scope in which the member variable is declared, and must lexically
17 precede all references the variable.
- 18 • A **groupprivate** variable must not have an incomplete type or a reference type.

C++

Fortran

- 19 • Each list item must be a named variable or a named common block; a named common block
20 must appear between slashes.
- 21 • The list argument must not include any coarrays or associate names.
- 22 • The **groupprivate** directive must appear in the declaration section of a scoping unit in
23 which the common block or variable is declared.
- 24 • If a **groupprivate** directive that specifies a common block name appears in one program
25 unit, then such a directive must also appear in every other program unit that contains a
26 **COMMON** statement that specifies the same name. Each such directive must appear after the
27 last such **COMMON** statement in that program unit.
- 28 • If a **groupprivate** variable or a **groupprivate** common block is declared with the **BIND**
29 attribute, the corresponding C entities must also be specified in a **groupprivate** directive
30 in the C program.
- 31 • A variable may only appear as an argument in a **groupprivate** directive in the scope in
32 which it is declared. It must not be an element of a common block or appear in an
33 **EQUIVALENCE** statement.

- A variable that appears as a list item in a **groupprivate** directive must be declared in the scope of a module or have the **SAVE** attribute, either explicitly or implicitly.
- The effect of an access to a groupprivate variable in a **DO CONCURRENT** construct is unspecified.

Fortran

Cross References

- **device_type** clause, see [Section 13.1](#)

5.13 local Clause

Name: <code>local</code>	Properties: data-environment attribute
---------------------------------	---

Arguments

Name	Type	Properties
<i>list</i>	list of variable list item type	<i>default</i>

Directives

[declare target](#)

Semantics

The **local** clause specifies that a reference to a list item on a given device will refer to a device local copy of the list item in memory associated with the device.

Cross References

- **declare target** directive, see [Section 7.8.1](#)

6 Memory Management

This chapter defines directives, clauses and related concepts for managing memory used by OpenMP programs.

6.1 Memory Spaces

OpenMP memory spaces represent storage resources where variables can be stored and retrieved. Table 6.1 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 6.1: 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
<code>omp_high_bw_mem_space</code>	Represents storage with high bandwidth
<code>omp_low_lat_mem_space</code>	Represents storage with low latency

Variables allocated in the `omp_const_mem_space` memory space may be initialized through the `firstprivate` clause or with compile time constants for static and constant variables. Implementation-defined mechanisms to provide the constant value of these variables may also be supported.

Restrictions

Restrictions to OpenMP memory spaces are as follows:

- Variables in the `omp_const_mem_space` memory space may not be written.

6.2 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 6.2 shows the allowed allocator traits, their possible values and the default value of each trait.

TABLE 6.2: 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>	Positive integer powers of 2	1 byte
<code>access</code>	<code>all</code> , <code>memspace</code> , <code>device</code> , <code>cgroup</code> , <code>pteam</code> , <code>thread</code>	<code>memspace</code>
<code>pool_size</code>	Any positive integer	Implementation de- fined
<code>fallback</code>	<code>default_mem_fb</code> , <code>null_fb</code> , <code>abort_fb</code> , <code>allocator_fb</code>	See below
<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>
<code>pin_device</code>	Conforming device number	(none)
<code>preferred_device</code>	Conforming device number	(none)
<code>target_access</code>	<code>single</code> , <code>multiple</code>	<code>single</code>
<code>atomic_scope</code>	<code>all</code> , <code>device</code>	<code>device</code>
<code>part_size</code>	Positive integer value	Implementation de- fined

The `sync_hint` trait describes the expected manner in which multiple threads may use the allocator. The values and their descriptions are:

- **contended:** high contention is expected on the allocator; that is, many tasks are expected

- 1 to request allocations simultaneously;
- 2 • **uncontended**: low contention is expected on the allocator; that is, few tasks are expected
3 to request allocations simultaneously;
- 4 • **serialized**: one task at a time will request allocations with the allocator. Requesting two
5 allocations simultaneously when specifying **serialized** results in unspecified behavior;
6 and
- 7 • **private**: the same thread will execute all tasks that request allocations with the allocator.
8 Requesting an allocation from tasks that different threads execute, simultaneously or not,
9 when specifying **private** results in unspecified behavior.

10 Allocated memory will be byte aligned to at least the value specified for the **alignment** trait of
11 the allocator. Some directives and API routines can specify additional requirements on alignment
12 beyond those described in this section.

13 The **access** trait defines the *access group* of tasks that may access memory that is allocated by an
14 allocator. If the value is **all**, the access group consists of all tasks that execute on all available
15 devices. If the value is **memspace**, the access group consists of all tasks that execute on all devices
16 that are associated with the allocator. If the value is **device**, the access group consists of all tasks
17 that execute on the device where the allocation was requested. If the value is **cgroup**, the access
18 group consists of all tasks in the same contention group as the task that requested the allocation. If
19 the value is **pteam**, the access group consists of all current team tasks of the innermost enclosing
20 parallel region in which the allocation was requested. If the value is **thread**, the access group
21 consists of all tasks that are executed by the same thread that executed the allocation request.
22 Memory returned by the allocator will be memory accessible by all tasks in the same access group
23 as the task that requested the allocation. Attempts to access this memory from a task that is not in
24 same access group results in unspecified behavior.

25 The total amount of storage in bytes that an allocator can use for allocation requests from tasks in
26 the same access group is limited by the **pool_size** trait. Requests that would result in using more
27 storage than **pool_size** will not be fulfilled by the allocator.

28 The **fallback** trait specifies how the allocator behaves when it cannot fulfill an allocation
29 request. If the **fallback** trait is set to **null_fb**, the allocator returns the value zero if it fails to
30 allocate the memory. If the **fallback** trait is set to **abort_fb**, the behavior is as if an **error**
31 directive for which *sev-level* is **fatal** and *action-time* is **execution** is encountered if the
32 allocation fails. If the **fallback** trait is set to **allocator_fb** then when an allocation fails the
33 request will be delegated to the allocator specified in the **fb_data** trait. If the **fallback** trait is
34 set to **default_mem_fb** then when an allocation fails another allocation will be tried in
35 **omp_default_mem_space**, which assumes all allocator traits to be set to their default values
36 except for **fallback** trait, which will be set to **null_fb**. The default value for the **fallback**
37 trait is **null_fb** for any allocator that is associated with a target memory space. Otherwise, the
38 default value is **default_mem_fb**.

39 All memory that is allocated with an allocator for which the **pinned** trait is specified as **true**

1 must remain in the same storage resource at the same location for its entire lifetime. If
2 **pin_device** is also specified then the allocation must be allocated in that device.

3 The **partition** trait describes the partitioning of allocated memory over the storage resources
4 represented by the memory space associated with the allocator. The partitioning will be done in
5 parts with a minimum size that is implementation defined. The values are:

- 6 • **environment**: the placement of allocated memory is determined by the execution
7 environment;
- 8 • **nearest**: allocated memory is placed in the storage resource that is nearest to the thread
9 that requests the allocation;
- 10 • **blocked**: allocated memory is partitioned into parts of approximately the same size with at
11 most one part per storage resource; and
- 12 • **interleaved**: allocated memory parts are distributed in a round-robin fashion across the
13 storage resources such that the size of each part is the value of the **part_size** trait except
14 possibly the last part, which can be smaller.

15 The **part_size** trait specifies the size of the parts allocated over the storage resources for some
16 of the partition trait policies. The actual value of the trait might be rounded up to an
17 implementation defined value to comply with hardware restrictions of the storage resources.

18 If the **preferred_device** trait is specified then storage resources of the specified device are
19 preferred to fulfill the allocation.

20 If the value of the **target_access** trait is **single** then data from this allocator cannot be
21 accessed on two different devices unless, for any given host access, the entry and exit of the target
22 region in which any accesses occur either both precede or both follow the host access in happens
23 before order. Additionally, for any two **target** regions that may access data from this allocator
24 and execute on distinct devices, the entry and exit of one of the regions must precede those of the
25 other in happens before order. If the value of the **target_access** trait is **multiple** then
26 accesses of data from this allocator from different devices may be arbitrarily interleaved, provided
27 that synchronization ensures data races do not occur.

28 If the value of the **atomic_scope** trait is **all** then all storage locations of data from this
29 allocator have an atomic scope that consists of all threads on the devices associated with the
30 allocator. If the value is **device** then all storage locations have an atomic scope that consists of all
31 threads on the device on which the atomic operation is performed.

32 Table 6.3 shows the list of predefined memory allocators and their associated memory spaces. The
33 predefined memory allocators have default values for their allocator traits unless otherwise
34 specified.

TABLE 6.3: Predefined Allocators

Allocator name	Associated memory space	Non-default trait values
<code>omp_default_mem_alloc</code>	<code>omp_default_mem_space</code>	<code>fallback:null_fb</code>
<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	<code>access:cgroup</code>
<code>omp_pteam_mem_alloc</code>	Implementation defined	<code>access:pteam</code>
<code>omp_thread_mem_alloc</code>	Implementation defined	<code>access:thread</code>

Fortran

1 If any operation of the base language causes a reallocation of a variable that is allocated with a
 2 memory allocator then that memory allocator will be used to deallocate the current memory and to
 3 allocate the new memory. For any allocatable subcomponents, the allocator that is used for the
 4 deallocation and allocation is unspecified.

Fortran

Restrictions

- If the `pin_device` trait is specified, its value must be the device number of a device associated with the memory allocator.
- If the `preferred_device` trait is specified, its value must be the device number of a device associated with the memory allocator.

6.3 align Clause

Name: <code>align</code>	Properties: unique
---------------------------------	---------------------------

Arguments

Name	Type	Properties
<code>alignment</code>	expression of integer type	constant, positive

Directives

`allocate`

Semantics

The **align** clause is used to specify the byte alignment to use for allocations associated with the construct on which the clause appears. Specifically, each allocation is byte aligned to at least the maximum of the value to which *alignment* evaluates, the **alignment** trait of the allocator being used for the allocation, and the alignment required by the base language for the type of the variable that is allocated. On constructs on which the clause may appear, if it is not specified then the effect is as if it was specified with the **alignment** trait of the allocator being used for the allocation.

Restrictions

Restrictions to the **align** clause are as follows:

- *alignment* must evaluate to a power of two.

Cross References

- **allocate** directive, see [Section 6.5](#)
- Memory Allocators, see [Section 6.2](#)

6.4 allocator Clause

Name: <code>allocator</code>	Properties: unique
-------------------------------------	---------------------------

Arguments

Name	Type	Properties
<i>allocator</i>	expression of <code>allocator_handle</code> type	<i>default</i>

Directives

[allocate](#)

Semantics

The **allocator** clause specifies the memory allocator to be used for allocations associated with the construct on which the clause appears. Specifically, the allocator to which *allocator* evaluates is used for the allocations. On constructs on which the clause may appear, if it is not specified then the effect is as if it was specified with the value of the *def-allocator-var* ICV.

Cross References

- **allocate** directive, see [Section 6.5](#)
- Memory Allocators, see [Section 6.2](#)
- *def-allocator-var* ICV, see [Table 2.1](#)

6.5 allocate Directive

Name: <code>allocate</code> Category: declarative	Association: none Properties: <i>default</i>
--	---

Arguments

Name	Type	Properties
<i>list</i>	list of variable list item type	<i>default</i>

Clauses

[align](#), [allocator](#)

Semantics

The storage for each list item that appears in the **allocate** directive is provided an allocation through the memory allocator as determined by the **allocator** clause with an alignment as determined by the **align** clause. 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.

For allocations that arise from this directive the **null_fb** value of the fallback allocator trait behaves as if the **abort_fb** had been specified.

Restrictions

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

- An **allocate** directive must appear in the same scope as the declarations of each of its list items and must follow all such declarations.
- A declared variable may appear as a list item in at most one **allocate** directive in a given compilation unit.
- **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 static storage duration, the **allocator** clause must be specified and the *allocator* expression in the clause must be a constant expression that evaluates to one of the predefined memory allocator values.
- A variable that is declared in a namespace or global scope may only appear as a list item in an **allocate** directive if an **allocate** directive that lists the variable follows a declaration that defines the variable and if all **allocate** directives that list it specify the same allocator.
- A list item must not be a function parameter.

C / C++

C

- After a list item has been allocated, the scope that contains the **allocate** directive must not end abnormally, such as through a call to the **longjmp** function.

C++

- After a list item has been allocated, the scope that contains the **allocate** directive must not end abnormally, such as through a call to the **longjmp** function, other than through C++ exceptions.
- A variable that has a reference type must not appear as a list item in an **allocate** directive.

C++

Fortran

- A list item that is specified in an **allocate** directive must not have the **ALLOCATABLE** or **POINTER** attribute.
- If a list item has the **SAVE** attribute, either explicitly or implicitly, or is a common block name then the **allocator** clause must be specified and only predefined memory allocator parameters can be used in the clause.
- A variable that is part of a common block must not be specified as a list item in an **allocate** directive, except implicitly via the named common block.
- A named common block may appear as a list item in at most one **allocate** directive in a given compilation unit.
- If a named common block appears as a list item in an **allocate** directive, it must appear as a list item in an **allocate** directive that specifies the same allocator in every compilation unit in which the common block is used.
- An associate name must not appear as a list item in an **allocate** directive.
- A list item must not be a dummy argument.

Fortran

Cross References

- **align** clause, see [Section 6.3](#)
- **allocator** clause, see [Section 6.4](#)
- Memory Allocators, see [Section 6.2](#)

6.6 allocate Clause

Name: <code>allocate</code>	Properties: <i>default</i>
------------------------------------	-----------------------------------

Arguments

Name	Type	Properties
<i>list</i>	list of variable list item type	<i>default</i>

1

Modifiers

Name	Modifies	Type	Properties
<i>allocator-simple-modifier</i>	<i>list</i>	expression of OpenMP <code>allocator_handle</code> type	exclusive, unique
<i>allocator-complex-modifier</i>	<i>list</i>	Complex, name: allocator Arguments: allocator expression of <code>allocator_handle</code> type (<i>default</i>)	unique
<i>align-modifier</i>	<i>list</i>	Complex, name: align Arguments: alignment expression of integer type (constant, positive)	unique

2

3

Directives

4

allocators, **distribute**, **do**, **for**, **parallel**, **scope**, **sections**, **single**, **target**, **task**, **taskgroup**, **taskloop**, **teams**

5

6

Semantics

7

The **allocate** clause specifies the memory allocator to be used to obtain storage for a list of variables. If a list item in the clause also appears in a data-sharing attribute clause on the same directive that privatizes the list item, allocations that arise from that list item in the clause will be provided by the memory allocator. If the *allocator-simple-modifier* is specified, the behavior is as if the *allocator-complex-modifier* is instead specified with *allocator-simple-modifier* as its *allocator* argument. The *allocator-complex-modifier* and *align-modifier* have the same syntax and semantics for the **allocate** clause as the **allocator** and **align** clauses have for the **allocate** directive.

8

9

10

11

12

13

14

15

For allocations that arise from this clause the **null_fb** value of the fallback allocator trait behaves as if the **abort_fb** had been specified.

16

17

Restrictions

18

Restrictions to the **allocate** clause are as follows:

19

20

21

- For any list item that is specified in the **allocate** clause on a directive other than the **allocators** 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.

22

23

- **allocate** clauses that appear on a **target** construct or on constructs in a **target** region must specify an *allocator-simple-modifier* or *allocator-complex-modifier* unless a **requires** directive with the **dynamic_allocators** clause is present in the same compilation unit.

Cross References

- **align** clause, see [Section 6.3](#)
- **allocator** clause, see [Section 6.4](#)
- **allocators** directive, see [Section 6.7](#)
- **distribute** directive, see [Section 11.6](#)
- **do** directive, see [Section 11.5.2](#)
- **for** directive, see [Section 11.5.1](#)
- **parallel** directive, see [Section 10.2](#)
- **scope** directive, see [Section 11.2](#)
- **sections** directive, see [Section 11.3](#)
- **single** directive, see [Section 11.1](#)
- **target** directive, see [Section 13.8](#)
- **task** directive, see [Section 12.5](#)
- **taskgroup** directive, see [Section 15.4](#)
- **taskloop** directive, see [Section 12.6](#)
- **teams** directive, see [Section 10.3](#)
- Memory Allocators, see [Section 6.2](#)

Fortran

6.7 allocators Construct

Name: <code>allocators</code>	Association: block (allocator structured block)
Category: executable	Properties: <i>default</i>

Clauses

[allocate](#)

Semantics

The **allocators** construct specifies that OpenMP memory allocators are used for certain variables that are allocated by the associated *allocate-stmt*. If a variable that is to be allocated appears as a list item in an **allocate** clause on the directive, an OpenMP allocator is used to allocate storage for the variable according to the semantics of the **allocate** clause. If a variable that is to be allocated does not appear as a list item in an **allocate** clause, the allocation is performed according to the base language implementation.

Restrictions

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

- A list item that appears in an **allocate** clause must appear as one of the variables that is allocated by the *allocate-stmt* in the associated allocator structured block.

Cross References

- **allocate** clause, see [Section 6.6](#)
- Memory Allocators, see [Section 6.2](#)
- OpenMP Allocator Structured Blocks, see [Section 4.3.1.1](#)

Fortran

6.8 uses_allocators Clause

Name: <code>uses_allocators</code>	Properties: data-environment attribute, data-sharing attribute
---	---

Arguments

Name	Type	Properties
<i>allocator</i>	expression of <code>allocator_handle</code> type	<i>default</i>

Modifiers

Name	Modifies	Type	Properties
<i>mem-space</i>	<i>Generic</i>	Complex, name: memspace Arguments: memspace-handle expression of <code>memspace_handle</code> type (<i>default</i>)	default
<i>traits-array</i>	<i>Generic</i>	Complex, name: traits Arguments: traits variable of alloctrait array type (<i>default</i>)	default

Directives

target

Semantics

The **uses_allocators** clause enables the use of the specified *allocator* in the region associated with the directive on which the clause appears. If *allocator* refers to a predefined allocator, that predefined allocator will be available for use in the region. If *allocator* does not refer to a predefined allocator, the effect is as if *allocator* is specified on a **private** clause. The resulting corresponding item is assigned the result of a call to **omp_init_allocator** at the beginning of the associated region with arguments *memspace-handle*, the number of traits in the *traits* array, and *traits*. If *mem-space* is not specified or **omp_null_mem_space** is specified, the effect is as if *memspace-handle* is specified as **omp_default_mem_space**. If *traits-array* is not specified, the effect is as if *traits* is specified as an empty array. Further, at the end of the associated region, the effect is as if this allocator is destroyed as if by a call to **omp_destroy_allocator**.

Restrictions

- The *allocator* expression must be a base language identifier.
- If *allocator* is a predefined allocator, no modifiers may be specified.
- If *allocator* is not a predefined allocator, it must be a variable.
- The *allocator* argument must not appear in other data-sharing attribute clauses or data-mapping attribute clauses on the same construct.

C / C++

- The *traits* argument for the *traits-array* modifier must be a constant array, have constant values and be defined in the same scope as the construct on which the clause appears.

C / C++

Fortran

- The *traits* argument for the *traits-array* modifier must be a a named constant of rank one.

Fortran

- The *memspace-handle* argument for the *mem-space* modifier must be an identifier that matches one of the predefined memory space names.

Cross References

- **target** directive, see [Section 13.8](#)
- Memory Allocators, see [Section 6.2](#)
- Memory Spaces, see [Section 6.1](#)
- **omp_destroy_allocator**, see [Section 18.13.5](#)
- **omp_init_allocator**, see [Section 18.13.3](#)

7 Variant Directives

This chapter defines directives and related concepts to support the seamless adaption of programs to OpenMP contexts.

7.1 OpenMP Contexts

At any point in a program, an OpenMP context exists that defines traits that describe the active OpenMP constructs, the execution devices, functionality supported by the implementation and available dynamic values. The traits are grouped into trait sets. The following trait sets exist: *construct*, *device*, *target_device*, *implementation* and *dynamic*. Traits are categorized as name-list traits, clause-list traits, non-property traits and extension traits. This categorization determines the syntax that is used to match the trait, as defined in [Section 7.2](#).

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. Whether the **dispatch** construct is added to the *construct* set is implementation defined. If it is added, it will only be added for the *target-call* of the associated code. 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:

1. For procedures 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 one.
2. For procedures that are determined to be function variants by a declare variant directive, the selectors c_1, \dots, c_M of the **construct** selector set are added in the same order to the beginning of the set as c_1, \dots, c_M so the total size of the set is increased by M .
3. For procedures that are determined to be target function variants by a declare target directive, the target trait is added to the beginning of the set as c_1 so the total size of the set is increased by one.

The *simd* trait is a clause-list trait that is defined with properties that match the clauses accepted by the **declare simd** directive with the same name and semantics. The *simd* trait defines at least the *simdlen* property and one of the *inbranch* or *notinbranch* properties. Traits in the *construct* set other than *simd* are non-property traits.

1 The *device* set includes traits that define the characteristics of the device being targeted by the
2 compiler at that point in the program. For each *target device* that the implementation supports, a
3 *target_device* set exists that defines the characteristics of that device. At least the following traits
4 must be defined for the *device* and all *target_device* sets:

- 5 • The *kind(kind-name-list)* trait specifies the general kind of the device. The following
6 *kind-name* values are defined:
 - 7 – *host*, which specifies that the device is the host device;
 - 8 – *nohost*, which specifies that the device is not the host device; and
 - 9 – the values defined in the *OpenMP Additional Definitions* document.
- 10 • The *isa(isa-name-list)* trait specifies the Instruction Set Architectures supported by the
11 device. The accepted *isa-name* values are implementation defined.
- 12 • The *arch(arch-name-list)* trait specifies the architectures supported by the device. The
13 accepted *arch-name* values are implementation defined.

14 The *kind*, *isa* and *arch* traits in the *device* and *target_device* sets are name-list traits.

15 Additionally, the *target_device* set defines the following trait:

- 16 • The *device_num* trait specifies the *device number* of the device.

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

- 19 • The *vendor(vendor-name-list)* trait, which specifies the vendor identifiers of the
20 implementation. OpenMP defined values for *vendor-name* are defined in the *OpenMP*
21 *Additional Definitions* document.
- 22 • The *extension(extension-name-list)* trait, which specifies vendor specific extensions to the
23 OpenMP specification. The accepted *extension-name* values are implementation defined.
- 24 • A *requires(requires-clause-list)* trait, which is a clause-list trait for which the properties are
25 the clauses that have been supplied to the **requires** directive prior to the program point as
26 well as implementation-defined implicit requirements.

27 The *vendor* and *extension* traits in the *implementation* set are name-list traits.

28 Implementations can define additional traits in the *device*, *target_device* and *implementation* sets;
29 these traits are extension traits.

30 The *dynamic* trait set includes traits that define the dynamic properties of a program at a point in its
31 execution. The *data state* trait in the *dynamic* trait set refers to the complete data state of the
32 program that may be accessed at runtime.

7.2 Context Selectors

Context selectors are used to define the properties that can match an OpenMP context. OpenMP defines different sets of selectors, each containing different selectors.

The syntax for a context selector is *context-selector-specification* as described in the following grammar:

```
context-selector-specification :  
    trait-set-selector[ , trait-set-selector[ , ...]]  
  
trait-set-selector :  
    trait-set-selector-name={trait-selector[ , trait-selector[ , ...]]}  
  
trait-selector :  
    trait-selector-name[ ([trait-score: ] trait-property[ , trait-property[ , ...]]) ]  
  
trait-property :  
    trait-property-name  
    trait-property-clause  
    trait-property-expression  
    trait-property-extension  
  
trait-property-clause :  
    clause  
  
trait-property-name :  
    identifier  
    string-literal  
  
trait-property-expression  
    scalar-expression (for C/C++)  
    scalar-logical-expression (for Fortran)  
    scalar-integer-expression (for Fortran)  
  
trait-score :  
    score (score-expression)  
  
trait-property-extension :  
    trait-property-name  
    identifier (trait-property-extension[ , trait-property-extension[ , ...]])  
    constant integer expression
```

For trait selectors that correspond to name-list traits, each *trait-property* should be *trait-property-name* and for any value that is a valid identifier both the identifier and the

1 corresponding string literal (for C/C++) and the corresponding *char-literal-constant* (for Fortran)
2 representation are considered representations of the same value.

3 For trait selectors that correspond to clause-list traits, each *trait-property* should be
4 *trait-property-clause*. The syntax is the same as for the matching OpenMP clause.

5 The **construct** selector set defines the *construct* traits that should be active in the OpenMP
6 context. Each selector that can be defined in the **construct** set is the *directive-name* of a
7 context-matching construct. Each *trait-property* of the **simd** selector is a *trait-property-clause*.
8 The syntax is the same as for a valid clause of the **declare simd** directive and the restrictions on
9 the clauses from that directive apply. The **construct** selector is an ordered list c_1, \dots, c_N .

10 The **device** and **implementation** selector sets define the traits that should be active in the
11 corresponding trait set of the OpenMP context. The **target_device** selector set defines the
12 traits that should be active in the *target_device* trait set for the device that the specified
13 **device_num** selector identifies. The same traits that are defined in the corresponding traits sets
14 can be used as selectors with the same properties. The **kind** selector of the **device** and
15 **target_device** selector sets can also specify the value **any**, which is as if no **kind** selector
16 was specified. If a **device_num** selector does not appear in the **target_device** selector set
17 then a **device_num** selector that specifies the value of the *default-device-var* ICV is implied. For
18 the **device_num** selector of the **target_device** selector set, a single
19 *trait-property-expression* must be specified. For the **atomic_default_mem_order** selector of
20 the **implementation** set, a single *trait-property* must be specified as an identifier equal to one
21 of the valid arguments to the **atomic_default_mem_order** clause on the **requires**
22 directive. For the **requires** selector of the **implementation** set, each *trait-property* is a
23 *trait-property-clause*. The syntax is the same as for a valid clause of the **requires** directive and
24 the restrictions on the clauses from that directive apply.

25 The **user** selector set defines the **condition** selector that provides additional user-defined
26 conditions.

27 The **condition** selector contains a single *trait-property-expression* that must evaluate to *true* for
28 the selector to be true.

29 Any non-constant expression that is evaluated to determine the suitability of a variant is evaluated
30 according to the *data state* trait in the *dynamic* trait set of the OpenMP context.

31 The **user** selector set is dynamic if the **condition** selector is present and the expression in the
32 **condition** selector is not a constant expression; otherwise, it is static.

33 All parts of a context selector define the static part of the context selector except the following
34 parts, which define the dynamic part of a context selector:

- 35 • Its **user** selector set if it is dynamic; and
- 36 • Its **target_device** selector set.

37 For the **match** clause of a **declare variant** directive, any argument of the base function that
38 is referenced in an expression that appears in the context selector is treated as a reference to the

1 expression that is passed into that argument at the call to the base function. Otherwise, a variable or
2 procedure reference in an expression that appears in a context selector is a reference to the variable
3 or procedure of that name that is visible at the location of the directive on which the selector
4 appears.

▼ C++ ▼

5 Each occurrence of the **this** pointer in an expression in a context selector that appears in the
6 **match** clause of a **declare variant** directive is treated as an expression that is the address of
7 the object on which the associated base function is invoked.

▲ C++ ▲

8 Implementations can allow further selectors to be specified. Each specified *trait-property* for these
9 implementation-defined selectors should be *trait-property-extension*. Implementations can ignore
10 specified selectors that are not those described in this section.

11 Restrictions

12 Restrictions to context selectors are as follows:

- 13 ● Each *trait-property* can only be specified once in a *trait-selector* other than the **construct**
14 selector set.
- 15 ● Each *trait-set-selector-name* can only be specified once.
- 16 ● Each *trait-selector-name* can only be specified once.
- 17 ● A *trait-score* cannot be specified in traits from the **construct**, **device** or
18 **target_device** *trait-selector-sets*.
- 19 ● A *score-expression* must be a non-negative constant integer expression.
- 20 ● The expression of a **device_num** trait must evaluate to a non-negative integer value that is
21 less than or equal to the value of **omp_get_num_devices()**.
- 22 ● A variable or procedure that is referenced in an expression that appears in a context selector
23 must be visible at the location of the directive on which the selector appears unless the
24 directive is a **declare variant** directive and the variable is an argument of the associated
25 base function.
- 26 ● If *trait-property any* is specified in the **kind** *trait-selector* of the **device** or
27 **target_device** selector set, no other *trait-property* may be specified in the same selector.
- 28 ● For a *trait-selector* that corresponds to a name-list trait, at least one *trait-property* must be
29 specified.
- 30 ● For a *trait-selector* that corresponds to a non-property trait, no *trait-property* may be
31 specified.
- 32 ● For the **requires** selector of the **implementation** selector set, at least one
33 *trait-property* must be specified.

7.3 Matching and Scoring Context Selectors

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

- All selectors in the **user** set of the context selector are true;
- All traits and trait properties that are defined by selectors in the **target_device** set of the context selector are active in the *target_device* trait set for the device that is identified by the **device_num** selector;
- All traits and trait properties that are defined by selectors in the **construct**, **device** and **implementation** sets of the context selector are active in the corresponding trait sets of the OpenMP context;
- For each selector in the context selector, its properties are a subset of the properties of the corresponding trait of the OpenMP context;
- Selectors in the **construct** set of the context selector appear in the same relative order as their corresponding traits in the *construct* trait set of the OpenMP context; and
- No specified implementation-defined selector is ignored by the implementation.

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

- The **simdlen**(N) property of the selector matches the *simdlen*(M) trait of the OpenMP context if M is a multiple of N ; and
- The **aligned**(*list*: N) property of the selector matches the *aligned*(*list*: M) trait of the OpenMP context if N is a multiple of M .

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

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

7.4 Metadirectives

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. A metadirective is replaced by a **nothing** directive or one of the directive variants specified by the **when** clauses or the **otherwise** clause. If no **otherwise** clause is specified the effect is as if one was specified without an associated directive variant.

The OpenMP context for a given metadirective is defined according to [Section 7.1](#). The order of clauses that appear on a metadirective is significant and **otherwise** must be the last clause specified on a metadirective.

Replacement candidates are ordered according to the following rules in decreasing precedence:

- A candidate is before another one if the score associated with the context selector of the corresponding **when** clause is higher.
- A candidate that was explicitly specified is before one that was implicitly specified.
- Candidates are ordered according to the order in which they lexically appear on the metadirective.

The list of dynamic replacement candidates is the prefix of the sorted list of replacement candidates up to and including the first candidate for which the corresponding **when** clause has a static context selector. The first dynamic replacement candidate for which the corresponding **when** clause has a compatible context selector, according to the matching rules defined in [Section 7.3](#), replaces the metadirective.

Restrictions

Restrictions to metadirectives are as follows:

- Replacement of the metadirective with the directive variant associated with any of the dynamic replacement candidates must result in a conforming OpenMP program.
- Insertion of user code at the location of a metadirective must be allowed if the first dynamic replacement candidate does not have a static context selector.
- All items must be executable directives if the first dynamic replacement candidate does not have a static context selector.

Fortran

- A metadirective that appears in the specification part of a subprogram must follow all *variant-generating* declarative directives that appear in the same specification part.
- All directive variants of a metadirective must be pure otherwise the metadirective is not pure.

Fortran

7.4.1 when Clause

Name: <code>when</code>	Properties: <i>default</i>
--------------------------------	-----------------------------------

Arguments

Name	Type	Properties
<i>directive-variant</i>	directive-specification	optional, unique

Modifiers

Name	Modifies	Type	Properties
<i>context-selector</i>	<i>directive-variant</i>	An OpenMP context-selector-specification	required, unique

Directives

`begin metadirective, metadirective`

Semantics

The directive variant specified by a **when** clause is a candidate to replace the metadirective on which the clause is specified if the static part of the corresponding context selector is compatible with the OpenMP context according to the matching rules defined in [Section 7.3](#). If a **when** clause does not explicitly specify a directive variant it implicitly specifies a **nothing** directive as the directive variant.

Expressions that appear in the context selector of a **when** clause are evaluated if no prior dynamic replacement candidate has a compatible context selector, and the number of times each expression is evaluated is implementation defined. All variables referenced by these expressions are considered to be referenced by the metadirective.

A directive variant that is associated with a **when** clause can only affect the program if the directive variant is a dynamic replacement candidate.

Restrictions

Restrictions to the **when** clause are as follows:

- *directive-variant* must not specify a metadirective.
- *context-selector* must not specify any properties for the **simd** selector.

• <i>directive-variant</i> must not specify a begin declare variant directive.

Cross References

- `begin metadirective` directive, see [Section 7.4.4](#)
- `metadirective` directive, see [Section 7.4.3](#)
- `nothing` directive, see [Section 8.4](#)
- Context Selectors, see [Section 7.2](#)

7.4.2 otherwise Clause

Name: <code>otherwise</code>	Properties: unique, ultimate
-------------------------------------	-------------------------------------

Arguments

Name	Type	Properties
<i>directive-variant</i>	directive-specification	optional, unique

Directives

[begin metadirective](#), [metadirective](#)


Semantics

The `otherwise` clause is treated as a `when` clause with the specified directive variant, if any, and an always compatible static context selector that has a score lower than the scores associated with any other clause.

Restrictions

Restrictions to the `otherwise` clause are as follows:

- *directive-variant* must not specify a metadirective.



Cross References

- `when` clause, see [Section 7.4.1](#)
- `begin metadirective` directive, see [Section 7.4.4](#)
- `metadirective` directive, see [Section 7.4.3](#)

7.4.3 metadirective

Name: <code>metadirective</code> Category: meta	Association: none Properties: pure
--	---

Clauses

[otherwise](#), [when](#)

Semantics

The **metadirective** specifies metadirective semantics.

Cross References

- **otherwise** clause, see [Section 7.4.2](#)
- **when** clause, see [Section 7.4.1](#)
- Metadirectives, see [Section 7.4](#)

7.4.4 begin metadirective

Name: <code>begin metadirective</code> Category: meta	Association: delimited Properties: pure
--	--

Clauses

[otherwise](#), [when](#)

Semantics

The **begin metadirective** is a metadirective for which the specified directive variants other than the **nothing** directive must accept a paired **end** directive. For any directive variant that is selected to replace the **begin metadirective** directive, the **end metadirective** directive is implicitly replaced by its paired **end** directive to demarcate the statements that are affected by or are associated with the directive variant. If the **nothing** directive is selected to replace the **begin metadirective** directive, the paired **end metadirective** is ignored.

Restrictions

The restrictions to **begin metadirective** are as follows:

- Any *directive-variant* that is specified by a **when** or **otherwise** clause must be an OpenMP directive that has a paired **end** directive or must be the **nothing** directive.

Cross References

- **otherwise** clause, see [Section 7.4.2](#)
- **when** clause, see [Section 7.4.1](#)
- **nothing** directive, see [Section 8.4](#)
- Metadirectives, see [Section 7.4](#)

7.5 Declare Variant Directives

Declare variant directives declare *base functions* to have the specified function variant. The context selector in the **match** clause is associated with the variant.

The OpenMP context for a direct call to a given base function is defined according to [Section 7.1](#). If a declare variant directive for the base function is visible at the call site and the static part of the context selector that is associated with the declared function variant is compatible with the OpenMP context of the call according to the matching rules defined in [Section 7.3](#) then the variant is a replacement candidate to be called instead of the base function. Replacement candidates are ordered in decreasing order of the score associated with the context selector. If two replacement candidates have the same score then their order is implementation defined.

The list of dynamic replacement candidates is the prefix of the sorted list of replacement candidates up to and including the first candidate for which the corresponding context selector is static.

The first dynamic replacement candidate for which the corresponding context selector is compatible, according to the matching rules defined in [Section 7.3](#), is called instead of the base function. If no compatible candidate exists then the base function is called.

Expressions that appear in the context selector of a **match** clause are evaluated if no prior dynamic replacement candidate has a compatible context selector, and the number of times each expression is evaluated is implementation defined. All variables referenced by these expressions are considered to be referenced at the call site.

▼ C++ ▲

For calls to **constexpr** base functions that are evaluated in constant expressions, whether any variant replacement occurs is implementation defined.

▲ C++ ▼

For indirect function calls that can be determined to call a particular base function, whether any variant replacement occurs is unspecified.

Any differences that the specific OpenMP context requires in the prototype of the variant from the base function prototype are implementation defined.

Different declare variant directives may be specified for different declarations of the same base function.

Restrictions

Restrictions to declare variant directives are as follows:

- Calling functions that a declare variant directive determined to be a function variant directly in an OpenMP context that is different from the one that the **construct** selector set of the context selector specifies is non-conforming.
- If a function is determined to be a function variant through more than one declare variant directive then the **construct** selector set of their context selectors must be the same.

- A function determined to be a function variant may not be specified as a base function in another `declare variant` directive.
- An `adjust_args` clause or `append_args` clause can only be specified if the `dispatch` selector of the `construct` selector set appears in the `match` clause.

C / C++

- The type of the function variant must be compatible with the type of the base function after the implementation-defined transformation for its OpenMP context.

C / C++

C++

- Declare variant directives cannot be specified for virtual, defaulted or deleted functions.
- Declare variant directives cannot be specified for constructors or destructors.
- Declare variant directives cannot be specified for immediate functions.
- The function that a declare variant directive determined to be a function variant may not be an immediate function.

C++

Cross References

- `begin declare variant` directive, see [Section 7.5.5](#)
- `declare variant` directive, see [Section 7.5.4](#)
- Context Selectors, see [Section 7.2](#)
- OpenMP Contexts, see [Section 7.1](#)

7.5.1 match Clause

Name: <code>match</code>	Properties: unique, required
---------------------------------	-------------------------------------

Arguments

Name	Type	Properties
<i>context-selector</i>	An OpenMP context-selector-specification	<i>default</i>

Directives

[begin declare variant](#), [declare variant](#)

Semantics

The `match` clause specifies the *context-selector* to use to determine if a specified variant function is a replacement candidate for the specified base function in a given context.

Restrictions

Restrictions to the **match** clause are as follows:

- All variables that are referenced in an expression that appears in the context selector of a **match** clause must be accessible at a call site to the base function according to the base language rules.

Cross References

- **begin declare variant** directive, see [Section 7.5.5](#)
- **declare variant** directive, see [Section 7.5.4](#)
- Context Selectors, see [Section 7.2](#)

7.5.2 adjust_args Clause

Name: adjust_args	Properties: <i>default</i>
--------------------------	----------------------------

Arguments

Name	Type	Properties
<i>parameter-list</i>	list of parameter list item type	<i>default</i>

Modifiers

Name	Modifies	Type	Properties
<i>adjust-op</i>	<i>parameter-list</i>	Keyword: need_device_ptr , nothing	required

Directives

declare variant

Semantics

The **adjust_args** clause specifies how to adjust the arguments of the base function when a specified variant function is selected for replacement. For each **adjust_args** clause that is present on the selected variant the adjustment operation specified by *adjust-op* is applied to each argument specified in the clause before being passed to the selected variant. If the *adjust-op* modifier is **nothing**, the argument is passed to the selected variant without being modified.

If the *adjust-op* modifier is **need_device_ptr**, the arguments are converted to corresponding device pointers of the default device. If an argument has the *is_device_ptr* property in its *interoperability requirement set* then the argument is not adjusted. Otherwise, the argument is converted in the same manner that a **use_device_ptr** clause on a **target data** construct converts its pointer list items into device pointers. If the argument cannot be converted into a device pointer then *NULL* is passed as the argument.

Restrictions

Fortran

- Each argument that appears in a **need_device_ptr** *adjust-op* must be of type **C_PTR** in the dummy argument declaration of the variant function.

Fortran

Cross References

- **declare variant** directive, see [Section 7.5.4](#)

7.5.3 **append_args** Clause

Name: <code>append_args</code>	Properties: unique
---------------------------------------	---------------------------

Arguments

Name	Type	Properties
<i>append-op-list</i>	list of OpenMP operation list item type	<i>default</i>

Directives

[declare variant](#)

Semantics

The **append_args** clause specifies additional arguments to pass in the call when a specified variant function is selected for replacement. If no **interop** clause is specified on an associated **dispatch** construct then the arguments are constructed according to each specified list item in *append-op-list*. If an **interop** clause is specified with *n* variables on an associated **dispatch** construct then the arguments are constructed in the same order in which they appear in the **interop** clause and the first *n* list items in the *append-op-list* are omitted. Any remaining list items in the *append-op-list* are used to construct additional arguments that follow the arguments that are constructed from the variables from the **interop** clause. In either case, the arguments are passed to the variant function after any named arguments of the base function in the same order in which they are constructed. If the base function is variadic, the constructed arguments are passed before any variadic arguments.

The supported OpenMP operations in *append-op-list* are:

interop

The **interop** operation accepts as its *operator-parameter-specification* any *modifier-specification-list* that is accepted by the **init** clause on the **interop** construct.

Each **interop** operation for an *append-op-list* list item that is not omitted constructs an argument of **interop** OpenMP type using the *interoperability requirement set* of the encountering task. The argument is constructed as if by an **interop** construct with an **init** clause that specifies the *modifier-specification-list* specified in the **interop** operation. If the *interoperability requirement set* contains one or more properties that could be used as clauses for an **interop** construct of

1 *interop-type*, the behavior is as if the corresponding clauses would also be part of the **interop**
2 construct and those properties are removed from the *interopability requirement set*.

3 This argument is destroyed after the call to the selected variant returns, as if an **interop** construct
4 with a **destroy** clause was used with the same clauses that were used to initialize the argument.

5 **Cross References**

- 6 • **init** clause, see [Section 14.1.2](#)
- 7 • **declare variant** directive, see [Section 7.5.4](#)
- 8 • **interop** directive, see [Section 14.1](#)
- 9 • Interoperability Requirement Set, see [Section 14.2](#)
- 10 • OpenMP Operations, see [Section 3.2.3](#)

11 **7.5.4 declare variant Directive**

12 Name: <code>declare variant</code>	Association: declaration
Category: declarative	Properties: pure

13 **Arguments**

14 **declare variant** (*[base-name:]variant-name*)

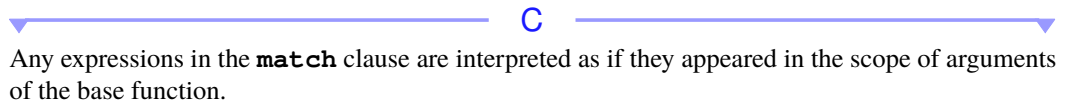
15 Name	Type	Properties
<i>base-name</i>	identifier of function type	optional
<i>variant-name</i>	identifier of function type	<i>default</i>

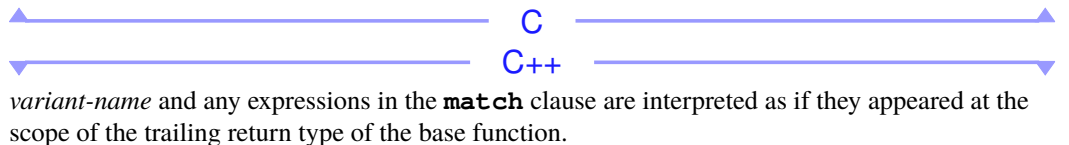
16 **Clauses**

17 [adjust_args](#), [append_args](#), [match](#)

18 **Semantics**

19 The **declare variant** specifies declare variant semantics for a single replacement candidate.
20 *variant-name* identifies the function variant while *base-name* identifies the base function.

21  **C** Any expressions in the **match** clause are interpreted as if they appeared in the scope of arguments
22 of the base function.

23  **C++** *variant-name* and any expressions in the **match** clause are interpreted as if they appeared at the
24 scope of the trailing return type of the base function.

25 The function variant is determined by base language standard name lookup rules ([basic.lookup])
26 of *variant-name* using the argument types at the call site after implementation-defined changes have
27 been made according to the OpenMP context.

 **C++**

Fortran

The procedure to which *base-name* refers is resolved at the location of the directive according to the establishment rules for procedure names in the base language.

If a **declare variant** directive appears in the specification part of a subprogram or an interface body, its bound procedure is this subprogram or the procedure defined by the interface body, respectively. Otherwise there is no bound procedure.

Fortran

Restrictions

C / C++

- If *base-name* is specified, it must match the name used in the associated declaration, if any declaration is associated.

C / C++

Fortran

- If the **declare variant** directive does not have a bound procedure or the base function is not the bound procedure, *base-name* must be specified.
- *base-name* must not be a generic name, an entry name, the name of a procedure pointer, a dummy procedure or a statement function.
- The procedure *base-name* must have an accessible explicit interface at the location of the directive.

Fortran

Cross References

- **adjust_args** clause, see [Section 7.5.2](#)
- **append_args** clause, see [Section 7.5.3](#)
- **match** clause, see [Section 7.5.1](#)
- Declare Variant Directives, see [Section 7.5](#)

C / C++

7.5.5 begin declare variant Directive

Name: `begin declare variant`

Association: delimited (declaration-definition-seq)

Category: declarative

Properties: *default*

Clauses

match

Semantics

The **begin declare variant** directive associates the context selector in the **match** clause with each function definition in *declaration-definition-seq*. For the purpose of call resolution, each function definition that appears between a **begin declare variant** directive and its paired **end** directive is a function variant for an assumed base function, with the same name and a compatible prototype, that is declared elsewhere without an associated declare variant directive.

If a declare variant directive appears between a **begin declare variant** directive and its paired **end** directive, the effective context selectors of the outer directive are appended to the context selector of the inner directive to form the effective context selector of the inner directive. If a *trait-set-selector* is present on both directives, the *trait-selector* list of the outer directive is appended to the *trait-selector* list of the inner directive after equivalent *trait-selectors* have been removed from the outer list. Restrictions that apply to explicitly specified context selectors also apply to effective context selectors constructed through this process.

The symbol name of a function definition that appears between a **begin declare variant** directive and its paired **end** directive is determined through the base language rules after the name of the function has been augmented with a string that is determined according to the effective context selector of the **begin declare variant** directive. The symbol names of two definitions of a function are considered to be equal if and only if their effective context selectors are equivalent.

If the context selector of a **begin declare variant** directive contains traits in the *device* or *implementation* set that are known never to be compatible with an OpenMP context during the current compilation, the preprocessed code that follows the **begin declare variant** directive up to its paired **end** directive is elided.

Any expressions in the **match** clause are interpreted at the location of the directive.

Restrictions

The restrictions to **begin declare variant** directive are as follows:

- **match** clause must not contain a **simd** *trait-selector-name*.
- Two **begin declare variant** directives and their paired **end** directives must either encompass disjoint source ranges or be perfectly nested.
- **match** clause must not contain a dynamic context selector that references the **this** pointer.
- If an expression in the context selector that appears in **match** clause references the **this** pointer, the base function must be a non-static member function.

Cross References

- **match** clause, see [Section 7.5.1](#)
- Declare Variant Directives, see [Section 7.5](#)

7.6 `dispatch` Construct

Name: <code>dispatch</code>	Association: block (function dispatch structured block)
Category: executable	Properties: context-matching

Clauses

`depend`, `device`, `interop`, `is_device_ptr`, `nocontext`, `novariants`, `nowait`

Binding

The binding task set for a `dispatch` region is the generating task. The `dispatch` region binds to the region of the generating task.

Semantics

The `dispatch` construct controls whether variant substitution occurs for *target-call* in the associated function dispatch structured block. The `dispatch` construct may also specify properties to be passed to the variant function if variant substitution occurs.

Properties added to the *interoperability requirement set* can be removed by the effect of other directives (see [Section 14.2](#)) before the `dispatch` region is executed. If one or more `depend` clauses are present on the `dispatch` construct, they are added as *depend* properties of the *interoperability requirement set*. If a `nowait` clause is present on the `dispatch` construct the *nowait* property is added to the *interoperability requirement set*. For each list item specified in an `is_device_ptr` clause, an *is_device_ptr* property for that list item is added to the *interoperability requirement set*.

If the *interoperability requirement set* contains one or more *depend* properties, the behavior is as if those properties were applied as `depend` clauses to a `taskwait` construct that is executed before the `dispatch` region is executed.

The presence of the `nowait` property in the *interoperability requirement set* has no effect on the `dispatch` construct.

If the `device` clause is present, the value of the *default-device-var* ICV is set to the value of the expression in the clause on entry to the `dispatch` region and is restored to its previous value at the end of the region.

If variant substitution occurs, the `interop` clause specifies additional arguments to pass to the variant function selected for replacement.

If the `interop` clause is present and has only one *interop-var*, and the `device` clause is not specified, the behavior is as if the `device` clause is present with a *device-description* equivalent to the `device_num` property of the *interop-var*.

Restrictions

Restrictions to the `dispatch` construct are as follows:

- If the `interop` clause is present and has more than one *interop-var* then the `device` clause must also be present.

Cross References

- **depend** clause, see [Section 15.9.5](#)
- **device** clause, see [Section 13.2](#)
- **interop** clause, see [Section 7.6.1](#)
- **is_device_ptr** clause, see [Section 5.4.7](#)
- **nocontext** clause, see [Section 7.6.3](#)
- **novariants** clause, see [Section 7.6.2](#)
- **nowait** clause, see [Section 15.6](#)
- Interoperability Requirement Set, see [Section 14.2](#)
- OpenMP Function Dispatch Structured Blocks, see [Section 4.3.1.2](#)

7.6.1 interop Clause

Name: <code>interop</code>	Properties: unique
-----------------------------------	---------------------------

Arguments

Name	Type	Properties
<i>interop-var-list</i>	list of variable of interop OpenMP type	<i>default</i>

Directives

[dispatch](#)

Semantics

The **interop** clause specifies additional arguments to pass to the variant function when variant substitution occurs for the *target-call* in a **dispatch** construct. The variables in the *interop-var-list* are passed in the same order in which they are specified in the **interop** clause.

Restrictions

Restrictions to the **interop** clause are as follows:

- If the **interop** clause is specified on a **dispatch** construct, the matching **declare variant** directive for the *target-call* must have an **append_args** clause with a number of list items that equals or exceeds the number of list items in the **interop** clause.

Cross References

- **dispatch** directive, see [Section 7.6](#)

7.6.2 novariants Clause

Name: <code>novariants</code>	Properties: unique
--------------------------------------	---------------------------

Arguments

Name	Type	Properties
<i>do-not-use-variant</i>	expression of logical type	<i>default</i>

Directives

[dispatch](#)

Semantics

If *do-not-use-variant* evaluates to *true*, no function variant is selected for the *target-call* of the **dispatch** region associated with the **novariants** clause even if one would be selected normally. The use of a variable in *do-not-use-variant* causes an implicit reference to the variable in all enclosing constructs. *do-not-use-variant* is evaluated in the enclosing context.

Cross References

- **dispatch** directive, see [Section 7.6](#)

7.6.3 nocontext Clause

Name: <code>nocontext</code>	Properties: unique
-------------------------------------	---------------------------

Arguments

Name	Type	Properties
<i>do-not-update-context</i>	expression of logical type	<i>default</i>

Directives

[dispatch](#)

Semantics

If *do-not-update-context* evaluates to *true*, the construct on which the **nocontext** clause appears is not added to the *construct* set of the OpenMP context. The use of a variable in *do-not-update-context* causes an implicit reference to the variable in all enclosing constructs. *do-not-update-context* is evaluated in the enclosing context.

Cross References

- **dispatch** directive, see [Section 7.6](#)

7.7 declare simd Directive

Name: <code>declare simd</code> Category: declarative	Association: declaration Properties: pure
--	--

Arguments

`declare simd[(proc-name)]`

Name	Type	Properties
<i>proc-name</i>	identifier of function type	optional

Clause groups

branch

Clauses

`aligned`, `linear`, `simdlen`, `uniform`

Semantics

The association of one or more `declare simd` directives with 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.

If a SIMD version is created and the `simdlen` clause is not specified, the number of concurrent arguments for the function is implementation defined.

For purposes of the `linear` clause, any integer-typed parameter that is specified in a `uniform` clause on the directive is considered to be constant and so may be used in *linear-step*.

▼ C / C++ ▼

The expressions that appear in the clauses of each directive are evaluated in the scope of the arguments of the function declaration or definition.

▲ C / C++ ▲

▼ C++ ▼

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

▲ C++ ▲

Restrictions

Restrictions to the `declare simd` directive are as follows:

- 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.

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

C / C++

C++

- The function cannot contain **throw** statements.

C++

Fortran

- *proc-name* must not be a generic name, procedure pointer, or entry name.
- If *proc-name* is omitted, the **declare simd** directive must appear in the specification part of a subroutine subprogram or a function subprogram for which creation of the SIMD versions is enabled.
- Any **declare simd** directive must appear in the specification part of a subroutine subprogram, function subprogram, or interface body to which it applies.
- If a **declare simd** directive is specified in an interface block for a procedure, it must match a **declare simd** directive in the definition of the procedure.
- If a procedure is declared via a procedure declaration statement, the procedure *proc-name* should appear in the same specification.
- If a **declare simd** directive is specified for a procedure name with explicit interface and a **declare simd** directive is also specified for the definition of the procedure then the two **declare simd** directives must match.
- Procedure pointers may not be used to access versions created by the **declare simd** directive.

Fortran

Cross References

- **aligned** clause, see [Section 5.11](#)
- **linear** clause, see [Section 5.4.6](#)
- **reduction** clause, see [Section 5.5.8](#)
- **simdlen** clause, see [Section 10.5.3](#)
- **uniform** clause, see [Section 5.10](#)

7.7.1 *branch* Clauses

Clause groups

Properties: unique, exclusive, inarguable	Members: <code>inbranch</code> , <code>notinbranch</code>
--	--

Directives

`declare simd`

Semantics

The *branch* clause grouping defines a set of clauses that indicate if a function can be assumed to be or not to be encountered in a branch. The `inbranch` clause specifies that the function will always be called from inside a conditional statement of the calling context. The `notinbranch` clause specifies that the function will never be called from inside a conditional statement of the calling context. If neither clause is specified, then the function may or may not be called from inside a conditional statement of the calling context.

Cross References

- `declare simd` directive, see [Section 7.7](#)

7.8 Declare Target Directives

Declare target directives apply to procedures and/or variables to ensure that they can be executed or accessed on a device. Variables are either replicated as device local variables for each device through a `local` clause, are mapped for all device executions through an `enter` clause, or are mapped for specific device executions through a `link` clause. An implementation may generate different versions of a procedure to be used for `target` regions that execute on different devices. Whether the same version is generated for different devices, or whether a version that is called in a `target` region differs from the version that is called outside a `target` region, is implementation defined.

To facilitate device usage, OpenMP defines rules that implicitly specify declare target directives for procedures and variables. The remainder of this section defines those rules as well as restrictions that apply to all declare target directives.

If a variable with static storage duration that is not `groupprivate` is declared in a device routine then the named variable is treated as if it had appeared in an `enter` clause on a declare target directive while a declared `groupprivate` variable is treated as if it had appeared in a `local` clause on a declare target directive.

In the following, a non-host declare target directive is one that does not specify a `device_type` clause with `host`. Further, a reverse-offload region is a region that is associated with a `target` construct that specifies a `device` clause with the *ancestor device-modifier*.

If a procedure is referenced outside of any reverse-offload region in a procedure that appears as a list item in an `enter` clause on a non-host declare target directive then the name of the referenced procedure is treated as if it had appeared in an `enter` clause on a declare target directive.

C / C++

1 If a variable with static storage duration or a function (except *lambda* for C++) is referenced in the
2 initializer expression list of a variable with static storage duration that appears as a list item in an
3 **enter** or **local** clause on a declare target directive then the name of the referenced variable or
4 function is treated as if it had appeared in an **enter** clause on a declare target directive.

C / C++

Fortran

5 If a **declare target** directive has a **device_type** clause then any enclosed internal
6 procedures cannot contain any **declare target** directives. The enclosing **device_type**
7 clause implicitly applies to internal procedures.

Fortran

8 Execution Model Events

9 The *target-global-data-op* event occurs when an original variable is associated with a
10 corresponding variable on a device as a result of a declare target directive; the event occurs before
11 the first access to the corresponding variable.

12 Tool Callbacks

13 A thread dispatches a registered **ompt_callback_target_data_op** callback, or a registered
14 **ompt_callback_target_data_op_emi** callback with **ompt_scope_beginend** as its
15 endpoint argument for each occurrence of a *target-global-data-op* event in that thread. These
16 callbacks have type signature **ompt_callback_target_data_op_t** or
17 **ompt_callback_target_data_op_emi_t**, respectively.

18 Restrictions

19 Restrictions to any declare target directive are as follows:

- 20 • A variable declared in the directive must have a mappable type.
- 21 • A variable declared in the directive must have static storage duration.
- 22 • The same list item must not explicitly appear in both an **enter** clause on one declare target
23 directive and a **link** or **local** clause on another declare target directive.
- 24 • The same list item must not explicitly appear in both a **link** clause on one declare target
25 directive and a **local** clause on another declare target directive.
- 26 • If a variable appears in a **enter** clause on the declare target directive, its initializer must not
27 refer to a variable that appears in a **link** clause on a declare target directive.

Cross References

- **enter** clause, see [Section 5.8.4](#)
- **link** clause, see [Section 5.8.5](#)
- **begin declare target** directive, see [Section 7.8.2](#)
- **declare target** directive, see [Section 7.8.1](#)
- **target** directive, see [Section 13.8](#)
- **ompt_callback_target_data_op_emi_t** and **ompt_callback_target_data_op_t**, see [Section 19.5.2.25](#)

7.8.1 declare target Directive

Name: declare target Category: declarative	Association: none Properties: device, declare target, pure
---	---

Arguments

declare target (*extended-list*)

Name	Type	Properties
<i>extended-list</i>	list of extended list item type	optional

Clauses

[device_type](#), [enter](#), [indirect](#), [link](#), [local](#)

Semantics

The **declare target** directive is a declare target directive. If the *extended-list* argument is specified, the effect is as if any list items from *extended-list* that are not groupprivate variables appear in the *extended-list* argument to an implicit **enter** clause and any list items that are groupprivate variables appear in the *list* argument to an implicit **local** clause.

▼ C++ ▼

If the **declare target** directive is specified as an attribute specifier with the **decl** attribute and a **decl** attribute is not used on the declaration to specify groupprivate variables, the effect is as if an **enter** clause is specified if a **link** clause or **local** clause is not specified.

If the **declare target** directive is specified as an attribute specifier with the **decl** attribute and a **decl** attribute is used on the declaration to specify groupprivate variables, the effect is as if a **local** clause is specified.

▲ C++ ▲
▼ Fortran ▼

If a **declare target** directive does not have any clauses and does not have an *extended-list* then an implicit **enter** clause with one item is formed from the name of the enclosing subroutine subprogram, function subprogram or interface body to which it applies.

▲ Fortran ▲

Restrictions

Restrictions to the **declare target** directive are as follows:

- If the *extended-list* argument is specified, no clauses may be specified.
- If the directive has a clause, it must contain at least one **enter** clause, **link** clause, or **local** clause.
- A variable for which **nohost** is specified may not appear in a **link** clause.
- A variable that is **groupprivate** may not appear in an **enter** clause or **link** clause.

Fortran

- If a list item is a procedure name, it must not be a generic name, procedure pointer, entry name, or statement function name.
- If no clauses are specified or if a **device_type** clause is specified, the directive must appear in a specification part of a subroutine subprogram, function subprogram or interface body.
- If a list item is a procedure name, the directive must be in the specification part of that subroutine or function subprogram or in the specification part of that subroutine or function in an interface body.
- If an extended list item is a variable name, the directive must appear in the specification part of a subroutine subprogram, function subprogram, program or module.
- If the directive is specified in an interface block for a procedure, it must match a **declare target** directive in the definition of the procedure, including the **device_type** clause if present.
- If an external procedure is a type-bound procedure of a derived type and the directive is specified in the definition of the external procedure, it must appear in the interface block that is accessible to the derived-type definition.
- If any procedure is declared via a procedure declaration statement that is not in the type-bound procedure part of a derived-type definition, any **declare target** with the procedure name must appear in the same specification part.
- The directive must appear in the declaration section of a scoping unit in which the common block or variable is declared.
- If a **declare target** 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, after the last such **COMMON** statement in the program unit.
- If a list item is declared with the **BIND** attribute, the corresponding C entities must also be specified in a **declare target** directive in the C program.

- A variable can only appear in a **declare target** 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 **declare target** directive must be declared in the Fortran scope of a module or have the **SAVE** attribute, either explicitly or implicitly.

Fortran

Cross References

- **device_type** clause, see [Section 13.1](#)
- **enter** clause, see [Section 5.8.4](#)
- **indirect** clause, see [Section 7.8.3](#)
- **link** clause, see [Section 5.8.5](#)
- **local** clause, see [Section 5.13](#)
- Declare Target Directives, see [Section 7.8](#)

C / C++

7.8.2 begin declare target Directive

Name: <code>begin declare target</code>	Association: delimited (declaration-definition-seq)
Category: declarative	Properties: device, declare target

Clauses

[device_type](#), [indirect](#)

Semantics

The **begin declare target** directive is a declare target directive. The directive and its paired **end** directive form a delimited code region that defines an implicit *extended-list* and implicit *local-list* that is converted to an implicit **enter** clause with the *extended-list* as its argument and an implicit **local** clause with the *local-list* as its argument, respectively.

The implicit *extended-list* consists of the variable and function names of any variable or function declarations at file scope that appear in the delimited code region, excluding declarations of groupprivate variables. If any groupprivate variables are declared in the delimited code region, the effect is as if the variables appear in the implicit *local-list*.

C++

1 Additionally, the implicit *extended-list* and *local-list* consist of the variable and function names of
2 any variable or function declarations at namespace or class scope that appear in the delimited code
3 region, including the **operator ()** member function of the resulting closure type of any lambda
4 expression that is defined in the delimited code region.

C++

5 The delimited code region may contain declare target directives. If a **device_type** clause is
6 present on the contained declare target directive, then its argument determines which versions are
7 made available. If a list item appears both in an implicit and explicit list, the explicit list determines
8 which versions are made available.

Restrictions

9 Restrictions to the **begin declare target** directive are as follows:

C++

- 11 • The function names of overloaded functions or template functions may only be specified
12 within an implicit *extended-list*.
- 13 • If a *lambda declaration and definition* appears between a **begin declare target**
14 directive and the paired **end** directive, all variables that are captured by the *lambda*
15 expression must also appear in an **enter** clause.
- 16 • A module *export* or *import* statement cannot appear between a declare target directive and
17 the paired **end** directive.

C++

Cross References

- 18 • **device_type** clause, see [Section 13.1](#)
- 19 • **enter** clause, see [Section 5.8.4](#)
- 20 • **indirect** clause, see [Section 7.8.3](#)
- 21 • **indirect** clause, see [Section 7.8.3](#)
- 22 • Declare Target Directives, see [Section 7.8](#)

C / C++

7.8.3 indirect Clause

Name: <code>indirect</code>	Properties: unique
------------------------------------	---------------------------

Arguments

Name	Type	Properties
<i>invoked-by-fptr</i>	expression of logical type	constant, optional

1 **Directives**
2 **begin declare target, declare target**

3 **Semantics**

4 If *invoked-by-fptr* evaluates to true, any procedures that appear in an **enter** clause on the directive
5 on which the **indirect** clause is specified may be called with an indirect device invocation. If the
6 *invoked-by-fptr* does not evaluate to true, any procedures that appear in an **enter** clause on the
7 directive may not be called with an indirect device invocation. Unless otherwise specified by an
8 **indirect** clause, procedures may not be called with an indirect device invocation. If the
9 **indirect** clause is specified and *invoked-by-fptr* is not specified, the effect of the clause is as if
10 *invoked-by-fptr* evaluates to true.



11 If a function appears in the implicit **enter** clause of a **begin declare target** directive and in
12 the **enter** clause of a **declare target** directive that is contained in the delimited code region of the
13 **begin declare target** directive, and if an **indirect** clause appears on both directives, then
14 the **indirect** clause on the **begin declare target** directive has no effect for that function.



15 **Restrictions**

16 Restrictions to the **indirect** clause are as follows:

- 17 • If *invoked-by-fptr* evaluates to true, a **device_type** clause must not appear on the same
18 directive unless it specifies **any.** for its *device-type-description*.

19 **Cross References**

- 20 • **begin declare target** directive, see [Section 7.8.2](#)
- 21 • **declare target** directive, see [Section 7.8.1](#)

8 Informational and Utility Directives

An informational directive conveys information about code properties to the compiler while a utility directive facilitates interactions with the compiler or supports code readability. A utility directive is informational unless the **at** clause implies it to be executable.

8.1 at Clause

Name: <code>at</code>	Properties: <code>unique</code>
------------------------------	--

Arguments

Name	Type	Properties
<i>action-time</i>	Keyword: compilation , execution	<i>default</i>

Directives

[error](#)

Semantics

The **at** clause determines when the implementation performs an action that is associated with a utility directive. If *action-time* is **compilation**, the action is performed during compilation if the directive appears in a declarative context or in an executable context that is reachable at runtime. If *action-time* is **compilation** and the directive appears in an executable context that is not reachable at runtime, the action may or may not be performed. If *action-time* is **execution**, the action is performed during program execution when a thread encounters the directive and the directive is considered to be an executable directive. If the **at** clause is not specified, the effect is as if *action-time* is **compilation**.

Cross References

- **error** directive, see [Section 8.5](#)

8.2 requires Directive

Name: <code>requires</code> Category: <code>informational</code>	Association: <code>none</code> Properties: <code>default</code>
---	--

Clause groups

requirement

Semantics

The **requires** directive specifies features that an implementation must support for correct execution and requirements for the execution of all code in the current compilation unit. 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 clauses of a **requires** directive are added to the *requires* trait in the OpenMP context for all program points that follow the directive.

Restrictions

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

- All **requires** directives in the same compilation unit that specify the **atomic_default_mem_order** requirement must specify the same argument.
- A **requires** directive may not appear lexically after a context selector in which any clause of the **requires** directive is used.
- Either all compilation units of a program that contain declare target directives, device constructs or device routines or none of them must specify a **requires** directive that specifies the **reverse_offload**, **unified_address** or **unified_shared_memory** requirement.
- A **requires** directive that specifies the **atomic_default_mem_order** requirement must not appear lexically after any **atomic** construct on which *memory-order-clause* is not specified.

C

- The **requires** directive may only appear at file scope.

C

C++

- The **requires** directive may only appear at file or namespace scope.

C++

C / C++

- Any **requires** directive that specifies a **reverse_offload**, **unified_address**, or **unified_shared_memory** requirement must appear lexically before any device constructs or device routines.

C / C++

- The **requires** directive must appear in the specification part of a program unit, either after all **USE** statements, **IMPORT** statements, and **IMPLICIT** statements or by referencing a module. Additionally, it may appear in the specification part of an internal or module subprogram that appears by referencing a module if each clause already appeared with the same arguments in the specification part of the program unit.

8.2.1 *requirement* Clauses

Clause groups

Properties: unique	Members: <code>atomic_default_mem_order</code> , <code>dynamic_allocators</code> , <code>reverse_offload</code> , <code>unified_address</code> , <code>unified_shared_memory</code>
---------------------------	---

Directives

`requires`

Semantics

The *requirement* clause grouping defines a set of clauses that indicate the requirement that a program requires the implementation to support. Other than `atomic_default_mem_order`, the members of the set are inarguable.

If an implementation supports a given *requirement* clause then the use of that clause on a **requires** directive will cause the implementation to ensure the enforcement of a guarantee represented by the specific member of the clause grouping. If the implementation does not support the requirement then it must perform compile-time error termination.

The `reverse_offload` clause requires an implementation to guarantee that if a **target** construct specifies a **device** clause in which the **ancestor** modifier appears, the **target** region can execute on the parent device of an enclosing **target** region.

The `unified_address` clause requires an implementation to guarantee that all devices accessible through OpenMP API routines and directives use a unified address space. In this address space, a pointer will always refer to the same location in memory from all devices accessible through OpenMP. Any OpenMP mechanism that returns a device pointer is guaranteed to return a device address that supports pointer arithmetic, and the `is_device_ptr` clause is not necessary to obtain device addresses from device pointers for use inside **target** regions. Host pointers may be passed as device pointer arguments to device memory routines and device pointers may be passed as host pointer arguments to device memory routines. Non-host devices may still have discrete memories and dereferencing a device pointer on the host device or a host pointer on a non-host device remains unspecified behavior. Memory local to a specific execution context may be

1 exempt from the **unified_address** requirement, following the restrictions of locality to a given
2 execution context, thread or contention group.

3 The **unified_shared_memory** clause implies the **unified_address** requirement,
4 inheriting all of its behaviors. The implementation must also guarantee that storage locations in
5 memory are accessible to threads on all available devices, except for memory that is local to a
6 specific execution context as defined in the description of **unified_address** above. Every
7 device address that refers to storage allocated through OpenMP routines is a valid host pointer that
8 may be dereferenced.

9 The **unified_shared_memory** clause makes **map** clauses optional on **target** constructs and
10 declare target directives optional for variables with static storage duration that are accessed inside
11 functions to which a declare target directive is applied. Scalar variables are still firstprivate by
12 default when referenced inside **target** constructs. Values stored into memory by one device may
13 not be visible to another device until those two devices synchronize with each other or both devices
14 synchronize with the host.

15 The **dynamic_allocators** clause removes certain restrictions on the use of memory allocators
16 in **target** regions. Specifically, allocators (including the default allocator that is specified by the
17 *def-allocator-var* ICV) may be used in a **target** region or in an **allocate** clause on a **target**
18 construct without specifying the **uses_allocators** clause on the **target** construct.
19 Additionally, the implementation must support calls to the **omp_init_allocator** and
20 **omp_destroy_allocator** API routines in **target** regions.

21 The **atomic_default_mem_order** clause specifies the default memory ordering behavior for
22 **atomic** constructs that an implementation must provide. The effect is as if its argument appears as
23 a clause on any **atomic** construct that does not specify a memory order clause.

24 **Cross References**

- 25 • **requires** directive, see [Section 8.2](#)

26 **8.3 Assumption Directives**

27 Assumption directives provide invariants that specify additional information about the expected
28 properties of the program that can optionally be used for optimization. An implementation may
29 ignore this information without altering the behavior of the program. Different assumption
30 directive formats facilitate definition of assumptions for a scope that is appropriate to each base
31 language. The scope of a particular format is its *assumption scope* and is defined in the section that
32 defines that format. If the invariants do not hold at runtime, the behavior is unspecified.

8.3.1 *assumption* Clauses

Clause groups

Properties:	Members: absent , contains , holds , no_openmp , no_openmp_routines , no_parallelism
-------------	--

Directives

assume, **assumes**, **begin assumes**

Semantics

The *assumption* clause grouping defines a set of clauses that indicate the assumptions that a program ensures the implementation can exploit. Other than **absent**, **contains** and **holds**, the members of the set are inarguable and unique.

The **no_openmp** clause guarantees that no OpenMP related code is executed in the assumption scope. The **no_openmp_routines** clause guarantees that no explicit OpenMP runtime library calls are executed in the assumption scope. The **no_parallelism** clause guarantees that no OpenMP tasks (explicit or implicit) will be generated and that no SIMD constructs will be executed in the assumption scope.

▼ C++ ▼

The **no_openmp** clause also guarantees that no thread will throw an exception in the assumption scope if it is contained in a region that arises from an exception-aborting directive.

▲ C++ ▲

The **absent** and **contains** clauses accept a *directive-name* list that may match a construct that is encountered within the assumption scope. An encountered construct matches the directive name if it or (if it is a combined or composite construct) one of its leaf constructs has the same *directive-name* as one of the members of the list. The **absent** clause specifies that the program guarantees that no constructs that match a listed directive name are encountered in the assumption scope. The **contains** clause specifies that constructs that match the listed directive names are likely to be encountered in the assumption scope.

When the **holds** clause appears on an assumption directive, the program guarantees that the listed expression evaluates to *true* in the assumption scope. The effect of the clause does not include an observable evaluation of the expression.

Restrictions

The restrictions to *assumption* clauses are as follows:

- A *directive-name* list member must not specify a combined or composite directive.
- A *directive-name* list member must not specify a directive that is a declarative directive, an informational directive, or a metadirective.

Cross References

- **assume** directive, see [Section 8.3.3](#)
- **assumes** directive, see [Section 8.3.2](#)
- **begin assumes** directive, see [Section 8.3.4](#)

8.3.2 **assumes** Directive

Name: assumes Category: informational	Association: none Properties: pure
---	---

Clause groups

assumption

Semantics

The assumption scope of the **assumes** directive is the code executed and reached from the current compilation unit.

Restrictions

The restrictions to the *assumes* directive are as follows:

C
• The assumes directive may only appear at file scope.
C
C++
• The assumes directive may only appear at file or namespace scope.
C++
Fortran
• The assumes directive may only appear in the specification part of a module or subprogram, after all USE statements, IMPORT statements, and IMPLICIT statements.
Fortran

8.3.3 **assume** Directive

Name: assume Category: informational	Association: block Properties: pure
--	--

Clause groups

assumption

Semantics

The assumption scope of the **assume** directive is the code executed in the corresponding region or in any region that is nested in the corresponding region.

C / C++

8.3.4 begin assumes Directive

Name: <code>begin assumes</code>	Association: delimited (declaration-definition-seq)
Category: informational	Properties: <i>default</i>

Clause groups

assumption

Semantics

The assumption scope of the **begin assumes** directive is the code that is executed and reached from any of the declared functions in the delimited code region.

C / C++

8.4 nothing Directive

Name: <code>nothing</code>	Association: none
Category: utility	Properties: pure

Semantics

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

Cross References

- Metadirectives, see [Section 7.4](#)

8.5 error Directive

Name: <code>error</code>	Association: none
Category: utility	Properties: pure

Clauses

at, message, severity

Semantics

The **error** directive instructs the compiler or runtime to perform an error action. The error action displays an implementation-defined message. The **severity** clause determines whether the error action is abortive following the display of the message. If *sev-level* is **fatal** and *action-time* is **compilation**, the message is displayed and compilation of the current compilation unit is aborted. If *sev-level* is **fatal** and *action-time* is **execution**, the message is displayed and program execution is aborted.

Execution Model Events

The *runtime-error* event occurs when a thread encounters an **error** directive for which the **at** clause specifies **execution**.

Tool Callbacks

A thread dispatches a registered **ompt_callback_error** callback for each occurrence of a *runtime-error* event in the context of the encountering task. This callback has the type signature **ompt_callback_error_t**.

Restrictions

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

- The directive is pure only if *action-time* is **compilation**.

Cross References

- **at** clause, see [Section 8.1](#)
- **message** clause, see [Section 8.5.2](#)
- **severity** clause, see [Section 8.5.1](#)
- **ompt_callback_error_t**, see [Section 19.5.2.30](#)

8.5.1 severity Clause

Name: severity	Properties: unique
------------------------------	---------------------------

Arguments

Name	Type	Properties
<i>sev-level</i>	Keyword: fatal , warning	<i>default</i>

Directives

error

Semantics

The **severity** clause determines the action that the implementation performs. If *sev-level* is **warning**, the implementation takes no action besides displaying the message that is associated with the directive. If *sev-level* is **fatal**, the implementation performs the abortive action associated with the directive on which the clause appears. If no **severity** clause is specified then the effect is as if *sev-level* is **fatal**.

Cross References

- `error` directive, see [Section 8.5](#)

8.5.2 message Clause

Name: <code>message</code>	Properties: unique
-----------------------------------	---------------------------

Arguments

Name	Type	Properties
<i>msg-string</i>	expression of string type	<i>default</i>

Directives

`error`

Semantics

The `message` clause specifies that *msg-string* is included in the implementation-defined message that is associated with the directive on which the clause appears.

Restrictions

▼	C / C++	▼
•	If the <i>action-time</i> is compilation , <i>msg-string</i> must be a constant string literal.	▲
▲	C / C++	▲
▼	Fortran	▼
•	If the <i>action-time</i> is compilation , <i>msg-string</i> must be a constant character expression.	▲
▲	Fortran	▲

Cross References

- `error` directive, see [Section 8.5](#)

9 Loop-Transforming Constructs

A loop-transforming construct replaces itself, including its associated loop nest (see [Section 4.4.1](#)), with a structured block that may be another loop nest. If the replacement of a loop-transforming construct is another loop nest, that loop nest, possibly as part of an enclosing loop nest, may be associated with another loop-associated directive. A nested loop-transforming construct and any loop-transforming constructs that result from its **apply** clauses are replaced before any enclosing loop-transforming construct.

All generated loops have canonical loop nest form, unless otherwise specified. Loop iteration variables of generated loops are always private in the enclosing parallelism-generating construct.

Restrictions

The following restrictions apply to loop-transforming constructs:

- The replacement of a loop-transforming construct with its generated loops must result in a conforming OpenMP program.

Cross References

- Canonical Loop Nest Form, see [Section 4.4.1](#)

9.1 `tile` Construct

Name: <code>tile</code>	Association: loop
Category: executable	Properties: pure, loop-transforming

Clauses

[apply](#), [sizes](#)

Loop Modifiers for the `apply` Clause

<i>loop-modifier</i>	Number of Generated Loops	Description
grid	n	the grid loops g_1, \dots, g_n
intratile	n	the intra-tile loops t_1, \dots, t_n

Semantics

The **tile** construct is associated with n loops, where n is the number of items in the **sizes** clause, which consists of items s_1, \dots, s_n . Let ℓ_1, \dots, ℓ_n be the associated loops, from outermost to innermost, which the construct replaces with a loop nest that consists of $2n$ perfectly nested loops. Let $g_1, \dots, g_n, t_1, \dots, t_n$ be the generated loops, from outermost to innermost. The loops g_1, \dots, g_n are the *grid loops* and the loops t_1, \dots, t_n are the *intra-tile loops*.

Let Ω be the *logical iteration vector space* of the associated loops. For any $(\alpha_1, \dots, \alpha_n) \in \mathbb{N}^n$, define the set of iterations $\{(i_1, \dots, i_n) \in \Omega \mid \forall k \in \{1, \dots, n\} : s_k \alpha_k \leq i_k < s_k \alpha_k + s_k\}$ to be tile $T_{\alpha_1, \dots, \alpha_n}$ and $G = \{T_{\alpha_1, \dots, \alpha_n} \mid T_{\alpha_1, \dots, \alpha_n} \neq \emptyset\}$ to be the set of tiles with at least one iteration. Tiles that contain $\prod_{k=1}^n s_k$ iterations are complete tiles. Otherwise, they are partial tiles.

The grid loops iterate over all tiles $\{T_{\alpha_1, \dots, \alpha_n} \in G\}$ in lexicographic order with respect to their indices $(\alpha_1, \dots, \alpha_n)$ and the intra-tile loops iterate over the iterations in $T_{\alpha_1, \dots, \alpha_n}$ in the lexicographic order of the corresponding iteration vectors. An implementation may reorder the sequential execution of two iterations if at least one is from a partial tile and if their respective logical iteration vectors in *loop-nest* do not have a product order relation.

Restrictions

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

- The depth of the associated loop nest must be greater than or equal to n .
- All loops that are associated with the construct must be perfectly nested.
- No loop that is associated with the construct may be a non-rectangular loop.
- A grid loop and an intra-tile loop that are generated from the same **tile** construct must not be associated with the same loop-associated directive.

Cross References

- **apply** clause, see [Section 9.5](#)
- **sizes** clause, see [Section 9.1.1](#)

9.1.1 sizes Clause

Name: sizes	Properties: unique, required
---------------------------	-------------------------------------

Arguments

Name	Type	Properties
<i>size-list</i>	list of OpenMP integer expression type	constant, positive

Directives

tile

Semantics

The **sizes** clause specifies a list of n compile-time constant, positive OpenMP integer expressions. The list items are not required to be unique.

Cross References

- **tile** directive, see [Section 9.1](#)

9.2 unroll Construct

Name: <code>unroll</code>	Association: loop
Category: executable	Properties: pure, loop-transforming

Clauses

`apply`, `full`, `partial`

Clause set

Properties: exclusive	Members: <code>full</code> , <code>partial</code>
------------------------------	--

Loop Modifiers for the `apply` Clause

<i>loop-modifier</i>	Number of Generated Loops	Description
<code>unrolled</code> (<i>default</i>)	1	the grid loop g_1 of the tiling step

The default *loop-modifier* is `unrolled`.

Semantics

The `unroll` construct is associated with one loop, which is unrolled according to its specified clause. If no clauses are specified, if and how the loop is unrolled is implementation defined. The `unroll` construct results in a generated loop that has canonical loop nest form if and only if the `partial` clause is specified.

Restrictions

Restrictions to the `unroll` directive are as follows:

- The `apply` clause can only be specified if the `partial` clause is specified.

Cross References

- `apply` clause, see [Section 9.5](#)
- `full` clause, see [Section 9.2.1](#)
- `partial` clause, see [Section 9.2.2](#)

9.2.1 full Clause

Name: <code>full</code>	Properties: unique
--------------------------------	---------------------------

Directives

[unroll](#)

Semantics

The **full** clause specifies that the associated loop is *fully unrolled*. The construct is replaced by a structured block that only contains n instances of its loop body, one for each of the n logical iterations of the associated loop and in their logical iteration order.

Restrictions

Restrictions to the **full** clause are as follows:

- The iteration count of the associated loop must be a compile-time constant.

Cross References

- **unroll** directive, see [Section 9.2](#)

9.2.2 partial Clause

Name: <code>partial</code>	Properties: unique
-----------------------------------	---------------------------

Arguments

Name	Type	Properties
<i>unroll-factor</i>	expression of integer type	optional, constant, positive

Directives

[unroll](#)

Semantics

The **partial** clause specifies that the associated loop is first tiled with a tile size of *unroll-factor*. Then, the generated intra-tile loop is fully unrolled. If the **partial** clause is used without an *unroll-factor* argument then the unroll factor is a positive integer that is implementation defined.

Cross References

- **unroll** directive, see [Section 9.2](#)

9.3 reverse Construct

Name: <code>reverse</code> Category: executable	Association: loop Properties: pure, loop-transforming
--	--

Clauses

[apply](#)

Loop Modifiers for the `apply` Clause

<i>loop-modifier</i>	Number of Generated Loops	Description
reversed (<i>default</i>)	1	the reversed loop

Semantics

The **reverse** construct is associated with one loop, the outermost loop, where $0, 1, \dots, n-2, n-1$ are the logical iteration numbers of that loop. The construct transforms that loop into a loop in which iterations occur in the order $n-1, n-2, \dots, 1, 0$.

Cross References

- **apply** clause, see [Section 9.5](#)

9.4 interchange Construct

Name: <code>interchange</code>	Association: loop
Category: executable	Properties: pure, loop-transforming

Clauses

[apply](#), [permutation](#)

Loop Modifiers for the `apply` Clause

<i>loop-modifier</i>	Number of Generated Loops	Description
interchanged (<i>default</i>)	n	the generated loops, in the new order

Semantics

The **interchange** construct is associated with n loops, where s_1, \dots, s_n are the n items in the *permutation-list* argument of the **permutation** clause. Let ℓ_1, \dots, ℓ_n be the associated loops, from outermost to innermost. The original associated loops are replaced with the loops in the order $\ell_{s_1}, \dots, \ell_{s_n}$.

If the **permutation** clause is not specified, the effect is as if **permutation**(**2, 1**) was specified.

Restrictions

Restrictions to the **interchange** clause are as follows:

- The associated loop nest must be rectangular.
- The associated loop nest must be perfectly nested.

Cross References

- **apply** clause, see [Section 9.5](#)
- **permutation** clause, see [Section 9.4.1](#)

9.4.1 permutation Clause

Name: <code>permutation</code>	Properties: <code>unique</code>
---------------------------------------	--

Arguments

Name	Type	Properties
<i>permutation-list</i>	list of OpenMP integer expression type	constant, positive

Directives

[interchange](#)

Semantics

The **permutation** clause specifies a list of n compile-time constant, positive OpenMP integer expressions.

Restrictions

Restrictions to the **permutation** clause are as follows:

- Every integer from 1 to n must appear exactly once in *permutation-list*.
- n must be at least 2.

Cross References

- **interchange** directive, see [Section 9.4](#)

9.5 apply Clause

Name: <code>apply</code>	Properties: <code>default</code>
---------------------------------	---

Arguments

Name	Type	Properties
<i>applied-directives</i>	list of directive specification list item type	<i>default</i>

Modifiers

Name	Modifies	Type	Properties
<i>loop-modifier</i>	<i>applied-directives</i>	Keyword: grid , interchanged , intratile , reversed , unrolled	default

Directives

[tile](#), [unroll](#)

Semantics

The **apply** clause applies loop-transforming directives, specified by the *applied-directives* list, to the generated loops of a loop-transforming construct. The *loop-modifier* specifies where the applied directives are inserted with respect to the generated loops. An applied loop-transforming directive may also specify **apply** clauses.

The valid *loop-modifier* keywords, the default *loop-modifier* if it exists, the number of *applied-directives* list items, and the target of each *applied-directives* list item is defined by the loop-transforming construct to which it applies. The directive specified by the *i*-th item in the *applied-directives* list is applied to the *i*-th generated loop according to the *loop-modifier* keyword description, unless the *applied-directives* list item is **nothing** in which case the *i*-th generated loop is not transformed. If the *loop-modifier* is omitted and a default *loop-modifier* exists for the **apply** clause on the construct, the behavior is as if the default *loop-modifier* is specified.

The list items of the **apply** clause arguments are not required to be directive-wide unique.

Restrictions

Restrictions to the **apply** clause are as follows:

- A list item in an **apply** clause must be **nothing** or the *directive-specification* of a loop-transforming directive.
- A given *loop-modifier* keyword must not appear in more than one **apply** *clause-argument-specification* on the same construct.
- If a directive does not define a default *loop-modifier* keyword, the *loop-modifier* must not be omitted.

Cross References

- **metadirective** directive, see [Section 7.4.3](#)
- **tile** directive, see [Section 9.1](#)
- **unroll** directive, see [Section 9.2](#)

10 Parallelism Generation and Control

This chapter defines constructs for generating and controlling parallelism.

10.1 `omp_curr_progress_width` Identifier

The `omp_curr_progress_width` identifier is a context-specific OpenMP constant that is an OpenMP integer expression. It evaluates to the maximum size, in terms of hardware threads, of a progress unit that is available to threads that are executing tasks in the current contention group.

10.2 `parallel` Construct

Name: <code>parallel</code> Category: executable	Association: block Properties: parallelism-generating, cancellable, thread-limiting, context-matching
---	--

Clauses

`allocate`, `copyin`, `default`, `firstprivate`, `if`, `num_threads`, `private`, `proc_bind`, `reduction`, `safesync`, `shared`

Binding

The binding thread set for a `parallel` region is the encountering thread. The encountering thread becomes the primary thread of the new team.

Semantics

When a thread encounters a `parallel` construct, a team of threads is created to execute the `parallel` region (see [Section 10.2.1](#) for more information about how the number of threads in the team is determined, including the evaluation of the `if` and `num_threads` clauses). The thread that encountered the `parallel` construct becomes the primary thread of the new team, with a thread number of zero for the duration of the new `parallel` region. All threads in the new team, including the primary thread, execute the region. Once the team is created, the number of threads in the team remains constant for the duration of that `parallel` region.

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

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

7 The implementation may cause any thread to suspend execution of its implicit task at a task
8 scheduling point, and to switch to execution of any explicit task generated by any of the threads in
9 the team, before eventually resuming execution of the implicit task (for more details see
10 [Chapter 12](#)).

11 An implicit barrier occurs at the end of a **parallel** region. After the end of a **parallel** region,
12 only the primary thread of the team resumes execution of the enclosing task region.

13 If a thread in a team that is executing a **parallel** region encounters another **parallel**
14 directive, it creates a new team, according to the rules in [Section 10.2.1](#), and it becomes the primary
15 thread of that new team.

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

20 Execution Model Events

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

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

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

28 Events associated with implicit barriers occur at the end of a **parallel** region. [Section 15.3.2](#)
29 describes events associated with implicit barriers.

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

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

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

1 Tool Callbacks

2 A thread dispatches a registered `ompt_callback_parallel_begin` callback for each
3 occurrence of a *parallel-begin* event in that thread. The callback occurs in the task that encounters
4 the `parallel` construct. This callback has the type signature
5 `ompt_callback_parallel_begin_t`. In the dispatched callback,
6 `(flags & ompt_parallel_team)` evaluates to *true*.

7 A thread dispatches a registered `ompt_callback_implicit_task` callback with
8 `ompt_scope_begin` as its *endpoint* argument for each occurrence of an *implicit-task-begin*
9 event in that thread. Similarly, a thread dispatches a registered
10 `ompt_callback_implicit_task` callback with `ompt_scope_end` as its *endpoint*
11 argument for each occurrence of an *implicit-task-end* event in that thread. The callbacks occur in
12 the context of the implicit task and have type signature `ompt_callback_implicit_task_t`.
13 In the dispatched callback, `(flags & ompt_task_implicit)` evaluates to *true*.

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

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

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

24 Cross References

- 25 • `allocate` clause, see [Section 6.6](#)
- 26 • `copyin` clause, see [Section 5.7.1](#)
- 27 • `default` clause, see [Section 5.4.1](#)
- 28 • `firstprivate` clause, see [Section 5.4.4](#)
- 29 • `if` clause, see [Section 3.4](#)
- 30 • `num_threads` clause, see [Section 10.2.2](#)
- 31 • `private` clause, see [Section 5.4.3](#)
- 32 • `proc_bind` clause, see [Section 10.2.4](#)
- 33 • `reduction` clause, see [Section 5.5.8](#)
- 34 • `safesync` clause, see [Section 10.2.5](#)
- 35 • `shared` clause, see [Section 5.4.2](#)

- 1 • `omp_get_thread_num`, see [Section 18.2.4](#)
- 2 • `ompt_callback_implicit_task_t`, see [Section 19.5.2.11](#)
- 3 • `ompt_callback_parallel_begin_t`, see [Section 19.5.2.3](#)
- 4 • `ompt_callback_parallel_end_t`, see [Section 19.5.2.4](#)
- 5 • `ompt_callback_thread_begin_t`, see [Section 19.5.2.1](#)
- 6 • `ompt_callback_thread_end_t`, see [Section 19.5.2.2](#)
- 7 • `ompt_scope_endpoint_t`, see [Section 19.4.4.11](#)
- 8 • Determining the Number of Threads for a **parallel** Region, see [Section 10.2.1](#)

Algorithm 10.1 Determine Number of Threads

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

if an **if** clause exists **then let** *IfClauseValue* be the value of *if-expression*;

else let *IfClauseValue* = *true*;

if a **num_threads** clause exists **then let** *ThreadsRequested* be the value of *nthreads*;

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 (*active-levels-var* \geq *max-active-levels-var*) **then** number of threads = 1;

else if (*dyn-var* = *true*) **and** (*ThreadsRequested* \leq *ThreadsAvailable*)

then $1 \leq$ number of threads \leq *ThreadsRequested*;

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

then $1 \leq$ number of threads \leq *ThreadsAvailable*;

else if (*dyn-var* = *false*) **and** (*ThreadsRequested* \leq *ThreadsAvailable*)

then number of threads = *ThreadsRequested*;

else if (*dyn-var* = *false*) **and** (*ThreadsRequested* $>$ *ThreadsAvailable*)

then behavior is implementation defined

10.2.1 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 10.1.

Cross References

- `if` clause, see [Section 3.4](#)
- `num_threads` clause, see [Section 10.2.2](#)
- `parallel` directive, see [Section 10.2](#)
- `dyn-var` ICV, see [Table 2.1](#)
- `max-active-levels-var` ICV, see [Table 2.1](#)
- `nthreads-var` ICV, see [Table 2.1](#)
- `thread-limit-var` ICV, see [Table 2.1](#)

10.2.2 `num_threads` Clause

Name: <code>num_threads</code>	Properties: unique
---------------------------------------	---------------------------

Arguments

Name	Type	Properties
<code>nthreads</code>	expression of integer type	positive

Modifiers

Name	Modifies	Type	Properties
<code>prescriptiveness</code>	<code>nthreads</code>	Keyword: strict	default

Directives

[parallel](#)

Semantics

The `num_threads` clause specifies the desired number of threads to execute a `parallel` region. Algorithm 10.1 determines the number of threads that execute the `parallel` region. If *prescriptiveness* is specified as `strict` and an implementation determines that Algorithm 10.1 would always result in a number of threads other than the value of the *nthreads* expression then compile-time error termination may be performed. Otherwise, if *prescriptiveness* is specified as `strict` and Algorithm 10.1 would result in a number of threads other than the value of the *nthreads* expression then runtime error termination is performed.

Cross References

- `parallel` directive, see [Section 10.2](#)

10.2.3 Controlling OpenMP Thread Affinity

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

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

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

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

The purpose of the `spread` thread affinity policy is to create a sparse distribution for a team of T threads among the P places of the parent's place partition. A sparse distribution is achieved by first subdividing the parent partition into T subpartitions if $T \leq P$, or P subpartitions if $T > P$. Then

1 one thread ($T \leq P$) or a set of threads ($T > P$) is assigned to each subpartition. The
2 *place-partition-var* ICV of each implicit task is set to its subpartition. The subpartitioning is not
3 only a mechanism for achieving a sparse distribution, it also defines a subset of places for a thread
4 to use when creating a nested **parallel** region. The assignment of threads to places is as follows:

- 5 • $T \leq P$: The parent thread's place partition is split into T subpartitions, where each
6 subpartition contains $\lfloor P/T \rfloor$ or $\lceil P/T \rceil$ consecutive places. A single thread is assigned to
7 each subpartition. The primary thread executes on the place of the parent thread and is
8 assigned to the subpartition that includes that place. The thread with the next smallest thread
9 number is assigned to the first place in the next subpartition, and so on, with wrap around
10 with respect to the original place partition of the primary thread.
- 11 • $T > P$: The parent thread's place partition is split into P subpartitions, each consisting of a
12 single place. Each subpartition is assigned S_p threads with consecutive thread numbers,
13 where $\lfloor T/P \rfloor \leq S_p \leq \lceil T/P \rceil$. The first S_0 threads (including the primary thread) are
14 assigned to the subpartition that contains the place of the parent thread. The next S_1 threads
15 are assigned to the next subpartition, and so on, with wrap around with respect to the original
16 place partition of the primary thread. When P does not divide T evenly, the exact number of
17 threads in a particular subpartition is implementation defined.

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

20
21 **Note** – Wrap around is needed if the end of a place partition is reached before all thread
22 assignments are done. For example, wrap around may be needed in the case of **close** and $T \leq P$,
23 if the primary thread is assigned to a place other than the first place in the place partition. In this
24 case, thread 1 is assigned to the place after the place of the primary thread, thread 2 is assigned to
25 the place after that, and so on. The end of the place partition may be reached before all threads are
26 assigned. In this case, assignment of threads is resumed with the first place in the place partition.
27

28 Cross References

- 29 • **proc_bind** clause, see [Section 10.2.4](#)
- 30 • **parallel** directive, see [Section 10.2](#)
- 31 • *bind-var* ICV, see [Table 2.1](#)
- 32 • *place-partition-var* ICV, see [Table 2.1](#)

10.2.4 `proc_bind` Clause

Name: <code>proc_bind</code>	Properties: unique
------------------------------	--------------------

Arguments

Name	Type	Properties
<i>affinity-policy</i>	Keyword: close , primary , spread	<i>default</i>

Directives

[parallel](#)

Semantics

The `proc_bind` clause specifies the mapping of OpenMP threads to places within the current place partition, that is, within the places listed in the *place-partition-var* ICV for the implicit task of the encountering thread. The effect of the possible values for *affinity-policy* are described in [Section 10.2.3](#)

Cross References

- `parallel` directive, see [Section 10.2](#)
- Controlling OpenMP Thread Affinity, see [Section 10.2.3](#)

10.2.5 `safesync` Clause

Name: <code>safesync</code>	Properties: unique
-----------------------------	--------------------

Arguments

Name	Type	Properties
<i>width</i>	expression of integer type	positive, optional

Directives

[parallel](#)

Semantics

The `safesync` clause specifies that threads in the new team are partitioned, in thread number order, into *progress groups* of size *width*, except for the last progress group, which may contain less than *width* threads. If the *width* argument is not specified, the behavior is as if the *width* argument is one.

Among threads that are executing tasks in the same contention group in parallel, only threads that are in the same progress group execute in the same progress unit.

Cross References

- `parallel` directive, see [Section 10.2](#)

10.3 teams Construct

Name: <code>teams</code> Category: executable	Association: block Properties: parallelism-generating, thread-limiting, context-matching
--	---

Clauses

`allocate`, `default`, `firstprivate`, `if`, `num_teams`, `private`, `reduction`, `shared`, `thread_limit`

Binding

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

Semantics

When a thread encounters a `teams` construct, a league of teams is created. Each team is an initial team, and the initial thread in each team executes the `teams` region. The number of teams created is determined by evaluating the `if` and `num_teams` clauses. Once the teams are created, the number of initial teams remains constant for the duration of the `teams` region. Within a `teams` region, initial team numbers uniquely identify each initial team. Initial team numbers are consecutive whole numbers ranging from zero to one less than the number of initial teams.

When an `if` clause is present on a `teams` construct and the `if` clause expression evaluates to *false*, the number of created teams is one. The use of a variable in an `if` clause expression of a `teams` construct causes an implicit reference to the variable in all enclosing constructs. The `if` clause expression is evaluated in the context outside of the `teams` construct.

If a `thread_limit` clause is not present on the `teams` construct, but the construct is closely nested inside a `target` construct on which the `thread_limit` clause is specified, the behavior is as if that `thread_limit` clause is also specified for the `teams` construct.

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

The place list, given by the *place-partition-var* ICV of the encountering thread, is split into subpartitions in an implementation-defined manner, and each team is assigned to a subpartition by setting the *place-partition-var* of its initial thread to the subpartition.

The `teams` construct sets the *default-device-var* ICV for each initial thread to an implementation-defined value.

After the teams have completed execution of the `teams` region, the encountering task resumes execution of the enclosing task region.

Execution Model Events

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

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

If the **teams** 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 *initial-task-begin* event.

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

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

If a native thread is destroyed at the end of a **teams** 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 *teams-begin* event in that thread. The callback occurs in the task that encounters the **teams** construct. This callback has the type signature **ompt_callback_parallel_begin_t**. In the dispatched callback, *(flags & ompt_parallel_league)* 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 *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, *(flags & ompt_task_initial)* evaluates to *true*.

A thread dispatches a registered **ompt_callback_parallel_end** callback for each occurrence of a *teams-end* event in that thread. The callback occurs in the task that encounters the **teams** 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**.

A thread dispatches a registered **ompt_callback_thread_end** callback for the *native-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**.

Restrictions

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

- If a *reduction-modifier* is specified in a **reduction** clause that appears on the directive then the reduction modifier must be **default**.
- 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 not contain any statements, declarations or directives outside of the corresponding **teams** construct.
- **distribute** regions, including any **distribute** regions arising from composite constructs, **parallel** regions, including any **parallel** regions arising from combined constructs, **loop** regions, **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

- **allocate** clause, see [Section 6.6](#)
- **default** clause, see [Section 5.4.1](#)
- **firstprivate** clause, see [Section 5.4.4](#)
- **if** clause, see [Section 3.4](#)
- **num_teams** clause, see [Section 10.3.1](#)
- **private** clause, see [Section 5.4.3](#)
- **reduction** clause, see [Section 5.5.8](#)
- **shared** clause, see [Section 5.4.2](#)
- **thread_limit** clause, see [Section 13.3](#)
- **distribute** directive, see [Section 11.6](#)
- **parallel** directive, see [Section 10.2](#)
- **target** directive, see [Section 13.8](#)
- **omp_get_num_teams**, see [Section 18.4.1](#)
- **omp_get_team_num**, see [Section 18.4.2](#)
- **ompt_callback_implicit_task_t**, see [Section 19.5.2.11](#)
- **ompt_callback_parallel_begin_t**, see [Section 19.5.2.3](#)
- **ompt_callback_parallel_end_t**, see [Section 19.5.2.4](#)
- **ompt_callback_thread_begin_t**, see [Section 19.5.2.1](#)
- **ompt_callback_thread_end_t**, see [Section 19.5.2.2](#)

10.3.1 num_teams Clause

Name: num_teams	Properties: unique
------------------------	--------------------

Arguments

Name	Type	Properties
<i>upper-bound</i>	expression of integer type	positive

Modifiers

Name	Modifies	Type	Properties
<i>lower-bound</i>	<i>Generic</i>	OpenMP integer expression	positive, ultimate, unique

Directives

teams

Semantics

The **num_teams** clause specifies the bounds on the number of teams created by the construct on which it appears. *lower-bound* specifies the lower bound and *upper-bound* specifies the upper bound on the number of teams requested. If *lower-bound* is not specified, the effect is as if *lower-bound* is specified as equal to *upper-bound*. The number of teams created is implementation defined, but it will be greater than or equal to the lower bound and less than or equal to the upper bound.

If the **num_teams** clause is not specified on a construct then the effect is as if *upper-bound* was specified as follows. If the value of the *nteam*s-var ICV is greater than zero, the effect is as if *upper-bound* was specified to an implementation-defined value greater than zero but less than or equal to the value of the *nteam*s-var ICV. Otherwise, the effect is as if *upper-bound* was specified as an implementation defined value greater than or equal to one.

Restrictions

- *lower-bound* must be less than or equal to *upper-bound*.

Cross References

- **teams** directive, see [Section 10.3](#)

10.4 order Clause

Name: order	Properties: unique
--------------------	--------------------

Arguments

Name	Type	Properties
<i>ordering</i>	Keyword: concurrent	<i>default</i>

Modifiers

Name	Modifies	Type	Properties
<i>order-modifier</i>	<i>ordering</i>	Keyword: reproducible , unconstrained	default

Directives

distribute, **do**, **for**, **loop**, **simd**

Semantics

The **order** clause specifies an *ordering* of execution for the iterations of the associated loops of a loop-associated directive. If *ordering* is **concurrent**, different logical iterations of the associated loops may execute in any order, including in parallel, as if by the binding thread set of the region. The binding thread set may recruit or create additional threads to participate in the parallel execution of any logical iterations.

The *order-modifier* on the **order** clause affects the schedule specification for the purpose of determining its consistency with other schedules (see [Section 4.4.5](#)). If *order-modifier* is **reproducible**, the loop schedule for the construct on which the clause appears is reproducible, whereas if *order-modifier* is **unconstrained**, the loop schedule is not reproducible.

Restrictions

Restrictions to the **order** clause are as follows:

- The only constructs that may be encountered inside a region that corresponds to a construct with an **order** clause that specifies **concurrent** are the **loop** construct, the **parallel** construct, the **simd** construct, and combined constructs for which the first construct is a **parallel** construct.
- A region that corresponds to a construct with an **order** clause that specifies **concurrent** may not contain calls to procedures that contain OpenMP directives.
- A region that corresponds to a construct with an **order** clause that specifies **concurrent** may not contain OpenMP runtime API calls.
- If a threadprivate variable is referenced inside a region that corresponds to a construct with an **order** clause that specifies **concurrent**, the behavior is unspecified.

Cross References

- **distribute** directive, see [Section 11.6](#)
- **do** directive, see [Section 11.5.2](#)
- **for** directive, see [Section 11.5.1](#)
- **loop** directive, see [Section 11.7](#)
- **simd** directive, see [Section 10.5](#)

10.5 simd Construct

Name: <code>simd</code> Category: executable	Association: loop Properties: parallelism-generating, context-matching, simdizable, pure
---	---

Separating directives

`scan`

Clauses

`aligned`, `collapse`, `if`, `lastprivate`, `linear`, `nontemporal`, `order`, `private`, `reduction`, `safelen`, `simdlen`

Binding

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

Semantics

The `simd` construct enables the execution of multiple iterations of the associated loops concurrently by using SIMD instructions. At the beginning of each logical iteration, the loop iteration variable or the variable declared by *range-decl* of each associated loop has the value that it would have if the set of the associated loops was executed sequentially. The number of iterations that are executed concurrently at any given time is implementation defined. Each concurrent iteration will be executed by a different SIMD lane. Each set of concurrent iterations is a SIMD chunk. Lexical forward dependences in the iterations of the original loop must be preserved within each SIMD chunk, unless an `order` clause that specifies `concurrent` is present.

When an `if` clause is present and evaluates to *false*, the preferred number of iterations to be executed concurrently is one, regardless of whether a `simdlen` clause is specified.

Restrictions

Restrictions to the `simd` construct are as follows:

- If both `simdlen` and `safelen` clauses are specified, the value of the `simdlen length` must be less than or equal to the value of the `safelen length`.
- Only simdizable constructs can be encountered during execution of a `simd` region.
- If an `order` clause that specifies `concurrent` appears on a `simd` directive, the `safelen` clause may not also appear.

C / C++

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

C / C++

- No exception can be raised in the **simd** region.
- The only random access iterator types that are allowed for the associated loops are pointer types.

Cross References

- **aligned** clause, see [Section 5.11](#)
- **collapse** clause, see [Section 4.4.3](#)
- **if** clause, see [Section 3.4](#)
- **lastprivate** clause, see [Section 5.4.5](#)
- **linear** clause, see [Section 5.4.6](#)
- **nontemporal** clause, see [Section 10.5.1](#)
- **order** clause, see [Section 10.4](#)
- **private** clause, see [Section 5.4.3](#)
- **reduction** clause, see [Section 5.5.8](#)
- **safelen** clause, see [Section 10.5.2](#)
- **simdlen** clause, see [Section 10.5.3](#)
- **scan** directive, see [Section 5.6](#)

10.5.1 nontemporal Clause

Name: <code>nontemporal</code>	Properties: <i>default</i>
---------------------------------------	-----------------------------------

Arguments

Name	Type	Properties
<i>list</i>	list of variable list item type	<i>default</i>

Directives

simd

Semantics

The **nontemporal** clause specifies that accesses to the storage locations to which the list items refer have low temporal locality across the iterations in which those storage locations are accessed. The list items of the **nontemporal** clause may also appear as list items of data-environment attribute clauses.

Cross References

- `simd` directive, see [Section 10.5](#)

10.5.2 safelen Clause

Name: <code>safelen</code>	Properties: unique
----------------------------	--------------------

Arguments

Name	Type	Properties
<i>length</i>	expression of integer type	positive, constant

Directives

`simd`

Semantics

The `safelen` clause specifies that no two concurrent iterations within a SIMD chunk can have a distance in the logical iteration space that is greater than or equal to the value given in the clause.

Cross References

- `simd` directive, see [Section 10.5](#)

10.5.3 simdlen Clause

Name: <code>simdlen</code>	Properties: unique
----------------------------	--------------------

Arguments

Name	Type	Properties
<i>length</i>	expression of integer type	positive, constant

Directives

`declare simd`, `simd`

Semantics

When the `simdlen` clause appears on a `simd` construct, *length* is treated as a hint that specifies the preferred number of iterations to be executed concurrently. When the `simdlen` clause appears on a `declare simd` construct, if a SIMD version of the associated function is created, *length* corresponds to the number of concurrent arguments of the function.

Cross References

- `declare simd` directive, see [Section 7.7](#)
- `simd` directive, see [Section 10.5](#)

10.6 masked Construct

Name: <code>masked</code> Category: executable	Association: block Properties: thread-limiting
---	---

Clauses

`filter`

Binding

The binding thread set for a **masked** region is the current team. A **masked** region binds to the innermost enclosing parallel region.

Semantics

The **masked** construct specifies a structured block that is executed by a subset of the threads of the current team. The **filter** clause selects a subset of the threads of the team that executes the binding parallel region to execute the structured block of the masked region. Other threads in the team do not execute the associated structured block. No implied barrier occurs either on entry to or exit from the **masked** construct. The result of evaluating the *thread_num* parameter of the **filter** clause may vary across threads.

If more than one thread in the team executes the structured block of a **masked** region, the structured block must include any synchronization required to ensure that data races do not occur.

Execution Model Events

The *masked-begin* event occurs in any thread of a team that executes the **masked** region on entry to the region.

The *masked-end* event occurs in any thread of a team that executes the **masked** region on exit from the region.

Tool Callbacks

A thread dispatches a registered `ompt_callback_masked` callback with `ompt_scope_begin` as its *endpoint* argument for each occurrence of a *masked-begin* event in that thread. Similarly, a thread dispatches a registered `ompt_callback_masked` callback with `ompt_scope_end` as its *endpoint* argument for each occurrence of a *masked-end* event in that thread. These callbacks occur in the context of the task executed by the current thread and have the type signature `ompt_callback_masked_t`.

Cross References

- **filter** clause, see [Section 10.6.1](#)
- `ompt_callback_masked_t`, see [Section 19.5.2.12](#)
- `ompt_scope_endpoint_t`, see [Section 19.4.4.11](#)

10.6.1 filter Clause

Name: filter	Properties: unique
---------------------	--------------------

Arguments

Name	Type	Properties
<i>thread_num</i>	expression of integer type	<i>default</i>

Directives

masked

Semantics

If *thread_num* specifies the thread number of the current thread in the current team then the **filter** clause selects the current thread. If the **filter** clause is not specified, the effect is as if the clause is specified with *thread_num* equal to zero, so that the **filter** clause selects the primary thread. The use of a variable in a *thread_num* clause expression causes an implicit reference to the variable in all enclosing constructs.

Cross References

- **masked** directive, see [Section 10.6](#)

11 Work-Distribution Constructs

A work-distribution construct distributes the execution of the corresponding region among the threads in its binding thread set. Threads execute portions of the region in the context of the implicit tasks that each one is executing.

A work-distribution construct is *worksharing* if the binding thread set is a thread team. 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.

Restrictions

The following restrictions apply to work-distribution constructs:

- Each work-distribution region must be encountered by all threads in the binding thread set or by none at all unless cancellation has been requested for the innermost enclosing parallel region.
- The sequence of encountered work-distribution regions that have the same binding thread set must be the same for every thread in the binding thread set.
- The sequence of encountered worksharing regions and **barrier** regions that bind to the same thread team must be the same for every thread in the team.

11.1 single Construct

Name: <code>single</code> Category: executable	Association: block Properties: work-distribution, worksharing, thread-limiting
---	---

Clauses

`allocate`, `copyprivate`, `firstprivate`, `nowait`, `private`

Binding

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.

Semantics

The **single** construct specifies that the associated structured block is executed by only one of the threads in the team (not necessarily the primary thread), in the context of its implicit task. The method of choosing a thread to execute the structured block each time the team encounters the construct is implementation defined. An implicit barrier occurs at the end of a **single** region if the **nowait** clause is not specified.

Execution Model Events

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.

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

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.

Cross References

- **allocate** clause, see [Section 6.6](#)
- **copyprivate** clause, see [Section 5.7.2](#)
- **firstprivate** clause, see [Section 5.4.4](#)
- **nowait** clause, see [Section 15.6](#)
- **private** clause, see [Section 5.4.3](#)
- **ompt_callback_work_t**, see [Section 19.5.2.5](#)
- **ompt_scope_endpoint_t**, see [Section 19.4.4.11](#)
- **ompt_work_t**, see [Section 19.4.4.16](#)

11.2 scope Construct

Name: <code>scope</code> Category: executable	Association: block Properties: work-distribution, worksharing, thread-limiting
--	---

Clauses

`allocate`, `firstprivate`, `nowait`, `private`, `reduction`

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.

Semantics

The **scope** construct specifies that all threads in a team execute the associated structured block and any additionally specified OpenMP operations. An implicit barrier occurs at the end of a **scope** region if the **nowait** clause is not 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 *work_type* 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 *work_type* 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`.

Cross References

- `allocate` clause, see [Section 6.6](#)
- `firstprivate` clause, see [Section 5.4.4](#)
- `nowait` clause, see [Section 15.6](#)
- `private` clause, see [Section 5.4.3](#)
- `reduction` clause, see [Section 5.5.8](#)
- `ompt_callback_work_t`, see [Section 19.5.2.5](#)

- `ompt_scope_endpoint_t`, see [Section 19.4.4.11](#)
- `ompt_work_t`, see [Section 19.4.4.16](#)

11.3 sections Construct

Name: <code>sections</code> Category: executable	Association: block Properties: work-distribution, worksharing, thread-limiting, cancellable
---	--

Separating directives

`section`

Clauses

`allocate`, `firstprivate`, `lastprivate`, `nowait`, `private`, `reduction`

Binding

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

Semantics

The `sections` construct is a non-iterative worksharing construct that contains a structured block that consists of a set of structured block sequences that are to be distributed among and executed by the threads in a team. Each structured block sequence is executed by one of the threads in the team in the context of its implicit task. An implicit barrier occurs at the end of a `sections` region if the `nowait` clause is not specified.

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.

Execution Model Events

The *sections-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.

1 Tool Callbacks

2 A thread dispatches a registered `ompt_callback_work` callback with `ompt_scope_begin`
3 as its *endpoint* argument and `ompt_work_sections` as its *work_type* argument for each
4 occurrence of a *sections-begin* event in that thread. Similarly, a thread dispatches a registered
5 `ompt_callback_work` callback with `ompt_scope_end` as its *endpoint* argument and
6 `ompt_work_sections` as its *work_type* argument for each occurrence of a *sections-end* event
7 in that thread. The callbacks occur in the context of the implicit task. The callbacks have type
8 signature `ompt_callback_work_t`.

9 Cross References

- 10 • `allocate` clause, see [Section 6.6](#)
- 11 • `firstprivate` clause, see [Section 5.4.4](#)
- 12 • `lastprivate` clause, see [Section 5.4.5](#)
- 13 • `nowait` clause, see [Section 15.6](#)
- 14 • `private` clause, see [Section 5.4.3](#)
- 15 • `reduction` clause, see [Section 5.5.8](#)
- 16 • `section` directive, see [Section 11.3.1](#)
- 17 • `ompt_callback_dispatch_t`, see [Section 19.5.2.6](#)
- 18 • `ompt_callback_work_t`, see [Section 19.5.2.5](#)
- 19 • `ompt_scope_endpoint_t`, see [Section 19.4.4.11](#)
- 20 • `ompt_work_t`, see [Section 19.4.4.16](#)

21 11.3.1 section Directive

22 Name: <code>section</code> Category: subsidiary	Association: separating Properties: <i>default</i>
---	---

23 Separated directives

24 [sections](#)

25 Semantics

26 The `section` directive splits a structured block sequence that is associated with a `sections`
27 construct into two structured block sequences.

28 Execution Model Events

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

1 Tool Callbacks

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

5 Cross References

- 6 • `sections` directive, see [Section 11.3](#)

Fortran

7 11.4 workshare Construct

8 Name: <code>workshare</code>	Association: block
Category: executable	Properties: work-distribution, worksharing

9 Clauses

10 `nowait`

11 Binding

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

16 Semantics

17 The `workshare` construct divides the execution of the associated structured block into separate
18 units of work and causes the threads of the team to share the work such that each unit is executed
19 only once by one thread, in the context of its implicit task. An implicit barrier occurs at the end of a
20 `workshare` region if a `nowait` clause is not specified.

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

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

- 27 • For array expressions within each statement, including transformational array intrinsic
28 functions that compute scalar values from arrays:
 - 29 – Evaluation of each element of the array expression, including any references to
30 elemental functions, is a unit of work.
 - 31 – Evaluation of transformational array intrinsic functions may be freely subdivided into
32 any number of units of work.

- 1 • For array assignment statements, assignment of each element is a unit of work.
- 2 • For scalar assignment statements, each assignment operation is a unit of work.
- 3 • For **WHERE** statements or constructs, evaluation of the mask expression and the masked
- 4 assignments are each a unit of work.
- 5 • For **FORALL** statements or constructs, evaluation of the mask expression, expressions
- 6 occurring in the specification of the iteration space, and the masked assignments are each a
- 7 unit of work.
- 8 • For **atomic** constructs, **critical** constructs, and **parallel** constructs, the construct is
- 9 a unit of work. A new thread team executes the statements contained in a **parallel**
- 10 construct.
- 11 • If none of the rules above apply to a portion of a statement in the structured block, then that
- 12 portion is a unit of work.

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

16 How units of work are assigned to the threads that execute a **workshare** region is unspecified.

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

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

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

24 Execution Model Events

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

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

29 Tool Callbacks

30 A thread dispatches a registered **ompt_callback_work** callback with **ompt_scope_begin**
 31 as its *endpoint* argument and **ompt_work_workshare** as its *work_type* argument for each
 32 occurrence of a *workshare-begin* event in that thread. Similarly, a thread dispatches a registered
 33 **ompt_callback_work** callback with **ompt_scope_end** as its *endpoint* argument and
 34 **ompt_work_workshare** as its *work_type* argument for each occurrence of a *workshare-end*
 35 event in that thread. The callbacks occur in the context of the implicit task. The callbacks have type
 36 signature **ompt_callback_work_t**.

Restrictions

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

- The only OpenMP constructs that may be closely nested inside a **workshare** construct are the **atomic**, **critical**, and **parallel** constructs.
- Base language statements that are encountered inside a **workshare** construct but that are not enclosed within a **parallel** or **atomic** construct that is nested inside the **workshare** construct must consist of only the following:
 - array assignments;
 - scalar assignments;
 - **FORALL** statements;
 - **FORALL** constructs;
 - **WHERE** statements;
 - **WHERE** constructs; and
 - **BLOCK** constructs that are strictly structured blocks associated with OpenMP directives.
- All array assignments, scalar assignments, and masked array assignments that are encountered inside a **workshare** construct but are not nested inside a **parallel** construct that is nested inside the **workshare** construct must be intrinsic assignments.
- The construct must not contain any user-defined function calls unless either the function is pure and elemental or the function call is contained inside a **parallel** construct that is nested inside the **workshare** construct.

Cross References

- **nowait** clause, see [Section 15.6](#)
- **atomic** directive, see [Section 15.8.5](#)
- **critical** directive, see [Section 15.2](#)
- **parallel** directive, see [Section 10.2](#)
- **ompt_callback_work_t**, see [Section 19.5.2.5](#)
- **ompt_scope_endpoint_t**, see [Section 19.4.4.11](#)
- **ompt_work_t**, see [Section 19.4.4.16](#)

Fortran

11.5 Worksharing-Loop Constructs

Binding

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

Semantics

The worksharing-loop construct is a worksharing construct that 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. Each thread executes its assigned chunks in the context of its implicit task. The iterations of a given chunk are executed in sequential order.

If specified, the **schedule** clause determines the schedule of the logical iterations associated with the construct. That is, it determines the division of iterations into chunks and how those chunks are assigned to the threads. If the **schedule** clause is not specified then the schedule is implementation defined.

At the beginning of each logical iteration, the loop iteration variable or the variable declared by *range-decl* of each associated loop has the value that it would have if the set of the associated loops was executed sequentially.

The schedule is reproducible if one of the following conditions is true:

- The **order** clause is specified with the **reproducible** *order-modifier*; or
- The **schedule** clause is specified with **static** as the *kind* argument but not the **simd** *ordering-modifier* and the **order** clause is not specified with the **unconstrained** *order-modifier*.

Programs can only depend on which thread executes a particular iteration if the schedule is reproducible. Schedule reproducibility also determines the consistency with the execution of constructs with the same schedule.

Execution Model Events

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

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

The *ws-loop-iteration-begin* event occurs at the beginning of each iteration of a worksharing-loop region. The *ws-loop-chunk-begin* event occurs for each scheduled chunk of a worksharing-loop region before the implicit task executes any of the associated iterations.

Tool Callbacks

A thread dispatches a registered `ompt_callback_work` callback with `ompt_scope_begin` as its *endpoint* argument for each occurrence of a *ws-loop-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 *ws-loop-end* event in that thread. The callbacks occur in the context of the implicit task. The callbacks have type signature `ompt_callback_work_t` and the *work_type* argument indicates the schedule as shown in Table 11.1.

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

TABLE 11.1: `ompt_callback_work` Callback Work Types for Worksharing-Loop

Value of <i>work_type</i>	If determined schedule is
<code>ompt_work_loop</code>	unknown at runtime
<code>ompt_work_loop_static</code>	<code>static</code>
<code>ompt_work_loop_dynamic</code>	<code>dynamic</code>
<code>ompt_work_loop_guided</code>	<code>guided</code>
<code>ompt_work_loop_other</code>	implementation specific

Restrictions

Restrictions to the worksharing-loop construct are as follows:

- The logical iteration space of the loops associated with the worksharing-loop construct must be the same for all threads in the team.
- The value of the *run-sched-var* ICV must be the same for all threads in the team.

Cross References

- `OMP_SCHEDULE`, see [Section 21.2.1](#)
- `nowait` clause, see [Section 15.6](#)
- `order` clause, see [Section 10.4](#)
- `schedule` clause, see [Section 11.5.3](#)
- `do` directive, see [Section 11.5.2](#)
- `for` directive, see [Section 11.5.1](#)
- Consistent Loop Schedules, see [Section 4.4.5](#)
- `ompt_callback_work_t`, see [Section 19.5.2.5](#)
- `ompt_scope_endpoint_t`, see [Section 19.4.4.11](#)
- `ompt_work_t`, see [Section 19.4.4.16](#)

11.5.1 `for` Construct

Name: <code>for</code> Category: executable	Association: loop Properties: work-distribution, workshar- ing, worksharing-loop, cancellable, context- matching
--	---

Separating directives

`scan`

Clauses

`allocate`, `collapse`, `firstprivate`, `lastprivate`, `linear`, `nowait`, `order`, `ordered`, `private`, `reduction`, `schedule`

Semantics

The `for` construct is a worksharing-loop construct.

Cross References

- `allocate` clause, see [Section 6.6](#)
- `collapse` clause, see [Section 4.4.3](#)
- `firstprivate` clause, see [Section 5.4.4](#)
- `lastprivate` clause, see [Section 5.4.5](#)
- `linear` clause, see [Section 5.4.6](#)
- `nowait` clause, see [Section 15.6](#)
- `order` clause, see [Section 10.4](#)
- `ordered` clause, see [Section 4.4.4](#)
- `private` clause, see [Section 5.4.3](#)
- `reduction` clause, see [Section 5.5.8](#)
- `schedule` clause, see [Section 11.5.3](#)
- `scan` directive, see [Section 5.6](#)
- Worksharing-Loop Constructs, see [Section 11.5](#)

11.5.2 do Construct

Name: do Category: executable	Association: loop Properties: work-distribution, workshar- ing, worksharing-loop, cancellable, context- matching
--	---

Separating directives

scan

Clauses

allocate, **collapse**, **firstprivate**, **lastprivate**, **linear**, **nowait**, **order**, **ordered**, **private**, **reduction**, **schedule**

Semantics

The **do** construct is a worksharing-loop construct.

Cross References

- **allocate** clause, see [Section 6.6](#)
- **collapse** clause, see [Section 4.4.3](#)
- **firstprivate** clause, see [Section 5.4.4](#)
- **lastprivate** clause, see [Section 5.4.5](#)
- **linear** clause, see [Section 5.4.6](#)
- **nowait** clause, see [Section 15.6](#)
- **order** clause, see [Section 10.4](#)
- **ordered** clause, see [Section 4.4.4](#)
- **private** clause, see [Section 5.4.3](#)
- **reduction** clause, see [Section 5.5.8](#)
- **schedule** clause, see [Section 11.5.3](#)
- **scan** directive, see [Section 5.6](#)
- Worksharing-Loop Constructs, see [Section 11.5](#)

11.5.3 schedule Clause

Name: schedule	Properties: unique
-----------------------	--------------------

Arguments

Name	Type	Properties
<i>kind</i>	Keyword: auto , dynamic , guided , runtime , static	<i>default</i>
<i>chunk_size</i>	expression of integer type	ultimate, optional, positive, region-invariant

Modifiers

Name	Modifies	Type	Properties
<i>ordering-modifier</i>	<i>kind</i>	Keyword: monotonic , nonmonotonic	unique
<i>chunk-modifier</i>	<i>kind</i>	Keyword: simd	unique

Directives

do, **for**

Semantics

The **schedule** clause specifies how iterations of associated loops of a worksharing-loop construct are divided into contiguous non-empty subsets, called chunks, and how these chunks are distributed among threads of the team. The *chunk_size* expression is evaluated using the original list items of any variables that are made private in the worksharing-loop construct. Whether, in what order, or how many times, any side effects of the evaluation of this expression occur is unspecified. 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.

If the *kind* argument is **static**, iterations are divided into chunks of size *chunk_size*, 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 *chunk_size* iterations, except for the chunk that contains the sequentially last iteration, which may have fewer iterations. If *chunk_size* is not specified, the logical iteration space is divided into chunks that are approximately equal in size, and at most one chunk is distributed to each thread.

If the *kind* argument is **dynamic**, each thread executes a chunk, then requests another chunk, until no chunks remain to be assigned. Each chunk contains *chunk_size* iterations, except for the chunk that contains the sequentially last iteration, which may have fewer iterations. If *chunk_size* is not specified, it defaults to 1.

If the *kind* argument is **guided**, each thread executes a chunk, then requests another chunk, until no chunks remain to be assigned. For a *chunk_size* of 1, the size of each chunk is proportional to the number of unassigned iterations divided by the number of threads in the team, decreasing to 1. For a *chunk_size* with value $k > 1$, the size of each chunk is determined in the same way, with the

1 restriction that the chunks do not contain fewer than k iterations (except for the chunk that contains
2 the sequentially last iteration, which may have fewer than k iterations). If *chunk_size* is not
3 specified, it defaults to 1.

4 If the *kind* argument is **auto**, the decision regarding scheduling is implementation defined.

5 If the *kind* argument is **runtime**, the decision regarding scheduling is deferred until runtime, and
6 the behavior is as if the clause specifies *kind*, *chunk-size* and *ordering-modifier* as set in the
7 *run-sched-var* ICV. If the **schedule** clause explicitly specifies any modifiers then they override
8 any corresponding modifiers that are specified in the *run-sched-var* ICV.

9 If the **simd chunk-modifier** is specified and the loop is associated with a SIMD construct,
10 $new_chunk_size = \lceil chunk_size / simd_width \rceil * simd_width$ is the *chunk_size* for all chunks
11 except the first and last chunks, where *simd_width* is an implementation-defined value. The first
12 chunk will have at least *new_chunk_size* iterations except if it is also the last chunk. The last chunk
13 may have fewer iterations than *new_chunk_size*. If the **simd** modifier is specified and the loop is
14 not associated with a SIMD construct, the modifier is ignored.

15
▼
16 **Note** – For a team of p threads and a loop of n iterations, let $\lceil n/p \rceil$ be the integer q that satisfies
17 $n = p * q - r$, with $0 \leq r < p$. One compliant implementation of the **static** schedule (with no
18 specified *chunk_size*) would behave as though *chunk_size* had been specified with value q . Another
19 compliant implementation would assign q iterations to the first $p - r$ threads, and $q - 1$ iterations to
20 the remaining r threads. This illustrates why a conforming program must not rely on the details of a
21 particular implementation.

22 A compliant implementation of the **guided** schedule with a *chunk_size* value of k would assign
23 $q = \lceil n/p \rceil$ iterations to the first available thread and set n to the larger of $n - q$ and $p * k$. It would
24 then repeat this process until q is greater than or equal to the number of remaining iterations, at
25 which time the remaining iterations form the final chunk. Another compliant implementation could
26 use the same method, except with $q = \lceil n/(2p) \rceil$, and set n to the larger of $n - q$ and $2 * p * k$.
27 ▲

28 If the **monotonic ordering-modifier** is specified then each thread executes the chunks that it is
29 assigned in increasing logical iteration order. When the **nonmonotonic ordering-modifier** is
30 specified then chunks may be assigned to threads in any order and the behavior of an application
31 that depends on any execution order of the chunks is unspecified. If an *ordering-modifier* is not
32 specified, the effect is as if the **monotonic** modifier is specified if the *kind* argument is **static**
33 or an **ordered** clause is specified on the construct; otherwise, the effect is as if the
34 **nonmonotonic** modifier is specified.

35 Restrictions

36 Restrictions to the **schedule** clause are as follows:

- 37
- The **schedule** clause cannot be specified if any of the associated loops are non-rectangular.
 - The value of the *chunk_size* expression must be the same for all threads in the team.
- 38

- If **runtime** or **auto** is specified for *kind*, *chunk_size* must not be specified.
- The **nonmonotonic ordering-modifier** cannot be specified if an **ordered** clause is specified on the same construct.

Cross References

- **ordered** clause, see [Section 4.4.4](#)
- **do** directive, see [Section 11.5.2](#)
- **for** directive, see [Section 11.5.1](#)
- *run-sched-var* ICV, see [Table 2.1](#)

11.6 distribute Construct

Name: <code>distribute</code>	Association: loop
Category: executable	Properties: work-distribution

Clauses

[allocate](#), [collapse](#), [dist_schedule](#), [firstprivate](#), [lastprivate](#), [order](#), [private](#)

Binding

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

Semantics

The **distribute** construct specifies that the iterations of one or more loops will be executed by the initial teams in the context of their implicit tasks. The iterations are distributed across the initial threads of all initial teams that execute the **teams** region to which the **distribute** region binds. No implicit barrier occurs at the end of a **distribute** region. To avoid data races the original list items that are modified due to **lastprivate** clauses should not be accessed between the end of the **distribute** construct and the end of the **teams** region to which the **distribute** binds.

If the **dist_schedule** clause is not specified, the schedule is implementation defined.

At the beginning of each logical iteration, the loop iteration variable or the variable declared by *range-decl* of each associated loop has the value that it would have if the set of the associated loops was executed sequentially.

The schedule is reproducible if one of the following conditions is true:

- The **order** clause is specified with the **reproducible** modifier; or
- The **dist_schedule** clause is specified with **static** as the *kind* parameter and the **order** clause is not specified with the **unconstrained order-modifier**.

1 Programs can only depend on which team executes a particular iteration if the schedule is
2 reproducible. Schedule reproducibility also determines the consistency with the execution of
3 constructs with the same schedule.

4 **Execution Model Events**

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

7 The *distribute-end* event occurs after an initial task finishes execution of a **distribute** region
8 but before it resumes execution of the enclosing context.

9 The *distribute-chunk-begin* event occurs for each scheduled chunk of a **distribute** region
10 before execution of any associated iteration.

11 **Tool Callbacks**

12 A thread dispatches a registered **ompt_callback_work** callback with **ompt_scope_begin**
13 as its *endpoint* argument and **ompt_work_distribute** as its *work_type* argument for each
14 occurrence of a *distribute-begin* event in that thread. Similarly, a thread dispatches a registered
15 **ompt_callback_work** callback with **ompt_scope_end** as its *endpoint* argument and
16 **ompt_work_distribute** as its *work_type* argument for each occurrence of a *distribute-end*
17 event in that thread. The callbacks occur in the context of the implicit task. The callbacks have type
18 signature **ompt_callback_work_t**.

19 A thread dispatches a registered **ompt_callback_dispatch** callback for each occurrence of a
20 *distribute-chunk-begin* event in that thread. The callback occurs in the context of the initial task.
21 The callback has type signature **ompt_callback_dispatch_t**.

22 **Restrictions**

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

- 24 • The logical iteration space of the loops associated with the **distribute** construct must be
25 the same for all teams in the league.
- 26 • The region that corresponds to the **distribute** construct must be strictly nested inside a
27 **teams** region.
- 28 • A list item may appear in a **firstprivate** or **lastprivate** clause, but not in both.
- 29 • The **conditional** *lastprivate-modifier* must not be specified.

30 **Cross References**

- 31 • **allocate** clause, see [Section 6.6](#)
- 32 • **collapse** clause, see [Section 4.4.3](#)
- 33 • **dist_schedule** clause, see [Section 11.6.1](#)
- 34 • **firstprivate** clause, see [Section 5.4.4](#)
- 35 • **lastprivate** clause, see [Section 5.4.5](#)

- **order** clause, see [Section 10.4](#)
- **private** clause, see [Section 5.4.3](#)
- **teams** directive, see [Section 10.3](#)
- Consistent Loop Schedules, see [Section 4.4.5](#)
- **ompt_callback_work_t**, see [Section 19.5.2.5](#)
- **ompt_work_t**, see [Section 19.4.4.16](#)

11.6.1 **dist_schedule** Clause

Name: <code>dist_schedule</code>	Properties: unique
---	---------------------------

Arguments

Name	Type	Properties
<i>kind</i>	Keyword: static	<i>default</i>
<i>chunk_size</i>	expression of integer type	ultimate, optional, positive, region-invariant

Directives

[distribute](#)

Semantics

The **dist_schedule** clause specifies how iterations of associated loops of a **distribute** construct are divided into contiguous non-empty subsets, called chunks, and how these chunks are distributed among the teams of the league. if *chunk_size* is not specified, the iteration space is divided into chunks that are approximately equal in size, and at most one chunk is distributed to each initial team of the league.

If the *chunk_size* argument is specified, iterations are divided into chunks of size *chunk_size*. The *chunk_size* expression is evaluated using the original list items of any variables that are made private in the **distribute** construct. Whether, in what order, or how many times, any side effects of the evaluation of this expression occur is unspecified. The use of a variable in a **dist_schedule** clause expression of a **distribute** construct causes an implicit reference to the variable in all enclosing constructs. These chunks are assigned to the initial teams of the league in a round-robin fashion in the order of the initial team number.

Restrictions

Restrictions to the **dist_schedule** clause are as follows:

- The value of the *chunk_size* expression must be the same for all teams in the league.
- The **dist_schedule** clause cannot be specified if any of the associated loops are non-rectangular.

Cross References

- **distribute** directive, see [Section 11.6](#)

11.7 loop Construct

Name: <code>loop</code> Category: executable	Association: loop Properties: work-distribution, worksharing, simdizable
---	---

Clauses

`bind`, `collapse`, `lastprivate`, `order`, `private`, `reduction`

Binding

The `bind` clause determines the binding region, which determines the binding thread set.

Semantics

A `loop` construct specifies that the logical iterations of the associated loops execute in the context of the binding thread set, in an order specified by the `order` clause. If the `order` clause is not specified, the behavior is as if the `order` clause is present and specifies the **concurrent ordering**. The logical iterations are executed as if by the binding thread set, once per instance of the `loop` region that is encountered by the binding thread set.

At the beginning of each logical iteration, the loop iteration variable or the variable declared by *range-decl* of each associated loop has the value that it would have if the set of the associated loops was executed sequentially.

The loop schedule for a `loop` construct is reproducible unless the `order` clause is present with the **unconstrained order-modifier**.

If the `loop` region binds to a `teams` region, the threads in the binding thread set may continue execution after the `loop` region without waiting for all logical iterations of the associated loops to complete. The iterations are guaranteed to complete before the end of the `teams` region. If the `loop` region does not bind to a `teams` region, all logical iterations of the associated loops must complete before the encountering threads continue execution after the `loop` region.

While a `loop` construct is always a work-distribution construct, it is a worksharing construct if and only if its binding region is the innermost enclosing parallel region.

Restrictions

Restrictions to the `loop` construct are as follows:

- A list item may not appear in a **lastprivate** clause unless it is the loop iteration variable of a loop that is associated with the construct.
- If a *reduction-modifier* is specified in a **reduction** clause that appears on the directive then the reduction modifier must be **default**.

- If a **loop** construct is not nested inside another OpenMP construct then the **bind** clause must be present.
- If a **loop** region binds to a **teams** or parallel region, it must be encountered by all threads in the binding thread set or by none of them.

Cross References

- **bind** clause, see [Section 11.7.1](#)
- **collapse** clause, see [Section 4.4.3](#)
- **lastprivate** clause, see [Section 5.4.5](#)
- **order** clause, see [Section 10.4](#)
- **private** clause, see [Section 5.4.3](#)
- **reduction** clause, see [Section 5.5.8](#)
- **teams** directive, see [Section 10.3](#)
- Consistent Loop Schedules, see [Section 4.4.5](#)

11.7.1 bind Clause

Name: <code>bind</code>	Properties: unique
--------------------------------	---------------------------

Arguments

Name	Type	Properties
<i>binding</i>	Keyword: parallel , teams , thread	<i>default</i>

Directives

[loop](#)

Semantics

The **bind** clause specifies the binding region of the construct on which it appears. Specifically, if *binding* is **teams** and an innermost enclosing **teams** region exists then the binding region is that **teams** region; if *binding* is **parallel** then the binding region is the innermost enclosing parallel region, which may be an implicit parallel region; and if *binding* is **thread** then the binding region is not defined. If the **bind** clause is not specified on a construct for which it may be specified and the construct is closely nested inside a **teams** or **parallel** construct, the effect is as if *binding* is **teams** or **parallel**. If none of those conditions hold, the binding region is not defined.

The specified binding region determines the binding thread set. Specifically, if the binding region is a **teams** region, then the binding thread set is the set of initial threads that are executing that region while if the binding region is a parallel region, then the binding thread set is the team of threads that are executing that region. If the binding region is not defined, then the binding thread set is the encountering thread.

1
2
3
4
5
6
7
8
9
10
11
12
13

Restrictions

Restrictions to the **bind** clause are as follows:

- If **teams** is specified as *binding* then the corresponding **loop** region must be strictly nested inside a **teams** region.
- If **teams** is specified as *binding* and the corresponding **loop** region executes on a non-host device then the behavior of a **reduction** clause that appears on the corresponding **loop** construct is unspecified if the construct is not nested inside a **teams** construct.
- If **parallel** is specified as *binding*, the behavior is unspecified if the corresponding **loop** region is closely nested inside a **simd** region.

Cross References

- **loop** directive, see [Section 11.7](#)
- **parallel** construct, see [Section 10.2](#)
- **teams** construct, see [Section 10.3](#).

12 Tasking Constructs

This chapter defines directives and concepts related to explicit tasks.

12.1 `untied` Clause

Name: <code>untied</code>	Properties: unique, inarguable
----------------------------------	---------------------------------------

Directives

`task`, `taskloop`

Semantics

The **`untied`** clause specifies that tasks generated by the construct on which it appears are untied, which means that any thread in the team can resume the **`task`** region after a suspension. If the **`untied`** clause is not specified on a construct on which it may appear, generated tasks are tied; if a tied task is suspended, its **`task`** region can only be resumed by the thread that started its execution. If a generated task is a final or an included task, the **`untied`** clause is ignored and the task is tied.

Cross References

- **`task`** directive, see [Section 12.5](#)
- **`taskloop`** directive, see [Section 12.6](#)

12.2 `mergeable` Clause

Name: <code>mergeable</code>	Properties: unique, inarguable
-------------------------------------	---------------------------------------

Directives

`task`, `taskloop`

Semantics

The **`mergeable`** clause specifies that tasks generated by the construct on which it appears are mergeable tasks.

Cross References

- **`task`** directive, see [Section 12.5](#)
- **`taskloop`** directive, see [Section 12.6](#)

12.3 final Clause

Name: final	Properties: unique
--------------------	--------------------

Arguments

Name	Type	Properties
<i>finalize</i>	expression of logical type	<i>default</i>

Directives

[task](#), [taskloop](#)

Semantics

The **final** clause specifies that tasks generated by the construct on which it appears are final tasks if the *finalize* expression evaluates to *true*. All **task** constructs that are encountered during execution of a final task generate final and included tasks. The use of a variable in a *finalize* expression causes an implicit reference to the variable in all enclosing constructs. The *finalize* expression is evaluated in the context outside of the construct on which the clause appears,

Cross References

- **task** directive, see [Section 12.5](#)
- **taskloop** directive, see [Section 12.6](#)

12.4 priority Clause

Name: priority	Properties: unique
-----------------------	--------------------

Arguments

Name	Type	Properties
<i>priority-value</i>	expression of integer type	constant, non-negative

Directives

[task](#), [taskloop](#)

Semantics

The **priority** clause specifies a hint for the task execution order of tasks generated by the construct on which it appears in the *priority-value* argument. Among all tasks ready to be executed, higher priority tasks (those with a higher numerical *priority-value*) are recommended to execute before lower priority ones. The default *priority-value* when no **priority** clause is specified is zero (the lowest priority). If a specified *priority-value* is higher than the *max-task-priority-var* ICV then the implementation will use the value of that ICV. A program that relies on the task execution order being determined by the *priority-value* may have unspecified behavior.

Cross References

- **task** directive, see [Section 12.5](#)
- **taskloop** directive, see [Section 12.6](#)
- *max-task-priority-var* ICV, see [Table 2.1](#)

12.5 task Construct

Name: task Category: executable	Association: block Properties: parallelism-generating, thread-limiting, task-generating
---	--

Clauses

[affinity](#), [allocate](#), [default](#), [depend](#), [detach](#), [final](#), [firstprivate](#), [if](#), [in_reduction](#), [mergeable](#), [priority](#), [private](#), [shared](#), [untied](#)

Clause set

Properties: exclusive	Members: detach , mergeable
------------------------------	--

Binding

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

Semantics

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 and clauses. If a **task** construct is encountered during execution of an outer task, the generated **task** region that corresponds to this construct is not a part of the outer task region unless the generated task is an included task.

A detachable task is completed when the execution of its associated structured block is completed and the *allow-completion* event is fulfilled. 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.

A thread that encounters a task scheduling point within the **task** region may temporarily suspend the **task** region.

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

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

4 When an **if** clause is present on a **task** construct and the **if** clause expression evaluates to *false*,
5 an undeferred task is generated, and the encountering thread must suspend the current task region,
6 for which execution cannot be resumed until execution of the structured block that is associated
7 with the generated task is completed. The use of a variable in an **if** clause expression of a **task**
8 construct causes an implicit reference to the variable in all enclosing constructs. The **if** clause
9 expression is evaluated in the context outside of the **task** construct.

10 Execution Model Events

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

13 Tool Callbacks

14 A thread dispatches a registered **ompt_callback_task_create** callback for each occurrence
15 of a *task-create* event in the context of the encountering task. This callback has the type signature
16 **ompt_callback_task_create_t** and the *flags* argument indicates the task types shown in
17 Table 12.1.

TABLE 12.1: **ompt_callback_task_create** Callback Flags Evaluation

Operation	Evaluates to true
<i>(flags & ompt_task_explicit)</i>	Always in the dispatched callback
<i>(flags & ompt_task_undeferred)</i>	If the task is an undeferred task
<i>(flags & ompt_task_final)</i>	If the task is a final task
<i>(flags & ompt_task_untied)</i>	If the task is an untied task
<i>(flags & ompt_task_mergeable)</i>	If the task is a mergeable task
<i>(flags & ompt_task_merged)</i>	If the task is a merged task

18 Cross References

- 19 • **affinity** clause, see [Section 12.5.1](#)
- 20 • **allocate** clause, see [Section 6.6](#)
- 21 • **default** clause, see [Section 5.4.1](#)
- 22 • **depend** clause, see [Section 15.9.5](#)
- 23 • **detach** clause, see [Section 12.5.2](#)
- 24 • **final** clause, see [Section 12.3](#)

- **firstprivate** clause, see [Section 5.4.4](#)
- **if** clause, see [Section 3.4](#)
- **in_reduction** clause, see [Section 5.5.10](#)
- **mergeable** clause, see [Section 12.2](#)
- **priority** clause, see [Section 12.4](#)
- **private** clause, see [Section 5.4.3](#)
- **shared** clause, see [Section 5.4.2](#)
- **untied** clause, see [Section 12.1](#)
- Task Scheduling, see [Section 12.9](#)
- **omp_fulfill_event**, see [Section 18.11.1](#)
- **ompt_callback_task_create_t**, see [Section 19.5.2.7](#)

12.5.1 affinity Clause

Name: <code>affinity</code>	Properties: unique
------------------------------------	---------------------------

Arguments

Name	Type	Properties
<i>locator-list</i>	list of locator list item type	<i>default</i>

Modifiers

Name	Modifies	Type	Properties
<i>iterator</i>	<i>locator-list</i>	Complex, name: iterator Arguments: <i>iterator-specifier</i> OpenMP expression (repeatable)	unique

Directives

task

Semantics

The **affinity** clause specifies a hint to indicate data affinity of tasks generated by the construct on which it appears. The hint recommends to execute generated tasks close to the location of the original list items. A program that relies on the task execution location being determined by this list may have unspecified behavior.

The list items that appear in the **affinity** clause may also appear in data-environment clauses. The list items may reference any *iterators-identifier* that is defined in the same clause and may include array sections.

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

Cross References

- **task** directive, see [Section 12.5](#)
- **iterator** modifier, see [Section 3.2.6](#)

12.5.2 detach Clause

Name: <code>detach</code>	Properties: unique
----------------------------------	---------------------------

Arguments

Name	Type	Properties
<i>event-handle</i>	variable of <code>event_handle</code> type	<i>default</i>

Directives

task

Semantics

The **detach** clause specifies that the task generated by the construct on which it appears is a detachable task. A new *allow-completion* event is created and connected to the completion of the associated **task** region. The original *event-handle* is updated to represent that *allow-completion* event before the task data environment is created. The *event-handle* is 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.

Restrictions

Restrictions to the **detach** clause are as follows:

- If a **detach** clause appears on a directive, then the encountering task must not be a final task.
- A variable that appears in a **detach** clause cannot appear as a list item on a data-environment attribute clause on the same construct.
- A variable that is part of another variable (as an array element or a structure element) cannot appear in a **detach** clause.

- *event-handle* must not have the **POINTER** attribute.
- If *event-handle* has the **ALLOCATABLE** attribute, the allocation status must be allocated when the **task** construct is encountered, and the allocation status must not be changed, either explicitly or implicitly, in the **task** region.

Cross References

- `firstprivate` clause, see [Section 5.4.4](#).
- `task` directive, see [Section 12.5](#)

12.6 `taskloop` Construct

Name: <code>taskloop</code> Category: executable	Association: loop Properties: parallelism-generating, task-generating
---	--

Clauses

[allocate](#), [collapse](#), [default](#), [final](#), [firstprivate](#), [grainsize](#), [if](#), [in_reduction](#), [lastprivate](#), [mergeable](#), [nogroup](#), [num_tasks](#), [priority](#), [private](#), [reduction](#), [shared](#), [untied](#)

Clause set synchronization-clause

Properties: exclusive	Members: nogroup , reduction
------------------------------	---

Clause set granularity-clause

Properties: exclusive	Members: grainsize , num_tasks
------------------------------	---

Binding

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

Semantics

When a thread encounters a `taskloop` construct, the construct partitions the iterations of the associated loops into chunks, each of which is assigned to an explicit task for parallel execution. The iteration count for each associated loop is computed before entry to the outermost loop. The data environment of each generated task is created according to the data-sharing attribute clauses on the `taskloop` construct, per-data environment ICVs, and any defaults that apply. The order of the creation of the loop tasks is unspecified. Programs that rely on any execution order of the logical iterations are non-conforming.

If the `nogroup` clause is not present, the `taskloop` construct executes as if it was enclosed in a `taskgroup` construct with no statements or directives outside of the `taskloop` construct. Thus, the `taskloop` construct creates an implicit `taskgroup` region. If the `nogroup` clause is present, no implicit `taskgroup` region is created.

If a `reduction` clause is present, the behavior is as if a `task_reduction` clause with the same reduction operator and list items was applied to the implicit `taskgroup` construct that encloses the `taskloop` construct. The `taskloop` construct executes as if each generated task was defined by a `task` construct on which an `in_reduction` clause with the same reduction

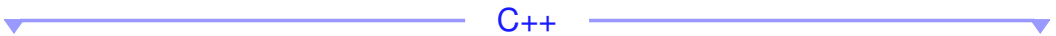

1 operator and list items is present. Thus, the generated tasks are participants of the reduction defined
2 by the **task_reduction** clause that was applied to the implicit **taskgroup** construct.

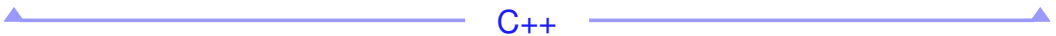

3 If an **in_reduction** clause is present, the behavior is as if each generated task was defined by a
4 **task** construct on which an **in_reduction** clause with the same reduction operator and list
5 items is present. Thus, the generated tasks are participants of a reduction previously defined by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

7 If no clause from the *granularity-clause* set is present, the number of loop tasks generated and the
8 number of logical iterations assigned to these tasks is implementation defined.

9 At the beginning of each logical iteration, the loop iteration variable or the variable declared by
10 *range-decl* of each associated loop has the value that it would have if the set of the associated loops
11 was executed sequentially.

12 When an **if** clause is present and the **if** clause expression evaluates to *false*, undeferred tasks are
13 generated. The use of a variable in an **if** clause expression causes an implicit reference to the
14 variable in all enclosing constructs.

15  C++ 
16 For **firstprivate** variables of class type, the number of invocations of copy constructors that
perform the initialization is implementation defined.

17  C++ 
18 When storage is shared by a **taskloop** region, the programmer must ensure, by adding proper
19 synchronization, that the storage does not reach the end of its lifetime before the **taskloop** region
and its descendent tasks complete their execution.

20 Execution Model Events

21 The *taskloop-begin* event occurs upon entering the **taskloop** region. A *taskloop-begin* will
22 precede any *task-create* events for the generated tasks. The *taskloop-end* event occurs upon
23 completion of the **taskloop** region.

24 Events for an implicit taskgroup region that surrounds the **taskloop** region are the same as for
25 the **taskgroup** construct.

26 The *taskloop-iteration-begin* event occurs at the beginning of each iteration of a **taskloop** region
27 before an explicit task executes the iteration. The *taskloop-chunk-begin* event occurs before an
28 explicit task executes any of its associated iterations in a **taskloop** region.

29 Tool Callbacks

30 A thread dispatches a registered **ompt_callback_work** callback for each occurrence of a
31 *taskloop-begin* and *taskloop-end* event in that thread. The callback occurs in the context of the
32 encountering task. The callback has type signature **ompt_callback_work_t**. The callback
33 receives **ompt_scope_begin** or **ompt_scope_end** as its *endpoint* argument, as appropriate,
34 and **ompt_work_taskloop** as its *work_type* argument.

1 A thread dispatches a registered `ompt_callback_dispatch` callback for each occurrence of a
2 *taskloop-iteration-begin* or *taskloop-chunk-begin* event in that thread.

3 The callback binds to the explicit task executing the iterations. The callback has type signature
4 `ompt_callback_dispatch_t`.

5 **Restrictions**

6 Restrictions to the `taskloop` construct are as follows:

- 7 • The *reduction-modifier* must be **default**.
- 8 • The **conditional** *lastprivate-modifier* must not be specified.

9 **Cross References**

- 10 • **allocate** clause, see [Section 6.6](#)
- 11 • **collapse** clause, see [Section 4.4.3](#)
- 12 • **default** clause, see [Section 5.4.1](#)
- 13 • **final** clause, see [Section 12.3](#)
- 14 • **firstprivate** clause, see [Section 5.4.4](#)
- 15 • **grainsize** clause, see [Section 12.6.1](#)
- 16 • **if** clause, see [Section 3.4](#)
- 17 • **in_reduction** clause, see [Section 5.5.10](#)
- 18 • **lastprivate** clause, see [Section 5.4.5](#)
- 19 • **mergeable** clause, see [Section 12.2](#)
- 20 • **nogroup** clause, see [Section 15.7](#)
- 21 • **num_tasks** clause, see [Section 12.6.2](#)
- 22 • **priority** clause, see [Section 12.4](#)
- 23 • **private** clause, see [Section 5.4.3](#)
- 24 • **reduction** clause, see [Section 5.5.8](#)
- 25 • **shared** clause, see [Section 5.4.2](#)
- 26 • **untied** clause, see [Section 12.1](#)
- 27 • **task** directive, see [Section 12.5](#)
- 28 • **taskgroup** directive, see [Section 15.4](#)
- 29 • Canonical Loop Nest Form, see [Section 4.4.1](#)
- 30 • `ompt_callback_dispatch_t`, see [Section 19.5.2.6](#)

- `ompt_callback_work_t`, see [Section 19.5.2.5](#)
- `ompt_scope_endpoint_t`, see [Section 19.4.4.11](#)
- `ompt_work_t`, see [Section 19.4.4.16](#)

12.6.1 grainsize Clause

Name: <code>grainsize</code>	Properties: unique
-------------------------------------	---------------------------

Arguments

Name	Type	Properties
<i>grain-size</i>	expression of integer type	positive

Modifiers

Name	Modifies	Type	Properties
<i>prescriptiveness</i>	<i>grain-size</i>	Keyword: strict	unique

Directives

[taskloop](#)

Semantics

The **grainsize** clause specifies the number of logical iterations, L_t , that are assigned to each generated task t . If *prescriptiveness* is not specified as **strict**, other than possibly for the generated task that contains the sequentially last iteration, L_t is greater than or equal to the minimum of the value of the *grain-size* expression and the number of logical iterations, but less than two times the value of the *grain-size* expression. If *prescriptiveness* is specified as **strict**, other than possibly for the generated task that contains the sequentially last iteration, L_t is equal to the value of the *grain-size* expression. In both cases, the generated task that contains the sequentially last iteration may have fewer iterations than the value of the *grain-size* expression.

Restrictions

Restrictions to the **grainsize** clause are as follows:

- None of the associated loops may be non-rectangular loops.

Cross References

- `taskloop` directive, see [Section 12.6](#)

12.6.2 num_tasks Clause

Name: <code>num_tasks</code>	Properties: unique
-------------------------------------	---------------------------

Arguments

Name	Type	Properties
<i>num-tasks</i>	expression of integer type	positive

Modifiers

Name	Modifies	Type	Properties
<i>prescriptiveness</i>	<i>num-tasks</i>	Keyword: strict	unique

Directives

taskloop

Semantics

The **num_tasks** clause specifies that the **taskloop** construct create as many tasks as the minimum of the *num-tasks* expression and the number of logical iterations. Each task must have at least one logical iteration. If *prescriptiveness* is specified as **strict** for a task loop with N logical iterations, the logical iterations are partitioned in a balanced manner and each partition is assigned, in order, to a generated task. The partition size is $\lceil N/num-tasks \rceil$ until the number of remaining iterations divides the number of remaining tasks evenly, at which point the partition size becomes $\lfloor N/num-tasks \rfloor$.

Restrictions

Restrictions to the **num_tasks** clause are as follows:

- None of the associated loops may be non-rectangular loops.

Cross References

- **taskloop** directive, see [Section 12.6](#)

12.7 taskyield Construct

Name: taskyield	Association: none
Category: executable	Properties: <i>default</i>

Binding

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

Semantics

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

Cross References

- Task Scheduling, see [Section 12.9](#)

12.8 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 execution.

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* 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 *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_callback_implicit_task_t`, see [Section 19.5.2.11](#)
- `ompt_callback_parallel_begin_t`, see [Section 19.5.2.3](#)
- `ompt_callback_parallel_end_t`, see [Section 19.5.2.4](#)
- `ompt_callback_thread_begin_t`, see [Section 19.5.2.1](#)
- `ompt_callback_thread_end_t`, see [Section 19.5.2.2](#)
- `ompt_task_flag_t`, see [Section 19.4.4.19](#)
- `ompt_thread_t`, see [Section 19.4.4.10](#)

12.9 Task Scheduling

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

- during the generation of an explicit task;
- the point immediately following the generation of an explicit task;
- after the point of completion of the structured block associated with a task;
- in a **taskyield** region;
- in a **taskwait** region;
- at the end of a **taskgroup** region;
- in an implicit barrier region;
- in an explicit **barrier** region;
- during the generation of a **target** region;
- the point immediately following the generation of a **target** region;
- at the beginning and end of a **target data** region;
- in a **target update** region;
- in a **target enter data** region;
- in a **target exit data** region;
- in the **omp_target_memcpy** routine;
- in the **omp_target_memcpy_async** routine;
- in the **omp_target_memcpy_rect** routine; and
- in the **omp_target_memcpy_rect_async** routine.

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

- begin execution of a tied task bound to the current team;
- resume any suspended task region, bound to the current team, to which it is tied;
- begin execution of an untied task bound to the current team; or
- resume any suspended untied task region bound to the current team.

If more than one of the above choices is available, which one is chosen is unspecified.

1 *Task Scheduling Constraints* are as follows:

- 2 1. If any suspended tasks are tied to the thread and are not suspended in a barrier region, a new
3 explicit tied task may be scheduled only if it is a descendent task of all of those suspended
4 tasks. Otherwise, any new explicit tied task may be scheduled.
- 5 2. A dependent task shall not start its execution until its task dependences are fulfilled.
- 6 3. A task shall not be scheduled while any task with which it is mutually exclusive has been
7 scheduled but has not yet completed.
- 8 4. When an explicit task is generated by a construct that contains an **if** clause for which the
9 expression evaluated to *false*, and the previous constraints are already met, the task is
10 executed immediately after generation of the task.

11 A program that relies on any other assumption about task scheduling is non-conforming.

12

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

17 A program must behave correctly and consistently with all conceivable scheduling sequences that
18 are compatible with the rules above.

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

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

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

31 **Execution Model Events**

32 The *task-schedule* event occurs in a thread when the thread switches tasks at a task scheduling
33 point; no event occurs when switching to or from a merged task.

34 **Tool Callbacks**

35 A thread dispatches a registered **ompt_callback_task_schedule** callback for each
36 occurrence of a *task-schedule* event in the context of the task that begins or resumes. This callback

1 has the type signature `ompt_callback_task_schedule_t`. The argument *prior_task_status*
2 is used to indicate the cause for suspending the prior task. This cause may be the completion of the
3 prior task region, the encountering of a `taskyield` construct, or the encountering of an active
4 cancellation point.

5 **Cross References**

- 6 • `ompt_callback_task_schedule_t`, see [Section 19.5.2.10](#)

13 Device Directives and Clauses

This chapter defines constructs and concepts related to device execution.

13.1 `device_type` Clause

Name: <code>device_type</code>	Properties: unique
---------------------------------------	---------------------------

Arguments

Name	Type	Properties
<i>device-type-description</i>	Keyword: any , host , nohost	<i>default</i>

Directives

`begin declare target`, `declare target`, `groupprivate`

Semantics

The `device_type` clause specifies if a version of the procedure or variable should be made available on the host device, non-host devices or both the host device and non-host devices. If **host** is specified then only a host device version of the procedure or variable is made available. If **any** is specified then both host device and non-host device versions of the procedure or variable are made available. If **nohost** is specified for a procedure then only non-host device versions of the procedure are made available. If **nohost** is specified for a variable then that variable is not available on the host device. If the `device_type` clause is not specified, the behavior is as if the `device_type` clause appears with **any** specified.

Cross References

- `begin declare target` directive, see [Section 7.8.2](#)
- `declare target` directive, see [Section 7.8.1](#)
- `groupprivate` directive, see [Section 5.12](#)

13.2 device Clause

Name: <code>device</code>	Properties: unique
----------------------------------	---------------------------

Arguments

Name	Type	Properties
<i>device-description</i>	expression of integer type	<i>default</i>

Modifiers

Name	Modifies	Type	Properties
<i>device-modifier</i>	<i>device-description</i>	Keyword: ancestor , device_num	default

Directives

`dispatch`, `interop`, `target`, `target data`, `target enter data`, `target exit data`, `target update`

Semantics

The **device** clause identifies the target device that is associated with a device construct.

If **device_num** is specified as the *device-modifier*, the *device-description* specifies the device number of the target device. If *device-modifier* does not appear in the clause, the behavior of the clause is as if *device-modifier* is **device_num**. If the *device-description* evaluates to **omp_invalid_device**, runtime error termination is performed.

If **ancestor** is specified as the *device-modifier*, the *device-description* specifies the number of target nesting level of the target device. Specifically, if the *device-description* evaluates to 1, the target device is the parent device of the enclosing **target** region. If the construct on which the **device** clause appears is not encountered in a **target** region, the current device is treated as the parent device.

Unless otherwise specified, for directives that accept the **device** clause, if no **device** clause is present, the behavior is as if the **device** clause appears without a *device-modifier* and with a *device-description* that evaluates to the value of the *default-device-var* ICV.

Restrictions

- The **ancestor** *device-modifier* must not appear on the **device** clause on any directive other than the **target** construct.
- If the **ancestor** *device-modifier* is specified, the **device-description** must evaluate to 1 and a **requires** directive with the **reverse_offload** clause must be specified;
- If the **device_num** *device-modifier* is specified and *target-offload-var* is not **mandatory**, *device-description* must evaluate to a conforming device number.

Cross References

- `dispatch` directive, see [Section 7.6](#)
- `interop` directive, see [Section 14.1](#)
- `target` directive, see [Section 13.8](#)
- `target data` directive, see [Section 13.5](#)
- `target enter data` directive, see [Section 13.6](#)
- `target exit data` directive, see [Section 13.7](#)
- `target update` directive, see [Section 13.9](#)
- `target-offload-var` ICV, see [Table 2.1](#)

13.3 `thread_limit` Clause

Name: <code>thread_limit</code>	Properties: unique
---------------------------------	--------------------

Arguments

Name	Type	Properties
<code>threadlim</code>	expression of integer type	positive

Directives

[target](#), [teams](#)

Semantics

As described in [Section 2.4](#), some constructs limit the number of threads that may participate in the parallel execution of tasks in a contention group initiated by each team by setting the value of the `thread-limit-var` ICV for the initial task to an implementation-defined value greater than zero. If the `thread_limit` clause is specified, the number of threads will be less than or equal to `threadlim`. Otherwise, if the `teams-thread-limit-var` ICV is greater than zero, the effect is as if the `thread_limit` clause was specified with a `threadlim` that evaluates to an implementation defined value less than or equal to the `teams-thread-limit-var` ICV.

Cross References

- `target` directive, see [Section 13.8](#)
- `teams` directive, see [Section 10.3](#)

13.4 Device Initialization

Execution Model Events

The *device-initialize* event occurs in a thread that begins initialization of OpenMP on the device, after 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.

The *device-finalize* event for a target device that has been initialized occurs in some thread before an OpenMP implementation shuts down.

Tool Callbacks

A thread dispatches a registered `ompt_callback_device_initialize` callback for each occurrence of a *device-initialize* event in that thread. This callback has type signature `ompt_callback_device_initialize_t`.

A thread dispatches a registered `ompt_callback_device_load` callback for each occurrence of a *device-load* event in that thread. This callback has type signature `ompt_callback_device_load_t`.

A thread dispatches a registered `ompt_callback_device_unload` callback for each occurrence of a *device-unload* event in that thread. This callback has type signature `ompt_callback_device_unload_t`.

A thread dispatches a registered `ompt_callback_device_finalize` callback for each occurrence of a *device-finalize* event in that thread. This callback has type signature `ompt_callback_device_finalize_t`.

Restrictions

Restrictions to OpenMP device initialization are as follows:

- No thread may offload execution of an OpenMP construct to a device until a dispatched `ompt_callback_device_initialize` callback completes.
- No thread may offload execution of an OpenMP construct to a device after a dispatched `ompt_callback_device_finalize` callback occurs.

Cross References

- `ompt_callback_device_finalize_t`, see [Section 19.5.2.20](#)
- `ompt_callback_device_initialize_t`, see [Section 19.5.2.19](#)
- `ompt_callback_device_load_t`, see [Section 19.5.2.21](#)
- `ompt_callback_device_unload_t`, see [Section 19.5.2.22](#)

13.5 target data Construct

Name: <code>target data</code> Category: executable	Association: block Properties: device, device-affecting, data-mapping, map-entering, map-exiting, mapping-only
--	---

Clauses

`device`, `if`, `map`, `use_device_addr`, `use_device_ptr`

Clause set data-environment-clause

Properties: required	Members: <code>map</code> , <code>use_device_addr</code> , <code>use_device_ptr</code>
-----------------------------	---

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.

Semantics

The `target data` construct maps variables to a device data environment. When a `target data` construct is encountered, the encountering task executes the region. When an `if` clause is present and the `if` clause expression evaluates to *false*, the target 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.

A list item that appears in a `map` clause may also appear in a `use_device_ptr` clause or a `use_device_addr` clause. If one or more `map` clauses are present, the list item conversions that are performed for any `use_device_ptr` or `use_device_addr` clause occur after all variables are mapped on entry to the region 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, as described in [Section 13.6](#).

The events associated with exiting a `target data` region are the same events as associated with a `target exit data` construct, as described in [Section 13.7](#).

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, as described in [Section 13.6](#).

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, as described in [Section 13.7](#).

Restrictions

Restrictions to the **target data** construct are as follows:

- A *map-type* in a **map** clause must be **to**, **from**, **tofrom** or **alloc**.

Cross References

- **device** clause, see [Section 13.2](#)
- **if** clause, see [Section 3.4](#)
- **map** clause, see [Section 5.8.3](#)
- **use_device_addr** clause, see [Section 5.4.10](#)
- **use_device_ptr** clause, see [Section 5.4.8](#)

13.6 target enter data Construct

Name: target enter data Category: executable	Association: none Properties: parallelism-generating, task-generating, device, device-affecting, data-mapping, map-entering, mapping-only
--	--

Clauses

[depend](#), [device](#), [if](#), [map](#), [nowait](#)

Binding

The binding task set for a **target enter data** region is the generating task, which is the *target task* generated by the **target enter data** construct. The **target enter data** region binds to the corresponding *target task* region.

Semantics

When a **target enter data** construct is encountered, the list items are mapped to the device data environment according to the **map** clause semantics. The **target enter data** construct generates a *target task*. The generated task region encloses the **target enter data** region. If a **depend** clause is present, it is associated with the *target task*. If the **nowait** clause is present, execution of the *target task* may be deferred. If the **nowait** clause is not present, the *target task* is an included task.

All clauses are evaluated when the **target enter data** construct is encountered. The data environment of the *target task* is created according to the data-mapping attribute clauses on the **target enter data** construct, per-data environment ICVs, and any default data-sharing attribute rules that apply to the **target enter data** construct. If a variable or part of a variable is mapped by the **target enter data** construct, the variable has a default data-sharing attribute of shared in the data environment of the *target task*.

1 Assignment operations associated with mapping a variable (see [Section 5.8.3](#)) occur when the
2 *target task* executes.

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

5 Execution Model Events

6 Events associated with a *target task* are the same as for the **task** construct defined in [Section 12.5](#).

7 The *target-enter-data-begin* event occurs after creation of the target task and completion of all
8 predecessor tasks that are not target tasks for the same device. The *target-enter-data-begin* event is
9 a *target-task-begin* event.

10 The *target-enter-data-end* event occurs after all other events associated with the
11 **target enter data** construct.

12 Tool Callbacks

13 Callbacks associated with events for *target tasks* are the same as for the **task** construct defined in
14 [Section 12.5](#); (*flags & ompt_task_target*) always evaluates to *true* in the dispatched callback.

15 A thread dispatches a registered **ompt_callback_target** or
16 **ompt_callback_target_emi** callback with **ompt_scope_begin** as its *endpoint*
17 argument and **ompt_target_enter_data** or **ompt_target_enter_data_nowait** if
18 the **nowait** clause is present as its *kind* argument for each occurrence of a *target-enter-data-begin*
19 event in that thread in the context of the target task on the host. Similarly, a thread dispatches a
20 registered **ompt_callback_target** or **ompt_callback_target_emi** callback with
21 **ompt_scope_end** as its *endpoint* argument and **ompt_target_enter_data** or
22 **ompt_target_enter_data_nowait** if the **nowait** clause is present as its *kind* argument
23 for each occurrence of a *target-enter-data-end* event in that thread in the context of the target task
24 on the host. These callbacks have type signature **ompt_callback_target_t** or
25 **ompt_callback_target_emi_t**, respectively.

26 Restrictions

27 Restrictions to the **target enter data** construct are as follows:

- 28 • At least one **map** clause must appear on the directive.
- 29 • All **map** clauses must be *map-entering*.

30 Cross References

- 31 • **depend** clause, see [Section 15.9.5](#)
- 32 • **device** clause, see [Section 13.2](#)
- 33 • **if** clause, see [Section 3.4](#)
- 34 • **map** clause, see [Section 5.8.3](#)
- 35 • **nowait** clause, see [Section 15.6](#)

- **task** directive, see [Section 12.5](#)
- **ompt_callback_target_emi_t** and **ompt_callback_target_t**, see [Section 19.5.2.26](#)

13.7 target exit data Construct

Name: <code>target exit data</code> Category: executable	Association: none Properties: parallelism-generating, task-generating, device, device-affecting, data-mapping, map-exiting, mapping-only
---	---

Clauses

[depend](#), [device](#), [if](#), [map](#), [nowait](#)

Binding

The binding task set for a **target exit data** region is the generating task, which is the *target task* generated by the **target exit data** construct. The **target exit data** region binds to the corresponding *target task* region.

Semantics

When a **target exit data** construct is encountered, the list items in the **map** clauses are unmapped from the device data environment according to the **map** clause semantics. The **target exit data** construct generates a *target task*. The generated task region encloses the **target exit data** region. If a **depend** clause is present, it is associated with the *target task*. If the **nowait** clause is present, execution of the *target task* may be deferred. If the **nowait** clause is not present, the *target task* is an included task.

All clauses are evaluated when the **target exit data** construct is encountered. The data environment of the *target task* is created according to the data-mapping attribute clauses on the **target exit data** construct, per-data environment ICVs, and any default data-sharing attribute rules that apply to the **target exit data** construct. If a variable or part of a variable is mapped by the **target exit data** construct, the variable has a default data-sharing attribute of shared in the data environment of the *target task*.

Assignment operations associated with mapping a variable (see [Section 5.8.3](#)) occur when the *target task* executes.

When an **if** clause is present and the **if** clause expression evaluates to *false*, the target 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 [Section 12.5](#).

3 The *target-exit-data-begin* event occurs after creation of the target task and completion of all
4 predecessor tasks that are not target tasks for the same device. The *target-exit-data-begin* event is a
5 *target-task-begin* event.

6 The *target-exit-data-end* event occurs after all other events associated with the
7 **target exit data** construct.

8 Tool Callbacks

9 Callbacks associated with events for *target tasks* are the same as for the **task** construct defined in
10 [Section 12.5](#); (*flags & ompt_task_target*) always evaluates to *true* in the dispatched callback.

11 A thread dispatches a registered **ompt_callback_target** or
12 **ompt_callback_target_emi** callback with **ompt_scope_begin** as its *endpoint*
13 argument and **ompt_target_exit_data** or **ompt_target_exit_data_nowait** if the
14 **nowait** clause is present as its *kind* argument for each occurrence of a *target-exit-data-begin*
15 event in that thread in the context of the target task on the host. Similarly, a thread dispatches a
16 registered **ompt_callback_target** or **ompt_callback_target_emi** callback with
17 **ompt_scope_end** as its *endpoint* argument and **ompt_target_exit_data** or
18 **ompt_target_exit_data_nowait** if the **nowait** clause is present as its *kind* argument for
19 each occurrence of a *target-exit-data-end* event in that thread in the context of the target task on the
20 host. These callbacks have type signature **ompt_callback_target_t** or
21 **ompt_callback_target_emi_t**, respectively.

22 Restrictions

23 Restrictions to the **target exit data** construct are as follows:

- 24 • At least one **map** clause must appear on the directive.
- 25 • All **map** clauses must be a *map-exiting*.

26 Cross References

- 27 • **depend** clause, see [Section 15.9.5](#)
- 28 • **device** clause, see [Section 13.2](#)
- 29 • **if** clause, see [Section 3.4](#)
- 30 • **map** clause, see [Section 5.8.3](#)
- 31 • **nowait** clause, see [Section 15.6](#)
- 32 • **task** directive, see [Section 12.5](#)
- 33 • **ompt_callback_target_emi_t** and **ompt_callback_target_t**, see
34 [Section 19.5.2.26](#)

13.8 target Construct

Name: <code>target</code> Category: executable	Association: block Properties: parallelism-generating, thread-limiting, exception-aborting, task-generating, device, device-affecting, data-mapping, map-entering, map-exiting, context-matching
---	---

Clauses

`allocate`, `defaultmap`, `depend`, `device`, `firstprivate`, `has_device_addr`, `if`, `in_reduction`, `is_device_ptr`, `map`, `nowait`, `private`, `thread_limit`, `uses_allocators`

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.

Semantics

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 on a device. The `target` construct generates a *target task*. The generated task region encloses the `target` region. If a `depend` clause is present, it is associated with the *target task*. The `device` clause determines the device on which the `target` region executes. If the `nowait` clause is present, execution of the *target task* may be deferred. If the `nowait` clause is not present, the *target task* is an included task.

All clauses are evaluated when the `target` construct is encountered. The data environment of the *target task* is created according to the data-sharing and data-mapping attribute clauses on the `target` construct, per-data environment ICVs, and any default data-sharing attribute rules that apply to the `target` construct. If a variable or part of a variable is mapped by the `target` construct and does not appear as a list item in an `in_reduction` clause on the construct, the variable has a default data-sharing attribute of shared in the data environment of the *target task*. Assignment operations associated with mapping a variable (see [Section 5.8.3](#)) occur when the *target task* executes.

If the `device` clause is specified with the *ancestor device-modifier*, the encountering thread waits for completion of the `target` region on the parent device before resuming. For any list item that appears in a `map` clause on the same construct, if the corresponding list item exists in the device data environment of the parent device, it is treated as if it has a reference count of positive infinity.

When an `if` clause is present and the `if` clause expression evaluates to *false*, the effect is as if a `device` clause that specifies `omp_initial_device` as the device number is present, regardless of any other `device` clause on the directive.

1 If a procedure is explicitly or implicitly referenced in a **target** construct that does not specify a
2 **device** clause in which the **ancestor** *device-modifier* appears then that procedure is treated as
3 if its name had appeared in an **enter** clause on a declare target directive.

4 If a variable with static storage duration is declared in a **target** construct that does not specify a
5 **device** clause in which the **ancestor** *device-modifier* appears then the named variable is treated
6 as if it had appeared in an **enter** clause on a declare target directive if it is not a groupprivate
7 variable and otherwise as if it had appeared in a **local** clause on a declare target directive.

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

8 If a list item in a **map** clause has a base pointer that is predetermined firstprivate (see [Section 5.1.1](#))
9 and on entry to the **target** region the list item is mapped, the firstprivate pointer is updated via
10 corresponding base pointer initialization.

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

▼────────────────────────────────── Fortran ───────────────────────────────────►

11 When an internal procedure is called in a **target** region, any references to variables that are host
12 associated in the procedure have unspecified behavior.

▲────────────────────────────────── Fortran ───────────────────────────────────►

13 Execution Model Events

14 Events associated with a *target task* are the same as for the **task** construct defined in [Section 12.5](#).

15 Events associated with the *initial task* that executes the **target** region are defined in [Section 12.8](#).

16 The *target-submit-begin* event occurs prior to initiating creation of an initial task on a target device
17 for a **target** region.

18 The *target-submit-end* event occurs after initiating creation of an initial task on a target device for a
19 **target** region.

20 The *target-begin* event occurs after creation of the target task and completion of all predecessor
21 tasks that are not target tasks for the same device. The *target-begin* event is a *target-task-begin*
22 event.

23 The *target-end* event occurs after all other events associated with the **target** construct.

24 Tool Callbacks

25 Callbacks associated with events for *target tasks* are the same as for the **task** construct defined in
26 [Section 12.5](#); (*flags & omp_target_target*) always evaluates to *true* in the dispatched callback.

27 A thread dispatches a registered **omp_callback_target** or
28 **omp_callback_target_emi** callback with **omp_scope_begin** as its *endpoint*
29 argument and **omp_target** or **omp_target_nowait** if the **nowait** clause is present as its
30 *kind* argument for each occurrence of a *target-begin* event in that thread in the context of the target
31 task on the host. Similarly, a thread dispatches a registered **omp_callback_target** or

1 **ompt_callback_target_emi** callback with **ompt_scope_end** as its *endpoint* argument
2 and **ompt_target** or **ompt_target_nowait** if the **nowait** clause is present as its *kind*
3 argument for each occurrence of a *target-end* event in that thread in the context of the target task on
4 the host. These callbacks have type signature **ompt_callback_target_t** or
5 **ompt_callback_target_emi_t**, respectively.

6 A thread dispatches a registered **ompt_callback_target_submit_emi** callback with
7 **ompt_scope_begin** as its endpoint argument for each occurrence of a *target-submit-begin*
8 event in that thread. Similarly, a thread dispatches a registered
9 **ompt_callback_target_submit_emi** callback with **ompt_scope_end** as its endpoint
10 argument for each occurrence of a *target-submit-end* event in that thread. These callbacks have type
11 signature **ompt_callback_target_submit_emi_t**.

12 A thread dispatches a registered **ompt_callback_target_submit** callback for each
13 occurrence of a *target-submit-begin* event in that thread. The callback occurs in the context of the
14 target task and has type signature **ompt_callback_target_submit_t**.

15 Restrictions

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

- 17 • Device-affecting constructs, other than **target** constructs for which the **ancestor**
18 *device-modifier* is specified, must not be encountered during execution of a **target** region.
- 19 • The result of an **omp_set_default_device**, **omp_get_default_device**, or
20 **omp_get_num_devices** routine called within a **target** region is unspecified.
- 21 • The effect of an access to a threadprivate variable in a **target** region is unspecified.
- 22 • If a list item in a **map** clause is a structure element, any other element of that structure that is
23 referenced in the **target** construct must also appear as a list item in a **map** clause.
- 24 • A list item in a data-sharing attribute clause that is specified on a **target** construct must not
25 have the same base variable as a list item in a **map** clause on the construct.
- 26 • A variable referenced in a **target** region but not the **target** construct that is not declared
27 in the **target** region must appear in a **declare target** directive.
- 28 • A *map-type* in a **map** clause must be **to**, **from**, **tofrom** or **alloc**.
- 29 • If a **device** clause is specified with the *ancestor device-modifier*, only the **device**,
30 **firstprivate**, **private**, **defaultmap**, and **map** clauses may appear on the construct
31 and no OpenMP constructs or calls to OpenMP API runtime routines are allowed inside the
32 corresponding **target** region.
- 33 • Memory allocators that do not appear in a **uses_allocators** clause cannot appear as an
34 allocator in an **allocate** clause or be used in the **target** region unless a **requires**
35 directive with the **dynamic_allocators** clause is present in the same compilation unit.
- 36 • Any IEEE floating-point exception status flag, halting mode, or rounding mode set prior to a
37 **target** region is unspecified in the region.

- 1 • Any IEEE floating-point exception status flag, halting mode, or rounding mode set in a
2 **target** region is unspecified upon exiting the region.
- 3 • A program must not rely on the value of a function address in a **target** region except for
4 assignments, comparisons to zero and indirect calls.
- ▼────────────────────────────────── C / C++ ───────────────────────────────────►
- 5 • Upon exit from a **target** region, the value of an attached pointer must not be different from
6 the value when entering the region.
- ▲────────────────────────────────── C / C++ ───────────────────────────────────►
- ▼────────────────────────────────── C++ ───────────────────────────────────►
- 7 • The run-time type information (RTTI) of an object can only be accessed from the device on
8 which it was constructed.
- 9 • Invoking a virtual member function of an object on a device other than the device on which
10 the object was constructed results in unspecified behavior, unless the object is accessible and
11 was constructed on the host device.
- 12 • If an object of polymorphic class type is destructed, virtual member functions of any
13 previously existing corresponding objects in other device data environments must not be
14 invoked.
- ▲────────────────────────────────── C++ ───────────────────────────────────►
- ▼────────────────────────────────── Fortran ───────────────────────────────────►
- 15 • An attached pointer that is associated with a given pointer target must not be associated with
16 a different pointer target upon exit from a **target** region.
- 17 • A reference to a coarray that is encountered on a non-host device must not be coindexed or
18 appear as an actual argument to a procedure where the corresponding dummy argument is a
19 coarray.
- 20 • If the allocation status of a mapped variable or a list item that appears in a
21 **has_device_addr** clause that has the **ALLOCATABLE** attribute is unallocated on entry to
22 a **target** region, the allocation status of the corresponding variable in the device data
23 environment must be unallocated upon exiting the region.
- 24 • If the allocation status of a mapped variable or a list item that appears in a
25 **has_device_addr** clause that has the **ALLOCATABLE** attribute is allocated on entry to a
26 **target** region, the allocation status and shape of the corresponding variable in the device
27 data environment may not be changed, either explicitly or implicitly, in the region after entry
28 to it.
- 29 • If the association status of a list item with the **POINTER** attribute that appears in a **map** or
30 **has_device_addr** clause on the construct is associated upon entry to the **target**
31 region, the list item must be associated with the same pointer target upon exit from the region.

- If the association status of a list item with the **POINTER** attribute that appears in a **map** or **has_device_addr** clause on the construct is disassociated upon entry to the **target** region, the list item must be disassociated upon exit from the region.
- A program must not rely on the association status of a procedure pointer in a **target** region except for calls to the **ASSOCIATED** inquiry function without the optional *proc-target* argument, pointer assignments and indirect calls.

Fortran

Cross References

- **allocate** clause, see [Section 6.6](#)
- **defaultmap** clause, see [Section 5.8.6](#)
- **depend** clause, see [Section 15.9.5](#)
- **device** clause, see [Section 13.2](#)
- **firstprivate** clause, see [Section 5.4.4](#)
- **has_device_addr** clause, see [Section 5.4.9](#)
- **if** clause, see [Section 3.4](#)
- **in_reduction** clause, see [Section 5.5.10](#)
- **is_device_ptr** clause, see [Section 5.4.7](#)
- **map** clause, see [Section 5.8.3](#)
- **nowait** clause, see [Section 15.6](#)
- **private** clause, see [Section 5.4.3](#)
- **thread_limit** clause, see [Section 13.3](#)
- **uses_allocators** clause, see [Section 6.8](#)
- **target data** directive, see [Section 13.5](#)
- **task** directive, see [Section 12.5](#)
- **ompt_callback_target_emi_t** and **ompt_callback_target_t**, see [Section 19.5.2.26](#)
- **ompt_callback_target_submit_emi_t** and **ompt_callback_target_submit_t**, see [Section 19.5.2.28](#)

13.9 target update Construct

Name: <code>target update</code> Category: executable	Association: none Properties: parallelism-generating, task-generating, device, device-affecting
--	--

Clauses

`depend`, `device`, `from`, `if`, `nowait`, `to`

Clause set

Properties: required	Members: <code>from</code> , <code>to</code>
-----------------------------	---

Binding

The binding task set for a **target update** region is the generating task, which is the *target task* generated by the **target update** construct. The **target update** region binds to the corresponding *target task* region.

Semantics

The **target update** directive makes the corresponding list items in the device data environment consistent with their original list items, according to the specified *data-motion-clauses*. The **target update** construct generates a *target task*. The generated task region encloses the **target update** region. If a **depend** clause is present, it is associated with the *target task*. If the **nowait** clause is present, execution of the *target task* may be deferred. If the **nowait** clause is not present, the *target task* is an included task.

All clauses are evaluated when the **target update** construct is encountered. The data environment of the *target task* is created according to *data-motion-clauses* on the **target update** construct, per-data environment ICVs, and any default data-sharing attribute rules that apply to the **target update** construct. If a variable or part of a variable is a list item in a *data-motion-clause* on the **target update** construct, the variable has a default data-sharing attribute of shared in the data environment of the *target task*.

Assignment operations associated with any motion clauses occur when the *target task* executes. When an **if** clause is present and the **if** clause expression evaluates to *false*, no assignments occur.

Execution Model Events

Events associated with a *target task* are the same as for the **task** construct defined in [Section 12.5](#).

The *target-update-begin* event occurs after creation of the target task and completion of all predecessor tasks that are not target tasks for the same device.

The *target-update-end* event occurs after all other events associated with the **target update** construct.

The *target-data-op-begin* event occurs in the **target update** region before a thread initiates a data operation on the target device.

The *target-data-op-end* event occurs in the **target update** region after a thread initiates a data operation on the target device.

1 Tool Callbacks

2 Callbacks associated with events for *target tasks* are the same as for the **task** construct defined in
3 [Section 12.5](#); (*flags & ompt_task_target*) always evaluates to *true* in the dispatched callback.

4 A thread dispatches a registered **ompt_callback_target** or
5 **ompt_callback_target_emi** callback with **ompt_scope_begin** as its *endpoint*
6 argument and **ompt_target_update** or **ompt_target_update_nowait** if the **nowait**
7 clause is present as its *kind* argument for each occurrence of a *target-update-begin* event in that
8 thread in the context of the target task on the host. Similarly, a thread dispatches a registered
9 **ompt_callback_target** or **ompt_callback_target_emi** callback with
10 **ompt_scope_end** as its *endpoint* argument and **ompt_target_update** or
11 **ompt_target_update_nowait** if the **nowait** clause is present as its *kind* argument for each
12 occurrence of a *target-update-end* event in that thread in the context of the target task on the host.
13 These callbacks have type signature **ompt_callback_target_t** or
14 **ompt_callback_target_emi_t**, respectively.

15 A thread dispatches a registered **ompt_callback_target_data_op_emi** callback with
16 **ompt_scope_begin** as its endpoint argument for each occurrence of a *target-data-op-begin*
17 event in that thread. Similarly, a thread dispatches a registered
18 **ompt_callback_target_data_op_emi** callback with **ompt_scope_end** as its endpoint
19 argument for each occurrence of a *target-data-op-end* event in that thread. These callbacks have
20 type signature **ompt_callback_target_data_op_emi_t**.

21 A thread dispatches a registered **ompt_callback_target_data_op** callback for each
22 occurrence of a *target-data-op-end* event in that thread. The callback occurs in the context of the
23 target task and has type signature **ompt_callback_target_data_op_t**.

24 Cross References

- 25 • **depend** clause, see [Section 15.9.5](#)
- 26 • **device** clause, see [Section 13.2](#)
- 27 • **from** clause, see [Section 5.9.2](#)
- 28 • **if** clause, see [Section 3.4](#)
- 29 • **nowait** clause, see [Section 15.6](#)
- 30 • **to** clause, see [Section 5.9.1](#)
- 31 • **task** directive, see [Section 12.5](#)
- 32 • **ompt_callback_target_emi_t** and **ompt_callback_target_t**, see
33 [Section 19.5.2.26](#)
- 34 • **ompt_callback_task_create_t**, see [Section 19.5.2.7](#)

14 Interoperability

An OpenMP implementation may interoperate with one or more foreign runtime environments through the use of the **interop** construct that is described in this chapter, the **interop** operation for a declared variant function and the interoperability routines that are available through the OpenMP Runtime API.

C / C++

The implementation must provide *foreign-runtime-id* values that are enumerators of type **omp_interop_fr_t** and that correspond to the supported foreign runtime environments.

C / C++

Fortran

The implementation must provide *foreign-runtime-id* values that are named integer constants with kind **omp_interop_fr_kind** and that correspond to the supported foreign runtime environments.

Fortran

Each *foreign-runtime-id* value provided by an implementation will be available as **omp_ifr_name**, where *name* is the name of the foreign runtime environment. Available names include those that are listed in the *OpenMP Additional Definitions* document; implementation-defined names may also be supported. The value of **omp_ifr_last** is defined as one greater than the value of the highest supported *foreign-runtime-id* value that is listed in the aforementioned document.

Cross References

- Interoperability Routines, see [Section 18.12](#)

14.1 interop Construct

Name: interop Category: executable	Association: none Properties: device
--	---

Clauses

depend, destroy, device, init, nowait, use

Clause set action-clause

Properties: required	Members: destroy, init, use
-----------------------------	---

1 Binding

2 The binding task set for an **interop** region is the generating task. The **interop** region binds to
3 the region of the generating task.

4 Semantics

5 The **interop** construct retrieves interoperability properties from the OpenMP implementation to
6 enable interoperability with foreign execution contexts. When an **interop** construct is
7 encountered, the encountering task executes the region.

8 For each *action-clause*, the *interop-type* set is the set of *interop-type* modifiers specified for the
9 clause if the clause is **init** or for the *init* clause that initialized the *interop-var* that is specified for
10 the clause if the clause is not **init**.

11 If the *interop-type* set includes **targetsync**, an empty *mergeable task* is generated. If the
12 **nowait** clause is not present on the construct then the task is also an *included task*. Any **depend**
13 clauses that are present on the construct apply to the generated task.

14 The **interop** construct ensures an ordered execution of the generated task relative to foreign tasks
15 executed in the foreign execution context through the foreign synchronization object that is
16 accessible through the **targetsync** property. When the creation of the foreign task precedes the
17 encountering of an **interop** construct in happens before order (see [Section 1.4.5](#)), the foreign
18 task must complete execution before the generated task begins execution. Similarly, when the
19 creation of a foreign task follows the encountering of an **interop** construct in happens before
20 order, the foreign task must not begin execution until the generated task completes execution. No
21 ordering is imposed between the encountering thread and either foreign tasks or OpenMP tasks by
22 the **interop** construct.

23 If the *interop-type* set does not include **targetsync**, the **nowait** clause has no effect.

24 Restrictions

25 Restrictions to the **interop** construct are as follows:

- 26 • A **depend** clause can only appear on the directive if the *interop-type* includes
27 **targetsync**.
- 28 • Each *interop-var* may be specified for at most one *action-clause* of each **interop** construct.

29 Cross References

- 30 • **depend** clause, see [Section 15.9.5](#)
- 31 • **destroy** clause, see [Section 3.5](#)
- 32 • **device** clause, see [Section 13.2](#)
- 33 • **init** clause, see [Section 14.1.2](#)
- 34 • **nowait** clause, see [Section 15.6](#)
- 35 • **use** clause, see [Section 14.1.3](#)
- 36 • Interoperability Routines, see [Section 18.12](#)

14.1.1 OpenMP Foreign Runtime Identifiers

An OpenMP foreign runtime identifier, *foreign-runtime-id*, is a base language string literal or a compile-time constant OpenMP integer expression. Allowed values for *foreign-runtime-id* include the names (as string literals) and integer values that the *OpenMP Additional Definitions* document specifies and the corresponding `omp_ifr_name` constants of OpenMP `interop_fr` type. Implementation-defined values for *foreign-runtime-id* may also be supported.

14.1.2 `init` Clause

Name: <code>init</code>	Properties: <i>default</i>
--------------------------------	-----------------------------------

Arguments

Name	Type	Properties
<i>interop-var</i>	variable of <code>omp_interop_t</code> type	<i>default</i>

Modifiers

Name	Modifies	Type	Properties
<i>interop-preference</i>	<i>Generic</i>	Complex, name: prefer_type Arguments: <i>preference_list</i> OpenMP foreign runtime preference list (<i>default</i>)	complex, unique
<i>interop-type</i>	<i>Generic</i>	Keyword: target , targetsync	repeatable, required

Directives

`interop`

Semantics

The `init` clause specifies that *interop-var* is initialized to refer to the list of properties associated with any *interop-type*. For any *interop-type*, the properties **type**, **type_name**, **vendor**, **vendor_name** and **device_num** will be available. If the implementation cannot initialize *interop-var*, it is initialized to the value of `omp_interop_none`, which is defined to be zero.

The **targetsync** *interop-type* will additionally provide the **targetsync** property, which is the handle to a foreign synchronization object for enabling synchronization between OpenMP tasks and foreign tasks that execute in the foreign execution context.

The **target** *interop-type* will additionally provide the following properties:

- **device**, which will be a foreign device handle;
- **device_context**, which will be a foreign device context handle; and
- **platform**, which will be a handle to a foreign platform of the device.

If the **prefer_type** *interop-modifier* clause is specified, the first supported *foreign-runtime-id* in *preference-list* in left-to-right order is used. The *foreign-runtime-id* that is used if the implementation does not support any of the items in *preference-list* is implementation defined.

Restrictions

Restrictions to the **init** clause are as follows:

- Each *interop-type* may be specified at most once.
- *interop-var* must be non-const.

Cross References

- **interop** directive, see [Section 14.1](#)
- OpenMP Foreign Runtime Identifiers, see [Section 14.1.1](#)

14.1.3 use Clause

Name: <code>use</code>	Properties: <i>default</i>
-------------------------------	-----------------------------------

Arguments

Name	Type	Properties
<i>interop-var</i>	variable of <code>omp_interop_t</code> type	<i>default</i>

Directives

[interop](#)

Semantics

The **use** clause specifies the *interop-var* that is used for the effects of the directive on which the clause appears. However, *interop-var* is not initialized, destroyed or otherwise modified. The *interop-type* is inferred based on the *interop-type* used to initialize *interop-var*.

Cross References

- **interop** directive, see [Section 14.1](#)

14.2 Interoperability Requirement Set

The *interoperability requirement set* of each task is a logical set of properties that can be added or removed by different directives. These properties can be queried by other constructs that have interoperability semantics.

A construct can add the following properties to the set:

- *depend*, which specifies that the construct requires enforcement of the synchronization relationship expressed by the *depend* clause;

- 1 • *nowait*, which specifies that the construct is asynchronous; and
- 2 • *is_device_ptr(list-item)*, which specifies that the *list-item* is a device pointer in the construct.

3 The following directives may add properties to the set:

- 4 • **dispatch**.

5 The following directives may remove properties from the set:

- 6 • **declare variant**.

7 **Cross References**

- 8 • **dispatch** directive, see [Section 7.6](#)
- 9 • Declare Variant Directives, see [Section 7.5](#)

15 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](#) and [Section 1.4.6](#). [Section 15.8.7](#) defines the behavior of synchronizing flush operations that are implied at various other locations in an OpenMP program.

15.1 Synchronization Hints

The programmer can provide hints about the expected dynamic behavior or suggested implementation of a lock by using `omp_init_lock_with_hint` or `omp_init_nest_lock_with_hint` to initialize it. Synchronization hints may also be provided for `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.

Cross References

- `hint` clause, see [Section 15.1.2](#)
- `omp_init_lock_with_hint` and `omp_init_nest_lock_with_hint`, see [Section 18.9.2](#)

15.1.1 Synchronization Hint Type

Synchronization hints are specified with an OpenMP `sync_hint` type. The C/C++ header file (`omp.h`) and the Fortran include file (`omp_lib.h`) and/or Fortran 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_sync_hint_uncontended = 0x1,

```



```

1  omp_sync_hint_contended = 0x2,
2  omp_sync_hint_nonspeculative = 0x4,
3  omp_sync_hint_speculative = 0x8,
4  } omp_sync_hint_t;

```

C / C++

Fortran

```

5  integer (kind=omp_sync_hint_kind), &
6     parameter :: omp_sync_hint_none = &
7         int(Z'0', kind=omp_sync_hint_kind)
8  integer (kind=omp_sync_hint_kind), &
9     parameter :: omp_sync_hint_uncontended = &
10         int(Z'1', kind=omp_sync_hint_kind)
11 integer (kind=omp_sync_hint_kind), &
12     parameter :: omp_sync_hint_contended = &
13         int(Z'2', kind=omp_sync_hint_kind)
14 integer (kind=omp_sync_hint_kind), &
15     parameter :: omp_sync_hint_nonspeculative = &
16         int(Z'4', kind=omp_sync_hint_kind)
17 integer (kind=omp_sync_hint_kind), &
18     parameter :: omp_sync_hint_speculative = &
19         int(Z'8', kind=omp_sync_hint_kind)

```

Fortran

20 The hints can be combined by using the + or | operators in C/C++ or the + operator in Fortran.
21 Combining `omp_sync_hint_none` with any other hint is equivalent to specifying the other hint.

22 The intended meaning of each hint is:

- 23 • **omp_sync_hint_uncontended**: low contention is expected in this operation, that is,
24 few threads are expected to perform the operation simultaneously in a manner that requires
25 synchronization;
- 26 • **omp_sync_hint_contended**: high contention is expected in this operation, that is,
27 many threads are expected to perform the operation simultaneously in a manner that requires
28 synchronization;
- 29 • **omp_sync_hint_speculative**: the programmer suggests that the operation should be
30 implemented using speculative techniques such as transactional memory; and
- 31 • **omp_sync_hint_nonspeculative**: the programmer suggests that the operation
32 should not be implemented using speculative techniques such as transactional memory.

Note – Future OpenMP specifications may add additional hints to the **sync_hint** type. Implementers are advised to add implementation-defined hints starting from the most significant bit of the type and to include the name of the implementation in the name of the added hint to avoid name conflicts with other OpenMP implementations.

Restrictions

Restrictions to the synchronization hints are as follows:

- The hints **omp_sync_hint_uncontended** and **omp_sync_hint_contended** cannot be combined.
- The hints **omp_sync_hint_nonspeculative** and **omp_sync_hint_speculative** cannot be combined.

15.1.2 hint Clause

Name: <code>hint</code>	Properties: unique
--------------------------------	---------------------------

Arguments

Name	Type	Properties
<i>hint-expr</i>	expression of <code>sync_hint</code> type	<i>default</i>

Directives

atomic, **critical**

Semantics

The **hint** clause gives the implementation additional information about the expected runtime properties of the region that corresponds to the construct on which it appears and that can optionally be used to optimize the implementation. The presence of a **hint** clause does not affect the semantics of the construct. If no **hint** clause is specified for a construct that accepts it, the effect is as if **hint (omp_sync_hint_none)** had been specified.

Restrictions

- *hint-expr* must evaluate to a valid synchronization hint.

Cross References

- **atomic** directive, see [Section 15.8.5](#)
- **critical** directive, see [Section 15.2](#)
- Synchronization Hint Type, see [Section 15.1.1](#)

15.2 critical Construct

Name: <code>critical</code> Category: executable	Association: block Properties: thread-limiting
---	---

Arguments

`critical` (*name*)

Name	Type	Properties
<i>name</i>	base language identifier	optional

Clauses

`hint`

Binding

The binding thread set for a `critical` region is all threads executing tasks in the contention group.

Semantics

The *name* argument is used to identify the `critical` construct. For any `critical` construct for which *name* is not specified, the effect is as if an identical (unspecified) name was specified. The region that corresponds to a `critical` construct of a given name is executed as if only a single thread at a time among all threads in the contention group executes the region, without regard to the teams to which the threads belong.

▼ 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 ▲

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

Restrictions to the `critical` construct are as follows:

- Unless `omp_sync_hint_none` is specified, the `critical` construct must specify a name.
- The *hint-expr* that is applied to each of the `critical` constructs with the same *name* must evaluate to the same value.

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

- `hint` clause, see [Section 15.1.2](#)
- `ompt_callback_mutex_acquire_t`, see [Section 19.5.2.14](#)
- `ompt_callback_mutex_t`, see [Section 19.5.2.15](#)
- `ompt_mutex_t`, see [Section 19.4.4.17](#)

15.3 Barriers

15.3.1 barrier Construct

Name: <code>barrier</code> Category: executable	Association: none Properties: <i>default</i>
--	---

1 Binding

2 The binding thread set for a **barrier** region is the current team. A **barrier** region binds to the
3 innermost enclosing **parallel** region.

4 Semantics

5 The **barrier** construct specifies an explicit barrier at the point at which the construct appears.
6 Unless the binding region is canceled, all threads of the team that executes that binding region must
7 enter the **barrier** region and complete execution of all explicit tasks bound to that binding region
8 before any of the threads continue execution beyond the barrier.

9 The **barrier** region includes an implicit task scheduling point in the current task region.

10 Execution Model Events

11 The *explicit-barrier-begin* event occurs in each thread that encounters the **barrier** construct on
12 entry to the **barrier** region.

13 The *explicit-barrier-wait-begin* event occurs when a task begins an interval of active or passive
14 waiting in a **barrier** region.

15 The *explicit-barrier-wait-end* event occurs when a task ends an interval of active or passive waiting
16 and resumes execution in a **barrier** region.

17 The *explicit-barrier-end* event occurs in each thread that encounters the **barrier** construct after
18 the barrier synchronization on exit from the **barrier** region.

19 A *cancellation* event occurs if cancellation is activated at an implicit cancellation point in a
20 **barrier** region.

21 Tool Callbacks

22 A thread dispatches a registered **ompt_callback_sync_region** callback with
23 **ompt_sync_region_barrier_explicit** as its *kind* argument and **ompt_scope_begin**
24 as its *endpoint* argument for each occurrence of an *explicit-barrier-begin* event. Similarly, a thread
25 dispatches a registered **ompt_callback_sync_region** callback with
26 **ompt_sync_region_barrier_explicit** as its *kind* argument and **ompt_scope_end** as
27 its *endpoint* argument for each occurrence of an *explicit-barrier-end* event. These callbacks occur
28 in the context of the task that encountered the **barrier** construct and have type signature
29 **ompt_callback_sync_region_t**.

30 A thread dispatches a registered **ompt_callback_sync_region_wait** callback with
31 **ompt_sync_region_barrier_explicit** as its *kind* argument and **ompt_scope_begin**
32 as its *endpoint* argument for each occurrence of an *explicit-barrier-wait-begin* event. Similarly, a
33 thread dispatches a registered **ompt_callback_sync_region_wait** callback with
34 **ompt_sync_region_barrier_explicit** as its *kind* argument and **ompt_scope_end** as
35 its *endpoint* argument for each occurrence of an *explicit-barrier-wait-end* event. These callbacks
36 occur in the context of the task that encountered the **barrier** construct and have type signature
37 **ompt_callback_sync_region_t**.

1 A thread dispatches a registered `ompt_callback_cancel` callback with
2 `ompt_cancel_detected` as its *flags* argument for each occurrence of a *cancellation* event in
3 that thread. The callback occurs in the context of the encountering task. The callback has type
4 signature `ompt_callback_cancel_t`.

5 **Restrictions**

6 Restrictions to the `barrier` construct are as follows:

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

11 **Cross References**

- 12 • `ompt_callback_cancel_t`, see [Section 19.5.2.18](#)
- 13 • `ompt_callback_sync_region_t`, see [Section 19.5.2.13](#)
- 14 • `ompt_scope_endpoint_t`, see [Section 19.4.4.11](#)
- 15 • `ompt_sync_region_t`, see [Section 19.4.4.14](#)

16 **15.3.2 Implicit Barriers**

17 This section describes the OMPT events and tool callbacks associated with implicit barriers, which
18 occur at the end of various regions as defined in the description of the constructs to which they
19 correspond. Implicit barriers are task scheduling points. For a description of task scheduling
20 points, associated events, and tool callbacks, see [Section 12.9](#).

21 **Execution Model Events**

22 The *implicit-barrier-begin* event occurs in each implicit task at the beginning of an implicit barrier
23 region.

24 The *implicit-barrier-wait-begin* event occurs when a task begins an interval of active or passive
25 waiting in an implicit barrier region.

26 The *implicit-barrier-wait-end* event occurs when a task ends an interval of active or waiting and
27 resumes execution of an implicit barrier region.

28 The *implicit-barrier-end* event occurs in each implicit task after the barrier synchronization on exit
29 from an implicit barrier region.

30 A *cancellation* event occurs if cancellation is activated at an implicit cancellation point in an
31 implicit barrier region.

1 Tool Callbacks

2 A thread dispatches a registered `ompt_callback_sync_region` callback for each implicit
3 barrier *begin* and *end* event. Similarly, a thread dispatches a registered
4 `ompt_callback_sync_region_wait` callback for each implicit barrier *wait-begin* and
5 *wait-end* event. All callbacks for implicit barrier events execute in the context of the encountering
6 task and have type signature `ompt_callback_sync_region_t`.

7 For the implicit barrier at the end of a worksharing construct, the *kind* argument is
8 `ompt_sync_region_barrier_implicit_workshare`. For the implicit barrier at the end
9 of a `parallel` region, the *kind* argument is
10 `ompt_sync_region_barrier_implicit_parallel`. For an extra barrier added by an
11 OpenMP implementation, the *kind* argument is
12 `ompt_sync_region_barrier_implementation`. For a barrier at the end of a `teams`
13 region, the *kind* argument is `ompt_sync_region_barrier_teams`.

14 A thread dispatches a registered `ompt_callback_cancel` callback with
15 `ompt_cancel_detected` as its *flags* argument for each occurrence of a *cancellation* event in
16 that thread. The callback occurs in the context of the encountering task. The callback has type
17 signature `ompt_callback_cancel_t`.

18 Restrictions

19 Restrictions to implicit barriers are as follows:

- 20 • If a thread is in the state `ompt_state_wait_barrier_implicit_parallel`, a call
21 to `ompt_get_parallel_info` may return a pointer to a copy of the data object
22 associated with the parallel region rather than a pointer to the associated data object itself.
23 Writing to the data object returned by `ompt_get_parallel_info` when a thread is in the
24 `ompt_state_wait_barrier_implicit_parallel` results in unspecified behavior.

25 Cross References

- 26 • `ompt_callback_cancel_t`, see [Section 19.5.2.18](#)
- 27 • `ompt_callback_sync_region_t`, see [Section 19.5.2.13](#)
- 28 • `ompt_cancel_flag_t`, see [Section 19.4.4.26](#)
- 29 • `ompt_scope_endpoint_t`, see [Section 19.4.4.11](#)
- 30 • `ompt_sync_region_t`, see [Section 19.4.4.14](#)

31 15.3.3 Implementation-Specific Barriers

32 An OpenMP implementation can execute implementation-specific barriers that the OpenMP
33 specification does not imply; therefore, no *execution model events* are bound to them. The
34 implementation can handle these barriers like implicit barriers and dispatch all events as for
35 implicit barriers. These callbacks use `ompt_sync_region_barrier_implementation`
36 — or `ompt_sync_region_barrier`, if the implementation cannot make a distinction — as
37 the *kind* argument when they are dispatched.

15.4 taskgroup Construct

Name: <code>taskgroup</code> Category: executable	Association: block Properties: cancellable
--	---

Clauses

`allocate`, `task_reduction`

Binding

The binding task set of a `taskgroup` region is all tasks of the current team that are generated in the region. A `taskgroup` region binds to the innermost enclosing `parallel` region.

Semantics

The `taskgroup` construct specifies a wait on completion of the *taskgroup set* associated with the `taskgroup` region. When a thread encounters a `taskgroup` construct, it starts executing the region.

An implicit task scheduling point occurs at the end of the `taskgroup` region. The current task is suspended at the task scheduling point until all tasks in the *taskgroup set* complete execution.

Execution Model Events

The *taskgroup-begin* event occurs in each thread that encounters the `taskgroup` construct on entry to the `taskgroup` region.

The *taskgroup-wait-begin* event occurs when a task begins an interval of active or passive waiting in a `taskgroup` region.

The *taskgroup-wait-end* event occurs when a task ends an interval of active or passive waiting and resumes execution in a `taskgroup` region.

The *taskgroup-end* event occurs in each thread that encounters the `taskgroup` construct after the `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

1 dispatches a registered `ompt_callback_sync_region_wait` callback with
2 `ompt_sync_region_taskgroup` as its *kind* argument and `ompt_scope_end` as its
3 *endpoint* argument for each occurrence of a *taskgroup-wait-end* event. These callbacks occur in the
4 context of the task that encounters the `taskgroup` construct and have type signature
5 `ompt_callback_sync_region_t`.

6 Cross References

- 7 • `allocate` clause, see [Section 6.6](#)
- 8 • `task_reduction` clause, see [Section 5.5.9](#)
- 9 • Task Scheduling, see [Section 12.9](#)
- 10 • `ompt_callback_sync_region_t`, see [Section 19.5.2.13](#)
- 11 • `ompt_scope_endpoint_t`, see [Section 19.4.4.11](#)
- 12 • `ompt_sync_region_t`, see [Section 19.4.4.14](#)

13 15.5 taskwait Construct

14 Name: <code>taskwait</code>	Association: none
Category: executable	Properties: <i>default</i>

15 Clauses

16 [depend](#), [nowait](#)

17 Binding

18 The `taskwait` region binds to the current task region. The binding thread set of the `taskwait`
19 region is the current team.

20 Semantics

21 The `taskwait` construct specifies a wait on the completion of child tasks of the current task.

22 If no `depend` clause is present on the `taskwait` construct, the current task region is suspended
23 at an implicit task scheduling point associated with the construct. The current task region remains
24 suspended until all child tasks that it generated before the `taskwait` region complete execution.

25 If one or more `depend` clauses are present on the `taskwait` construct and the `nowait` clause is
26 not also present, the behavior is as if these clauses were applied to a `task` construct with an empty
27 associated structured block that generates a *mergeable* and *included task*. Thus, the current task
28 region is suspended until the *predecessor tasks* of this task complete execution.

29 If one or more `depend` clauses are present on the `taskwait` construct and the `nowait` clause is
30 also present, the behavior is as if these clauses were applied to a `task` construct with an empty
31 associated structured block that generates a task for which execution may be deferred. Thus, all
32 *predecessor tasks* of this task must complete execution before any subsequently generated task that
33 depends on this task starts its execution.

Execution Model Events

The *taskwait-begin* event occurs in a thread when it encounters a **taskwait** construct with no **depend** clause 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 region corresponding to a **taskwait** construct with no **depend** clause.

The *taskwait-wait-end* event occurs when a task ends an interval of active or passive waiting and resumes execution from a region corresponding to a **taskwait** construct with no **depend** clause.

The *taskwait-end* event occurs in a thread when it encounters a **taskwait** construct with no **depend** clause after the taskwait synchronization on exit from the **taskwait** region.

The *taskwait-init* event occurs in a thread when it encounters a **taskwait** construct with one or more **depend** clauses on entry to the **taskwait** region.

The *taskwait-complete* event occurs on completion of the dependent task that results from a **taskwait** construct with one or more **depend** clauses, in the context of the thread that executes the dependent task and before any subsequently generated task that depends on the dependent task starts its execution.

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**.

A thread dispatches a registered **ompt_callback_task_create** callback for each occurrence of a *taskwait-init* event in the context of the encountering task. This callback has the type signature **ompt_callback_task_create_t**. In the dispatched callback, (*flags & ompt_task_taskwait*) always evaluates to *true*. If the **nowait** clause is not present, (*flags & ompt_task_underrred*) also evaluates to *true*.

1 A thread dispatches a registered `ompt_callback_task_schedule` callback for each
2 occurrence of a *taskwait-complete* event. This callback has the type signature
3 `ompt_callback_task_schedule_t` with `ompt_taskwait_complete` as its
4 *prior_task_status* argument.

5 Restrictions

6 Restrictions to the `taskwait` construct are as follows:

- 7 • The `mutexinoutset` *dependence-type* may not appear in a `depend` clause on a
8 `taskwait` construct.
- 9 • If the *dependence-type* of a `depend` clause is `depobj` then the dependence objects cannot
10 represent dependences of the `mutexinoutset` dependence type.
- 11 • The `nowait` clause may only appear on a `taskwait` directive if the `depend` clause is
12 present.

13 Cross References

- 14 • `depend` clause, see [Section 15.9.5](#)
- 15 • `nowait` clause, see [Section 15.6](#)
- 16 • `task` directive, see [Section 12.5](#)
- 17 • `ompt_callback_sync_region_t`, see [Section 19.5.2.13](#)
- 18 • `ompt_scope_endpoint_t`, see [Section 19.4.4.11](#)
- 19 • `ompt_sync_region_t`, see [Section 19.4.4.14](#)

20 15.6 nowait Clause

21 Name: <code>nowait</code>	Properties: unique, end-clause
------------------------------	--------------------------------

22 Directives

23 [dispatch](#), [do](#), [for](#), [interop](#), [scope](#), [sections](#), [single](#), [target](#), [target enter](#)
24 [data](#), [target exit data](#), [target update](#), [taskwait](#), [workshare](#)

25 Semantics

26 The `nowait` clause overrides any synchronization that would otherwise occur at the end of a
27 construct. It can also specify that an *interoperability requirement set* includes the *nowait* property.
28 If the construct includes an implicit barrier, the `nowait` clause specifies that the barrier will not
29 occur. For constructs that generate a task, the `nowait` clause specifies that the generated task may
30 be deferred. If the `nowait` clause is not present on the directive then the generated task is an
31 included task (so it executes synchronously in the context of the encountering task). For constructs
32 that generate an *interoperability requirement set*, the `nowait` clause adds the *nowait* property to
33 the set.

Cross References

- `dispatch` directive, see [Section 7.6](#)
- `do` directive, see [Section 11.5.2](#)
- `for` directive, see [Section 11.5.1](#)
- `interop` directive, see [Section 14.1](#)
- `scope` directive, see [Section 11.2](#)
- `sections` directive, see [Section 11.3](#)
- `single` directive, see [Section 11.1](#)
- `target` directive, see [Section 13.8](#)
- `target enter data` directive, see [Section 13.6](#)
- `target exit data` directive, see [Section 13.7](#)
- `target update` directive, see [Section 13.9](#)
- `taskwait` directive, see [Section 15.5](#)
- `workshare` directive, see [Section 11.4](#)

15.7 nogroup Clause

Name: <code>nogroup</code>	Properties: unique
----------------------------	--------------------

Directives

[taskloop](#)

Semantics

The `nogroup` clause overrides any implicit `taskgroup` that would otherwise enclose the construct.

Cross References

- `taskloop` directive, see [Section 12.6](#)

15.8 OpenMP Memory Ordering

This sections describes constructs and clauses in OpenMP that support ordering of memory operations.

15.8.1 *memory-order* Clauses

Clause groups

Properties: unique, exclusive, inarguable	Members: acq_rel , acquire , relaxed , release , seq_cst
--	---

Directives

[atomic](#), [flush](#)

Semantics

The *memory-order* clause grouping defines a set of clauses that indicate the memory ordering requirements for the visibility of the effects of the constructs on which they may be specified.

Cross References

- [atomic](#) directive, see [Section 15.8.5](#)
- [flush](#) directive, see [Section 15.8.6](#)

15.8.2 *atomic* Clauses

Clause groups

Properties: unique, exclusive, inarguable	Members: read , update , write
--	---

Directives

[atomic](#)

Semantics

The *atomic* clause grouping defines a set of clauses that defines the semantics for which a directive enforces atomicity. If a construct accepts the *atomic* clause grouping and no member of the grouping is specified, the effect is as if the **update** clause is specified.

Cross References

- [atomic](#) directive, see [Section 15.8.5](#)

15.8.3 *extended-atomic* Clauses

Clause groups

Properties: unique	Members: capture , compare , fail , weak
---------------------------	---

Directives

`atomic`

Semantics

The *extended-atomic* clause grouping defines a set of clauses that extend the atomicity semantics specified by members of the *atomic* clause grouping. Other than the **fail** clause, they are inarguable; the **fail** clause takes a member of the *memory-order* clause grouping as an argument.

The **capture** clause extends the semantics to capture the value of the variable being updated atomically. The **compare** clause extends the semantics to perform the atomic update conditionally. The **fail** clause extends the semantics to specify the memory ordering requirements for any comparison performed by any atomic conditional update that fails. Its argument overrides any other specified memory ordering. If the **fail** clause is not specified on an atomic conditional update the effect is as if the **fail** clause is specified with a default argument that depends on the effective memory ordering. If the effective memory ordering is **acq_rel**, the default argument is **acquire**. If the effective memory ordering is **release**, the default argument is **relaxed**. For any other effective memory ordering, the default argument is equal to that effective memory ordering. The **weak** clause specifies that the comparison performed by a conditional atomic update may spuriously fail, evaluating to not equal even when the values are equal.

Note – Allowing for spurious failure by specifying a **weak** clause can result in performance gains on some systems when using compare-and-swap in a loop. For cases where a single compare-and-swap would otherwise be sufficient, using a loop over a **weak** compare-and-swap is unlikely to improve performance.

Restrictions

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

- **acq_rel** and **release** cannot be specified as arguments to the **fail** clause.

Cross References

- *atomic* Clauses, see [Section 15.8.2](#)
- *memory-order* Clauses, see [Section 15.8.1](#)
- **atomic** directive, see [Section 15.8.5](#)

15.8.4 memscope Clause

Name: memscope	Properties: unique
-----------------------	--------------------

Arguments

Name	Type	Properties
<i>scope-specifier</i>	Keyword: all , cgroup , device	<i>default</i>

Directives

[atomic](#), [flush](#)

Semantics

The **memscope** clause determines the binding thread set of the region that corresponds to the construct on which it is specified.

If the *scope-specifier* is **device**, the binding thread set consists of all threads on the device. If the *scope-specifier* is **cgroup**, the binding thread set consists of all threads that are executing tasks in the contention group. If the *scope-specifier* is **all**, the binding thread set consists of all threads on all devices.

Unless otherwise stated, the thread-set of any flush operations that are performed in an **atomic** or **flush** region is the same as the binding thread set of the region, as determined by the **memscope** clause.

Restrictions

The restrictions for the **memscope** clause are as follows:

- The binding thread set defined by the *scope-specifier* of the **memscope** clause on an **atomic** construct must be a subset of the atomic scope of the atomically accessed memory.
- The binding thread set defined by the *scope-specifier* of the **memscope** clause on an **atomic** construct must be a subset of all threads that are executing tasks in the contention group if the size of the atomically accessed storage location is not 8, 16, 32, or 64 bits.

Cross References

- **atomic** directive, see [Section 15.8.5](#)
- **flush** directive, see [Section 15.8.6](#)

15.8.5 atomic Construct

Name: atomic Category: executable	Association: block (atomic structured block) Properties: simdizable
---	--

Clause groups

[atomic](#), [extended-atomic](#), [memory-order](#)

Clauses

[hint](#), [memscope](#)

This section uses the terminology and symbols defined for OpenMP Atomic Structured Blocks (see [Section 4.3.1.3](#)).

1 Binding

2 The **memscope** clause determines the binding thread set for an **atomic** region. If the **memscope**
3 clause is not present the behavior is as if the **memscope** clause appeared on the construct with the
4 **cgroup scope-specifier** or, if the size of the atomically accessed storage location is 8, 16, 32, or 64
5 bits, the **device scope-specifier**.

6 Semantics

7 The **atomic** construct ensures that a specific storage location is accessed atomically so that
8 possible simultaneous reads and writes by multiple threads do not result in indeterminate values.
9 An **atomic** region enforces exclusive access with respect to other **atomic** regions that access the
10 same storage location x among all threads in the binding thread set without regard to the teams to
11 which the threads belong.

12 The **atomic** construct with the **read** clause results in an atomic read of the location designated
13 by x . The **atomic** construct with the **write** clause results in an atomic write of the location
14 designated by x . The **atomic** construct with the **update** clause results in an atomic update of the
15 location designated by x using the designated operator or intrinsic. Only the read and write of the
16 location designated by x are performed mutually atomically. The evaluation of $expr$ or $expr-list$
17 need not be atomic with respect to the read or write of the location designated by x . No task
18 scheduling points are allowed between the read and the write of the location designated by x .

19 If the **capture** clause is present, the atomic update is an atomic captured update — an atomic
20 update to the location designated by x using the designated operator or intrinsic while also
21 capturing the original or final value of the location designated by x with respect to the atomic
22 update. The original or final value of the location designated by x is written in the location
23 designated by v based on the base language semantics of structured block or statements of the
24 **atomic** construct. Only the read and write of the location designated by x are performed mutually
25 atomically. Neither the evaluation of $expr$ or $expr-list$, nor the write to the location designated by v ,
26 need be atomic with respect to the read or write of the location designated by x .

27 If the **compare** clause is present, the atomic update is an atomic conditional update. For forms
28 that use an equality comparison, the operation is an atomic compare-and-swap. It atomically
29 compares the value of x to e and writes the value of d into the location designated by x if they are
30 equal. Based on the base language semantics of the associated structured block, the original or final
31 value of the location designated by x is written to the location designated by v , which is allowed to
32 be the same location as designated by e , or the result of the comparison is written to the location
33 designated by r . Only the read and write of the location designated by x are performed mutually
34 atomically. Neither the evaluation of either e or d nor writes to the locations designated by v and r
35 need be atomic with respect to the read or write of the location designated by x .

▼ C / C++ ▼

36 If the **compare** clause is present, forms that use *ordop* are logically an atomic maximum or
37 minimum, but they may be implemented with a compare-and-swap loop with short-circuiting. For
38 forms where *statement* is *cond-expr-stmt*, if the result of the condition implies that the value of x
39 does not change then the update may not occur.

▲ C / C++ ▲

1 If a *memory-order* clause is present, or implicitly provided by a **requires** directive, it specifies
2 the effective memory ordering. Otherwise the effect is as if the **relaxed** memory ordering clause
3 is specified.

4 The **atomic** construct may be used to enforce memory consistency between threads, based on the
5 guarantees provided by Section 1.4.6. A strong flush on the location designated by x is performed
6 on entry to and exit from the atomic operation, ensuring that the set of all atomic operations applied
7 to the same location in a race-free program has a total completion order. If the **write** or **update**
8 clause is specified, the atomic operation is not an atomic conditional update for which the
9 comparison fails, and the effective memory ordering is **release**, **acq_rel**, or **seq_cst**, the
10 strong flush on entry to the atomic operation is also a release flush. If the **read** or **update** clause
11 is specified and the effective memory ordering is **acquire**, **acq_rel**, or **seq_cst** then the
12 strong flush on exit from the atomic operation is also an acquire flush. Therefore, if the effective
13 memory ordering is not **relaxed**, release and/or acquire flush operations are implied and permit
14 synchronization between the threads without the use of explicit **flush** directives.

15 For all forms of the **atomic** construct, any combination of two or more of these **atomic**
16 constructs enforces mutually exclusive access to the locations designated by x among threads in the
17 binding thread set. To avoid data races, all accesses of the locations designated by x that could
18 potentially occur in parallel must be protected with an **atomic** construct.

19 **atomic** regions do not guarantee exclusive access with respect to any accesses outside of
20 **atomic** regions to the same storage location x even if those accesses occur during a **critical**
21 or **ordered** region, while an OpenMP lock is owned by the executing task, or during the
22 execution of a **reduction** clause.

23 However, other OpenMP synchronization can ensure the desired exclusive access. For example, a
24 barrier that follows a series of atomic updates to x guarantees that subsequent accesses do not form
25 a race with the atomic accesses.

26 A compliant implementation may enforce exclusive access between **atomic** regions that update
27 different storage locations. The circumstances under which this occurs are implementation defined.

28 If the storage location designated by x is not size-aligned (that is, if the byte alignment of x is not a
29 multiple of the size of x), then the behavior of the **atomic** region is implementation defined.

30 **Execution Model Events**

31 The *atomic-acquiring* event occurs in the thread that encounters the **atomic** construct on entry to
32 the atomic region before initiating synchronization for the region.

33 The *atomic-acquired* event occurs in the thread that encounters the **atomic** construct after it
34 enters the region, but before it executes the structured block of the **atomic** region.

35 The *atomic-released* event occurs in the thread that encounters the **atomic** construct after it
36 completes any synchronization on exit from the **atomic** region.

1 Tool Callbacks

2 A thread dispatches a registered `ompt_callback_mutex_acquire` callback for each
3 occurrence of an *atomic-acquiring* event in that thread. This callback has the type signature
4 `ompt_callback_mutex_acquire_t`.

5 A thread dispatches a registered `ompt_callback_mutex_acquired` callback for each
6 occurrence of an *atomic-acquired* event in that thread. This callback has the type signature
7 `ompt_callback_mutex_t`.

8 A thread dispatches a registered `ompt_callback_mutex_released` callback with
9 `ompt_mutex_atomic` as the *kind* argument if practical, although a less specific *kind* may be
10 used, for each occurrence of an *atomic-released* event in that thread. This callback has the type
11 signature `ompt_callback_mutex_t` and occurs in the task that encounters the atomic
12 construct.

13 Restrictions

14 Restrictions to the `atomic` construct are as follows:

- 15 • OpenMP constructs may not be encountered during execution of an `atomic` region.
- 16 • If a `capture` or `compare` clause is specified, the *atomic* clause must be `update`.
- 17 • If a `capture` clause is specified but the `compare` clause is not specified, an
18 *update-capture-atomic* structured block must be associated with the construct.
- 19 • If both `capture` and `compare` clauses are specified, a *conditional-update-capture-atomic*
20 structured block must be associated with the construct.
- 21 • If a `compare` clause is specified but the `capture` clause is not specified, a
22 *conditional-update-atomic* structured block must be associated with the construct.
- 23 • If a `write` clause is specified, a *write-atomic* structured block must be associated with the
24 construct.
- 25 • If a `read` clause is specified, a *read-atomic* structured block must be associated with the
26 construct.
- 27 • If the *atomic* clause is `read` then the *memory-order* clause must not be `release`.
- 28 • If the *atomic* clause is `write` then the *memory-order* clause must not be `acquire`.
- 29 • The `weak` clause may only appear if the resulting atomic operation is an atomic conditional
30 update for which the comparison tests for equality.

▼ C / C++ ▼

- 31 • All atomic accesses to the storage locations designated by *x* throughout the program are
32 required to have a compatible type.
- 33 • The `fail` clause may only appear if the resulting atomic operation is an atomic conditional
34 update.

▲ C / C++ ▲

Fortran

- All atomic accesses to the storage locations designated by x throughout the program are required to have the same type and type parameters.
- The **fail** clause may only appear if the resulting atomic operation is an atomic conditional update or an atomic update where *intrinsic-procedure-name* is either **MAX** or **MIN**.

Fortran

Cross References

- **hint** clause, see [Section 15.1.2](#)
- **memscope** clause, see [Section 15.8.4](#)
- **barrier** directive, see [Section 15.3.1](#)
- **critical** directive, see [Section 15.2](#)
- **flush** directive, see [Section 15.8.6](#)
- **requires** directive, see [Section 8.2](#)
- Lock Routines, see [Section 18.9](#)
- OpenMP Atomic Structured Blocks, see [Section 4.3.1.3](#)
- Synchronization Hints, see [Section 15.1](#)
- **ompt_callback_mutex_acquire_t**, see [Section 19.5.2.14](#)
- **ompt_callback_mutex_t**, see [Section 19.5.2.15](#)
- **ompt_mutex_t**, see [Section 19.4.4.17](#)
- **ordered** Construct, see [Section 15.10](#)

15.8.6 flush Construct

Name: <code>flush</code>	Association: none
Category: executable	Properties: <i>default</i>

Arguments

flush (*list*)

Name	Type	Properties
<i>list</i>	list of variable list item type	optional

Clause groups

[memory-order](#)

Clauses

[memscope](#)

1 Binding

2 The **memscope** clause determines the binding thread set for a **flush** region. If the **memscope**
3 clause is not present the behavior is as if the **memscope** clause appeared on the construct with the
4 **device scope-specifier**.

5 Semantics

6 The **flush** construct executes the OpenMP flush operation. This operation makes a thread's
7 temporary view of memory consistent with memory and enforces an order on the memory
8 operations of the variables explicitly specified or implied. Execution of a **flush** region affects the
9 memory and it affects the temporary view of memory of the encountering thread. It does not affect
10 the temporary view of other threads. Other threads in the *thread-set* must themselves execute a
11 flush operation in order to be guaranteed to observe the effects of the flush operation of the
12 encountering thread. See the memory model description in [Section 1.4](#) and the **memscope** clause
13 description in [Section 15.8.4](#) for more details on thread-sets.

14 If neither a *memory-order* clause nor a *list* argument appears on a **flush** construct then the
15 behavior is as if the *memory-order* clause is **seq_cst**.

16 A **flush** construct with the **seq_cst** clause, executed on a given thread, operates as if all data
17 storage blocks that are accessible to the thread are flushed by a strong flush operation. A **flush**
18 construct with a list applies a strong flush operation to the items in the list, and the flush operation
19 does not complete until the operation is complete for all specified list items. An implementation
20 may implement a **flush** construct with a list by ignoring the list and treating it the same as a
21 **flush** construct with the **seq_cst** clause.

22 If no list items are specified, the flush operation has the release and/or acquire flush properties:

- 23 • If the *memory-order* clause is **seq_cst** or **acq_rel**, the flush operation is both a release
24 flush and an acquire flush.
- 25 • If the *memory-order* clause is **release**, the flush operation is a release flush.
- 26 • If the *memory-order* clause is **acquire**, the flush operation is an acquire flush.

▼ C / C++ ▼

27 If a pointer is present in the list, the pointer itself is flushed, not the memory block to which the
28 pointer refers.

29 A **flush** construct without a list corresponds to a call to **atomic_thread_fence**, where the
30 argument is given by the identifier that results from prefixing **memory_order_** to the
31 *memory-order* clause name.

32 For a **flush** construct without a list, the generated **flush** region implicitly performs the
33 corresponding call to **atomic_thread_fence**. The behavior of an explicit call to
34 **atomic_thread_fence** that occurs in the program and does not have the argument
35 **memory_order_consume** is as if the call is replaced by its corresponding **flush** construct.

▲ C / C++ ▲

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 of
 3 type **C_PTR**, the variable is flushed, but the storage that corresponds to that address is not flushed.
 4 If the list item or the subobject of the list item has the **ALLOCATABLE** attribute and has an
 5 allocation status of allocated, the allocated variable is flushed; otherwise the allocation status is
 6 flushed.

7 **Execution Model Events**

8 The *flush* event occurs in a thread that encounters the **flush** construct.

9 **Tool Callbacks**

10 A thread dispatches a registered **ompt_callback_flush** callback for each occurrence of a
 11 *flush* event in that thread. This callback has the type signature **ompt_callback_flush_t**.

12 **Restrictions**

13 Restrictions to the **flush** construct are as follows:

- 14 • If a *memory-order* clause is specified, the *list* argument must not be specified.
- 15 • The *memory-order* clause must not be **relaxed**.

16 **Cross References**

- 17 • **memscope** clause, see [Section 15.8.4](#)
- 18 • **ompt_callback_flush_t**, see [Section 19.5.2.17](#)

19 **15.8.7 Implicit Flushes**

20 Flush operations implied when executing an **atomic** region are described in [Section 15.8.5](#).

21 A **flush** region that corresponds to a **flush** directive with the **release** clause present is
 22 implied at the following locations:

- 23 • During a barrier region;
- 24 • At entry to a **parallel** region;
- 25 • At entry to a **teams** region;
- 26 • At exit from a **critical** region;
- 27 • During an **omp_unset_lock** region;
- 28 • During an **omp_unset_nest_lock** region;
- 29 • During an **omp_fulfill_event** region;


- 1 • Immediately before every task scheduling point;
- 2 • At exit from the task region of each implicit task;
- 3 • At exit from an **ordered** region, if a **threads** clause or a **doacross** clause with a
- 4 **source** dependence type is present, or if no clauses are present; and
- 5 • During a **cancel** region, if the *cancel-var* ICV is *true*.

6 For a **target** construct, the *thread-set* of an implicit release flush that is performed in a target task
7 during the generation of the **target** region and that is performed on exit from the initial task
8 region that implicitly encloses the **target** region consists of the thread that executes the target
9 task and the initial thread that executes the **target** region.

10 A **flush** region that corresponds to a **flush** directive with the **acquire** clause present is
11 implied at the following locations:

- 12 • During a barrier region;
- 13 • At exit from a **teams** region;
- 14 • At entry to a **critical** region;
- 15 • If the region causes the lock to be set, during:
 - 16 – an **omp_set_lock** region;
 - 17 – an **omp_test_lock** region;
 - 18 – an **omp_set_nest_lock** region; and
 - 19 – an **omp_test_nest_lock** region;
- 20 • Immediately after every task scheduling point;
- 21 • At entry to the task region of each implicit task;
- 22 • At entry to an **ordered** region, if a **threads** clause or a **doacross** clause with a **sink**
- 23 dependence type is present, or if no clauses are present; and
- 24 • Immediately before a cancellation point, if the *cancel-var* ICV is *true* and cancellation has
- 25 been activated.

26 For a **target** construct, the *thread-set* of an implicit acquire flush that is performed in a target
27 task following the generation of the **target** region or that is performed on entry to the initial task
28 region that implicitly encloses the **target** region consists of the thread that executes the target
29 task and the initial thread that executes the **target** region.



31 Note – A **flush** region is not implied at the following locations:

- 32 • At entry to worksharing regions; and

- At entry to or exit from **masked** regions.

The synchronization behavior of implicit flushes is as follows:

- When a thread executes an **atomic** region for which the corresponding construct has the **release**, **acq_rel**, or **seq_cst** clause and specifies an atomic operation that starts a given release sequence, the release flush that is performed on entry to the atomic operation synchronizes with an acquire flush that is performed by a different thread and has an associated atomic operation that reads a value written by a modification in the release sequence.
- When a thread executes an **atomic** region for which the corresponding construct has the **acquire**, **acq_rel**, or **seq_cst** clause and specifies an atomic operation that reads a value written by a given modification, a release flush that is performed by a different thread and has an associated release sequence that contains that modification synchronizes with the acquire flush that is performed on exit from the atomic operation.
- When a thread executes a **critical** region that has a given name, the behavior is as if the release flush performed on exit from the region synchronizes with the acquire flush performed on entry to the next **critical** region with the same name that is performed by a different thread, if it exists.
- When a thread team executes a **barrier** region, the behavior is as if the release flush performed by each thread within the region, and the release flush performed by any other thread upon fulfilling the *allow-completion* event for a detachable task bound to the binding parallel region of the region, synchronizes with the acquire flush performed by all other threads within the region.
- When a thread executes a **taskwait** region that does not result in the creation of a dependent task and the task that encounters the corresponding **taskwait** construct has at least one child task, the behavior is as if each thread that executes a child task that is generated before the **taskwait** region performs a release flush upon completion of the associated structured block of the child task that synchronizes with an acquire flush performed in the **taskwait** region. If the child task is detachable, the thread that fulfills its *allow-completion* event performs a release flush upon fulfilling the event that synchronizes with the acquire flush performed in the **taskwait** region.
- When a thread executes a **taskgroup** region, the behavior is as if each thread that executes a remaining descendent task performs a release flush upon completion of the associated structured block of the descendent task that synchronizes with an acquire flush performed on exit from the **taskgroup** region. If the descendent task is detachable, the thread that fulfills its *allow-completion* event performs a release flush upon fulfilling the event that synchronizes with the acquire flush performed in the **taskgroup** region.
- When a thread executes an **ordered** region that does not arise from a stand-alone **ordered** directive, the behavior is as if the release flush performed on exit from the region

1 synchronizes with the acquire flush performed on entry to an **ordered** region encountered
2 in the next logical iteration to be executed by a different thread, if it exists.

- 3 • When a thread executes an **ordered** region that arises from a stand-alone **ordered**
4 directive, the behavior is as if the release flush performed in the **ordered** region from a
5 given source iteration synchronizes with the acquire flush performed in all **ordered** regions
6 executed by a different thread that are waiting for dependences on that iteration to be satisfied.
- 7 • When a thread team begins execution of a **parallel** region, the behavior is as if the release
8 flush performed by the primary thread on entry to the **parallel** region synchronizes with
9 the acquire flush performed on entry to each implicit task that is assigned to a different thread.
- 10 • When an initial thread begins 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 generating
12 thread in the target task synchronizes with the acquire flush performed by the initial thread on
13 entry to its initial task region.
- 14 • When an initial thread completes execution of a **target** region that is generated by a
15 different thread from a target task, the behavior is as if the release flush performed by the
16 initial thread on exit from its initial task region synchronizes with the acquire flush performed
17 by the generating thread in the target task.
- 18 • When a thread encounters a **teams** construct, the behavior is as if the release flush
19 performed by the thread on entry to the **teams** region synchronizes with the acquire flush
20 performed on entry to each initial task that is executed by a different initial thread that
21 participates in the execution of the **teams** region.
- 22 • When a thread that encounters a **teams** construct reaches the end of the **teams** region, the
23 behavior is as if the release flush performed by each different participating initial thread at
24 exit from its initial task synchronizes with the acquire flush performed by the thread at exit
25 from the **teams** region.
- 26 • When a task generates an explicit task that begins execution on a different thread, the
27 behavior is as if the thread that is executing the generating task performs a release flush that
28 synchronizes with the acquire flush performed by the thread that begins to execute the
29 explicit task.
- 30 • When an undeferred task completes execution on a given thread that is different from the
31 thread on which its generating task is suspended, the behavior is as if a release flush
32 performed by the thread that completes execution of the associated structured block of the
33 undeferred task synchronizes with an acquire flush performed by the thread that resumes
34 execution of the generating task.
- 35 • When a dependent task with one or more predecessor tasks begins execution on a given
36 thread, the behavior is as if each release flush performed by a different thread on completion
37 of the associated structured block of a predecessor task synchronizes with the acquire flush
38 performed by the thread that begins to execute the dependent task. If the predecessor task is
39 detachable, the thread that fulfills its *allow-completion* event performs a release flush upon

- 1 fulfilling the event that synchronizes with the acquire flush performed when the dependent
2 task begins to execute.
- 3 • When a task begins execution on a given thread and it is mutually exclusive with respect to
4 another sibling task that is executed by a different thread, the behavior is as if each release
5 flush performed on completion of the sibling task synchronizes with the acquire flush
6 performed by the thread that begins to execute the task.
 - 7 • When a thread executes a **cancel** region, the *cancel-var* ICV is *true*, and cancellation is not
8 already activated for the specified region, the behavior is as if the release flush performed
9 during the **cancel** region synchronizes with the acquire flush performed by a different
10 thread immediately before a cancellation point in which that thread observes cancellation was
11 activated for the region.
 - 12 • When a thread executes an **omp_unset_lock** region that causes the specified lock to be
13 unset, the behavior is as if a release flush is performed during the **omp_unset_lock**
14 region that synchronizes with an acquire flush that is performed during the next
15 **omp_set_lock** or **omp_test_lock** region to be executed by a different thread that
16 causes the specified lock to be set.
 - 17 • When a thread executes an **omp_unset_nest_lock** region that causes the specified
18 nested lock to be unset, the behavior is as if a release flush is performed during the
19 **omp_unset_nest_lock** region that synchronizes with an acquire flush that is performed
20 during the next **omp_set_nest_lock** or **omp_test_nest_lock** region to be
21 executed by a different thread that causes the specified nested lock to be set.

22 15.9 OpenMP Dependences

23 This section describes constructs and clauses in OpenMP that support the specification and
24 enforcement of dependences. OpenMP supports two kinds of dependences: *task dependences*,
25 which enforce orderings between tasks; and *cross-iteration dependences*, which enforce orderings
26 between loop iterations.

27 15.9.1 *task-dependence-type* Modifier

28 Modifiers

Name	Modifies	Type	Properties
29 <i>task-dependence-type</i>	<i>locator-list</i>	Keyword: depobj , in , inout , inoutset , mutexinoutset , out	required, ultimate

30 Clauses

31 **depend**, **update**

Semantics

OpenMP clauses that are related to task dependences use the *task-dependence-type* modifier to identify the type of dependence relevant to that clause. The effect of the type of dependence is associated with locator list items as described with the **depend** clause, see [Section 15.9.5](#).

Cross References

- **depend** clause, see [Section 15.9.5](#)
- **update** clause, see [Section 15.9.3](#)

15.9.2 Depend Objects

OpenMP depend objects 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.

15.9.3 update Clause

Name: <code>update</code>	Properties: unique
----------------------------------	---------------------------

Arguments

Name	Type	Properties
<i>task-dependence-type</i>	Keyword: depobj , in , inout , inoutset , mutexinoutset , out	<i>default</i>

Directives

[depobj](#)

Semantics

The **update** clause sets the dependence type of an OpenMP depend object to *task-dependence-type*.

Restrictions

Restrictions to the **update** clause are as follows:

- *task-dependence-type* must not be **depobj**.

Cross References

- **depobj** directive, see [Section 15.9.4](#)
- **task-dependence-type** modifier, see [Section 15.9.1](#)

15.9.4 depobj Construct

Name: <code>depobj</code> Category: executable	Association: none Properties: <i>default</i>
---	---

Arguments

`depobj` (*depend-object*)

Name	Type	Properties
<i>depend-object</i>	variable of depend type	<i>default</i>

Clauses

`depend`, `destroy`, `update`

Clause set

Properties: unique, required, exclusive	Members: <code>depend</code> , <code>destroy</code> , <code>update</code>
--	--

Binding

The binding thread set for a `depobj` region is the encountering thread.

Semantics

The `depobj` construct initializes, updates or destroys an OpenMP depend object. If a `depend` clause is specified, the state of *depend-object* is set to initialized and *depend-object* is set to represent the dependence that the `depend` clause specifies. If an `update` clause is specified, *depend-object* is updated to represent the new dependence type. If a `destroy` clause is specified, the state of *depend-object* is set to uninitialized.

Restrictions

Restrictions to the `depobj` construct are as follows:

- A `depend` clause on a `depobj` construct must only specify one locator.
- The state of *depend-object* must be uninitialized if a `depend` clause is specified.
- The state of *depend-object* must be initialized if a `destroy` clause or `update` clause is specified.
- If the *depend-object* represents a dependence for the `omp_all_memory` locator, an `update` clause must specify either an `out` or `inout` *task-dependence-type*.

Cross References

- `depend` clause, see [Section 15.9.5](#)
- `destroy` clause, see [Section 3.5](#)
- `update` clause, see [Section 15.9.3](#)
- `task-dependence-type` modifier, see [Section 15.9.1](#)

15.9.5 depend Clause

Name: <code>depend</code>	Properties: <i>default</i>
----------------------------------	-----------------------------------

Arguments

Name	Type	Properties
<i>locator-list</i>	list of locator list item type	<i>default</i>

Modifiers

Name	Modifies	Type	Properties
<i>task-dependence-type</i>	<i>locator-list</i>	Keyword: depobj , in , inout , inoutset , mutexinoutset , out	required, ultimate
<i>iterator</i>	<i>locator-list</i>	Complex, name: iterator Arguments: <i>iterator-specifier</i> OpenMP expression (repeatable)	unique

Directives

`depobj`, `interop`, `target`, `target enter data`, `target exit data`, `target update`, `task`, `taskwait`

Semantics

The **depend** clause enforces additional constraints on the scheduling of tasks. These constraints establish dependences only between sibling tasks. Task dependences are derived from the *task-dependence-type* and the list items.

The storage location of a list item matches the storage location of another list item if they have the same storage location, or if any of the list items is **omp_all_memory**.

For the **in** *task-dependence-type*, if the storage location of at least one of the list items matches the storage location of a list item appearing in a **depend** clause with an **out**, **inout**, **mutexinoutset**, or **inoutset** *task-dependence-type* on a construct from which a sibling task was previously generated, then the generated task will be a dependent task of that sibling task.

For the **out** and **inout** *task-dependence-types*, if the storage location of at least one of the list items matches the storage location of a list item appearing in a **depend** clause with an **in**, **out**, **inout**, **mutexinoutset**, or **inoutset** *task-dependence-type* on a construct from which a sibling task was previously generated, then the generated task will be a dependent task of that sibling task.

For the **mutexinoutset** *task-dependence-type*, if the storage location of at least one of the list items matches the storage location of a list item appearing in a **depend** clause with an **in**, **out**, **inout**, or **inoutset** *task-dependence-type* on a construct from which a sibling task was previously generated, then the generated task will be a dependent task of that sibling task.

1 If a list item appearing in a **depend** clause with a **mutexinoutset** *task-dependence-type* on a
2 task generating construct matches a list item appearing in a **depend** clause with a
3 **mutexinoutset** *task-dependence-type* on a different task generating construct, and both
4 constructs generate sibling tasks, the sibling tasks will be mutually exclusive tasks.

5 For the **inoutset** *task-dependence-type*, if the storage location of at least one of the list items
6 matches the storage location of a list item appearing in a **depend** clause with an **in**, **out**, **inout**,
7 or **mutexinoutset** *task-dependence-type* on a construct from which a sibling task was
8 previously generated, then the generated task will be a dependent task of that sibling task.

9 When the *task-dependence-type* is **depobj**, the task dependences are derived from the
10 dependences represented by the depend objects specified in the **depend** clause as if the **depend**
11 clauses of the **depobj** constructs were specified in the current construct.

12 The list items that appear in the **depend** clause may reference any *iterators-identifier* defined in its
13 *iterator* modifier.

14 The list items that appear in the **depend** clause may include array sections or the
15 **omp_all_memory** reserved locator.

▼ Fortran ▼

16 If a list item has the **ALLOCATABLE** attribute and its allocation status is unallocated, the behavior
17 is unspecified. If a list item has the **POINTER** attribute and its association status is disassociated or
18 undefined, the behavior is unspecified.

▲ Fortran ▲

▼ C / C++ ▼

19 The list items that appear in a **depend** clause may use shape-operators.

▲ C / C++ ▲

20 ▼

21 **Note** – The enforced task dependence establishes a synchronization of memory accesses
22 performed by a dependent task with respect to accesses performed by the predecessor tasks.
23 However, the programmer must properly synchronize with respect to other concurrent accesses that
24 occur outside of those tasks.

▲

26 Execution Model Events

27 The *task-dependences* event occurs in a thread that encounters a task generating construct or a
28 **taskwait** construct with a **depend** clause immediately after the *task-create* event for the new
29 task or the *taskwait-init* event.

30 The *task-dependence* event indicates an unfulfilled dependence for the generated task. This event
31 occurs in a thread that observes the unfulfilled dependence before it is satisfied.

1 Tool Callbacks

2 A thread dispatches the `ompt_callback_dependences` callback for each occurrence of the
3 *task-dependences* event to announce its dependences with respect to the list items in the `depend`
4 clause. This callback has type signature `ompt_callback_dependences_t`.

5 A thread dispatches the `ompt_callback_task_dependence` callback for a *task-dependence*
6 event to report a dependence between a predecessor task (*src_task_data*) and a dependent task
7 (*sink_task_data*). This callback has type signature `ompt_callback_task_dependence_t`.

8 Restrictions

9 Restrictions to the `depend` clause are as follows:

- 10 • List items, other than reserved locators, used in `depend` clauses of the same task or sibling
11 tasks must indicate identical storage locations or disjoint storage locations.
- 12 • List items used in `depend` clauses cannot be zero-length array sections.
- 13 • The `omp_all_memory` reserved locator can only be used in a `depend` clause with an `out`
14 or `inout` *task-dependence-type*.
- 15 • Array sections cannot be specified in `depend` clauses with the `depobj`
16 *task-dependence-type*.
- 17 • List items used in `depend` clauses with the `depobj` *task-dependence-type* must be
18 expressions of the OpenMP `depend` type that correspond to depend objects in the initialized
19 state.
- 20 • List items that are expressions of the OpenMP `depend` type can only be used in `depend`
21 clauses with the `depobj` *task-dependence-type*.

▼ Fortran ▼

- 22 • A common block name cannot appear in a `depend` clause.

▲ Fortran ▲

▼ C / C++ ▼

- 23 • A bit-field cannot appear in a `depend` clause.

▲ C / C++ ▲

Cross References

- `depobj` directive, see [Section 15.9.4](#)
- `interop` directive, see [Section 14.1](#)
- `target` directive, see [Section 13.8](#)
- `target enter data` directive, see [Section 13.6](#)
- `target exit data` directive, see [Section 13.7](#)
- `target update` directive, see [Section 13.9](#)
- `task` directive, see [Section 12.5](#)
- `taskwait` directive, see [Section 15.5](#)
- Array Sections, see [Section 3.2.5](#)
- Array Shaping, see [Section 3.2.4](#)
- `iterator` modifier, see [Section 3.2.6](#)
- `task-dependence-type` modifier, see [Section 15.9.1](#)
- `ompt_callback_dependences_t`, see [Section 19.5.2.8](#)
- `ompt_callback_task_dependence_t`, see [Section 19.5.2.9](#)

15.9.6 doacross Clause

Name: <code>doacross</code>	Properties: required
------------------------------------	-----------------------------

Arguments

Name	Type	Properties
<i>vector</i>	loop-iteration vector	<i>default</i>

Modifiers

Name	Modifies	Type	Properties
<i>dependence-type</i>	<i>vector</i>	Keyword: sink, source	required

Directives

[ordered-standalone](#)

Semantics

The **doacross** clause identifies cross-iteration dependences that imply additional constraints on the scheduling of loop iterations. These constraints establish dependences only between loop iterations.

The **source dependence-type** specifies the satisfaction of cross-iteration dependences that arise from the current iteration. If the **source dependence-type** is specified then the *vector* argument is optional; if *vector* is omitted, it is assumed to be **omp_cur_iteration**.

The **sink dependence-type** specifies a cross-iteration dependence, where *vector* indicates the iteration that satisfies the dependence. If *vector* does not occur in the iteration space, the **doacross** clause is ignored. If all **doacross** clauses on an **ordered** construct are ignored then the construct is ignored.

Note – If the **sink dependence-type** is specified for a *vector* that does not indicate an earlier iteration of the logical iteration space, deadlock may occur.

Restrictions

Restrictions to the **doacross** clause are as follows:

- If *vector* is specified without the **omp_cur_iteration** keyword and it has n dimensions, the innermost loop-associated construct that encloses the construct on which the clause appears must specify an **ordered** clause for which the parameter value equals n .
- If *vector* is specified with the **omp_cur_iteration** keyword and with **sink** as the *dependence-type* then it must be **omp_cur_iteration - 1**.
- If *vector* is specified with **source** as the *dependence-type* then it must be **omp_cur_iteration**.
- For each element of *vector* for which the **sink dependence-type** is specified, if the loop iteration variable var_i has an integral or pointer type, the i^{th} expression of *vector* must be computable without overflow in that type for any value of var_i that can encounter the construct on which the **doacross** clause appears.

C++

- For each element of *vector* for which the **sink dependence-type** is specified, if the loop iteration variable var_i is of a random access iterator type other than pointer type, the i^{th} expression of *vector* must be computable without overflow in the type that would be used by **std::distance** applied to variables of the type of var_i for any value of var_i that can encounter the construct on which the **doacross** clause appears.

C++

Cross References

- **ordered** clause, see [Section 4.4.4](#)
- **ordered** directive, see [Section 15.10.1](#)
- OpenMP Loop-Iteration Spaces and Vectors, see [Section 4.4.2](#)

15.10 ordered Construct

This section describes two forms for the **ordered** construct, the stand-alone **ordered** construct and the block-associated **ordered** construct. Both forms include the execution model events, tool callbacks, and restrictions listed in this section.

Execution Model Events

The *ordered-acquiring* event occurs in the task that encounters the **ordered** construct on entry to the ordered region before it initiates synchronization for the region.

The *ordered-acquired* event occurs in the task that encounters the **ordered** construct after it enters the region, but before it executes the structured block of the **ordered** region.

The *ordered-released* event occurs in the task that encounters the **ordered** construct after it completes any synchronization on exit from the **ordered** region.

Tool Callbacks

A thread dispatches a registered **ompt_callback_mutex_acquire** callback for each occurrence of an *ordered-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 an *ordered-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 with **ompt_mutex_ordered** as the *kind* argument if practical, although a less specific kind may be used, for each occurrence of an *ordered-released* event in that thread. This callback has the type signature **ompt_callback_mutex_t** and occurs in the task that encounters the construct.

Restrictions

- The construct that corresponds to the binding region of an **ordered** region must specify an **ordered** clause.
- The construct that corresponds to the binding region of an **ordered** region must not specify a **reduction** clause with the **inscan** modifier.
- The regions of a stand-alone **ordered** construct and a block-associated **ordered** construct must not have the same binding region.

Cross References

- `ompt_callback_mutex_acquire_t`, see [Section 19.5.2.14](#)
- `ompt_callback_mutex_t`, see [Section 19.5.2.15](#)

15.10.1 Stand-alone ordered Construct

Name: <code>ordered</code>	Association: none
Category: executable	Properties: <i>default</i>

Clauses

[doacross](#)

Binding

The binding thread set for a stand-alone **ordered** region is the current team. A stand-alone **ordered** region binds to the innermost enclosing worksharing-loop region.

Semantics

The stand-alone **ordered** construct specifies that execution must not violate cross-iteration dependences as specified in the **doacross** clauses that appear on the construct. When a thread that is executing an iteration encounters a **ordered** construct with one or more **doacross** clauses for which the **sink dependence-type** is specified, the thread waits until its dependences on all valid iterations specified by the **doacross** clauses are satisfied before it continues execution. A specific dependence is satisfied when a thread that is executing the corresponding iteration encounters an **ordered** construct with a **doacross** clause for which the **source dependence-type** is specified.

Execution Model Events

The *doacross-sink* event occurs in the task that encounters an **ordered** construct for each **doacross** clause for which the **sink dependence-type** is specified after the dependence is fulfilled.

The *doacross-source* event occurs in the task that encounters an **ordered** construct with a **doacross** clause for which the **source dependence-type** is specified before signaling that the dependence has been fulfilled.

Tool Callbacks

A thread dispatches a registered **ompt_callback_dependences** callback with all vector entries listed as **ompt_dependence_type_sink** in the *deps* argument for each occurrence of a *doacross-sink* event in that thread. A thread dispatches a registered **ompt_callback_dependences** callback with all vector entries listed as **ompt_dependence_type_source** in the *deps* argument for each occurrence of a *doacross-source* event in that thread. These callbacks have the type signature **ompt_callback_dependences_t**.

Restrictions

Additional restrictions to the stand-alone **ordered** construct are as follows:

- At most one **doacross** clause may appear on the construct with **source** as the *dependence-type*.
- All **doacross** clauses that appear on the construct must specify the same *dependence-type*.
- The construct must not be an orphaned construct.

Cross References

- **doacross** clause, see [Section 15.9.6](#)
- Worksharing-Loop Constructs, see [Section 11.5](#)
- `ompt_callback_dependences_t`, see [Section 19.5.2.8](#)

15.10.2 Block-associated **ordered** Construct

Name: <code>ordered</code> Category: executable	Association: block Properties: simdizable, thread-limiting
--	---

Clause groups

[parallelization-level](#)

Binding

The binding thread set for a block-associated **ordered** region is the current team. A block-associated **ordered** region binds to the innermost enclosing worksharing-loop, **simd** or worksharing-loop SIMD region.

Semantics

If no clauses are specified, the effect is as if the **threads** *parallelization-level* clause was 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 the **simd** *parallelization-level* clause is specified, the **ordered** regions encountered by any thread will execute one at a time in the order of the loop iterations. With either *parallelization-level*, execution of code outside the region for different iterations can run in parallel; execution of that code within the same iteration must observe any constraints imposed by the base-language semantics.

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 a block-associated **ordered** construct, it waits at the beginning of the **ordered** region until execution of all **ordered** regions that belong to all previous iterations has completed. **ordered** regions that bind to different regions execute independently of each other.

Restrictions

Additional restrictions to the block-associated **ordered** construct are as follows:

- The construct is simdizable only if the **simd** *parallelization-level* is specified.
- If the **simd** *parallelization-level* is specified, the binding region must be a **simd** region or one that corresponds to a combined or composite construct for which the **simd** construct is a leaf construct.
- If the **threads** *parallelization-level* is specified, the binding region must be a *worksharing-loop* region or one that corresponds to a combined or composite construct for which the *worksharing-loop* is a leaf construct.
- If the **threads** *parallelization-level* is specified and the binding region corresponds to a combined or composite construct then **simd** construct must not be a leaf construct unless the **simd** *parallelization-level* is also specified.
- During execution of the logical iteration of a loop-associated construct, a thread must not execute more than one block-associated **ordered** region that binds to the corresponding region of the loop-associated construct.
- An **ordered** clause with a parameter value equal to one must appear on the construct that corresponds to the binding region.

Cross References

- *parallelization-level* Clauses, see [Section 15.10.3](#)
- **ordered** clause, see [Section 4.4.4](#)
- **simd** directive, see [Section 10.5](#)
- Worksharing-Loop Constructs, see [Section 11.5](#)

15.10.3 *parallelization-level* Clauses

Clause groups

Properties: unique, inarguable	Members: simd , threads
--------------------------------	---------------------------------------

Directives

ordered-blockassoc

Semantics

The *parallelization-level* clause grouping defines a set of clauses that indicate the level of parallelization with which to associate a construct.

Cross References

- **ordered** directive, see [Section 15.10.2](#)

16 Cancellation Constructs

This chapter defines constructs related to cancellation of OpenMP regions.

16.1 `cancel` Construct

Name: <code>cancel</code> Category: executable	Association: none Properties: <i>default</i>
---	---

Clauses

`if`, `do`, `for`, `parallel`, `sections`, `taskgroup`

Additional information

The *cancel-directive-name* clause set consists of the *directive-name* of each directive that has the cancellable property (i.e., *directive-name* for the worksharing-loop construct, `parallel`, `sections` and `taskgroup`). This clause set has the required, unique and exclusive properties.

Binding

The binding thread set of the `cancel` region is the current team. The binding region of the `cancel` region is the innermost enclosing region of the type that corresponds to *cancel-directive-name*.

Semantics

The `cancel` construct activates cancellation of the innermost enclosing region of the type specified by *cancel-directive-name*, which must be the *directive-name* of a cancellable construct. Cancellation of the binding region is activated only if the *cancel-var* ICV is *true*, in which case the `cancel` construct causes the encountering task to continue execution at the end of the binding region if *cancel-directive-name* is not `taskgroup`. If the *cancel-var* ICV is *true* and *cancel-directive-name* is `taskgroup`, the encountering task continues execution at the end of the current task region. If the *cancel-var* ICV is *false*, the `cancel` construct is ignored.

Threads check for active cancellation only at cancellation points that are implied at the following locations:

- `cancel` regions;
- `cancellation point` regions;
- `barrier` regions;

- 1 • at the end of a worksharing-loop construct with a **nowait** clause and for which the same list
- 2 item appears in both **firstprivate** and **lastprivate** clauses; and
- 3 • implicit barrier regions.

4 When a thread reaches one of the above cancellation points and if the *cancel-var* ICV is *true*, then:

- 5 • If the thread is at a **cancel** or **cancellation point** region and *cancel-directive-name*
- 6 is not **taskgroup**, the thread continues execution at the end of the canceled region if
- 7 cancellation has been activated for the innermost enclosing region of the type specified.
- 8 • If the thread is at a **cancel** or **cancellation point** region and *cancel-directive-name*
- 9 is **taskgroup**, the encountering task checks for active cancellation of all of the *taskgroup*
- 10 *sets* to which the encountering task belongs, and continues execution at the end of the current
- 11 task region if cancellation has been activated for any of the *taskgroup sets*.
- 12 • If the encountering task is at a barrier region or at the end of a worksharing-loop construct
- 13 with a **nowait** clause and for which the same list item appears in both **firstprivate**
- 14 and **lastprivate** clauses, the encountering task checks for active cancellation of the
- 15 innermost enclosing **parallel** region. If cancellation has been activated, then the
- 16 encountering task continues execution at the end of the canceled region.

17 When cancellation of tasks is activated through a **cancel** construct with **taskgroup** for

18 *cancel-directive-name*, the tasks that belong to the *taskgroup set* of the innermost enclosing

19 **taskgroup** region will be canceled. The task that encountered that construct continues execution

20 at the end of its task region, which implies completion of that task. Any task that belongs to the

21 innermost enclosing **taskgroup** and has already begun execution must run to completion or until

22 a cancellation point is reached. Upon reaching a cancellation point and if cancellation is active, the

23 task continues execution at the end of its task region, which implies the completion of the task. Any

24 task that belongs to the innermost enclosing **taskgroup** and that has not begun execution may be

25 discarded, which implies its completion.

26 When cancellation of tasks is activated through a **cancel** construct with *cancel-directive-name*

27 other than **taskgroup**, each thread of the binding thread set resumes execution at the end of the

28 canceled region if a cancellation point is encountered. If the canceled region is a parallel region,

29 any tasks that have been created by a **task** or a **taskloop** construct and their descendent tasks

30 are canceled according to the above **taskgroup** cancellation semantics. If the canceled region is

31 not a parallel region, no task cancellation occurs.

32  C++ 

The usual C++ rules for object destruction are followed when cancellation is performed.

33  C++ 

34  Fortran 

All private objects or subobjects with **ALLOCATABLE** attribute that are allocated inside the

canceled construct are deallocated.

 Fortran 

1 If the canceled construct contains a reduction-scoping or **lastprivate** clause, the final values of
2 the list items that appeared in those clauses are undefined.

3 When an **if** clause is present on a **cancel** construct and the **if** expression evaluates to *false*, the
4 **cancel** construct does not activate cancellation. The cancellation point associated with the
5 **cancel** construct is always encountered regardless of the value of the **if** expression.

6
7 **Note** – The programmer is responsible for releasing locks and other synchronization data
8 structures that might cause a deadlock when a **cancel** construct is encountered and blocked
9 threads cannot be canceled. The programmer is also responsible for ensuring proper
10 synchronizations to avoid deadlocks that might arise from cancellation of OpenMP regions that
11 contain OpenMP synchronization constructs.
12

13 Execution Model Events

14 If a task encounters a **cancel** construct that will activate cancellation then a *cancel* event occurs.

15 A *discarded-task* event occurs for any discarded tasks.

16 Tool Callbacks

17 A thread dispatches a registered **ompt_callback_cancel** callback for each occurrence of a
18 *cancel* event in the context of the encountering task. This callback has type signature
19 **ompt_callback_cancel_t; (flags & ompt_cancel_activated)** always evaluates to
20 *true* in the dispatched callback; **(flags & ompt_cancel_parallel)** evaluates to *true* in the
21 dispatched callback if *cancel-directive-name* is **parallel**;
22 **(flags & ompt_cancel_sections)** evaluates to *true* in the dispatched callback if
23 *cancel-directive-name* is **sections**; **(flags & ompt_cancel_loop)** evaluates to *true* in the
24 dispatched callback if *cancel-directive-name* is **for** or **do**; and
25 **(flags & ompt_cancel_taskgroup)** evaluates to *true* in the dispatched callback if
26 *cancel-directive-name* is **taskgroup**.

27 A thread dispatches a registered **ompt_callback_cancel** callback with the *ompt_data_t*
28 associated with the discarded task as its *task_data* argument and
29 **ompt_cancel_discarded_task** as its *flags* argument for each occurrence of a
30 *discarded-task* event. The callback occurs in the context of the task that discards the task and has
31 type signature **ompt_callback_cancel_t**.

32 Restrictions

33 Restrictions to the **cancel** construct are as follows:

- 34 • The behavior for concurrent cancellation of a region and a region nested within it is
35 unspecified.
- 36 • If *cancel-directive-name* is **taskgroup**, the **cancel** construct must be closely nested
37 inside a **task** or a **taskloop** construct and the **cancel** region must be closely nested
38 inside a **taskgroup** region.

- 1 • If *cancel-directive-name* is not **taskgroup**, the **cancel** construct must be closely nested
2 inside an OpenMP construct that matches *cancel-directive-name*.
- 3 • A worksharing construct that is canceled must not have a **nowait** clause or a **reduction**
4 clause with a user-defined reduction that uses **omp_orig** in the *initializer-expr* of the
5 corresponding **declare reduction** directive.
- 6 • A worksharing-loop construct that is canceled must not have an **ordered** clause or a
7 **reduction** clause with the **inscan** modifier.
- 8 • When cancellation is active for a **parallel** region, a thread in the team that binds to that
9 region may not be executing or encounter a worksharing construct with an **ordered** clause,
10 a **reduction** clause with the **inscan** modifier or a **reduction** clause with a
11 user-defined reduction that uses **omp_orig** in the *initializer-expr* of the corresponding
12 **declare reduction** directive.
- 13 • When cancellation is active for a **parallel** region, a thread in the team that binds to that
14 region may not be executing or encounter a **scope** construct with a **reduction** clause
15 with a user-defined reduction that uses **omp_orig** in the *initializer-expr* of the
16 corresponding **declare reduction** directive.
- 17 • During execution of a construct that may be subject to cancellation, a thread must not
18 encounter an orphaned cancellation point. That is, a cancellation point must only be
19 encountered within that construct and must not be encountered elsewhere in its region.

20 Cross References

- 21 • **firstprivate** clause, see [Section 5.4.4](#)
- 22 • **if** clause, see [Section 3.4](#)
- 23 • **nowait** clause, see [Section 15.6](#)
- 24 • **ordered** clause, see [Section 4.4.4](#)
- 25 • **private** clause, see [Section 5.4.3](#)
- 26 • **reduction** clause, see [Section 5.5.8](#)
- 27 • **barrier** directive, see [Section 15.3.1](#)
- 28 • **cancellation point** directive, see [Section 16.2](#)
- 29 • **declare reduction** directive, see [Section 5.5.11](#)
- 30 • **do** directive, see [Section 11.5.2](#)
- 31 • **for** directive, see [Section 11.5.1](#)
- 32 • **parallel** directive, see [Section 10.2](#)
- 33 • **sections** directive, see [Section 11.3](#)
- 34 • **task** directive, see [Section 12.5](#)

- `taskgroup` directive, see [Section 15.4](#)
- `cancel-var` ICV, see [Table 2.1](#)
- `omp_get_cancellation`, see [Section 18.2.8](#)
- `ompt_callback_cancel_t`, see [Section 19.5.2.18](#)
- `ompt_cancel_flag_t`, see [Section 19.4.4.26](#)

16.2 cancellation point Construct

Name: <code>cancellation point</code> Category: executable	Association: none Properties: <i>default</i>
---	---

Clauses

[do](#), [for](#), [parallel](#), [sections](#), [taskgroup](#)

Additional information

The *cancel-directive-name* clause set consists of the *directive-name* of each directive that has the cancellable property (i.e., *directive-name* for the worksharing-loop construct, **parallel**, **sections** and **taskgroup**). This clause set has the required, unique and exclusive properties.

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 that corresponds to *cancel-directive-name*.

Semantics

The **cancellation point** construct 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 by *cancel-directive-name*, which must be the *directive-name* of a cancellable construct, has been activated. 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 *cancel-directive-name* of the encountered **cancellation point** construct is not **taskgroup**, 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 *cancel-directive-name* 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.

1 Execution Model Events

2 The *cancellation* event occurs if a task encounters a cancellation point and detects the activation of
3 cancellation.

4 Tool Callbacks

5 A thread dispatches a registered `ompt_callback_cancel` callback for each occurrence of a
6 *cancel* event in the context of the encountering task. This callback has type signature
7 `ompt_callback_cancel_t; (flags & ompt_cancel_detected)` always evaluates to *true*
8 in the dispatched callback; `(flags & ompt_cancel_parallel)` evaluates to *true* in the
9 dispatched callback if *cancel-directive-name* of the encountered **cancellation point**
10 construct is **parallel**; `(flags & ompt_cancel_sections)` evaluates to *true* in the
11 dispatched callback if *cancel-directive-name* of the encountered **cancellation point**
12 construct is **sections**; `(flags & ompt_cancel_loop)` evaluates to *true* in the dispatched
13 callback if *cancel-directive-name* of the encountered **cancellation point** construct is **for**
14 or **do**; and `(flags & ompt_cancel_taskgroup)` evaluates to *true* in the dispatched callback if
15 *cancel-directive-name* of the encountered **cancellation point** construct is **taskgroup**.

16 Restrictions

17 Restrictions to the **cancellation point** construct are as follows:

- 18 • A **cancellation point** construct for which *cancel-directive-name* is **taskgroup** must
19 be closely nested inside a **task** or **taskloop** construct, and the **cancellation point**
20 region must be closely nested inside a **taskgroup** region.
- 21 • A **cancellation point** construct for which *cancel-directive-name* is not **taskgroup**
22 must be closely nested inside an OpenMP construct that matches *cancel-directive-name*.

23 Cross References

- 24 • **do** directive, see [Section 11.5.2](#)
- 25 • **for** directive, see [Section 11.5.1](#)
- 26 • **parallel** directive, see [Section 10.2](#)
- 27 • **sections** directive, see [Section 11.3](#)
- 28 • **taskgroup** directive, see [Section 15.4](#)
- 29 • *cancel-var* ICV, see [Table 2.1](#)
- 30 • `omp_get_cancellation`, see [Section 18.2.8](#)
- 31 • `ompt_callback_cancel_t`, see [Section 19.5.2.18](#)

17 Composition of Constructs

This chapter defines rules and mechanisms for nesting regions and for combining constructs.

17.1 Nesting of Regions

This section describes a set of restrictions on the nesting of regions. The restrictions on nesting are as follows:

- A worksharing region may not be closely nested inside a worksharing, **task**, **taskloop**, **critical**, **ordered**, **atomic**, or **masked** region.
- A **barrier** region may not be closely nested inside a worksharing, **task**, **taskloop**, **critical**, **ordered**, **atomic**, or **masked** region.
- A **masked** region may not be closely nested inside a worksharing, **atomic**, **task**, or **taskloop** region.
- An **ordered** region that corresponds 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 that corresponds to an **ordered** construct without the **simd** clause specified must be closely nested inside a worksharing-loop region.
- An **ordered** region that corresponds to an **ordered** construct with the **simd** clause specified must be closely nested inside a **simd** or worksharing-loop SIMD region.
- An **ordered** region that corresponds 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 without a defined binding region, the **simd** construct and the **ordered** construct with the **simd** clause.

- 1 • If a **target update**, **target data**, **target enter data**, or **target exit data**
2 construct is encountered during execution of a **target** region, the behavior is unspecified.
- 3 • If a **target** construct is encountered during execution of a **target** region and a **device**
4 clause in which the **ancestor device-modifier** appears is not present on the construct, the
5 behavior is unspecified.
- 6 • A **teams** region must be strictly nested either within the implicit parallel region that
7 surrounds the whole OpenMP program or within a **target** region. If a **teams** construct is
8 nested within a **target** construct, that **target** construct must contain no statements,
9 declarations or directives outside of the **teams** construct.
- 10 • **distribute** regions, including any **distribute** regions arising from composite
11 constructs, **parallel** regions, including any **parallel** regions arising from combined
12 constructs, **loop** regions, **omp_get_num_teams()** regions, and
13 **omp_get_team_num()** regions are the only OpenMP regions that may be strictly nested
14 inside the **teams** region.
- 15 • A **loop** region that binds to a **teams** region must be strictly nested inside a **teams** region.
- 16 • A **distribute** region must be strictly nested inside a **teams** region.
- 17 • If *cancel-directive-name* is **taskgroup**, the **cancel** construct must be closely nested
18 inside a **task** construct and the **cancel** region must be closely nested inside a
19 **taskgroup** region. Otherwise, the **cancel** construct must be closely nested inside an
20 OpenMP construct for which *directive-name* is *cancel-directive-name*.
- 21 • A **cancellation point** construct for which *cancel-directive-name* is **taskgroup** must
22 be closely nested inside a **task** construct, and the **cancellation point** region must be
23 closely nested inside a **taskgroup** region. Otherwise, a **cancellation point**
24 construct must be closely nested inside an OpenMP construct for which *directive-name* is
25 *cancel-directive-name*.
- 26 • The only constructs that may be encountered inside a region that corresponds to a construct
27 with an **order** clause that specifies **concurrent** are the **loop**, **parallel** and **simd**
28 constructs, and combined constructs for which *directive-name-A* is **parallel**.
- 29 • A region that corresponds to a construct with an **order** clause that specifies **concurrent**
30 may not contain calls to the OpenMP Runtime API or to procedures that contain OpenMP
31 directives.

32 17.2 Clauses on Combined and Composite 33 Constructs

34 This section specifies the handling of clauses on combined or composite constructs and the
35 handling of implicit clauses from variables with predetermined data sharing if they are not
36 predetermined only on a particular construct. Some clauses are permitted only on a single leaf

1 construct of the combined or composite construct, in which case the effect is as if the clause is
2 applied to that specific construct. Other clauses that are permitted on more than one leaf construct
3 have the effect as if they are applied to a subset of those constructs, as detailed in this section.

4 The **collapse** clause is applied once to the combined or composite construct.

5 The effect of the **private** clause is as if it is applied only to the innermost leaf construct that
6 permits it.

7 The effect of the **firstprivate** clause is as if it is applied to one or more leaf constructs as
8 follows:

- 9 • To the **distribute** construct if it is among the constituent constructs;
- 10 • To the **teams** construct if it is among the constituent constructs and the **distribute**
11 construct is not;
- 12 • To a worksharing construct that accepts the clause if one is among the constituent constructs;
- 13 • To the **taskloop** construct if it is among the constituent constructs;
- 14 • To the **parallel** construct if it is among the constituent constructs and neither a
15 **taskloop** construct nor a worksharing construct that accepts the clause is among them;
- 16 • To the **target** construct if it is among the constituent constructs and the same list item
17 neither appears in a **lastprivate** clause nor is the base variable or base pointer of a list
18 item that appears in a **map** clause.

19 If the **parallel** construct is among the constituent constructs and the effect is not as if the
20 **firstprivate** clause is applied to it by the above rules, then the effect is as if the **shared**
21 clause with the same list item is applied to the **parallel** construct. If the **teams** construct is
22 among the constituent constructs and the effect is not as if the **firstprivate** clause is applied to
23 it by the above rules, then the effect is as if the **shared** clause with the same list item is applied to
24 the **teams** construct.

25 The effect of the **lastprivate** clause is as if it is applied to all leaf constructs that permit the
26 clause. If the **parallel** construct is among the constituent constructs and the list item is not also
27 specified in the **firstprivate** clause, then the effect of the **lastprivate** clause is as if the
28 **shared** clause with the same list item is applied to the **parallel** construct. If the **teams**
29 construct is among the constituent constructs and the list item is not also specified in the
30 **firstprivate** clause, then the effect of the **lastprivate** clause is as if the **shared** clause
31 with the same list item is applied to the **teams** construct. If the **target** construct is among the
32 constituent constructs and the list item is not the base variable or base pointer of a list item that
33 appears in a **map** clause, the effect of the **lastprivate** clause is as if the same list item appears
34 in a **map** clause with a *map-type* of **tofrom**.

35 The effect of the **shared**, **default**, **thread_limit**, or **order** clause is as if it is applied to
36 all leaf constructs that permit the clause.

1 The effect of the **allocate** clause is as if it is applied to all leaf constructs that permit the clause
2 and to which a data-sharing attribute clause that may create a private copy of the same list item is
3 applied.

4 The effect of the **reduction** clause is as if it is applied to all leaf constructs that permit the
5 clause, except for the following constructs:

- 6 • The **parallel** construct, when combined with the **sections**, worksharing-loop, **loop**,
7 or **taskloop** construct; and
- 8 • The **teams** construct, when combined with the **loop** construct.

9 For the **parallel** and **teams** constructs above, the effect of the **reduction** clause instead is as
10 if each list item or, for any list item that is an array item, its corresponding base array or base
11 pointer appears in a **shared** clause for the construct. If the **task reduction-modifier** is specified,
12 the effect is as if it only modifies the behavior of the **reduction** clause on the innermost leaf
13 construct that accepts the modifier (see [Section 5.5.8](#)). If the **inscan reduction-modifier** is
14 specified, the effect is as if it modifies the behavior of the **reduction** clause on all constructs of
15 the combined construct to which the clause is applied and that accept the modifier. If a list item in a
16 **reduction** clause on a combined target construct does not have the same base variable or base
17 pointer as a list item in a **map** clause on the construct, then the effect is as if the list item in the
18 **reduction** clause appears as a list item in a **map** clause with a *map-type* of **tofrom**.

19 The effect of the **if** clause is described in [Section 3.4](#).

20 The effect of the **linear** clause is as if it is applied to the innermost leaf construct. Additionally,
21 if the list item is not the iteration variable of a **simd** or worksharing-loop SIMD construct, the
22 effect on the outer leaf constructs is as if the list item was specified in **firstprivate** and
23 **lastprivate** clauses on the combined or composite construct, with the rules specified above
24 applied. If a list item of the **linear** clause is the iteration variable of a **simd** or worksharing-loop
25 SIMD construct and it is not declared in the construct, the effect on the outer leaf constructs is as if
26 the list item was specified in a **lastprivate** clause on the combined or composite construct with
27 the rules specified above applied.

28 The effect of the **nowait** clause is as if it is applied to the outermost leaf construct that permits it.

29 If the clauses have expressions on them, such as for various clauses where the argument of the
30 clause is an expression, or *lower-bound*, *length*, or *stride* expressions inside array sections (or
31 *subscript* and *stride* expressions in *subscript-triplet* for Fortran), or *linear-step* or *alignment*
32 expressions, the expressions are evaluated immediately before the construct to which the clause has
33 been split or duplicated per the above rules (therefore inside of the outer leaf constructs). However,
34 the expressions inside the **num_teams** and **thread_limit** clauses are always evaluated before
35 the outermost leaf construct.

36 The restriction that a list item may not appear in more than one data sharing clause with the
37 exception of specifying a variable in both **firstprivate** and **lastprivate** clauses applies
38 after the clauses are split or duplicated per the above rules.

Restrictions

Restrictions to clauses on combined and composite constructs are as follows:

- A clause that appears on a combined or composite construct must apply to at least one of the leaf constructs per the rules defined in this section.

17.3 Combined and Composite Directive Names

Combined constructs are shortcuts for specifying one construct immediately nested inside another construct. Composite constructs are also shortcuts for specifying the effect of one construct immediately following the effect of another construct. However, composite constructs define semantics to combine constructs that cannot otherwise be immediately nested.

For all combined and composite constructs, *directive-name* concatenates *directive-name-A*, the directive name of the enclosing construct, with an intervening space followed by *directive-name-B*, the directive name of the nested construct. If *directive-name-A* and *directive-name-B* both correspond to loop-associated constructs then *directive-name* is a composite construct. Otherwise *directive-name* is a combined construct.

If *directive-name-A* is **taskloop**, **for** or **do** then *directive-name-B* may be **simd**.

If *directive-name-A* is **masked** then *directive-name-B* may be **taskloop** or the directive name of a combined or composite construct for which *directive-name-A* is **taskloop**.

If *directive-name-A* is **parallel** then *directive-name-B* may be **loop**, **sections**, **workshare**, **masked**, **for**, **do** or the directive name of a combined or composite construct for which *directive-name-A* is **masked**, **for** or **do**.

If *directive-name-A* is **distribute** then *directive-name-B* may be **simd** or the directive name of a combined or composite construct for which *directive-name-A* is **parallel** and **for** or **do** is a leaf construct.

If *directive-name-A* is **teams** then *directive-name-B* may be **loop**, **distribute** or the directive name of a combined or composite construct for which *directive-name-A* is **distribute**.

If *directive-name-A* is **target** then *directive-name-B* may be **simd**, **parallel**, **teams**, the directive name of a combined or composite construct for which *directive-name-A* is **teams** or the directive name of a combined or composite construct for which *directive-name-A* is **parallel** and **loop**, **for** or **do** is a leaf construct.

Cross References

- **distribute** directive, see [Section 11.6](#)
- **do** directive, see [Section 11.5.2](#)
- **for** directive, see [Section 11.5.1](#)
- **loop** directive, see [Section 11.7](#)

- 1 • **masked** directive, see [Section 10.6](#)
- 2 • **parallel** directive, see [Section 10.2](#)
- 3 • **sections** directive, see [Section 11.3](#)
- 4 • **target** directive, see [Section 13.8](#)
- 5 • **taskloop** directive, see [Section 12.6](#)
- 6 • **teams** directive, see [Section 10.3](#)
- 7 • **workshare** directive, see [Section 11.4](#)

8 17.4 Combined Construct Semantics

9 The semantics of the combined constructs are identical to that of explicitly specifying the first
10 construct containing one instance of the second construct and no other statements. All combined
11 and composite directives for which a loop-associated construct is a leaf construct are themselves
12 loop-associated constructs. For combined constructs, tool callbacks are invoked as if the constructs
13 were explicitly nested.

14 Restrictions

15 Restrictions to combined constructs are as follows:

- 16 • The restrictions of *directive-name-A* and *directive-name-B* apply.
- 17 • If *directive-name-A* is **parallel**, the **in_reduction** clause must not be specified.
- 18 • If *directive-name-A* is **parallel** and **target** is not among the constituent constructs, the
19 **nowait** clause must not be specified.
- 20 • If *directive-name-A* is **target**, the **copyin** clause must not be specified.

21 Cross References

- 22 • **copyin** clause, see [Section 5.7.1](#)
- 23 • **in_reduction** clause, see [Section 5.5.10](#)
- 24 • **nowait** clause, see [Section 15.6](#)
- 25 • **parallel** directive, see [Section 10.2](#)
- 26 • **target** directive, see [Section 13.8](#)

17.5 Composite Construct Semantics

Composite constructs combine constructs that otherwise cannot be immediately nested. Specifically, composite constructs apply multiple loop-associated constructs to the same canonical loop nest. The semantics of each composite construct first apply the semantics of the enclosing construct as specified by *directive-name-A* and any clauses that apply to it. For each task (possibly implicit, possibly initial) as appropriate for the semantics of *directive-name-A*, the application of its semantics yields a nested loop of depth two in which the outer loop iterates over the chunks assigned to that task and the inner loop iterates over the logical iterations of each chunk. The semantics of *directive-name-B* and any clauses that apply to it are then applied to that inner loop. For composite constructs, tool callbacks are invoked as if the constructs were explicitly nested.

If *directive-name-A* is **taskloop** and *directive-name-B* is **simd** then for the application of the **simd** construct, the effect of any **in_reduction** clause is as if a **reduction** clause with the same reduction operator and list items is present.

Restrictions

Restrictions to composite constructs are as follows:

- The restrictions of *directive-name-A* and *directive-name-B* apply.
- If *directive-name-A* is **distribute**, the **linear** clause may only be specified for loop iteration variables of loops that are associated with the construct.
- If *directive-name-A* is **distribute**, the **ordered** clause must not be specified.

Cross References

- **in_reduction** clause, see [Section 5.5.10](#)
- **linear** clause, see [Section 5.4.6](#)
- **ordered** clause, see [Section 4.4.4](#)
- **reduction** clause, see [Section 5.5.8](#)
- **distribute** directive, see [Section 11.6](#)
- **simd** directive, see [Section 10.5](#)
- **taskloop** directive, see [Section 12.6](#)

18 Runtime Library Routines

This chapter describes the OpenMP API runtime library routines and queryable runtime states. All OpenMP Runtime API names have an `omp_` prefix. Names that begin with the `omp_x_` prefix are reserved for implementation-defined extensions to the OpenMP Runtime API. In this chapter, *true* and *false* are used as generic terms to simplify the description of the routines.

C / C++

true means a non-zero 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 restrictions apply to all OpenMP runtime library routines:

- OpenMP runtime library routines may not be called from **PURE** or **ELEMENTAL** procedures.
- OpenMP runtime library routines may not be called in **DO CONCURRENT** constructs.

Fortran

18.1 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.

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 (for
5 C++ possibly scoped) enum type with at least the **omp_null_allocator** enumerator
6 with the value zero and an enumerator for each predefined memory allocator in Table 6.3;
- 7 • **omp_atv_default**, which is an instance of a type compatible with **omp_uintptr_t**
8 with the value -1;
- 9 • The type **omp_control_tool_result_t**;
- 10 • The type **omp_control_tool_t**;
- 11 • The type **omp_depend_t**;
- 12 • The type **omp_event_handle_t**, which must be an implementation-defined (for C++
13 possibly scoped) enum type;
- 14 • The enumerator **omp_initial_device** with value -1;
- 15 • The type **omp_interop_t**, which must be an implementation-defined integral or pointer
16 type;
- 17 • The type **omp_interop_fr_t**, which must be an implementation-defined enum type with
18 enumerators named **omp_ifr_name** where *name* is a foreign runtime name that is defined
19 in the *OpenMP Additional Definitions* document;
- 20 • The type **omp_intptr_t**, which is a signed integer type that is at least the size of a pointer
21 on any device;
- 22 • The enumerator **omp_invalid_device** with an implementation-defined value less than
23 -1;
- 24 • The type **omp_lock_t**;
- 25 • The type **omp_memspace_handle_t**, which must be an implementation-defined (for
26 C++ possibly scoped) enum type with at least the **omp_null_mem_space** enumerator
27 with the value zero and an enumerator for each predefined memory space in Table 6.1;
- 28 • The type **omp_nest_lock_t**;
- 29 • The type **omp_pause_resource_t**;
- 30 • The type **omp_proc_bind_t**;
- 31 • The type **omp_sched_t**;
- 32 • The type **omp_sync_hint_t**; and

- 1 • The type `omp_uintptr_t`, which is an unsigned integer type capable of holding a pointer
2 on any device.



3 The OpenMP enumeration types provided in the `omp.h` header file shall not be scoped
4 enumeration types unless explicitly allowed.

5 The `omp.h` header file also defines a class template that models the **Allocator** concept in the
6 `omp::allocator` namespace for each predefined memory allocator in Table 6.3 for which the
7 name includes neither the `omp_` prefix nor the `_alloc` suffix.



8 The OpenMP Fortran API runtime library routines are external procedures. The return values of
9 these routines are of default kind, unless otherwise specified.

10 Interface declarations for the OpenMP Fortran runtime library routines described in this chapter
11 shall be provided in the form of a Fortran **module** named `omp_lib` or a Fortran **include** file
12 named `omp_lib.h`. Whether the `omp_lib.h` file provides derived-type definitions or those
13 routines that require an explicit interface is implementation defined. Whether the **include** file or
14 the **module** file (or both) is provided is also implementation defined.

15 These files also define the following:

- 16 • The default integer named constant `omp_allocator_handle_kind`;
- 17 • An integer named constant of kind `omp_allocator_handle_kind` for each predefined
18 memory allocator in Table 6.3;
- 19 • The integer named constant `omp_null_allocator` of kind
20 `omp_allocator_handle_kind`;
- 21 • The default integer named constant `omp_alloctrait_key_kind`;
- 22 • The default integer named constant `omp_alloctrait_val_kind`;
- 23 • The default integer named constant `omp_control_tool_kind`;
- 24 • The default integer named constant `omp_control_tool_result_kind`;
- 25 • The default integer named constant `omp_depend_kind`;
- 26 • The default integer named constant `omp_event_handle_kind`;
- 27 • The default integer named constant `omp_initial_device` with value -1;
- 28 • The default integer named constant `omp_interop_kind`;
- 29 • The default integer named constant `omp_interop_fr_kind`;

- 1 • An integer named constant `omp_ifr_name` of kind `omp_interop_fr_kind` for each
- 2 *name* that is a foreign runtime name that is defined in the *OpenMP Additional Definitions*
- 3 document;
- 4 • The default integer named constant `omp_invalid_device` with an
- 5 implementation-defined value less than -1;
- 6 • The default integer named constant `omp_lock_kind`;
- 7 • The default integer named constant `omp_memspace_handle_kind`;
- 8 • An integer named constant of kind `omp_memspace_handle_kind` for each predefined
- 9 memory space in Table 6.1;
- 10 • The integer named constant `omp_null_mem_space` of kind
- 11 `omp_memspace_handle_kind`;
- 12 • The default integer named constant `omp_nest_lock_kind`;
- 13 • The default integer named constant `omp_pause_resource_kind`;
- 14 • The default integer named constant `omp_proc_bind_kind`;
- 15 • The default integer named constant `omp_sched_kind`;
- 16 • The default integer named constant `omp_sync_hint_kind`; and
- 17 • The default integer named constant `openmp_version` with a value `yyyymm` where `yyyy`
- 18 and `mm` are the year and month designations of the version of the OpenMP Fortran API that
- 19 the implementation supports; this value matches that of the C preprocessor macro `_OPENMP`,
- 20 when a macro preprocessor is supported (see Section 3.3).

21 Whether any of the OpenMP runtime library routines that take an argument are extended with a
 22 generic interface so arguments of different **KIND** type can be accommodated is implementation
 23 defined.

24 18.2 Thread Team Routines

25 This section describes routines that affect and monitor thread teams that execute tasks in the current
 26 contention group.

27 18.2.1 `omp_set_num_threads`

28 Summary

29 The `omp_set_num_threads` routine affects the number of threads to be used for subsequent
 30 **parallel** regions that do not specify a **num_threads** clause, by setting the value of the first
 31 element of the *nthreads-var* ICV of the current task.

Format

C / C++

```
void omp_set_num_threads(int num_threads);
```

C / C++

Fortran

```
subroutine omp_set_num_threads(num_threads)  
integer num_threads
```

Fortran

Constraints on Arguments

The value of the argument passed to this routine must evaluate to a positive integer, or else the behavior of this routine is implementation defined.

Binding

The binding task set for an `omp_set_num_threads` region is the generating task.

Effect

The effect of this routine is to set the value of the first element of the *nthreads-var* ICV of the current task to the value specified in the argument.

Cross References

- `num_threads` clause, see [Section 10.2.2](#)
- `parallel` directive, see [Section 10.2](#)
- *nthreads-var* ICV, see [Table 2.1](#)
- Determining the Number of Threads for a `parallel` Region, see [Section 10.2.1](#)

18.2.2 omp_get_num_threads

Summary

The `omp_get_num_threads` routine returns the number of threads in the current team.

Format

C / C++

```
int omp_get_num_threads(void);
```

C / C++

Fortran

```
integer function omp_get_num_threads()
```

Fortran

1 **Binding**
2 The binding region for an `omp_get_num_threads` region is the innermost enclosing parallel
3 region.

4 **Effect**
5 The `omp_get_num_threads` routine returns the number of threads in the team that is executing
6 the parallel region to which the routine region binds.

7 18.2.3 `omp_get_max_threads`

8 **Summary**
9 The `omp_get_max_threads` routine returns an upper bound on the number of threads that
10 could be used to form a new team if a `parallel` construct without a `num_threads` clause is
11 encountered after execution returns from this routine.

12 **Format**

	C / C++
13	<code>int omp_get_max_threads(void);</code>
	C / C++
	Fortran
14	<code>integer function omp_get_max_threads()</code>
	Fortran

15 **Binding**
16 The binding task set for an `omp_get_max_threads` region is the generating task.

17 **Effect**
18 The value returned by `omp_get_max_threads` is the value of the first element of the
19 *nthreads-var* ICV of the current task. This value is also an upper bound on the number of threads
20 that could be used to form a new team if a parallel region without a `num_threads` clause is
21 encountered after execution returns from this routine.

- 22 **Cross References**
- 23 • `num_threads` clause, see [Section 10.2.2](#)
 - 24 • `parallel` directive, see [Section 10.2](#)
 - 25 • *nthreads-var* ICV, see [Table 2.1](#)
 - 26 • Determining the Number of Threads for a `parallel` Region, see [Section 10.2.1](#)

18.2.4 `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++
`int omp_get_thread_num(void);`

C / C++

Fortran
`integer function omp_get_thread_num()`

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 primary thread of the team is 0.

Cross References

- `omp_get_num_threads`, see [Section 18.2.2](#)

18.2.5 `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*.

Format

C / C++
`int omp_in_parallel(void);`

C / C++

Fortran
`logical function omp_in_parallel()`

Fortran

Binding

The binding task set for an `omp_in_parallel` region is the generating task.

Effect

The effect of the `omp_in_parallel` routine is to return *true* if the current task is enclosed by an active `parallel` region, and the `parallel` region is enclosed by the outermost initial task region on the device; otherwise it returns *false*.

Cross References

- `parallel` directive, see [Section 10.2](#)
- *active-levels-var* ICV, see [Table 2.1](#)

18.2.6 `omp_set_dynamic`

Summary

The `omp_set_dynamic` routine enables or disables dynamic adjustment of the number of threads available for the execution of subsequent `parallel` regions by setting the value of the *dyn-var* ICV.

Format

C / C++
`void omp_set_dynamic(int dynamic_threads);`

C / C++

Fortran
`subroutine omp_set_dynamic(dynamic_threads)
logical dynamic_threads`

Fortran

Binding

The binding task set for an `omp_set_dynamic` region is the generating task.

Effect

For implementations that support dynamic adjustment of the number of threads, if the argument to `omp_set_dynamic` evaluates to *true*, dynamic adjustment is enabled for the current task; otherwise, dynamic adjustment is disabled for the current task. For implementations that do not support dynamic adjustment of the number of threads, this routine has no effect: the value of *dyn-var* remains *false*.

Cross References

- *dyn-var* ICV, see [Table 2.1](#)

18.2.7 `omp_get_dynamic`

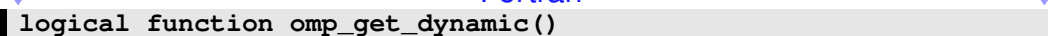

Summary

The `omp_get_dynamic` routine returns the value of the *dyn-var* ICV, which determines whether dynamic adjustment of the number of threads is enabled or disabled.

Format

 C / C++
`int omp_get_dynamic(void);`

 C / C++
 Fortran

`logical function omp_get_dynamic()`
 Fortran

Binding

The binding task set for an `omp_get_dynamic` region is the generating task.

Effect

This routine returns *true* if dynamic adjustment of the number of threads is enabled for the current task; otherwise, it returns *false*. If an implementation does not support dynamic adjustment of the number of threads, then this routine always returns *false*.

Cross References

- *dyn-var* ICV, see [Table 2.1](#)

18.2.8 `omp_get_cancellation`

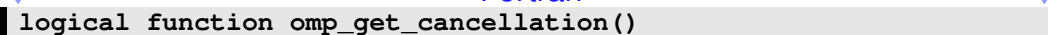

Summary

The `omp_get_cancellation` routine returns the value of the *cancel-var* ICV, which determines if cancellation is enabled or disabled.

Format

 C / C++
`int omp_get_cancellation(void);`

 C / C++
 Fortran

`logical function omp_get_cancellation()`
 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 [Table 2.1](#)

18.2.9 omp_set_schedule

Summary

The `omp_set_schedule` routine affects the schedule that is applied when `runtime` is used as schedule kind, by setting the value of the *run-sched-var* ICV.

Format

C / C++

```
void omp_set_schedule(omp_sched_t kind, int chunk_size);
```

C / C++

Fortran

```
subroutine omp_set_schedule(kind, chunk_size)
integer (kind=omp_sched_kind) kind
integer chunk_size
```

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 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++

```
typedef enum omp_sched_t {
```

```
    // schedule kinds
    omp_sched_static = 0x1,
    omp_sched_dynamic = 0x2,
    omp_sched_guided = 0x3,
    omp_sched_auto = 0x4,

    // schedule modifier
    omp_sched_monotonic = 0x80000000u
} omp_sched_t;
```

C / C++

Fortran

```
1  ! schedule kinds
2  integer(kind=omp_sched_kind), &
3     parameter :: omp_sched_static = &
4         int(Z'1', kind=omp_sched_kind)
5  integer(kind=omp_sched_kind), &
6     parameter :: omp_sched_dynamic = &
7         int(Z'2', kind=omp_sched_kind)
8  integer(kind=omp_sched_kind), &
9     parameter :: omp_sched_guided = &
10         int(Z'3', kind=omp_sched_kind)
11 integer(kind=omp_sched_kind), &
12     parameter :: omp_sched_auto = &
13         int(Z'4', kind=omp_sched_kind)
14
15 ! schedule modifier
16 integer(kind=omp_sched_kind), &
17     parameter :: omp_sched_monotonic = &
18         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 [Table 2.1](#)

18.2.10 `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 18.2.9](#), or any implementation-specific schedule kind. If the returned schedule kind is **static**, **dynamic**, or **guided**, 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. The value returned by the second argument is implementation defined for any other schedule kinds.

Cross References

- *run-sched-var* ICV, see [Table 2.1](#)

18.2.11 `omp_get_thread_limit`

Summary

The `omp_get_thread_limit` routine returns the maximum number of OpenMP threads available to execute tasks in the current contention group.

Format

C / C++

```
int omp_get_thread_limit(void);
```

C / C++

Fortran

```
integer function omp_get_thread_limit()
```

Fortran

Binding

The binding task set for an `omp_get_thread_limit` region is the generating task.

Effect

The `omp_get_thread_limit` routine returns the value of the *thread-limit-var* ICV.

Cross References

- *thread-limit-var* ICV, see [Table 2.1](#)

18.2.12 `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 cannot have a value that is greater than this number. The value that the `omp_get_supported_active_levels` routine returns is implementation defined, but it must be greater than 0.

Cross References

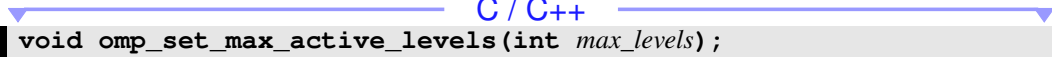
- *max-active-levels-var* ICV, see [Table 2.1](#)

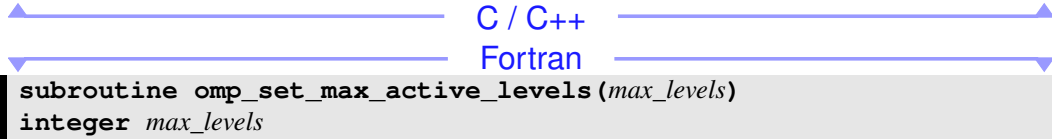
18.2.13 `omp_set_max_active_levels`


Summary

The `omp_set_max_active_levels` routine limits the number of nested active parallel regions when a new nested parallel region is generated by the current task by setting the *max-active-levels-var* ICV.

1 **Format**

2  **C / C++**

3  **Fortran**

4 

5 **Constraints on Arguments**

6 The value of the argument passed to this routine must evaluate to a non-negative integer, otherwise

7 the behavior of this routine is implementation defined.

8 **Binding**

9 The binding task set for an `omp_set_max_active_levels` region is the generating task.

10 **Effect**

11 The effect of this routine is to set the value of the *max-active-levels-var* ICV to the value specified

12 in the argument.

13 If the number of active levels requested exceeds the number of active levels of parallelism supported

14 by the implementation, the value of the *max-active-levels-var* ICV will be set to the number of

15 active levels supported by the implementation. If the number of active levels requested is less than

16 the value of the *active-levels-var* ICV, the value of the *max-active-levels-var* ICV will be set to an

17 implementation-defined value between the requested number and *active-levels-var*, inclusive.

18 **Cross References**

19

- *max-active-levels-var* ICV, see [Table 2.1](#)

20 18.2.14 `omp_get_max_active_levels`

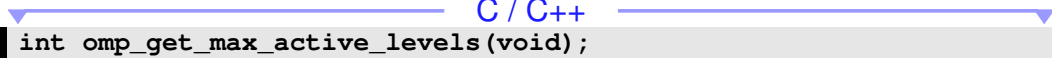
21 **Summary**

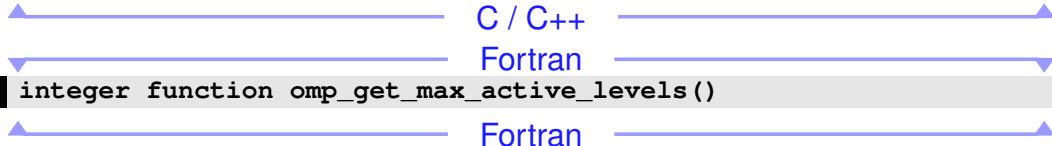
22 The `omp_get_max_active_levels` routine returns the value of the *max-active-levels-var*

23 ICV, which determines the maximum number of nested active parallel regions when the innermost

24 parallel region is generated by the current task.

25 **Format**

26  **C / C++**

27  **Fortran**

Binding

The binding task set for an `omp_get_max_active_levels` region is the generating task.

Effect

The `omp_get_max_active_levels` routine returns the value of the *max-active-levels-var* ICV. The current task may only generate an active parallel region if the returned value is greater than the value of the *active-levels-var* ICV.

Cross References

- *max-active-levels-var* ICV, see [Table 2.1](#)

18.2.15 `omp_get_level`

Summary

The `omp_get_level` routine returns the value of the *levels-var* ICV.

Format

<code>int omp_get_level(void);</code>	C / C++
<code>integer function omp_get_level()</code>	Fortran

Binding

The binding task set for an `omp_get_level` region is the generating task.

Effect

The effect of the `omp_get_level` routine is to return the number of nested `parallel` regions (whether active or inactive) that enclose the current task such that all of the `parallel` regions are enclosed by the outermost initial task region on the current device.

Cross References

- `parallel` directive, see [Section 10.2](#)
- *levels-var* ICV, see [Table 2.1](#)

18.2.16 `omp_get_ancestor_thread_num`

Summary

The `omp_get_ancestor_thread_num` routine returns, for a given nested level of the current thread, the thread number of the ancestor of the current thread.

Format

C / C++
`int omp_get_ancestor_thread_num(int level);`

C / C++

Fortran

Fortran
`integer function omp_get_ancestor_thread_num(level)`
`integer level`

Fortran

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` directive, see [Section 10.2](#)
- `omp_get_level`, see [Section 18.2.15](#)
- `omp_get_thread_num`, see [Section 18.2.4](#)

18.2.17 omp_get_team_size

Summary

The `omp_get_team_size` routine returns, for a given nested level of the current thread, the size of the thread team to which the ancestor or the current thread belongs.

Format

C / C++
`int omp_get_team_size(int level);`

C / C++

Fortran

```
integer function omp_get_team_size(level)
integer level
```

Fortran

Binding

The binding thread set for an `omp_get_team_size` region is the encountering thread. The binding region for an `omp_get_team_size` region is the innermost enclosing `parallel` region.

Effect

The `omp_get_team_size` routine returns the size of the thread team to which the ancestor or the current thread belongs. If the requested nested level is outside the range of 0 and the nested level of the current thread, as returned by the `omp_get_level` routine, the routine returns -1. Inactive parallel regions are regarded as 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

- `parallel` directive, see [Section 10.2](#)
- `omp_get_level`, see [Section 18.2.15](#)
- `omp_get_num_threads`, see [Section 18.2.2](#)

18.2.18 omp_get_active_level

Summary

The `omp_get_active_level` routine returns the value of the *active-levels-var* ICV.

Format

C / C++

```
int omp_get_active_level(void);
```

C / C++

Fortran

```
integer function omp_get_active_level()
```

Fortran

Binding

The binding task set for 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

- `parallel` directive, see [Section 10.2](#)
- `active-levels-var` ICV, see [Table 2.1](#)

18.3 Thread Affinity Routines

This section describes routines that affect and access thread affinity policies that are in effect.

18.3.1 `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.

Format

```

C / C++
omp_proc_bind_t omp_get_proc_bind(void);
C / C++
Fortran
integer (kind=omp_proc_bind_kind) function omp_get_proc_bind()
Fortran
```

Constraints on Arguments

The value returned by this routine must be one of the valid affinity policy kinds. The C/C++ header file (`omp.h`) and the Fortran include file (`omp_lib.h`) and/or Fortran module file (`omp_lib`) define the valid constants. The valid constants must include the following:

```

C / C++
typedef enum omp_proc_bind_t {
    omp_proc_bind_false = 0,
    omp_proc_bind_true = 1,
    omp_proc_bind_primary = 2,
    omp_proc_bind_close = 3,
    omp_proc_bind_spread = 4
} omp_proc_bind_t;
C / C++
```

Fortran

```
1 integer (kind=omp_proc_bind_kind), &  
2     parameter :: omp_proc_bind_false = 0  
3 integer (kind=omp_proc_bind_kind), &  
4     parameter :: omp_proc_bind_true = 1  
5 integer (kind=omp_proc_bind_kind), &  
6     parameter :: omp_proc_bind_primary = 2  
7 integer (kind=omp_proc_bind_kind), &  
8     parameter :: omp_proc_bind_close = 3  
9 integer (kind=omp_proc_bind_kind), &  
10    parameter :: omp_proc_bind_spread = 4
```

Fortran

Binding

The binding task set for an `omp_get_proc_bind` region is the generating task.

Effect

The effect of this routine is to return the value of the first element of the *bind-var* ICV of the current task. See [Section 10.2.3](#) for the rules that govern the thread affinity policy.

Cross References

- `parallel` directive, see [Section 10.2](#)
- Controlling OpenMP Thread Affinity, see [Section 10.2.3](#)
- *bind-var* ICV, see [Table 2.1](#)

18.3.2 `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++

```
int omp_get_num_places(void);
```

C / C++

Fortran

```
integer function omp_get_num_places()
```

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 [Table 2.1](#)

18.3.3 `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 omp_get_place_num_procs(int place_num);</code>		
	C / C++	
	Fortran	
<code>integer function omp_get_place_num_procs(place_num)</code> <code>integer 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

- `omp_get_num_places`, see [Section 18.3.2](#)

18.3.4 `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

```
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

- `OMP_PLACES`, see [Section 21.1.5](#)
- `omp_get_num_places`, see [Section 18.3.2](#)
- `omp_get_place_num_procs`, see [Section 18.3.3](#)

18.3.5 `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      int omp_get_place_num(void);
3      integer function omp_get_place_num()

```

C / C++
C / C++
Fortran
Fortran

Binding
The binding thread set for an `omp_get_place_num` region is the encountering thread.

Effect
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.

- Cross References**
- `omp_get_num_places`, see [Section 18.3.2](#)

18.3.6 `omp_get_partition_num_places`

Summary
The `omp_get_partition_num_places` routine returns the number of places in the place partition of the innermost implicit task.

```

17     Format
18     int omp_get_partition_num_places(void);
19     integer function omp_get_partition_num_places()

```

C / C++
C / C++
Fortran
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 [Table 2.1](#)

18.3.7 `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

C / C++

```
void omp_get_partition_place_nums(int *place_nums);
```

C / C++

Fortran

```
subroutine omp_get_partition_place_nums(place_nums)  
integer place_nums(*)
```

Fortran

Binding

The binding task set for an `omp_get_partition_place_nums` region is the encountering implicit task.

Effect

The `omp_get_partition_place_nums` routine returns the list of place numbers 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 [Table 2.1](#)
- `omp_get_partition_num_places`, see [Section 18.3.6](#)

18.3.8 `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++

```
void omp_set_affinity_format(const char *format);
```

C / C++

Fortran

```
1 | subroutine omp_set_affinity_format(format)  
2 | character(len=*), intent(in) :: format
```

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.

Restrictions

Restrictions to the `omp_set_affinity_format` routine are as follows.

- When called from within a `target` region the effect is unspecified.

Cross References

- `OMP_AFFINITY_FORMAT`, see [Section 21.2.5](#)
- `OMP_DISPLAY_AFFINITY`, see [Section 21.2.4](#)
- Controlling OpenMP Thread Affinity, see [Section 10.2.3](#)
- `omp_capture_affinity`, see [Section 18.3.11](#)
- `omp_display_affinity`, see [Section 18.3.10](#)
- `omp_get_affinity_format`, see [Section 18.3.9](#)

18.3.9 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++

```
29 | 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.

Restrictions

Restrictions to the **omp_get_affinity_format** routine are as follows.

- When called from within a **target** region the effect is unspecified.

Cross References

- **parallel** directive, see [Section 10.2](#)
- **teams** directive, see [Section 10.3](#)
- *affinity-format-var* ICV, see [Table 2.1](#)

18.3.10 `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
`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.

Restrictions

Restrictions to the `omp_display_affinity` routine are as follows.

- When called from within a `target` region the effect is unspecified.

Cross References

- `affinity-format-var` ICV, see [Table 2.1](#)

18.3.11 `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

```
1 integer function omp_capture_affinity(buffer,format)
2 character(len=*), intent(out) :: buffer
3 character(len=*), intent(in) :: format
```

Fortran

Binding

The binding thread set for an **omp_capture_affinity** region is the encountering thread.

Effect

C / C++

The **omp_capture_affinity** routine returns the number of characters in the entire thread affinity information string excluding the terminating null byte (' \0 '). If *size* is non-zero, it 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.

Restrictions

Restrictions to the **omp_capture_affinity** routine are as follows.

- When called from within a **target** region the effect is unspecified.

Cross References

- *affinity-format-var* ICV, see [Table 2.1](#)

18.4 Teams Region Routines

This section describes routines that affect and monitor the league of teams that may execute a **teams** region.

18.4.1 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` directive, see [Section 10.3](#)

18.4.2 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

Binding

The binding task set for an `omp_get_team_num` region is the generating task.

Effect

The `omp_get_team_num` routine returns the initial team number of the calling thread. The initial team number is an integer between 0 and one less than the value returned by `omp_get_num_teams()`, inclusive. The routine returns 0 if it is called outside of a `teams` region.

Cross References

- **teams** directive, see [Section 10.3](#)
- **omp_get_num_teams**, see [Section 18.4.1](#)

18.4.3 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 *ntteams-var* ICV of the current device.

Format

	C / C++	
<pre>void omp_set_num_teams(int num_teams);</pre>		
	C / C++	
	Fortran	
<pre>subroutine omp_set_num_teams(num_teams) integer num_teams</pre>		
	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 *ntteams-var* ICV of the current device to the value specified in the argument.

Restrictions

Restrictions to the **omp_set_num_teams** routine are as follows:

- The routine may not be called from within a parallel region that is not the implicit parallel region that surrounds the whole OpenMP program.

Cross References

- **num_teams** clause, see [Section 10.3.1](#)
- **teams** directive, see [Section 10.3](#)
- *ntteams-var* ICV, see [Table 2.1](#)

18.4.4 `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

```
int omp_get_max_teams(void);
```

C / C++

```
integer function omp_get_max_teams()
```

Fortran

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 device. 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

- `num_teams` clause, see [Section 10.3.1](#)
- `teams` directive, see [Section 10.3](#)
- `nteamsvar` ICV, see [Table 2.1](#)

18.4.5 `omp_set_teams_thread_limit`

Summary

The `omp_set_teams_thread_limit` routine defines the maximum number of OpenMP threads that can execute tasks in each contention group that a `teams` construct creates.

Format

```
void omp_set_teams_thread_limit(int thread_limit);
```

C / C++

```
subroutine omp_set_teams_thread_limit(thread_limit)  
integer thread_limit
```

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_teams_thread_limit` region is the generating task.

Effect

The `omp_set_teams_thread_limit` routine sets the value of the *teams-thread-limit-var* ICV to the value of the *thread_limit* argument. If the value of *thread_limit* exceeds the number of OpenMP threads that an implementation supports for each contention group created by a **teams** construct, the value of the *teams-thread-limit-var* ICV will be set to the number that is supported by the implementation.

Restrictions

Restrictions to the `omp_set_teams_thread_limit` routine are as follows:

- The routine may not be called from within a parallel region other than the implicit parallel region that surrounds the whole OpenMP program.

Cross References

- `thread_limit` clause, see [Section 13.3](#)
- **teams** directive, see [Section 10.3](#)
- *teams-thread-limit-var* ICV, see [Table 2.1](#)

18.4.6 `omp_get_teams_thread_limit`

Summary

The `omp_get_teams_thread_limit` routine returns the maximum number of OpenMP threads available to execute tasks in each contention group that a **teams** construct creates.

Format

	C / C++	
<code>int omp_get_teams_thread_limit(void);</code>		
	C / C++	
	Fortran	
<code>integer function omp_get_teams_thread_limit()</code>		
	Fortran	

Binding

The binding task set for an `omp_get_teams_thread_limit` region is the generating task.

Effect

The `omp_get_teams_thread_limit` routine returns the value of the *teams-thread-limit-var* ICV.

Cross References

- `teams` directive, see [Section 10.3](#)
- *teams-thread-limit-var* ICV, see [Table 2.1](#)

18.5 Tasking Routines

This section describes routines that pertain to OpenMP explicit tasks.

18.5.1 `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	

Binding

The binding thread set for an `omp_get_max_task_priority` region is all threads on the device. The effect of executing this routine is not related to any specific region that corresponds to any construct or API routine.

Effect

The `omp_get_max_task_priority` routine returns the value of the *max-task-priority-var* ICV, which determines the maximum value that can be specified in the `priority` clause.

Cross References

- `priority` clause, see [Section 12.4](#)
- *max-task-priority-var* ICV, see [Table 2.1](#)

18.5.2 `omp_in_explicit_task`

Summary

The `omp_in_explicit_task` routine returns the value of the *explicit-task-var* ICV.

Format

C / C++

```
int omp_in_explicit_task(void);
```

C / C++

Fortran

```
logical function omp_in_explicit_task()
```

Fortran

Binding

The binding task set for an `omp_in_explicit_task` region is the generating task.

Effect

The `omp_in_explicit_task` routine returns the value of the *explicit-task-var* ICV, which indicates whether the encountering region is an explicit task region.

Cross References

- `task` directive, see [Section 12.5](#)
- *explicit-task-var* ICV, see [Table 2.1](#)

18.5.3 `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++

```
int omp_in_final(void);
```

C / C++

Fortran

```
logical function omp_in_final()
```

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*.

18.6 Resource Relinquishing Routines

This section describes routines that relinquish resources used by the OpenMP runtime.

18.6.1 `omp_pause_resource`

Summary

The `omp_pause_resource` routine allows the runtime to relinquish resources used by OpenMP on the specified device.

Format

C / C++

```
int omp_pause_resource(omp_pause_resource_t kind, int device_num);
```

C / C++

Fortran

```
integer function omp_pause_resource(kind, device_num)  
integer (kind=omp_pause_resource_kind) kind  
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 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++

```
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 a conforming device number. If the device number has the value `omp_invalid_device`, runtime error termination is performed.

1 Binding



2 The binding task set for an `omp_pause_resource` region is the whole program.

3 Effect

4 The `omp_pause_resource` routine allows the runtime to relinquish resources used by OpenMP
5 on the specified device.

6 If successful, the `omp_pause_hard` value results in a hard pause for which the OpenMP state is
7 not guaranteed to persist across the `omp_pause_resource` call. A hard pause may relinquish
8 any data allocated by OpenMP on a given device, including data allocated by memory routines for
9 that device as well as data present on the device as a result of a declare target directive or
10 `target data` construct. A hard pause may also relinquish any data associated with a
11 `threadprivate` directive. When relinquished and when applicable, base language appropriate
12 deallocation/finalization is performed. When relinquished and when applicable, mapped data on a
13 device will not be copied back from the device to the host.

14 If successful, the `omp_pause_soft` value results in a soft pause for which the OpenMP state is
15 guaranteed to persist across the call, with the exception of any data associated with a
16 `threadprivate` directive, which may be relinquished across the call. When relinquished and
17 when applicable, base language appropriate deallocation/finalization is performed.

18 
19 **Note** – A hard pause may relinquish more resources, but may resume processing OpenMP regions
20 more slowly. A soft pause allows OpenMP regions to restart more quickly, but may relinquish fewer
21 resources. An OpenMP implementation will reclaim resources as needed for OpenMP regions
22 encountered after the `omp_pause_resource` region. Since a hard pause may unmap data on the
23 specified device, appropriate data mapping is required before using data on the specified device
24 after the `omp_pause_region` region.
25 

26 The routine returns zero in case of success, and non-zero otherwise.

27 Tool Callbacks

28 If the tool is not allowed to interact with the specified device after encountering this call, then the
29 runtime must call the tool finalizer for that device.

30 Restrictions

31 Restrictions to the `omp_pause_resource` routine are as follows:

- 32 • The `omp_pause_resource` region may not be nested in any explicit OpenMP region.
- 33 • The routine may only be called when all explicit tasks have finalized execution.

34 Cross References

- 35 • `target data` directive, see [Section 13.5](#)
- 36 • `threadprivate` directive, see [Section 5.2](#)
- 37 • Declare Target Directives, see [Section 7.8](#)

18.6.2 `omp_pause_resource_all`

Summary

The `omp_pause_resource_all` routine allows the runtime to relinquish resources used by OpenMP on all devices.

Format

C / C++
`int omp_pause_resource_all(omp_pause_resource_t kind);`

C / C++
Fortran
`integer function omp_pause_resource_all(kind)`
`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 18.6.1](#), 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

Restrictions to the `omp_pause_resource_all` routine are as follows:

- 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.

Cross References

- `omp_pause_resource`, see [Section 18.6.1](#)

18.7 Device Information Routines

This section describes routines that pertain to the set of devices that are available to an OpenMP program.

18.7.1 `omp_get_num_procs`

Summary

The `omp_get_num_procs` routine returns the number of processors available to the device.

Format

C / C++
`int omp_get_num_procs(void);`

C / C++
Fortran
integer function `omp_get_num_procs()`
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.

18.7.2 `omp_get_max_progress_width`

Summary

The `omp_get_max_progress_width` routine returns the maximum size of progress units on the specified device.

Format

C / C++
`int omp_get_max_progress_width(int device_num);`

C / C++
Fortran
integer function `omp_get_max_progress_width(device_num)`
integer *device_num*
Fortran

Constraints on Arguments

The `device_num` argument must be a conforming device number.

Binding

The binding task set for an `omp_get_max_progress_width` region is the generating task.

Effect

The effect of the `omp_get_progress_max_width` routine is to return the maximum size, in terms of hardware threads, of progress units on the device specified by *device_num*.

Cross References

- `parallel` directive, see [Section 10.2](#)

18.7.3 `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

- `target` directive, see [Section 13.8](#)
- *default-device-var* ICV, see [Table 2.1](#)

18.7.4 `omp_get_default_device`

Summary

The `omp_get_default_device` routine returns the default target device.

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24

Format

```
int omp_get_default_device(void);  
  
integer function omp_get_default_device()
```

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

- `target` directive, see [Section 13.8](#)
- *default-device-var* ICV, see [Table 2.1](#)

18.7.5 `omp_get_num_devices`

Summary

The `omp_get_num_devices` routine returns the number of non-host devices available for offloading code or data.

Format

```
int omp_get_num_devices(void);  
  
integer function omp_get_num_devices()
```

Binding

The binding task set for an `omp_get_num_devices` region is the generating task.

Effect

The `omp_get_num_devices` routine returns the number of available non-host devices onto which code or data may be offloaded. When called from within a `target` region the effect of this routine is unspecified.

Cross References









- `target` directive, see [Section 13.8](#)

18.7.6 `omp_get_device_num`

Summary

The `omp_get_device_num` routine returns the device number of the device on which the calling thread is executing.

Format

		C / C++	
8	<code>int omp_get_device_num(void);</code>		
		C / C++	
		Fortran	
9	<code>integer function omp_get_device_num()</code>		
		Fortran	

Binding

The binding task set for an `omp_get_device_num` region is the generating task.

Effect









The `omp_get_device_num` routine returns the device number of the device on which the calling thread is executing. When called on the host device, it will return the same value as the `omp_get_initial_device` routine.

18.7.7 `omp_is_initial_device`

Summary

The `omp_is_initial_device` routine returns *true* if the current task is executing on the host device; otherwise, it returns *false*.

Format

		C / C++	
21	<code>int omp_is_initial_device(void);</code>		
		C / C++	
		Fortran	
22	<code>logical function omp_is_initial_device()</code>		
		Fortran	

Binding

The binding task set for an `omp_is_initial_device` region is the generating task.

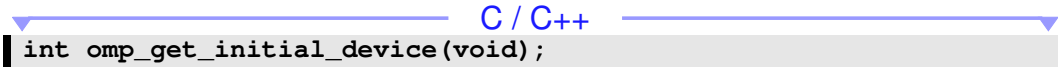
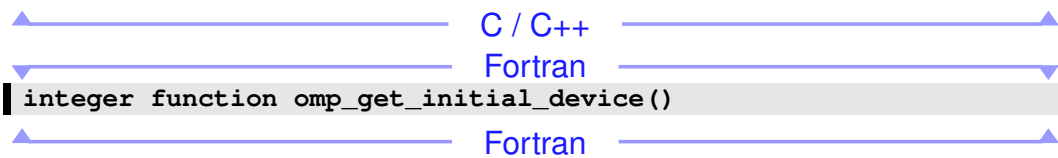
1 **Effect**
2 The effect of this routine is to return *true* if the current task is executing on the host device;
3 otherwise, it returns *false*.

4 **18.7.8 omp_get_initial_device**

5 **Summary**

6 The **omp_get_initial_device** routine returns a device number that represents the host
7 device.

8 **Format**

9 
10 

11 **Binding**

12 The binding task set for an **omp_get_initial_device** region is the generating task.

13 **Effect**

14 The effect of this routine is to return the device number of the host device. The value of the device
15 number is the value returned by the **omp_get_num_devices** routine. When called from within
16 a **target** region the effect of this routine is unspecified.

17 **Cross References**

- 18 • **target** directive, see [Section 13.8](#)

19 **18.8 Device Memory Routines**

20 This section describes routines that support allocation of memory and management of pointers in
21 the data environments of target devices.

22 If the *device_num*, *src_device_num*, or *dst_device_num* argument of a device memory routine has
23 the value **omp_invalid_device**, runtime error termination is performed.

24 **18.8.1 omp_target_alloc**

25 **Summary**

26 The **omp_target_alloc** routine allocates memory in a device data environment and returns a
27 device pointer to that memory.

Format

C / C++
`void* omp_target_alloc(size_t size, int device_num);`

C / C++

Fortran

Fortran
`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

Constraints on Arguments

The `device_num` argument must be a conforming device number.

Binding

The binding task set for an `omp_target_alloc` region is the generating task, which is the *target task* generated by the call to the `omp_target_alloc` routine.

Effect

The `omp_target_alloc` routine returns a device pointer that references 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`. The `omp_target_alloc` routine executes as if part of a target task that is generated by the call to the routine and that is an included task. The `omp_target_alloc` routine returns `NULL` if it cannot dynamically allocate the memory in the device data environment. The device pointer returned by `omp_target_alloc` can be used in an `is_device_ptr` clause (see [Section 5.4.7](#)).

Fortran

The `omp_target_alloc` routine requires an explicit interface and so might not be provided in `omp_lib.h`.

Fortran

Execution Model Events

The *target-data-allocation-begin* event occurs before a thread initiates a data allocation on a target device.

The *target-data-allocation-end* event occurs after a thread initiates a data allocation on a target device.

Tool Callbacks

A thread dispatches a registered `ompt_callback_target_data_op_emi` callback with `ompt_scope_begin` as its endpoint argument for each occurrence of a *target-data-allocation-begin* event in that thread. Similarly, a thread dispatches a registered `ompt_callback_target_data_op_emi` callback with `ompt_scope_end` as its endpoint argument for each occurrence of a *target-data-allocation-end* event in that thread. These callbacks have type signature `ompt_callback_target_data_op_emi_t`.

1 A thread dispatches a registered `ompt_callback_target_data_op` callback for each
2 occurrence of a *target-data-allocation-end* event in that thread. The callback occurs in the context
3 of the target task and has type signature `ompt_callback_target_data_op_t`.

4 Restrictions

5 Restrictions to the `omp_target_alloc` routine are as follows.

- 6 • Freeing the storage returned by `omp_target_alloc` with any routine other than
7 `omp_target_free` results in unspecified behavior.
- 8 • When called from within a `target` region the effect is unspecified.

▼ C / C++ ▼

- 9 • Unless the `unified_address` clause appears on a `requires` directive in the
10 compilation unit, pointer arithmetic is not supported on the device pointer returned by
11 `omp_target_alloc`.

▲ C / C++ ▲

12 Cross References

- 13 • `is_device_ptr` clause, see [Section 5.4.7](#)
- 14 • `target` directive, see [Section 13.8](#)
- 15 • `omp_target_free`, see [Section 18.8.2](#)
- 16 • `ompt_callback_target_data_op_emi_t` and
17 `ompt_callback_target_data_op_t`, see [Section 19.5.2.25](#)

18 18.8.2 `omp_target_free`

19 Summary

20 The `omp_target_free` routine frees the device memory allocated by the
21 `omp_target_alloc` routine.

22 Format

▼ C / C++ ▼

```
23 void omp_target_free(void *device_ptr, int device_num);
```

▲ C / C++ ▲

▼ Fortran ▼

```
24 subroutine omp_target_free(device_ptr, device_num) bind(c)  
25 use, intrinsic :: iso_c_binding, only : c_ptr, c_int  
26 type(c_ptr), value :: device_ptr  
27 integer(c_int), value :: device_num
```

▲ Fortran ▲

28 Constraints on Arguments

29 A program that calls `omp_target_free` with a non-null pointer that does not have a value
30 returned from `omp_target_alloc` is non-conforming. The `device_num` argument must be a
31 conforming device number.

Binding

The binding task set for an `omp_target_free` region is the generating task, which is the *target task* generated by the call to the `omp_target_free` routine.

Effect

The `omp_target_free` routine frees the memory in the device data environment associated with `device_ptr`. If `device_ptr` is `NULL`, the operation is ignored. The `omp_target_free` routine executes as if part of a target task that is generated by the call to the routine and that is an included task. Synchronization must be inserted to ensure that all accesses to `device_ptr` are completed before the call to `omp_target_free`.

Fortran

The `omp_target_free` routine requires an explicit interface and so might not be provided in `omp_lib.h`.

Fortran

Execution Model Events

The *target-data-free-begin* event occurs before a thread initiates a data free on a target device.

The *target-data-free-end* event occurs after a thread initiates a data free on a target device.

Tool Callbacks

A thread dispatches a registered `ompt_callback_target_data_op_emi` callback with `ompt_scope_begin` as its endpoint argument for each occurrence of a *target-data-free-begin* event in that thread. Similarly, a thread dispatches a registered `ompt_callback_target_data_op_emi` callback with `ompt_scope_end` as its endpoint argument for each occurrence of a *target-data-free-end* event in that thread. These callbacks have type signature `ompt_callback_target_data_op_emi_t`.

A thread dispatches a registered `ompt_callback_target_data_op` callback for each occurrence of a *target-data-free-begin* event in that thread. The callback occurs in the context of the target task and has type signature `ompt_callback_target_data_op_t`.

Restrictions

Restrictions to the `omp_target_free` routine are as follows.

- When called from within a `target` region the effect is unspecified.

Cross References

- `target` directive, see [Section 13.8](#)
- `omp_target_alloc`, see [Section 18.8.1](#)
- `ompt_callback_target_data_op_emi_t` and `ompt_callback_target_data_op_t`, see [Section 19.5.2.25](#)

18.8.3 omp_target_is_present

Summary

The `omp_target_is_present` routine tests whether a host pointer refers to storage that is mapped to a given device.

Format

C / C++

```
int omp_target_is_present(const void *ptr, int device_num);
```

C / C++

Fortran

```
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*. The *device_num* argument must be a conforming device number.

Binding

The binding task set for an `omp_target_is_present` region is the encountering task.

Effect

The `omp_target_is_present` routine returns a non-zero value if *device_num* refers to the host device or if *ptr* refers to storage that has corresponding storage in the device data environment of device *device_num*. Otherwise, the routine returns zero.

Fortran

The `omp_target_is_present` routine requires an explicit interface and so might not be provided in `omp_lib.h`.

Fortran

Restrictions

Restrictions to the `omp_target_is_present` routine are as follows.

- When called from within a `target` region the effect is unspecified.

Cross References

- `target` directive, see [Section 13.8](#)

18.8.4 omp_target_is_accessible

Summary

The `omp_target_is_accessible` routine tests whether memory is accessible from a given device.

Format

C / C++

```
int omp_target_is_accessible( const void *ptr, size_t size,  
                             int device_num);
```

C / C++

Fortran

```
integer(c_int) function omp_target_is_accessible( &  
        ptr, size, device_num) bind(c)  
use, intrinsic :: iso_c_binding, only : c_ptr, c_size_t, c_int  
type(c_ptr), value :: ptr  
integer(c_size_t), value :: size  
integer(c_int), value :: device_num
```

Fortran

Constraints on Arguments

The value of *size* must be positive. The *device_num* argument must be a conforming device number.

Binding

The binding task set for an `omp_target_is_accessible` region is the encountering task.

Effect

This routine returns a non-zero value if the storage of *size* bytes starting at the address given by *ptr* is accessible from device *device_num*. Otherwise, it returns zero. The value of *ptr* is interpreted as an address in the address space of the specified device.

Fortran

The `omp_target_is_accessible` routine requires an explicit interface and so might not be provided in `omp_lib.h`.

Fortran

Restrictions

Restrictions to the `omp_target_is_accessible` routine are as follows.

- When called from within a `target` region the effect is unspecified.

Cross References

- `target` directive, see [Section 13.8](#)

18.8.5 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 conforming device numbers.

Binding

The binding task set for an `omp_target_memcpy` region is the generating task, which is the *target task* generated by the call to the `omp_target_memcpy` routine.

Effect

This routine copies `length` bytes of memory at offset `src_offset` from `src` in the device data environment of device `src_device_num` to `dst` starting at offset `dst_offset` in the device data environment of device `dst_device_num`. The `omp_target_memcpy` routine executes as if part of a target task that is generated by the call to the routine and that is an included task. The return value is zero on success and non-zero on failure. This routine contains a task scheduling point.

Fortran

The `omp_target_memcpy` routine requires an explicit interface and so might not be provided in `omp_lib.h`.

Fortran

Execution Model Events

The *target-data-op-begin* event occurs before a thread initiates a data transfer in the `omp_target_memcpy` region.

The *target-data-op-end* event occurs after a thread initiates a data transfer in the `omp_target_memcpy` region.

Tool Callbacks

A thread dispatches a registered `ompt_callback_target_data_op_emi` callback with `ompt_scope_begin` as its endpoint argument for each occurrence of a *target-data-op-begin* event in that thread. Similarly, a thread dispatches a registered `ompt_callback_target_data_op_emi` callback with `ompt_scope_end` as its endpoint argument for each occurrence of a *target-data-op-end* event in that thread. These callbacks have type signature `ompt_callback_target_data_op_emi_t`.

A thread dispatches a registered `ompt_callback_target_data_op` callback for each occurrence of a *target-data-op-end* event in that thread. The callback occurs in the context of the target task and has type signature `ompt_callback_target_data_op_t`.

Restrictions

Restrictions to the `omp_target_memcpy` routine are as follows.

- When called from within a `target` region the effect is unspecified.

Cross References

- `target` directive, see [Section 13.8](#)
- `ompt_callback_target_data_op_emi_t` and `ompt_callback_target_data_op_t`, see [Section 19.5.2.25](#)

18.8.6 `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++

```
int omp_target_memcpy_rect (  
    void *dst,  
    const void *src,  
    size_t element_size,  
    int num_dims,  
    const size_t *volume,  
    const size_t *dst_offsets,
```

```

1  const size_t *src_offsets,
2  const size_t *dst_dimensions,
3  const size_t *src_dimensions,
4  int dst_device_num,
5  int src_device_num
6  );

```

▲ C / C++ ▲

▼ Fortran ▼

```

7  integer(c_int) function omp_target_memcpy_rect(dst,src,element_size, &
8     num_dims, volume, dst_offsets, src_offsets, dst_dimensions, &
9     dst_device_num, src_device_num) bind(c)
10 use, intrinsic :: iso_c_binding, only : c_ptr, c_int, c_size_t
11 type(c_ptr), value :: dst, src
12 integer(c_size_t), value :: element_size
13 integer(c_int), value :: num_dims, dst_device_num, src_device_num
14 integer(c_size_t), intent(in) :: volume(*), dst_offsets(*), &
15     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 conforming device numbers. 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.

▼ Fortran ▼

Because the interface binds directly to a C language function the function assumes C memory ordering.

▲ Fortran ▲

Binding

The binding task set for an **omp_target_memcpy_rect** region is the generating task, which is the *target task* generated by the call to the **omp_target_memcpy_rect** routine.

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 the 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*).

1 The `omp_target_memcpy_rect` routine executes as if part of a target task that is generated by
2 the call to the routine and that is an included task. The routine returns zero if successful.
3 Otherwise, it returns a non-zero value. The routine contains a task scheduling point.

4 An application can determine the inclusive number of dimensions supported by an implementation
5 by passing `NULL` for both `dst` and `src`. The routine returns the number of dimensions supported by
6 the implementation for the specified device numbers. No copy operation is performed.

Fortran

7 The `omp_target_memcpy_rect` routine requires an explicit interface and so might not be
8 provided in `omp_lib.h`.

Fortran

9 Execution Model Events

10 The *target-data-op-begin* event occurs before a thread initiates a data transfer in the
11 `omp_target_memcpy_rect` region.

12 The *target-data-op-end* event occurs after a thread initiates a data transfer in the
13 `omp_target_memcpy_rect` region.

14 Tool Callbacks

15 A thread dispatches a registered `ompt_callback_target_data_op_emi` callback with
16 `ompt_scope_begin` as its endpoint argument for each occurrence of a *target-data-op-begin*
17 event in that thread. Similarly, a thread dispatches a registered
18 `ompt_callback_target_data_op_emi` callback with `ompt_scope_end` as its endpoint
19 argument for each occurrence of a *target-data-op-end* event in that thread. These callbacks have
20 type signature `ompt_callback_target_data_op_emi_t`.

21 A thread dispatches a registered `ompt_callback_target_data_op` callback for each
22 occurrence of a *target-data-op-end* event in that thread. The callback occurs in the context of the
23 target task and has type signature `ompt_callback_target_data_op_t`.

24 Restrictions

25 Restrictions to the `omp_target_memcpy_rect` routine are as follows.

- 26 • When called from within a `target` region the effect is unspecified.

27 Cross References

- 28 • `target` directive, see [Section 13.8](#)
- 29 • `ompt_callback_target_data_op_emi_t` and
30 `ompt_callback_target_data_op_t`, see [Section 19.5.2.25](#)

31 18.8.7 `omp_target_memcpy_async`

32 Summary

33 The `omp_target_memcpy_async` routine asynchronously performs a copy between any
34 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 conforming device numbers.

Binding

The binding task set for an **omp_target_memcpy_async** region is the generating task, which is the *target task* generated by the call to the **omp_target_memcpy_async** routine.

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*. The **omp_target_memcpy_async** routine executes as if part of a target task that is generated by the call to the routine and for which execution may be deferred. Task dependences are expressed with zero or more OpenMP depend objects. The dependences are specified by passing the number of depend objects followed by an array of the 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.

1 The routine returns zero if successful. Otherwise, it returns a non-zero value. The routine contains
2 a task scheduling point.

Fortran

3 The `omp_target_memcpy_async` routine requires an explicit interface and so might not be
4 provided in `omp_lib.h`.

Fortran

5 Execution Model Events

6 The *target-data-op-begin* event occurs before a thread initiates a data transfer in the
7 `omp_target_memcpy_async` region.

8 The *target-data-op-end* event occurs after a thread initiates a data transfer in the
9 `omp_target_memcpy_async` region.

10 Tool Callbacks

11 A thread dispatches a registered `ompt_callback_target_data_op_emi` callback with
12 `ompt_scope_begin` as its endpoint argument for each occurrence of a *target-data-op-begin*
13 event in that thread. Similarly, a thread dispatches a registered
14 `ompt_callback_target_data_op_emi` callback with `ompt_scope_end` as its endpoint
15 argument for each occurrence of a *target-data-op-end* event in that thread. These callbacks have
16 type signature `ompt_callback_target_data_op_emi_t`.

17 A thread dispatches a registered `ompt_callback_target_data_op` callback for each
18 occurrence of a *target-data-op-end* event in that thread. The callback occurs in the context of the
19 target task and has type signature `ompt_callback_target_data_op_t`.

20 Restrictions

21 Restrictions to the `omp_target_memcpy_async` routine are as follows.

- 22 • When called from within a `target` region the effect is unspecified.

23 Cross References

- 24 • `target` directive, see [Section 13.8](#)
- 25 • Depend Objects, see [Section 15.9.2](#)
- 26 • `ompt_callback_target_data_op_emi_t` and
27 `ompt_callback_target_data_op_t`, see [Section 19.5.2.25](#)

28 18.8.8 `omp_target_memcpy_rect_async`

29 Summary

30 The `omp_target_memcpy_rect_async` routine asynchronously performs a copy between
31 any combination of host and device pointers.

Format

C / C++

```
1 int omp_target_memcpy_rect_async(  
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    int depobj_count,  
14    omp_depend_t *depobj_list  
15 );  
16
```

C / C++

Fortran

```
17 integer(c_int) function omp_target_memcpy_rect_async(dst, src, &  
18     element_size, num_dims, volume, dst_offsets, src_offsets, &  
19     dst_dimensions, src_dimensions, dst_device_num, src_device_num, &  
20     depobj_count, depobj_list) bind(c)  
21 use, intrinsic :: iso_c_binding, only : c_ptr, c_int, c_size_t  
22 type(c_ptr), value :: dst, src  
23 integer(c_size_t), value :: element_size  
24 integer(c_int), value :: num_dims, dst_device_num, src_device_num, &  
25     depobj_count  
26 integer(c_size_t), intent(in) :: volume(*), dst_offsets(*), &  
27     src_offsets(*), dst_dimensions(*), src_dimensions(*)  
28 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 conforming device numbers. 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.

Fortran

Because the interface binds directly to a C language function the function assumes C memory ordering.

Fortran

Binding

The binding task set for an `omp_target_memcpy_rect_async` region is the generating task, which is the *target task* generated by the call to the `omp_target_memcpy_rect_async` routine.

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 the 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 `omp_target_memcpy_rect_async` routine executes as if part of a target task that is generated by the call to the routine and for which execution may be deferred. Task dependences are expressed with zero or more OpenMP depend objects. The dependences are specified by passing the number of depend objects followed by an array of the objects. The generated target task is not a dependent task if the program passes in a count of zero for *depopbj_count*. *depopbj_list* is ignored if the value of *depopbj_count* is zero.

The routine returns zero if successful. Otherwise, it returns a non-zero value. The routine contains a task scheduling point.

An application can determine the number of inclusive dimensions supported by an implementation by passing *NULL* 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.

Fortran

The `omp_target_memcpy_rect_async` routine requires an explicit interface and so might not be provided in `omp_lib.h`.

Fortran

Execution Model Events

The *target-data-op-begin* event occurs before a thread initiates a data transfer in the `omp_target_memcpy_rect_async` region.

The *target-data-op-end* event occurs after a thread initiates a data transfer in the `omp_target_memcpy_rect_async` region.

1 Tool Callbacks

2 A thread dispatches a registered `ompt_callback_target_data_op_emi` callback with
3 `ompt_scope_begin` as its endpoint argument for each occurrence of a *target-data-op-begin*
4 event in that thread. Similarly, a thread dispatches a registered
5 `ompt_callback_target_data_op_emi` callback with `ompt_scope_end` as its endpoint
6 argument for each occurrence of a *target-data-op-end* event in that thread. These callbacks have
7 type signature `ompt_callback_target_data_op_emi_t`.

8 A thread dispatches a registered `ompt_callback_target_data_op` callback for each
9 occurrence of a *target-data-op-end* event in that thread. The callback occurs in the context of the
10 target task and has type signature `ompt_callback_target_data_op_t`.

11 Restrictions

12 Restrictions to the `omp_target_memcpy_rect_async` routine are as follows.

- 13 • When called from within a `target` region the effect is unspecified.

14 Cross References

- 15 • `target` directive, see [Section 13.8](#)
- 16 • Depend Objects, see [Section 15.9.2](#)
- 17 • `ompt_callback_target_data_op_emi_t` and
18 `ompt_callback_target_data_op_t`, see [Section 19.5.2.25](#)

19 18.8.9 `omp_target_associate_ptr`

20 Summary

21 The `omp_target_associate_ptr` routine maps a device pointer, which may be returned
22 from `omp_target_alloc` or implementation-defined runtime routines, to a host pointer.

23 Format

```
24 int omp_target_associate_ptr(  
25     const void *host_ptr,  
26     const void *device_ptr,  
27     size_t size,  
28     size_t device_offset,  
29     int device_num  
30 );
```

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 a conforming device number.

Binding

The binding task set for an **omp_target_associate_ptr** region is the generating task, which is the *target task* generated by the call to the **omp_target_associate_ptr** routine.

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 using the **omp_target_disassociate_ptr** routine to disassociate it .

The **omp_target_associate_ptr** routine executes as if part of a target task that is generated by the call to the routine and that is an included task. 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.

Fortran

The **omp_target_associate_ptr** routine requires an explicit interface and so might not be provided in **omp_lib.h**.

Fortran

Execution Model Events

The *target-data-associate* event occurs before a thread initiates a device pointer association on a target device.

Tool Callbacks

A thread dispatches a registered `ompt_callback_target_data_op` callback, or a registered `ompt_callback_target_data_op_emi` callback with `ompt_scope_beginend` as its endpoint argument for each occurrence of a *target-data-associate* event in that thread. These callbacks have type signature `ompt_callback_target_data_op_t` or `ompt_callback_target_data_op_emi_t`, respectively.

Restrictions

Restrictions to the `omp_target_associate_ptr` routine are as follows.

- When called from within a `target` region the effect is unspecified.

Cross References

- `target` directive, see [Section 13.8](#)
- `omp_target_alloc`, see [Section 18.8.1](#)
- `omp_target_disassociate_ptr`, see [Section 18.8.10](#)
- `omp_target_is_present`, see [Section 18.8.3](#)
- `ompt_callback_target_data_op_emi_t` and `ompt_callback_target_data_op_t`, see [Section 19.5.2.25](#)

18.8.10 `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* argument must be a conforming device number.

Binding

The binding task set for an `omp_target_disassociate_ptr` region is the generating task, which is the *target task* generated by the call to the `omp_target_disassociate_ptr` routine.

Effect

The `omp_target_disassociate_ptr` removes the associated device data on device *device_num* from the presence table for host pointer *ptr*. A call to this routine on a pointer that is not *NULL* and does not have associated data on the given device results in unspecified behavior. The reference count of the mapping is reduced to zero, regardless of its current value. The `omp_target_disassociate_ptr` routine executes as if part of a target task that is generated by the call to the routine and that is an included task. 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.

Fortran

The `omp_target_disassociate_ptr` routine requires an explicit interface and so might not be provided in `omp_lib.h`.

Fortran

Execution Model Events

The *target-data-disassociate* event occurs before a thread initiates a device pointer disassociation on a target device.

Tool Callbacks

A thread dispatches a registered `ompt_callback_target_data_op` callback, or a registered `ompt_callback_target_data_op_emi` callback with `ompt_scope_beginend` as its endpoint argument for each occurrence of a *target-data-disassociate* event in that thread. These callbacks have type signature `ompt_callback_target_data_op_t` or `ompt_callback_target_data_op_emi_t`, respectively.

Restrictions

Restrictions to the `omp_target_disassociate_ptr` routine are as follows.

- When called from within a `target` region the effect is unspecified.

Cross References

- `target` directive, see [Section 13.8](#)
- `ompt_callback_target_data_op_emi_t` and `ompt_callback_target_data_op_t`, see [Section 19.5.2.25](#)

18.8.11 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++

```
void * omp_get_mapped_ptr(const void *ptr, int device_num);
```

C / C++

Fortran

```
type(c_ptr) function omp_get_mapped_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` argument must be a conforming device number.

Binding

The binding task set for an `omp_get_mapped_ptr` region is the encountering task.

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` and does not have an associated pointer on the given device will return `NULL`. The routine returns `NULL` if unsuccessful. Otherwise it returns the device pointer, which is `ptr` if `device_num` is the value returned by `omp_get_initial_device()`.

Fortran

The `omp_get_mapped_ptr` routine requires an explicit interface and so might not be provided in `omp_lib.h`.

Fortran

Execution Model Events

No events are associated with this routine.

Restrictions

Restrictions to the `omp_get_mapped_ptr` routine are as follows.

- When called from within a `target` region the effect is unspecified.

Cross References

- `omp_get_initial_device`, see [Section 18.7.8](#)

18.9 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. An OpenMP program does not need to include explicit **flush** directives to ensure that the lock variable's value is consistent among different tasks.

Binding

The binding task set for all lock routine regions is all tasks in the contention group.

Simple Lock Routines

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++

C / C++

Fortran

1 For the following routines, a simple lock variable must be an integer variable of
2 **kind=omp_lock_kind**.

Fortran

3 The simple lock routines are as follows:

- 4 • The **omp_init_lock** routine initializes a simple lock;
- 5 • The **omp_init_lock_with_hint** routine initializes a simple lock and attaches a hint to
6 it;
- 7 • The **omp_destroy_lock** routine uninitialized a simple lock;
- 8 • The **omp_set_lock** routine waits until a simple lock is available and then sets it;
- 9 • The **omp_unset_lock** routine unsets a simple lock; and
- 10 • The **omp_test_lock** routine tests a simple lock and sets it if it is available.

Nestable Lock Routines

C / C++

12 The type **omp_nest_lock_t** represents a nestable lock. For the following routines, a nestable
13 lock variable must be of **omp_nest_lock_t** type. All nestable lock routines require an
14 argument that is a pointer to a variable of type **omp_nest_lock_t**.

C / C++

Fortran

15 For the following routines, a nestable lock variable must be an integer variable of
16 **kind=omp_nest_lock_kind**.

Fortran

17 The nestable lock routines are as follows:

- 18 • The **omp_init_nest_lock** routine initializes a nestable lock;
- 19 • The **omp_init_nest_lock_with_hint** routine initializes a nestable lock and attaches
20 a hint to it;
- 21 • The **omp_destroy_nest_lock** routine uninitialized a nestable lock;
- 22 • The **omp_set_nest_lock** routine waits until a nestable lock is available and then sets it;
- 23 • The **omp_unset_nest_lock** routine unsets a nestable lock; and
- 24 • The **omp_test_nest_lock** routine tests a nestable lock and sets it if it is available.

Restrictions

25 Restrictions to OpenMP lock routines are as follows:

- 27 • The use of the same OpenMP lock in different contention groups results in unspecified
28 behavior.

18.9.1 `omp_init_lock` and `omp_init_nest_lock`

Summary

These routines initialize an OpenMP lock without a hint.

Format

C / C++

```
void omp_init_lock(omp_lock_t *lock);  
void omp_init_nest_lock(omp_nest_lock_t *lock);
```

C / C++

Fortran

```
subroutine omp_init_lock(svar)  
integer (kind=omp_lock_kind) svar  
  
subroutine omp_init_nest_lock(nvar)  
integer (kind=omp_nest_lock_kind) nvar
```

Fortran

Constraints on Arguments

A program that accesses a lock that is not in the uninitialized state through either routine is non-conforming.

Effect

The effect of these routines is to initialize the lock to the unlocked state; that is, 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` 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 19.5.2.14](#)

18.9.2 `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

```
subroutine omp_init_lock_with_hint(svar, hint)  
    integer (kind=omp_lock_kind) svar  
    integer (kind=omp_sync_hint_kind) hint  
  
subroutine omp_init_nest_lock_with_hint(nvar, hint)  
    integer (kind=omp_nest_lock_kind) nvar  
    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 15.1](#).

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-with-hint* 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.

1 Tool Callbacks

2 A thread dispatches a registered `ompt_callback_lock_init` callback with the same value
3 for its *hint* argument as the *hint* argument of the call to `omp_init_lock_with_hint` and
4 `ompt_mutex_lock` as the *kind* argument for each occurrence of a *lock-init-with-hint* event in
5 that thread. Similarly, a thread dispatches a registered `ompt_callback_lock_init` callback
6 with the same value for its *hint* argument as the *hint* argument of the call to
7 `omp_init_nest_lock_with_hint` and `ompt_mutex_nest_lock` as the *kind* argument
8 for each occurrence of a *nest-lock-init-with-hint* event in that thread. These callbacks have the type
9 signature `ompt_callback_mutex_acquire_t` and occur in the task that encounters the
10 routine.

11 Cross References

- 12 • Synchronization Hints, see [Section 15.1](#)
- 13 • `ompt_callback_mutex_acquire_t`, see [Section 19.5.2.14](#)

14 18.9.3 omp_destroy_lock and omp_destroy_nest_lock

15 Summary

16 These routines ensure that the OpenMP lock is uninitialized.

17 Format

C / C++

```
18 void omp_destroy_lock(omp_lock_t *lock);  
19 void omp_destroy_nest_lock(omp_nest_lock_t *lock);
```

C / C++

Fortran

```
20 subroutine omp_destroy_lock(svar)  
21 integer (kind=omp_lock_kind) svar  
22  
23 subroutine omp_destroy_nest_lock(nvar)  
24 integer (kind=omp_nest_lock_kind) nvar
```

Fortran

25 Constraints on Arguments

26 A program that accesses a lock that is not in the unlocked state through either routine is
27 non-conforming.

28 Effect

29 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* 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_t` and occur in the task that encounters the routine.

Cross References

- `ompt_callback_mutex_t`, see [Section 19.5.2.15](#)

18.9.4 `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 19.5.2.14](#)
- `ompt_callback_mutex_t`, see [Section 19.5.2.15](#)
- `ompt_callback_nest_lock_t`, see [Section 19.5.2.16](#)

18.9.5 `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 19.5.2.15](#)
- 13 • `ompt_callback_nest_lock_t`, see [Section 19.5.2.16](#)

14 18.9.6 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

C / C++

```
19 int omp_test_lock(omp_lock_t *lock);  
20 int omp_test_nest_lock(omp_nest_lock_t *lock);
```

C / C++

Fortran

```
21 logical function omp_test_lock(svar)  
22 integer (kind=omp_lock_kind) svar  
23  
24 integer function omp_test_nest_lock(nvar)  
25 integer (kind=omp_nest_lock_kind) nvar
```

Fortran

26 Constraints on Arguments

27 A program that accesses a lock that is in the uninitialized state through either routine is
28 non-conforming. The behavior is unspecified if a simple lock accessed by `omp_test_lock` is in
29 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

- `ompt_callback_mutex_acquire_t`, see [Section 19.5.2.14](#)
- `ompt_callback_mutex_t`, see [Section 19.5.2.15](#)
- `ompt_callback_nest_lock_t`, see [Section 19.5.2.16](#)

18.10 Timing Routines

This section describes routines that support a portable wall clock timer.

18.10.1 `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.

18.10.2 `omp_get_wtick`

Summary

The `omp_get_wtick` routine returns the precision of the timer used by `omp_get_wtime`.

Format

	C / C++
<code>double omp_get_wtick(void);</code>	
	C / C++
	Fortran
<code>double precision function omp_get_wtick()</code>	
	Fortran

1 **Binding**

2 The binding thread set for an `omp_get_wtick` region is the encountering thread. The routine’s
3 return value is not guaranteed to be consistent across any set of threads.

4 **Effect**

5 The `omp_get_wtick` routine returns a value equal to the number of seconds between successive
6 clock ticks of the timer used by `omp_get_wtime`.

7 **18.11 Event Routine**

8 This section describes a routine that supports OpenMP event objects.

9 **Binding**

10 The binding thread set for all event routine regions is the encountering thread.

11 **18.11.1 omp_fulfill_event**

12 **Summary**

13 This routine fulfills and destroys an OpenMP event.

14 **Format**

```

15       | void omp_fulfill_event(omp_event_handle_t event);
16       | subroutine omp_fulfill_event(event)
17       | integer (kind=omp_event_handle_kind) event

```

18 **Constraints on Arguments**

19 A program that calls this routine on an event that was already fulfilled is non-conforming. A
20 program that calls this routine with an event handle that was not created by the `detach` clause is
21 non-conforming.

22 **Effect**

23 The effect of this routine is to fulfill the event associated with the event handle argument. The effect
24 of fulfilling the event will depend on how the event was created. The event is destroyed and cannot
25 be accessed after calling this routine, and the event handle becomes unassociated with any event.

26 **Execution Model Events**

27 The *task-fulfill* event occurs in a thread that executes an `omp_fulfill_event` region before the
28 event is fulfilled if the OpenMP event object was created by a `detach` clause on a task.

1 Tool Callbacks

2 A thread dispatches a registered `ompt_callback_task_schedule` callback with `NULL` as its
3 `next_task_data` argument while the argument `prior_task_data` binds to the detachable task for each
4 occurrence of a `task-fulfill` event. If the `task-fulfill` event occurs before the detachable task finished
5 the execution of the associated `structured-block`, the callback has
6 `ompt_task_early_fulfill` as its `prior_task_status` argument; otherwise the callback has
7 `ompt_task_late_fulfill` as its `prior_task_status` argument. This callback has type
8 signature `ompt_callback_task_schedule_t`.

9 Restrictions

10 Restrictions to the `omp_fulfill_event` routine are as follows:

- 11 • The event handler passed to the routine must have been created by a thread in the same device
12 as the thread that invoked the routine.

13 Cross References

- 14 • `detach` clause, see [Section 12.5.2](#)
- 15 • `ompt_callback_task_schedule_t`, see [Section 19.5.2.10](#)

▼ C / C++ ▼

16 18.12 Interoperability Routines

17 The interoperability routines provide mechanisms to inspect the properties associated with an
18 `omp_interop_t` object. Such objects may be initialized, destroyed or otherwise used by an
19 `interop` construct. Additionally, an `omp_interop_t` object can be initialized to
20 `omp_interop_none`, which is defined to be zero. An `omp_interop_t` object may only be
21 accessed or modified through OpenMP directives and API routines.

22 An `omp_interop_t` object can be copied without affecting, or copying, the underlying state.
23 Destruction of an `omp_interop_t` object destroys the state to which all copies of the object refer.

24 OpenMP reserves all negative values for properties, as listed in [Table 18.1](#); implementation-defined
25 properties may use zero and positive values. The special property, `omp_ipr_first`, will always
26 have the lowest property value, which may change in future versions of this specification. Valid
27 values and types for the properties that [Table 18.1](#) lists are specified in the *OpenMP Additional*
28 *Definitions* document or are implementation defined unless otherwise specified.

29 [Table 18.2](#) lists the return codes used by routines that take an `int* ret_code` argument.

30 Binding

31 The binding task set for all interoperability routine regions is the generating task.

▲ C / C++ ▲

TABLE 18.1: Required Values of the `omp_interop_property_t` enum Type

Enum Name	Contexts	Name	Property
<code>omp_ipr_fr_id = -1</code>	all	<code>fr_id</code>	An <code>intptr_t</code> value that represents the foreign runtime id of context
<code>omp_ipr_fr_name = -2</code>	all	<code>fr_name</code>	C string value that represents the foreign runtime name of context
<code>omp_ipr_vendor = -3</code>	all	<code>vendor</code>	An <code>intptr_t</code> that represents the vendor of context
<code>omp_ipr_vendor_name = -4</code>	all	<code>vendor_name</code>	C string value that represents the vendor of context
<code>omp_ipr_device_num = -5</code>	all	<code>device_num</code>	The OpenMP device ID for the device in the range 0 to <code>omp_get_num_devices()</code> inclusive
<code>omp_ipr_platform = -6</code>	<i>target</i>	<code>platform</code>	A foreign platform handle usually spanning multiple devices
<code>omp_ipr_device = -7</code>	<i>target</i>	<code>device</code>	A foreign device handle
<code>omp_ipr_device_context = -8</code>	<i>target</i>	<code>device_context</code>	A handle to an instance of a foreign device context
<code>omp_ipr_targetsync = -9</code>	<i>targetsync</i>	<code>targetsync</code>	A handle to a synchronization object of a foreign execution context
<code>omp_ipr_first = -9</code>			

TABLE 18.2: Required Values for the `omp_interop_rc_t` enum Type

Enum Name	Description
<code>omp_irc_no_value = 1</code>	Parameters valid, no meaningful value available
<code>omp_irc_success = 0</code>	Successful, value is usable
<code>omp_irc_empty = -1</code>	The object provided is equal to <code>omp_interop_none</code>
<code>omp_irc_out_of_range = -2</code>	Property ID is out of range, see Table 18.1
<code>omp_irc_type_int = -3</code>	Property type is int; use <code>omp_get_interop_int</code>
<code>omp_irc_type_ptr = -4</code>	Property type is pointer; use <code>omp_get_interop_ptr</code>
<code>omp_irc_type_str = -5</code>	Property type is string; use <code>omp_get_interop_str</code>
<code>omp_irc_other = -6</code>	Other error; use <code>omp_get_interop_rc_desc</code>

18.12.1 `omp_get_num_interop_properties`

Summary

The `omp_get_num_interop_properties` routine retrieves the number of implementation-defined properties available for an `omp_interop_t` object.

Format

```
int omp_get_num_interop_properties(const omp_interop_t interop);
```

Effect

The `omp_get_num_interop_properties` routine returns the number of implementation-defined properties available for *interop*. The total number of properties available for *interop* is the returned value minus `omp_ipr_first`.

C / C++

C / C++

18.12.2 `omp_get_interop_int`

Summary

The `omp_get_interop_int` routine retrieves an integer property from an `omp_interop_t` object.

Format

```
omp_intptr_t omp_get_interop_int(const omp_interop_t interop,
                                omp_interop_property_t property_id,
                                int *ret_code);
```

Effect

The `omp_get_interop_int` routine returns the requested integer property, if available, and zero if an error occurs or no value is available. If the *interop* is `omp_interop_none`, an empty error occurs. If the *property_id* is less than `omp_ipr_first` or greater than or equal to `omp_get_num_interop_properties(interop)`, an out of range error occurs. If the requested property value is not convertible into an integer value, a type error occurs.

If a non-null pointer is passed to *ret_code*, an `omp_interop_rc_t` value that indicates the return code is stored in the object to which *ret_code* points. If an error occurred, the stored value will be negative and it will match the error as defined in Table 18.2. On success, zero will be stored. If no error occurred but no meaningful value can be returned, `omp_irc_no_value`, which is one, will be stored.

Restrictions

Restrictions to the `omp_get_interop_int` routine are as follows:

- The behavior of the routine is unspecified if an invalid `omp_interop_t` object is provided.

Cross References

- `omp_get_num_interop_properties`, see [Section 18.12.1](#)

C / C++

C / C++

18.12.3 `omp_get_interop_ptr`

Summary

The `omp_get_interop_ptr` routine retrieves a pointer property from an `omp_interop_t` object.

Format

```
void* omp_get_interop_ptr(const omp_interop_t interop,  
                          omp_interop_property_t property_id,  
                          int *ret_code);
```

Effect

The `omp_get_interop_ptr` routine returns the requested pointer property, if available, and `NULL` if an error occurs or no value is available. If the `interop` is `omp_interop_none`, an empty error occurs. If the `property_id` is less than `omp_ipr_first` or greater than or equal to `omp_get_num_interop_properties(interop)`, an out of range error occurs. If the requested property value is not convertible into a pointer value, a type error occurs.

If a non-null pointer is passed to `ret_code`, an `omp_interop_rc_t` value that indicates the return code is stored in the object to which the `ret_code` points. If an error occurred, the stored value will be negative and it will match the error as defined in Table 18.2. On success, zero will be stored. If no error occurred but no meaningful value can be returned, `omp_irc_no_value`, which is one, will be stored.

Restrictions

Restrictions to the `omp_get_interop_ptr` routine are as follows:

- The behavior of the routine is unspecified if an invalid `omp_interop_t` object is provided.
- Memory referenced by the pointer returned by the `omp_get_interop_ptr` routine is managed by the OpenMP implementation and should not be freed or modified.

Cross References

- `omp_get_num_interop_properties`, see [Section 18.12.1](#)

C / C++

18.12.4 `omp_get_interop_str`

Summary

The `omp_get_interop_str` routine retrieves a string property from an `omp_interop_t` object.

Format

```
const char* omp_get_interop_str(const omp_interop_t interop,
                               omp_interop_property_t property_id,
                               int *ret_code);
```

Effect

The `omp_get_interop_str` routine returns the requested string property as a C string, if available, and `NULL` if an error occurs or no value is available. If the `interop` is `omp_interop_none`, an empty error occurs. If the `property_id` is less than `omp_ipr_first` or greater than or equal to `omp_get_num_interop_properties(interop)`, an out of range error occurs. If the requested property value is not convertible into a string value, a type error occurs.

If a non-null pointer is passed to `ret_code`, an `omp_interop_rc_t` value that indicates the return code is stored in the object to which the `ret_code` points. If an error occurred, the stored value will be negative and it will match the error as defined in Table 18.2. On success, zero will be stored. If no error occurred but no meaningful value can be returned, `omp_irc_no_value`, which is one, will be stored.

Restrictions

Restrictions to the `omp_get_interop_str` routine are as follows:

- The behavior of the routine is unspecified if an invalid `omp_interop_t` object is provided.
- Memory referenced by the pointer returned by the `omp_get_interop_str` routine is managed by the OpenMP implementation and should not be freed or modified.

Cross References

- `omp_get_num_interop_properties`, see [Section 18.12.1](#)

18.12.5 `omp_get_interop_name`

Summary

The `omp_get_interop_name` routine retrieves a property name from an `omp_interop_t` object.

Format

```
const char* omp_get_interop_name(const omp_interop_t interop,  
                                omp_interop_property_t property_id)  
    ;
```

Effect

The `omp_get_interop_name` routine returns the name of the property identified by `property_id` as a C string. Property names for non-implementation defined properties are listed in Table 18.1. If the `property_id` is less than `omp_ipr_first` or greater than or equal to `omp_get_num_interop_properties (interop)`, `NULL` is returned.


Restrictions

Restrictions to the `omp_get_interop_name` routine are as follows:

- The behavior of the routine is unspecified if an invalid object is provided.
- Memory referenced by the pointer returned by the `omp_get_interop_name` routine is managed by the OpenMP implementation and should not be freed or modified.

Cross References

- `omp_get_num_interop_properties`, see [Section 18.12.1](#)



18.12.6 omp_get_interop_type_desc

Summary

The `omp_get_interop_type_desc` routine retrieves a description of the type of a property associated with an `omp_interop_t` object.

Format

```
const char* omp_get_interop_type_desc(const omp_interop_t interop,  
                                      omp_interop_property_t  
    property_id);
```

Effect

The `omp_get_interop_type_desc` routine returns a C string that describes the type of the property identified by `property_id` in human-readable form. That may contain a valid C type declaration possibly followed by a description or name of the type. If `interop` has the value `omp_interop_none`, `NULL` is returned. If the `property_id` is less than `omp_ipr_first` or greater than or equal to `omp_get_num_interop_properties (interop)`, `NULL` is returned.

Restrictions

Restrictions to the `omp_get_interop_type_desc` routine are as follows:

- The behavior of the routine is unspecified if an invalid object is provided.
- Memory referenced by the pointer returned from the `omp_get_interop_type_desc` routine is managed by the OpenMP implementation and should not be freed or modified.

Cross References

- `omp_get_num_interop_properties`, see [Section 18.12.1](#)

▲ C / C++ ▲

▼ C / C++ ▼

18.12.7 `omp_get_interop_rc_desc`

Summary

The `omp_get_interop_rc_desc` routine retrieves a description of the return code associated with an `omp_interop_t` object.

Format

```
const char* omp_get_interop_rc_desc(const omp_interop_t interop,  
                                   omp_interop_rc_t ret_code);
```

Effect

The `omp_get_interop_rc_desc` routine returns a C string that describes the return code `ret_code` in human-readable form.

Restrictions

Restrictions to the `omp_get_interop_rc_desc` routine are as follows:

- The behavior of the routine is unspecified if an invalid object is provided or if `ret_code` was not last written by an interoperability routine invoked with the `omp_interop_t` object `interop`.
- Memory referenced by the pointer returned by the `omp_get_interop_rc_desc` routine is managed by the OpenMP implementation and should not be freed or modified.

▲ C / C++ ▲

18.13 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.

Restrictions

C / C++

For all routines in this section that allocate memory, the following restrictions apply:

- Unless the **unified_address** requirement is specified or the current device is an associated device of the allocator, pointer arithmetic is not supported on the returned pointers.

C / C++

18.13.1 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_atk_pin_device = 9,
    omp_atk_preferred_device = 10,
    omp_atk_device_access = 11,
    omp_atk_target_access = 12,
    omp_atk_atomic_scope = 13,
    omp_atk_part_size = 14
} omp_alloctrail_key_t;

typedef enum omp_alloctrail_value_t {
    omp_atv_false = 0,
    omp_atv_true = 1,
    omp_atv_contended = 3,
    omp_atv_uncontended = 4,
    omp_atv_serialized = 5,
    omp_atv_private = 6,
    omp_atv_device = 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,
```



```

1      omp_atv_allocator_fb = 14,
2      omp_atv_environment = 15,
3      omp_atv_nearest = 16,
4      omp_atv_blocked = 17,
5      omp_atv_interleaved = 18,
6      omp_atv_all = 19,
7      omp_atv_single = 20,
8      omp_atv_multiple = 21,
9      omp_atv_memspace = 22
10     } omp_alloctrail_value_t;
11
12     typedef struct omp_alloctrail_t {
13         omp_alloctrail_key_t key;
14         omp_uintptr_t value;
15     } omp_alloctrail_t;

```

▲ C / C++ ▲

▼ Fortran ▼

```

16     integer(kind=omp_alloctrail_key_kind), &
17         parameter :: omp_atk_sync_hint = 1
18     integer(kind=omp_alloctrail_key_kind), &
19         parameter :: omp_atk_alignment = 2
20     integer(kind=omp_alloctrail_key_kind), &
21         parameter :: omp_atk_access = 3
22     integer(kind=omp_alloctrail_key_kind), &
23         parameter :: omp_atk_pool_size = 4
24     integer(kind=omp_alloctrail_key_kind), &
25         parameter :: omp_atk_fallback = 5
26     integer(kind=omp_alloctrail_key_kind), &
27         parameter :: omp_atk_fb_data = 6
28     integer(kind=omp_alloctrail_key_kind), &
29         parameter :: omp_atk_pinned = 7
30     integer(kind=omp_alloctrail_key_kind), &
31         parameter :: omp_atk_partition = 8
32     integer(kind=omp_alloctrail_key_kind), &
33         parameter :: omp_atk_pin_device = 9
34     integer(kind=omp_alloctrail_key_kind), &
35         parameter :: omp_atk_preferred_device = 10
36     integer(kind=omp_alloctrail_key_kind), &
37         parameter :: omp_atk_device_access = 11
38     integer(kind=omp_alloctrail_key_kind), &
39         parameter :: omp_atk_target_access = 12
40     integer(kind=omp_alloctrail_key_kind), &
41         parameter :: omp_atk_atomic_scope = 13

```

```

1  integer(kind=omp_alloctrail_key_kind), &
2      parameter :: omp_atk_part_size = 14
3
4
5  integer(kind=omp_alloctrail_val_kind), &
6      parameter :: omp_atv_default = -1
7  integer(kind=omp_alloctrail_val_kind), &
8      parameter :: omp_atv_false = 0
9  integer(kind=omp_alloctrail_val_kind), &
10     parameter :: omp_atv_true = 1
11 integer(kind=omp_alloctrail_val_kind), &
12     parameter :: omp_atv_contended = 3
13 integer(kind=omp_alloctrail_val_kind), &
14     parameter :: omp_atv_uncontended = 4
15 integer(kind=omp_alloctrail_val_kind), &
16     parameter :: omp_atv_serialized = 5
17 integer(kind=omp_alloctrail_val_kind), &
18     parameter :: omp_atv_private = 6
19 integer(kind=omp_alloctrail_val_kind), &
20     parameter :: omp_atv_device = 7
21 integer(kind=omp_alloctrail_val_kind), &
22     parameter :: omp_atv_thread = 8
23 integer(kind=omp_alloctrail_val_kind), &
24     parameter :: omp_atv_pteam = 9
25 integer(kind=omp_alloctrail_val_kind), &
26     parameter :: omp_atv_cgroup = 10
27 integer(kind=omp_alloctrail_val_kind), &
28     parameter :: omp_atv_default_mem_fb = 11
29 integer(kind=omp_alloctrail_val_kind), &
30     parameter :: omp_atv_null_fb = 12
31 integer(kind=omp_alloctrail_val_kind), &
32     parameter :: omp_atv_abort_fb = 13
33 integer(kind=omp_alloctrail_val_kind), &
34     parameter :: omp_atv_allocator_fb = 14
35 integer(kind=omp_alloctrail_val_kind), &
36     parameter :: omp_atv_environment = 15
37 integer(kind=omp_alloctrail_val_kind), &
38     parameter :: omp_atv_nearest = 16
39 integer(kind=omp_alloctrail_val_kind), &
40     parameter :: omp_atv_blocked = 17
41 integer(kind=omp_alloctrail_val_kind), &
42     parameter :: omp_atv_interleaved = 18
43 integer(kind=omp_alloctrail_val_kind), &

```

```

1      parameter :: omp_atv_all = 19
2      integer(kind=omp_alloctrail_val_kind), &
3      parameter :: omp_atv_single = 20
4      integer(kind=omp_alloctrail_val_kind), &
5      parameter :: omp_atv_multiple = 21
6      integer(kind=omp_alloctrail_val_kind), &
7      parameter :: omp_atv_memspace = 22
8
9      ! omp_alloctrail might not be provided in omp_lib.h.
10     type omp_alloctrail
11         integer(kind=omp_alloctrail_key_kind) key
12         integer(kind=omp_alloctrail_val_kind) value
13     end type omp_alloctrail
14
15     integer(kind=omp_memspace_handle_kind), &
16     parameter :: omp_null_mem_space = 0
17
18     integer(kind=omp_allocator_handle_kind), &
19     parameter :: omp_null_allocator = 0

```

Fortran

18.13.2 Routines to get memory spaces

Summary

The following routines return a memory space that represents a set of resources accessible by one or more devices.

Format

C / C++

```

25     omp_memspace_handle_t
26     omp_get_devices_memspace ( int ndevs, const int *devs,
27                               omp_memspace_handle_t memspace );
28
29     omp_memspace_handle_t
30     omp_get_device_memspace ( int dev,
31                               omp_memspace_handle_t memspace );
32
33     omp_memspace_handle_t
34     omp_get_devices_and_host_memspace ( int ndevs,
35                                         const int *devs,
36                                         omp_memspace_handle_t memspace );
37
38     omp_memspace_handle_t

```

```

1  omp_get_device_and_host_memspace ( int dev,
2                                     omp_memspace_handle_t memspace );
3
4  omp_memspace_handle_t
5  omp_get_devices_all_memspace ( omp_memspace_handle_t memspace );

```

▲ C / C++ ▲

▼ Fortran ▼

```

6  integer(kind=omp_memspace_handle_kind) &
7  function omp_get_devices_memspace ( ndevs, devs, memspace )
8  integer, intent(in) :: ndevs
9  integer, intent(in) :: devs(*)
10 integer(kind=omp_memspace_handle_kind), intent(in) :: memspace
11
12 integer(kind=omp_memspace_handle_kind) &
13 function omp_get_device_memspace ( dev, memspace )
14 integer, intent(in) :: dev
15 integer(kind=omp_memspace_handle_kind), intent(in) :: memspace
16
17 integer(kind=omp_memspace_handle_kind) &
18 function omp_get_devices_and_host_memspace ( ndevs, devs, memspace )
19 integer, intent(in) :: ndevs
20 integer, intent(in) :: devs(*)
21 integer(kind=omp_memspace_handle_kind), intent(in) :: memspace
22
23 integer(kind=omp_memspace_handle_kind) &
24 function omp_get_device_and_host_memspace ( dev, memspace )
25 integer, intent(in) :: dev
26 integer(kind=omp_memspace_handle_kind), intent(in) :: memspace
27
28 integer(kind=omp_memspace_handle_kind) &
29 function omp_get_devices_all_memspace ( memspace )
30 integer(kind=omp_memspace_handle_kind), intent(in) :: memspace

```

▲ Fortran ▲

Constraints on Arguments

The *memspace* argument must be one of the predefined memory spaces.

The *ndevs* argument to `omp_get_devices_memspace` and `omp_get_target_devices_and_host_memspace` must be greater than zero. The *devs* argument to `omp_get_devices_memspace` and `omp_get_devices_and_host_memspace` must point to an array that contains at least *ndevs* values. Each value must be a conforming device number. If there are more than *ndevs* values, the additional values will be ignored.

1 The *dev* argument to `omp_get_device_memspace` and
2 `omp_get_device_and_host_memspace` must be a conforming device number.

3 **Binding**

4 The binding thread set for these routines region is all threads on the device.

5 **Effect**

6 The effect of these routines is to return a handle to a memory space that represents a set of storage
7 resources such that for each storage resource the following requirements are true:

- 8 • The storage resource is accessible by each of the devices selected by the routine; and
- 9 • The storage resource is part of *memspace* in each of the devices selected by the routine.

10 If no set of storage resources matches the above requirements, then the special value
11 `omp_null_mem_space` is returned.

12 The devices selected by `omp_get_devices_memspace` are those specified in the *devs*
13 argument.

14 The device selected by `omp_get_device_memspace` is the device specified in the *dev*
15 argument.

16 The devices selected by `omp_get_devices_and_host_memspace` are those specified in the
17 *devs* argument and the initial device.

18 The device selected by `omp_get_device_and_host_memspace` are the device specified in
19 the *dev* argument and the initial device.

20 The devices selected by `omp_get_devices_all_memspace` are all available devices.

21 The memory spaces returned by these routine are target memory spaces if any of the selected
22 devices is not the current device.

23 **Restrictions**

24 The restrictions to these routines are as follows:

- 25 • These routines must only be invoked on the initial device.

26 **Cross References**

- 27 • `requires` directive, see [Section 8.2](#)
- 28 • `target` directive, see [Section 13.8](#)
- 29 • Memory Spaces, see [Section 6.1](#)

30 **18.13.3 omp_init_allocator**

31 **Summary**

32 The `omp_init_allocator` routine initializes an allocator and associates it with a memory
33 space.

Format

C / C++

```
omp_allocator_handle_t omp_init_allocator (
    omp_memspace_handle_t memspace,
    int ntraits,
    const omp_alloctrail_t traits[]
);
```

C / C++

Fortran

```
integer(kind=omp_allocator_handle_kind) &
function omp_init_allocator ( memspace, ntraits, traits )
integer(kind=omp_memspace_handle_kind), intent(in) :: memspace
integer, intent(in) :: ntraits
type(omp_alloctrail), intent(in) :: traits(*)
```

Fortran

Constraints on Arguments

The *memspace* argument must be a valid memory space handle or the value **omp_null_mem_space**. 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.

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 6.2 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.

If *memspace* has the value **omp_null_mem_space** the effect of this routine will be as if the value of *memspace* was **omp_default_mem_space**.

Restrictions

The restrictions to the `omp_init_allocator` routine are as follows:

- 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.
- 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.
- If `memspace` is a target memory space, the values `device`, `cgroup`, `pteam` or `thread` must not be specified for the `access` allocator trait.

Cross References

- `requires` directive, see [Section 8.2](#)
- `target` directive, see [Section 13.8](#)
- Memory Allocators, see [Section 6.2](#)
- Memory Spaces, see [Section 6.1](#)

18.13.4 Routines to Get Allocators

Summary

These routines return the default memory allocator for a given device for a certain kind of memory.

Format

C / C++

```
omp_allocator_handle_t
omp_get_devices_allocator ( int ndevs, const int *devs,
                           omp_memspace_handle_t memspace );

omp_allocator_handle_t
omp_get_device_allocator ( int dev,
                           omp_memspace_handle_t memspace );

omp_allocator_handle_t
omp_get_devices_and_host_allocator ( int ndevs,
                                     const int *devs,
                                     omp_memspace_handle_t memspace );

omp_allocator_handle_t
omp_get_device_and_host_allocator ( int dev,
                                    omp_memspace_handle_t memspace );

omp_allocator_handle_t
omp_get_devices_all_allocator ( omp_memspace_handle_t memspace );
```

C / C++

Fortran

```
1 integer(kind=omp_allocator_handle_kind) &
2 function omp_get_devices_allocator ( ndevs, devs, memspace )
3 integer, intent(in) :: ndevs
4 integer, intent(in) :: devs(*)
5 integer(kind=omp_memspace_handle_kind), intent(in) :: memspace
6
7 integer(kind=omp_allocator_handle_kind) &
8 function omp_get_device_allocator ( dev, memspace )
9 integer, intent(in) :: dev
10 integer(kind=omp_memspace_handle_kind), intent(in) :: memspace
11
12 integer(kind=omp_allocator_handle_kind) &
13 function omp_get_devices_and_host_allocator ( ndevs, devs, memspace
14 )
15 integer, intent(in) :: ndevs
16 integer, intent(in) :: devs(*)
17 integer(kind=omp_memspace_handle_kind), intent(in) :: memspace
18
19 integer(kind=omp_allocator_handle_kind) &
20 function omp_get_device_and_host_allocator ( dev, memspace )
21 integer, intent(in) :: dev
22 integer(kind=omp_memspace_handle_kind), intent(in) :: memspace
23
24 integer(kind=omp_allocator_handle_kind) &
25 function omp_get_devices_all_allocator ( memspace )
26 integer(kind=omp_memspace_handle_kind), intent(in) :: memspace
```

Fortran

27 Constraints on Arguments

28 The *memspace* argument must be one of the predefined memory spaces. The *ndevs* argument to
29 **omp_get_devices_allocator** and **omp_get_devices_and_host_allocator** must
30 be greater than zero. The *devs* argument to **omp_get_devices_allocator** and
31 **omp_get_devices_and_host_allocator** must point to an array that contains at least
32 *ndevs* values. Each value must be a conforming device number. If there are more than *ndevs* values,
33 the additional values will be ignored.

34 The *dev* argument to **omp_get_device_allocator** and
35 **omp_get_device_and_host_allocator** must be a conforming device number.

36 Binding

37 The binding thread set for these routines region is all threads on a device. The effect of executing
38 this routine is not related to any specific region that corresponds to any construct or API routine.

Effect

The effect of these routines is to return the predefined allocator for memory of kind *memspace* for the selected devices. If the implementation does not have a predefined allocator that satisfies the request, then the special value `omp_null_allocator` is returned.

The selected devices for `omp_get_devices_allocator` are those specified in the *devs* argument.

The selected device for `omp_get_device_allocator` is the device specified in the *dev* argument.

The selected devices for `omp_get_devices_and_host_allocator` are those specified in the *devs* argument and the initial device.

The selected devices for `omp_get_device_and_host_allocator` are the device specified in the *dev* argument and the initial device.

The selected devices for `omp_get_devices_all_allocator` are all available devices.

Each of these routines returns an allocator that may be used anywhere that requires a predefined allocator specified in Table 6.3. The allocator is associated with a target memory space if any of the selected devices is not the current device.

Restrictions

The restrictions to these routines are as follows:

- These routines can only be invoked on the initial device.

Cross References

- `requires` directive, see [Section 8.2](#)
- `target` directive, see [Section 13.8](#)
- Memory Allocators, see [Section 6.2](#)
- Memory Spaces, see [Section 6.1](#)

18.13.5 `omp_destroy_allocator`

Summary

The `omp_destroy_allocator` routine releases all resources used by the allocator handle.

Format

```
void omp_destroy_allocator (omp_allocator_handle_t allocator);  
  
subroutine omp_destroy_allocator ( allocator )  
integer(kind=omp_allocator_handle_kind), intent(in) :: allocator
```

Constraints on Arguments

The *allocator* argument must not represent a predefined memory allocator.

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. If *allocator* is **omp_null_allocator** then this routine will have no effect.

Restrictions

The restrictions to the **omp_destroy_allocator** routine are as follows:

- Accessing any memory allocated by the *allocator* after this call results in unspecified behavior.
- 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.

Cross References

- **requires** directive, see [Section 8.2](#)
- **target** directive, see [Section 13.8](#)
- Memory Allocators, see [Section 6.2](#)

18.13.6 omp_set_default_allocator

Summary

The **omp_set_default_allocator** routine sets the default memory allocator to be used by allocation calls, **allocate** clauses and **allocate** and **allocators** directives 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

- `allocate` clause, see [Section 6.6](#)
- `allocate` directive, see [Section 6.5](#)
- `allocators` directive, see [Section 6.7](#)
- Memory Allocators, see [Section 6.2](#)
- *def-allocator-var* ICV, see [Table 2.1](#)

18.13.7 `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` clauses and `allocate` and `allocators` directives that do not specify an allocator.

Format

```

C / C++
omp_allocator_handle_t omp_get_default_allocator (void);

C / C++
integer(kind=omp_allocator_handle_kind) &
function omp_get_default_allocator ()

Fortran
```

Binding

The binding task set for an `omp_get_default_allocator` region is the binding implicit task.

Effect

The effect of this routine is to return the value of the *def-allocator-var* ICV of the binding implicit task.

Cross References

- `allocate` clause, see [Section 6.6](#)
- `allocate` directive, see [Section 6.5](#)
- `allocators` directive, see [Section 6.7](#)
- Memory Allocators, see [Section 6.2](#)
- `def-allocator-var` ICV, see [Table 2.1](#)

18.13.8 `omp_alloc` and `omp_aligned_alloc`

Summary

The `omp_alloc` and `omp_aligned_alloc` routines request a memory allocation from a memory allocator.

Format

C

```
void *omp_alloc(size_t size, omp_allocator_handle_t allocator);
void *omp_aligned_alloc(
    size_t alignment,
    size_t size,
    omp_allocator_handle_t allocator);
```

C

C++

```
void *omp_alloc(
    size_t size,
    omp_allocator_handle_t allocator=omp_null_allocator
);
void *omp_aligned_alloc(
    size_t alignment,
    size_t size,
    omp_allocator_handle_t allocator=omp_null_allocator
);
```

C++

Fortran

```
1  type(c_ptr) function omp_alloc(size, allocator) bind(c)
2  use, intrinsic :: iso_c_binding, only : c_ptr, c_size_t
3  integer(c_size_t), value :: size
4  integer(omp_allocator_handle_kind), value :: allocator
5
6  type(c_ptr) function omp_aligned_alloc(alignment, &
7  size, allocator) bind(c)
8  use, intrinsic :: iso_c_binding, only : c_ptr, c_size_t
9  integer(c_size_t), value :: alignment, size
10 integer(omp_allocator_handle_kind), value :: allocator
```

Fortran

Constraints on Arguments

Unless **dynamic_allocators** appears on a **requires** directive in the same compilation unit, **omp_alloc** and **omp_aligned_alloc** invocations that appear in **target** regions must not pass **omp_null_allocator** as the *allocator* argument, which must be a constant expression that evaluates to one of the predefined memory allocator values. The *alignment* argument to **omp_aligned_alloc** must be a power of two and the *size* argument must be a multiple of *alignment*.

Binding

The binding task set for an **omp_alloc** or **omp_aligned_alloc** region is the generating task.

Effect

The **omp_alloc** and **omp_aligned_alloc** routines request 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 routines will be the one specified by the *def-allocator-var* ICV of the binding implicit task. Upon success they return a pointer to the allocated memory. Otherwise, the behavior that the **fallback** trait of the allocator specifies will be followed. If *size* is 0, **omp_alloc** and **omp_aligned_alloc** will return *NULL*.

Memory allocated by **omp_alloc** will be byte-aligned to at least the maximum of the alignment required by **malloc** and the **alignment** trait of the allocator. Memory allocated by **omp_aligned_alloc** will be byte-aligned to at least the maximum of the alignment required by **malloc**, the **alignment** trait of the allocator and the *alignment* argument value.

Pointers returned by these routines are considered device pointers if at least one of the devices associated with the allocator is not the current device.

Fortran

33 The **omp_alloc** and **omp_aligned_alloc** routines require an explicit interface and so might
34 not be provided in **omp_lib.h**.

Fortran

Cross References

- **requires** directive, see [Section 8.2](#)
- **target** directive, see [Section 13.8](#)
- Memory Allocators, see [Section 6.2](#)
- *def-allocator-var* ICV, see [Table 2.1](#)

18.13.9 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++

Fortran

```
subroutine omp_free(ptr, allocator) bind(c)  
use, intrinsic :: iso_c_binding, only : c_ptr  
type(c_ptr), value :: ptr  
integer(omp_allocator_handle_kind), value :: allocator
```

Fortran

Binding

The binding task set for an **omp_free** region is the generating task.

Effect

The **omp_free** routine deallocates the memory to which *ptr* points. The *ptr* argument must have been returned by an OpenMP allocation routine. 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. If *ptr* is *NULL*, no operation is performed.

Fortran

The **omp_free** routine requires an explicit interface and so might not be provided in **omp_lib.h**.

Fortran

Restrictions

The restrictions to the `omp_free` routine are as follows:

- 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 6.2](#)
- `omp_destroy_allocator`, see [Section 18.13.5](#)

18.13.10 `omp_calloc` and `omp_aligned_calloc`

Summary

The `omp_calloc` and `omp_aligned_calloc` routines request a zero initialized memory allocation from a memory allocator.

Format

C

```
void *omp_calloc(  
    size_t nmemb,  
    size_t size,  
    omp_allocator_handle_t allocator  
);  
void *omp_aligned_calloc(  
    size_t alignment,  
    size_t nmemb,  
    size_t size,  
    omp_allocator_handle_t allocator  
);
```

C

C++

```
void *omp_calloc(  
    size_t nmemb,  
    size_t size,  
    omp_allocator_handle_t allocator=omp_null_allocator  
);  
void *omp_aligned_calloc(  
    size_t alignment,  
    size_t nmemb,  
    size_t size,  
    omp_allocator_handle_t allocator=omp_null_allocator  
);
```

C++

Fortran

```
1 type(c_ptr) function omp_malloc(nmemb, size, allocator) bind(c)
2 use, intrinsic :: iso_c_binding, only : c_ptr, c_size_t
3 integer(c_size_t), value :: nmemb, size
4 integer(omp_allocator_handle_kind), value :: allocator
5
6 type(c_ptr) function omp_aligned_malloc(alignment, nmemb, size, &
7     allocator) bind(c)
8 use, intrinsic :: iso_c_binding, only : c_ptr, c_size_t
9 integer(c_size_t), value :: alignment, nmemb, size
10 integer(omp_allocator_handle_kind), value :: allocator
```

Fortran

Constraints on Arguments

Unless **dynamic_allocators** appears on a **requires** directive in the same compilation unit, **omp_malloc** and **omp_aligned_malloc** invocations that appear in **target** regions must not pass **omp_null_allocator** as the *allocator* argument, which must be a constant expression that evaluates to one of the predefined memory allocator values. The *alignment* argument to **omp_aligned_malloc** must be a power of two and the *size* argument must be a multiple of *alignment*.

Binding

The binding task set for an **omp_malloc** or **omp_aligned_malloc** region is the generating task.

Effect

The **omp_malloc** and **omp_aligned_malloc** routines request a memory allocation from the specified memory allocator for an array of *nmemb* elements each of which has a size of *size* bytes. If the *allocator* argument is **omp_null_allocator** the memory allocator used by the routines will be the one specified by the *def-allocator-var* ICV of the binding implicit task. Upon success they return a pointer to the allocated memory. Otherwise, the behavior that the **fallback** trait of the allocator specifies will be followed. Any memory allocated by these routines will be set to zero before returning. If either *nmemb* or *size* is 0, **omp_malloc** will return *NULL*.

Memory allocated by **omp_malloc** will be byte-aligned to at least the maximum of the alignment required by **malloc** and the **alignment** trait of the allocator. Memory allocated by **omp_aligned_malloc** will be byte-aligned to at least the maximum of the alignment required by **malloc**, the **alignment** trait of the allocator and the *alignment* argument value.

Fortran

The **omp_malloc** and **omp_aligned_malloc** routines require an explicit interface and so might not be provided in **omp_lib.h**.

Fortran

Cross References

- `requires` directive, see [Section 8.2](#)
- `target` directive, see [Section 13.8](#)
- Memory Allocators, see [Section 6.2](#)
- `def-allocator-var` ICV, see [Table 2.1](#)

18.13.11 `omp_realloc`

Summary

The `omp_realloc` routine deallocates previously allocated memory and requests a memory allocation from a memory allocator.

Format

C

```
void *omp_realloc(  
    void *ptr,  
    size_t size,  
    omp_allocator_handle_t allocator,  
    omp_allocator_handle_t free_allocator  
);
```

C

C++

```
void *omp_realloc(  
    void *ptr,  
    size_t size,  
    omp_allocator_handle_t allocator=omp_null_allocator,  
    omp_allocator_handle_t free_allocator=omp_null_allocator  
);
```

C++

Fortran

```
type(c_ptr) &  
function omp_realloc(ptr, size, allocator, free_allocator) bind(c)  
use, intrinsic :: iso_c_binding, only : c_ptr, c_size_t  
type(c_ptr), value :: ptr  
integer(c_size_t), value :: size  
integer(omp_allocator_handle_kind), value :: allocator, free_allocator
```

Fortran

Constraints on Arguments

Unless a **dynamic_allocators** clause appears on a **requires** directive in the same compilation unit, **omp_realloc** invocations that appear in **target** regions must not pass **omp_null_allocator** as the *allocator* or *free_allocator* argument, which must be constant expressions that evaluate to one of the predefined memory allocator values.

Binding

The binding task set for an **omp_realloc** region is the generating task.

Effect

The **omp_realloc** routine deallocates the memory to which *ptr* points and requests a new memory allocation of *size* bytes from the specified memory allocator. If the *free_allocator* argument is specified, it must be the memory allocator to which the previous allocation request was made. If the *free_allocator* argument is **omp_null_allocator** the implementation will determine that value automatically. If the *allocator* argument is **omp_null_allocator** the behavior is as if the memory allocator that allocated the memory to which *ptr* argument points is passed to the *allocator* argument. Upon success it returns a (possibly moved) pointer to the allocated memory and the contents of the new object shall be the same as that of the old object prior to deallocation, up to the minimum size of old allocated size and *size*. Any bytes in the new object beyond the old allocated size will have unspecified values. If the allocation failed, the behavior that the **fallback** trait of the *allocator* specifies will be followed. If *ptr* is *NULL*, **omp_realloc** will behave the same as **omp_alloc** with the same *size* and *allocator* arguments. If *size* is 0, **omp_realloc** will return *NULL* and the old allocation will be deallocated. If *size* is not 0, the old allocation will be deallocated if and only if the function returns a non-null value.

Memory allocated by **omp_realloc** will be byte-aligned to at least the maximum of the alignment required by **malloc** and the **alignment** trait of the allocator.

Fortran

The **omp_realloc** routine requires an explicit interface and so might not be provided in **omp_lib.h**.

Fortran

Restrictions

The restrictions to the **omp_realloc** routine are as follows:

- The *ptr* argument must have been returned by an OpenMP allocation routine.
- Using **omp_realloc** 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

- `requires` directive, see [Section 8.2](#)
- `target` directive, see [Section 13.8](#)
- Memory Allocators, see [Section 6.2](#)
- `omp_alloc` and `omp_aligned_alloc`, see [Section 18.13.8](#)
- `omp_destroy_allocator`, see [Section 18.13.5](#)

18.13.12 `omp_get_memspace_num_resources`

Summary

The `omp_get_memspace_num_resources` routine returns the number of resources associated with the specified memory space.

Format

C / C++

```
int omp_get_memspace_num_resources (
    omp_memspace_handle_t memspace
);
```

C / C++

Fortran

```
integer &
function omp_get_memspace_num_resources ( memspace )
integer(kind=omp_memspace_handle_kind), intent(in) :: memspace
```

Fortran

Constraints on Arguments

The *memspace* argument must be a valid memory space.

Binding

The binding thread set for an `omp_get_memspace_num_resources` 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_get_memspace_num_resources` returns the number of distinct storage resources that are associated with the memory space represented by the *memspace* handle.

Cross References

- Memory Spaces, see [Section 6.1](#)

18.13.13 `omp_get_submemspace`

Summary

The `omp_get_submemspace` routine returns a new memory space that contains a subset of the resources of the original memory space.

Format

C / C++

```
omp_memspace_handle_t omp_get_submemspace (  
    omp_memspace_handle_t memspace,  
    int num_resources,  
    int *resources  
);
```

C / C++

Fortran

```
integer(kind=omp_memspace_handle_kind) &  
function omp_get_submemspace ( memspace, num_resources, resources )  
integer(kind=omp_memspace_handle_kind), intent(in) :: memspace  
integer, intent(in):: num_resources  
integer, intent(in):: resources(*)
```

Fortran

Constraints on Arguments

The *memspace* argument must be a valid memory space.

The *num_resources* argument must be a non-negative value.

The *resources* array must contain at least as many entries as specified by the *num_resources* argument. Each entry value must be a value between 0 and the number of resources associated with *memspace* minus 1.

Binding

The binding thread set for an `omp_get_submemspace` 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_get_submemspace` returns a new memory space that represents only the resources of *memspace* that are specified by the *resources* argument.

If *num_resources* is zero or a memory space cannot be created for the requested resources the special value `omp_null_mem_space` is returned.

Cross References

- Memory Spaces, see [Section 6.1](#)

18.14 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`

Constraints on Arguments

The following enumeration type defines four standard commands. Table 18.3 describes the actions that these commands request from a tool.

C / C++
`typedef enum omp_control_tool_t {
 omp_control_tool_start = 1,
 omp_control_tool_pause = 2,
 omp_control_tool_flush = 3,
 omp_control_tool_end = 4
} omp_control_tool_t;`

C / C++
Fortran
`integer (kind=omp_control_tool_kind), &
 parameter :: omp_control_tool_start = 1
integer (kind=omp_control_tool_kind), &
 parameter :: omp_control_tool_pause = 2
integer (kind=omp_control_tool_kind), &
 parameter :: omp_control_tool_flush = 3
integer (kind=omp_control_tool_kind), &
 parameter :: omp_control_tool_end = 4`

Tool-specific values for *command* must be greater or equal to 64. Tools must ignore *command* values that they are not explicitly designed to handle. Other values accepted by a tool for *command*, and any values for *modifier* and *arg* are tool-defined.

TABLE 18.3: 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 **Binding**

2 The binding task set for an `omp_control_tool` region is the generating task.

3 **Effect**

4 An OpenMP program may use `omp_control_tool` to pass commands to a tool. An application
 5 can use `omp_control_tool` to request that a tool starts or restarts data collection when a code
 6 region of interest is encountered, that a tool pauses data collection when leaving the region of
 7 interest, that a tool flushes any data that it has collected so far, or that a tool ends data collection.
 8 Additionally, `omp_control_tool` can be used to pass tool-specific commands to a particular
 9 tool. The following types correspond to return values from `omp_control_tool`:

10

```

11 typedef enum omp_control_tool_result_t {
12     omp_control_tool_notool = -2,
13     omp_control_tool_nocallback = -1,
14     omp_control_tool_success = 0,
15     omp_control_tool_ignored = 1
16 } omp_control_tool_result_t;

```

15

Fortran

```
1 integer (kind=omp_control_tool_result_kind), &  
2     parameter :: omp_control_tool_notool = -2  
3 integer (kind=omp_control_tool_result_kind), &  
4     parameter :: omp_control_tool_nocallback = -1  
5 integer (kind=omp_control_tool_result_kind), &  
6     parameter :: omp_control_tool_success = 0  
7 integer (kind=omp_control_tool_result_kind), &  
8     parameter :: omp_control_tool_ignored = 1
```

Fortran

9 If the OMPT interface state is inactive, the OpenMP implementation returns
10 **omp_control_tool_notool**. If the OMPT interface state is active, but no callback is
11 registered for the *tool-control* event, the OpenMP implementation returns
12 **omp_control_tool_nocallback**. An OpenMP implementation may return other
13 implementation-defined negative values strictly smaller than -64; an application may assume that
14 any negative return value indicates that a tool has not received the command. A return value of
15 **omp_control_tool_success** indicates that the tool has performed the specified command. A
16 return value of **omp_control_tool_ignored** indicates that the tool has ignored the specified
17 command. A tool may return other positive values strictly greater than 64 that are tool-defined.

18 Execution Model Events

19 The *tool-control* event occurs in the thread that encounters a call to **omp_control_tool** at a
20 point inside its corresponding OpenMP region.

21 Tool Callbacks

22 A thread dispatches a registered **ompt_callback_control_tool** callback for each
23 occurrence of a *tool-control* event. The callback executes in the context of the call that occurs in the
24 user program and has type signature **ompt_callback_control_tool_t**. The callback may
25 return any non-negative value, which will be returned to the application by the OpenMP
26 implementation as the return value of the **omp_control_tool** call that triggered the callback.

27 Arguments passed to the callback are those passed by the user to **omp_control_tool**. If the
28 call is made in Fortran, the tool will be passed *NULL* as the third argument to the callback. If any of
29 the four standard commands is presented to a tool, the tool will ignore the *modifier* and *arg*
30 argument values.

31 Restrictions

32 Restrictions on access to the state of an OpenMP first-party tool are as follows:

- 33 • An application may access the tool state modified by an OMPT callback only by using
34 **omp_control_tool**.

Cross References

- OMPT Interface, see [Chapter 19](#)
- `ompt_callback_control_tool_t`, see [Section 19.5.2.29](#)

18.15 Environment Display Routine

Summary

The `omp_display_env` routine displays the OpenMP version number and the initial values of ICVs associated with the environment variables described in [Chapter 21](#).

Format

```
void omp_display_env(int verbose);
```

C / C++

```
subroutine omp_display_env(verbose)
logical, intent(in) :: verbose
```

Fortran

Binding

The binding thread set for an `omp_display_env` region is the encountering thread.

Effect

Each time the `omp_display_env` routine is invoked, the runtime system prints the OpenMP version number and the initial values of the ICVs associated with the environment variables described in [Chapter 21](#). The displayed values are the values of the ICVs after they have been modified according to the environment variable settings and before the execution of any OpenMP construct or API routine.

The display begins with "OPENMP DISPLAY ENVIRONMENT BEGIN", followed by the `_OPENMP` version macro (or the `openmp_version` named constant for Fortran) and ICV values, in the format `NAME '=' VALUE`. `NAME` corresponds to the macro or environment variable name, optionally prepended with a bracketed `DEVICE`. `VALUE` corresponds to the value of the macro or ICV associated with this environment variable. Values are enclosed in single quotes. `DEVICE` corresponds to the device on which the value of the ICV is applied. The display is terminated with "OPENMP DISPLAY ENVIRONMENT END".

If the `verbose` argument evaluates to *false*, the runtime displays the OpenMP version number defined by the `_OPENMP` version macro (or the `openmp_version` named constant for Fortran) value and the initial ICV values for the environment variables listed in [Chapter 21](#). If the `verbose` argument evaluates to *true*, the runtime may also display the values of vendor-specific ICVs that may be modified by vendor-specific environment variables.

Example output:


```
1 OPENMP DISPLAY ENVIRONMENT BEGIN
2   _OPENMP='202111'
3   [host] OMP_SCHEDULE='GUIDED,4'
4   [host] OMP_NUM_THREADS='4,3,2'
5   [device] OMP_NUM_THREADS='2'
6   [host,device] OMP_DYNAMIC='TRUE'
7   [host] OMP_PLACES='{0:4},{4:4},{8:4},{12:4}'
8   ...
9 OPENMP DISPLAY ENVIRONMENT END
```

10 Restrictions

11 Restrictions to the `omp_display_env` routine are as follows.

- 12 • When called from within a `target` region the effect is unspecified.

13 Cross References

- 14 • `OMP_DISPLAY_ENV`, see [Section 21.7](#)

19 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.

19.1 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 `omp-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++

19.2 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.

19.2.1 `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

```
ompt_start_tool_result_t *ompt_start_tool(  
    unsigned int omp_version,  
    const char *runtime_version  
);
```

Semantics

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

- Tool Initialization and Finalization, see [Section 19.4.1](#)

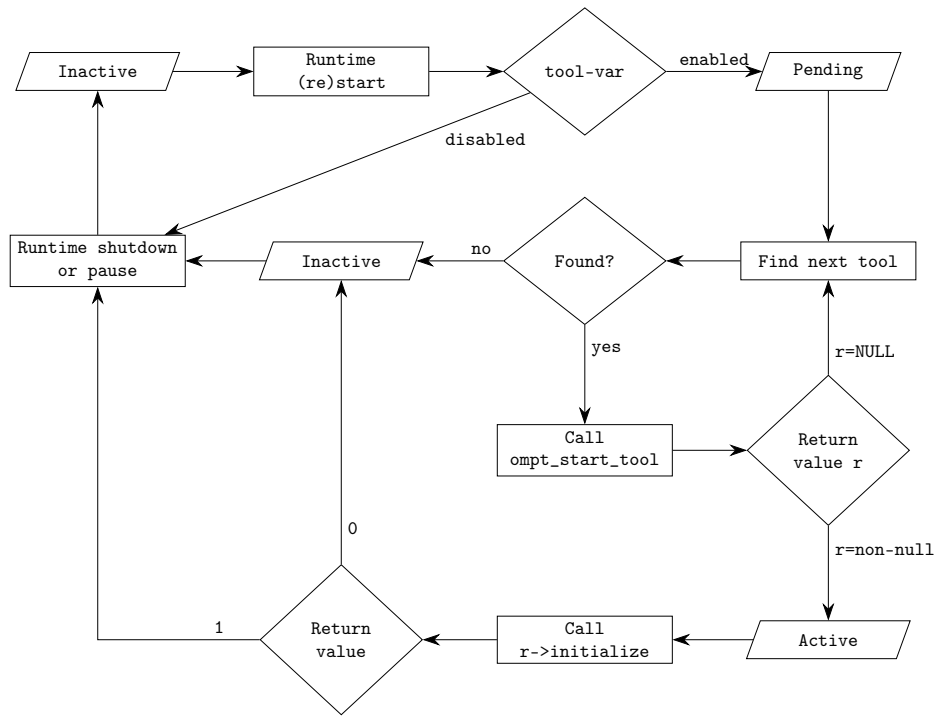


FIGURE 19.1: First-Party Tool Activation Flow Chart

19.2.2 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

1 includes a 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 21.3.2](#), 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 Initialization and Finalization, see [Section 19.4.1](#)
- 23 • *tool-libraries-var* ICV, see [Table 2.1](#)
- 24 • *tool-var* ICV, see [Table 2.1](#)
- 25 • `ompt_start_tool`, see [Section 19.2.1](#)

26 19.2.3 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 19.5.1.1](#), uses the function specified in its *lookup* argument
32 to look up pointers to OMPT interface runtime entry points that the OpenMP implementation
33 provides; this process is described in [Section 19.2.3.1](#). Typically, a tool initializer obtains a pointer
34 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 19.2.4](#).

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 • Tool Initialization and Finalization, see [Section 19.4.1](#)
- 10 • `ompt_enumerate_mutex_impls_t`, see [Section 19.6.1.2](#)
- 11 • `ompt_enumerate_states_t`, see [Section 19.6.1.1](#)
- 12 • `ompt_set_callback_t`, see [Section 19.6.1.3](#)
- 13 • `ompt_start_tool`, see [Section 19.2.1](#)

14 **19.2.3.1 Binding Entry Points in the OMPT Callback Interface**

15 Functions that an OpenMP implementation provides to support the OMPT interface are not defined
16 as global function symbols. Instead, they are defined as runtime entry points that a tool can only
17 identify through the *lookup* function that is provided as an argument with type signature
18 `ompt_function_lookup_t` to the tool initializer. A tool can use this function to obtain a
19 pointer to each of the runtime entry points that an OpenMP implementation provides to support the
20 OMPT interface. Once a tool has obtained a *lookup* function, it may employ it at any point in the
21 future.

22 For each runtime entry point in the OMPT interface for the host device, Table 19.1 provides the
23 string name by which it is known and its associated type signature. Implementations can provide
24 additional implementation-specific names and corresponding entry points. Any names that begin
25 with `ompt_` are reserved names.

26 During initialization, a tool should look up each runtime entry point in the OMPT interface by
27 name and bind a pointer maintained by the tool that can later be used to invoke the entry point. The
28 entry points described in Table 19.1 enable a tool to assess the thread states and mutual exclusion
29 implementations that an OpenMP implementation supports to register tool callbacks, to inspect
30 registered callbacks, to introspect OpenMP state associated with threads, and to use tracing to
31 monitor computations that execute on target devices.

32 Detailed information about each runtime entry point listed in Table 19.1 is included as part of the
33 description of its type signature.

TABLE 19.1: OMPT Callback Interface Runtime Entry Point Names and Their Type Signatures

Entry Point String Name	Type signature
"ompt_enumerate_states"	ompt_enumerate_states_t
"ompt_enumerate_mutex_impls"	ompt_enumerate_mutex_impls_t
"ompt_set_callback"	ompt_set_callback_t
"ompt_get_callback"	ompt_get_callback_t
"ompt_get_thread_data"	ompt_get_thread_data_t
"ompt_get_num_places"	ompt_get_num_places_t
"ompt_get_place_proc_ids"	ompt_get_place_proc_ids_t
"ompt_get_place_num"	ompt_get_place_num_t
"ompt_get_partition_place_nums"	ompt_get_partition_place_nums_t
"ompt_get_proc_id"	ompt_get_proc_id_t
"ompt_get_state"	ompt_get_state_t
"ompt_get_parallel_info"	ompt_get_parallel_info_t
"ompt_get_task_info"	ompt_get_task_info_t
"ompt_get_task_memory"	ompt_get_task_memory_t
"ompt_get_num_devices"	ompt_get_num_devices_t
"ompt_get_num_procs"	ompt_get_num_procs_t
"ompt_get_target_info"	ompt_get_target_info_t
"ompt_get_unique_id"	ompt_get_unique_id_t
"ompt_finalize_tool"	ompt_finalize_tool_t

Cross References

1
2
3
4
5
6
7
8
9
10
11
12
13
14

- `ompt_enumerate_mutex_impls_t`, see [Section 19.6.1.2](#)
- `ompt_enumerate_states_t`, see [Section 19.6.1.1](#)
- Lookup Entry Points: `ompt_function_lookup_t`, see [Section 19.6.3](#)
- `ompt_get_callback_t`, see [Section 19.6.1.4](#)
- `ompt_get_num_devices_t`, see [Section 19.6.1.17](#)
- `ompt_get_num_places_t`, see [Section 19.6.1.7](#)
- `ompt_get_num_procs_t`, see [Section 19.6.1.6](#)
- `ompt_get_parallel_info_t`, see [Section 19.6.1.13](#)
- `ompt_get_partition_place_nums_t`, see [Section 19.6.1.10](#)
- `ompt_get_place_num_t`, see [Section 19.6.1.9](#)
- `ompt_get_place_proc_ids_t`, see [Section 19.6.1.8](#)
- `ompt_get_proc_id_t`, see [Section 19.6.1.11](#)
- `ompt_get_state_t`, see [Section 19.6.1.12](#)

- 1 • `ompt_get_target_info_t`, see [Section 19.6.1.16](#)
- 2 • `ompt_get_task_info_t`, see [Section 19.6.1.14](#)
- 3 • `ompt_get_task_memory_t`, see [Section 19.6.1.15](#)
- 4 • `ompt_get_thread_data_t`, see [Section 19.6.1.5](#)
- 5 • `ompt_get_unique_id_t`, see [Section 19.6.1.18](#)
- 6 • `ompt_set_callback_t`, see [Section 19.6.1.3](#)

19.2.4 Monitoring Activity on the Host with OMPT

To monitor the execution of an OpenMP program on the host device, a tool initializer must register to receive notification of events that occur as an OpenMP program executes. A tool can use the `ompt_set_callback` runtime entry point to register callbacks for OpenMP events. The return codes for `ompt_set_callback` use the `ompt_set_result_t` enumeration type. If the `ompt_set_callback` runtime entry point is called outside a tool initializer, registration of supported callbacks may fail with a return value of `ompt_set_error`.

All callbacks registered with `ompt_set_callback` or returned by `ompt_get_callback` use the dummy type signature `ompt_callback_t`.

For callbacks listed in Table 19.2, `ompt_set_always` is the only registration return code that is allowed. An OpenMP implementation must guarantee that the callback will be invoked every time that a runtime event that is associated with it occurs. Support for such callbacks is required in a minimal implementation of the OMPT interface.

For callbacks listed in Table 19.3, the `ompt_set_callback` runtime entry may return any non-error code. Whether an OpenMP implementation invokes a registered callback never, sometimes, or always is implementation defined. If registration for a callback allows a return code of `ompt_set_never`, support for invoking such a callback may not be present in a minimal implementation of the OMPT interface. The return code from registering a callback indicates the implementation-defined level of support for the callback.

Two techniques reduce the size of the OMPT interface. First, in cases where events are naturally paired, for example, the beginning and end of a region, and the arguments needed by the callback at each endpoint are identical, a tool registers a single callback for the pair of events, with `ompt_scope_begin` or `ompt_scope_end` provided as an argument to identify for which endpoint the callback is invoked. Second, when a class of events is amenable to uniform treatment, OMPT provides a single callback for that class of events, for example, an `ompt_callback_sync_region_wait` callback is used for multiple kinds of synchronization regions, such as barrier, taskwait, and taskgroup regions. Some events, for example, `ompt_callback_sync_region_wait`, use both techniques.

Cross References

- `ompt_get_callback_t`, see [Section 19.6.1.4](#)

TABLE 19.2: Callbacks for which `ompt_set_callback` Must Return `ompt_set_always`

Callback Name
<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_emi</code>
<code>ompt_callback_target_data_op</code>
<code>ompt_callback_target_data_op_emi</code>
<code>ompt_callback_target_submit</code>
<code>ompt_callback_target_submit_emi</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>

- 1 • `ompt_set_callback_t`, see [Section 19.6.1.3](#)
- 2 • `ompt_set_result_t`, see [Section 19.4.4.2](#)

3 19.2.5 Tracing Activity on Target Devices with OMPT

4 A target device may or may not initialize a full OpenMP runtime system. Unless it does,
5 monitoring activity on a device using a tool interface based on callbacks may not be possible. To
6 accommodate such cases, the OMPT interface defines a monitoring interface for tracing activity on
7 target devices. Tracing activity on a target device involves the following steps:

- 8 • To prepare to trace device activity, a tool must register for an
9 `ompt_callback_device_initialize` callback. A tool may also register for an
10 `ompt_callback_device_load` callback to be notified when code is loaded onto a
11 target device or an `ompt_callback_device_unload` callback to be notified when
12 code is unloaded from a target device. A tool may also optionally register an
13 `ompt_callback_device_finalize` callback.
- 14 • When an OpenMP implementation initializes a target device, the OpenMP implementation
15 dispatches the device initialization callback of the tool on the host device. If the OpenMP
16 implementation or target device does not support tracing, the OpenMP implementation

TABLE 19.3: Callbacks for which `ompt_set_callback` May Return Any Non-Error Code

Callback Name
<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_masked</code>
<code>ompt_callback_target_map</code>
<code>ompt_callback_target_map_emi</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>

1 passes *NULL* to the device initializer of the tool for its *lookup* argument; otherwise, the
2 OpenMP implementation passes a pointer to a device-specific runtime entry point with type
3 signature `ompt_function_lookup_t` to the device initializer of the tool.

4 • If a non-null *lookup* pointer is provided to the device initializer of the tool, the tool may use it
5 to determine the runtime entry points in the tracing interface that are available for the device
6 and may bind the returned function pointers to tool variables. Table 19.4 indicates the names
7 of runtime entry points that may be available for a device; an implementation may provide
8 additional implementation-defined names and corresponding entry points. The driver for the
9 device provides the runtime entry points that enable a tool to control the trace collection
10 interface of the device. The *native* trace format that the interface uses may be device specific
11 and the available kinds of trace records are implementation defined. Some devices may allow
12 a tool to collect traces of records in a standard format known as OMPT trace records. Each
13 OMPT trace record serves as a substitute for an OMPT callback that is not appropriate to be
14 dispatched on the device. The fields in each trace record type are defined in the description of
15 the callback that the record represents. If this type of record is provided then the *lookup*
16 function returns values for the runtime entry points `ompt_set_trace_ompt` and
17 `ompt_get_record_ompt`, which support collecting and decoding OMPT traces. If the
18 native tracing format for a device is the OMPT format then tracing can be controlled using
19 the runtime entry points for native or OMPT tracing.

TABLE 19.4: OMPT Tracing Interface Runtime Entry Point Names and Their Type Signatures

Entry Point String Name	Type Signature
"ompt_get_device_num_procs"	ompt_get_device_num_procs_t
"ompt_get_device_time"	ompt_get_device_time_t
"ompt_translate_time"	ompt_translate_time_t
"ompt_set_trace_ompt"	ompt_set_trace_ompt_t
"ompt_set_trace_native"	ompt_set_trace_native_t
"ompt_start_trace"	ompt_start_trace_t
"ompt_pause_trace"	ompt_pause_trace_t
"ompt_flush_trace"	ompt_flush_trace_t
"ompt_stop_trace"	ompt_stop_trace_t
"ompt_advance_buffer_cursor"	ompt_advance_buffer_cursor_t
"ompt_get_record_type"	ompt_get_record_type_t
"ompt_get_record_ompt"	ompt_get_record_ompt_t
"ompt_get_record_native"	ompt_get_record_native_t
"ompt_get_record_abstract"	ompt_get_record_abstract_t

- 1 • The tool uses the **ompt_set_trace_native** and/or the **ompt_set_trace_ompt**
2 runtime entry point to specify what types of events or activities to monitor on the device. The
3 return codes for **ompt_set_trace_ompt** and **ompt_set_trace_native** use the
4 **ompt_set_result_t** enumeration type. If the **ompt_set_trace_native** or the
5 **ompt_set_trace_ompt** runtime entry point is called outside a device initializer,
6 registration of supported callbacks may fail with a return code of **ompt_set_error**.
- 7 • The tool initiates tracing of device activity by invoking **ompt_start_trace**. Arguments
8 to **ompt_start_trace** include two tool callbacks through which the OpenMP
9 implementation can manage traces associated with the device. One callback allocates a buffer
10 in which device activity can be deposited. The second callback processes a buffer of trace
11 events from the device.
- 12 • If the OpenMP implementation requires a trace buffer for device activity, the OpenMP
13 implementation invokes the tool-supplied callback function on the host device to request a
14 new buffer.
- 15 • The OpenMP implementation monitors the execution of OpenMP constructs on the device
16 and records a trace of events or activities into a trace buffer. If possible, device trace records
17 are marked with a *host_op_id*—an identifier that associates device activities with the target
18 operation that the host initiated to cause these activities. To correlate activities on the host
19 with activities on a device, a tool can register a
20 **ompt_callback_target_submit_emi** callback. Before and after the host initiates
21 creation of an initial task on a device associated with a structured block for a **target**
22 construct, the OpenMP implementation dispatches the
23 **ompt_callback_target_submit_emi** callback on the host in the thread that is
24 executing the task that encounters the **target** construct. This callback provides the tool

1 with a pair of identifiers: one that identifies the **target** region and a second that uniquely
2 identifies the initial task associated with that region. These identifiers help the tool correlate
3 activities on the target device with their **target** region.

- 4 • When appropriate, for example, when a trace buffer fills or needs to be flushed, the OpenMP
5 implementation invokes the tool-supplied buffer completion callback to process a non-empty
6 sequence of records in a trace buffer that is associated with the device.
- 7 • The tool-supplied buffer completion callback may return immediately, ignoring records in the
8 trace buffer, or it may iterate through them using the **ompt_advance_buffer_cursor**
9 entry point to inspect each record. A tool may use the **ompt_get_record_type** runtime
10 entry point to inspect the type of the record at the current cursor position. Three runtime
11 entry points (**ompt_get_record_ompt**, **ompt_get_record_native**, and
12 **ompt_get_record_abstract**) allow tools to inspect the contents of some or all
13 records in a trace buffer. The **ompt_get_record_native** runtime entry point uses the
14 native trace format of the device. The **ompt_get_record_abstract** runtime entry
15 point decodes the contents of a native trace record and summarizes them as an
16 **ompt_record_abstract_t** record. The **ompt_get_record_ompt** runtime entry
17 point can only be used to retrieve records in OMPT format.
- 18 • Once device tracing has been started, a tool may pause or resume device tracing at any time
19 by invoking **ompt_pause_trace** with an appropriate flag value as an argument.
- 20 • A tool may invoke the **ompt_flush_trace** runtime entry point for a device at any time
21 between device initialization and finalization to cause the pending trace records for that
22 device to be flushed.
- 23 • At any time, a tool may use the **ompt_start_trace** runtime entry point to start or the
24 **ompt_stop_trace** runtime entry point to stop device tracing. When device tracing is
25 stopped, the OpenMP implementation eventually gathers all trace records already collected
26 from device tracing and presents them to the tool using the buffer completion callback.
- 27 • An OpenMP implementation can be shut down while device tracing is in progress.
- 28 • When an OpenMP implementation is shut down, it finalizes each device. Device finalization
29 occurs in three steps. First, the OpenMP implementation halts any tracing in progress for the
30 device. Second, the OpenMP implementation flushes all trace records collected for the
31 device and uses the buffer completion callback associated with that device to present them to
32 the tool. Finally, the OpenMP implementation dispatches any
33 **ompt_callback_device_finalize** callback registered for the device.

34 **Restrictions**

35 Restrictions on tracing activity on devices are as follows:

- 36 • Implementation-defined names must not start with the prefix **ompt_**, which is reserved for
37 the OpenMP specification.

Cross References

- `ompt_advance_buffer_cursor_t`, see [Section 19.6.2.11](#)
- `ompt_callback_device_finalize_t`, see [Section 19.5.2.20](#)
- `ompt_callback_device_initialize_t`, see [Section 19.5.2.19](#)
- `ompt_flush_trace_t`, see [Section 19.6.2.9](#)
- `ompt_get_device_num_procs_t`, see [Section 19.6.2.1](#)
- `ompt_get_device_time_t`, see [Section 19.6.2.2](#)
- `ompt_get_record_abstract_t`, see [Section 19.6.2.15](#)
- `ompt_get_record_native_t`, see [Section 19.6.2.14](#)
- `ompt_get_record_ompt_t`, see [Section 19.6.2.13](#)
- `ompt_get_record_type_t`, see [Section 19.6.2.12](#)
- `ompt_pause_trace_t`, see [Section 19.6.2.8](#)
- `ompt_set_trace_native_t`, see [Section 19.6.2.5](#)
- `ompt_set_trace_ompt_t`, see [Section 19.6.2.4](#)
- `ompt_start_trace_t`, see [Section 19.6.2.7](#)
- `ompt_stop_trace_t`, see [Section 19.6.2.10](#)
- `ompt_translate_time_t`, see [Section 19.6.2.3](#)

19.3 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`, see [Section 19.5.1.2](#)

19.4 OMPT Data Types

The C/C++ header file (`omp-tools.h`) provides the definitions of the types that are specified throughout this subsection.

19.4.1 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

Restrictions to the `ompt_start_tool_result_t` type are as follows:

- 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_data_t`, see [Section 19.4.4.4](#)
- `ompt_finalize_t`, see [Section 19.5.1.2](#)
- `ompt_initialize_t`, see [Section 19.5.1.1](#)
- `ompt_start_tool`, see [Section 19.2.1](#)

19.4.2 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_masked                 = 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_callback_target_emi             = 33,
36    ompt_callback_target_data_op_emi     = 34,
37    ompt_callback_target_submit_emi      = 35,
38    ompt_callback_target_map_emi         = 36,
39    ompt_callback_error                  = 37
40 } ompt_callbacks_t;
```

C/C++

19.4.3 Tracing

OpenMP provides type definitions that support tracing with OMPT.

19.4.3.1 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++

19.4.3.2 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++

19.4.3.3 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++

Semantics

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`.

19.4.3.4 Standard Trace Record Type

Summary

The `ompt_record_ompt_t` type provides a 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;
        ompt_record_dependences_t dependences;
    };
};
```

```

1  ompt_record_task_dependence_t task_dependence;
2  ompt_record_task_schedule_t task_schedule;
3  ompt_record_implicit_task_t implicit_task;
4  ompt_record_masked_t masked;
5  ompt_record_sync_region_t sync_region;
6  ompt_record_mutex_acquire_t mutex_acquire;
7  ompt_record_mutex_t mutex;
8  ompt_record_nest_lock_t nest_lock;
9  ompt_record_flush_t flush;
10 ompt_record_cancel_t cancel;
11 ompt_record_target_t target;
12 ompt_record_target_data_op_t target_data_op;
13 ompt_record_target_map_t target_map;
14 ompt_record_target_kernel_t target_kernel;
15 ompt_record_control_tool_t control_tool;
16 ompt_record_error_t error;
17 } record;
18 } ompt_record_ompt_t;

```

C / C++

Semantics

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

Restrictions to the `ompt_record_ompt_t` type are as follows:

- If *type* is set to `ompt_callback_thread_end_t` then the value of *record* is undefined.

19.4.4 Miscellaneous Type Definitions

This section describes miscellaneous types and enumerations used by the tool interface.

19.4.4.1 `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);
```

19.4.4.2 ompt_set_result_t

Summary

The `ompt_set_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;
```

Semantics

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

- `ompt_set_callback_t`, see [Section 19.6.1.3](#)
- `ompt_set_trace_native_t`, see [Section 19.6.2.5](#)
- `ompt_set_trace_ompt_t`, see [Section 19.6.2.4](#)

19.4.4.3 `ompt_id_t`

Summary

The `ompt_id_t` type is used to provide various identifiers to tools.

Format

```
typedef uint64_t ompt_id_t;
```

Semantics

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

Restrictions to the `ompt_id_t` type are as follows:

- 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 data 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.

19.4.4.4 `ompt_data_t`

Summary

The `ompt_data_t` type represents data associated with threads and with parallel and task regions.

Format

```
typedef union ompt_data_t {  
    uint64_t value;  
    void *ptr;  
} ompt_data_t;
```

Semantics

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, `teams`, task, or target 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.

1 19.4.4.5 ompt_device_t

2 Summary

3 The `ompt_device_t` opaque object type represents a device.

4 Format

5 `typedef void ompt_device_t;`
6 C / C++

6 19.4.4.6 ompt_device_time_t

7 Summary

8 The `ompt_device_time_t` type represents raw device time values.

9 Format

10 `typedef uint64_t ompt_device_time_t;`
11 C / C++

11 Semantics

12 The `ompt_device_time_t` opaque object type represents raw device time values.

13 `ompt_time_none` refers to an unknown or unspecified time and is defined as an instance of type
14 `ompt_device_time_t` with the value 0.

15 19.4.4.7 ompt_buffer_t

16 Summary

17 The `ompt_buffer_t` opaque object type is a handle for a target buffer.

18 Format

19 `typedef void ompt_buffer_t;`
20 C / C++

20 19.4.4.8 ompt_buffer_cursor_t

21 Summary

22 The `ompt_buffer_cursor_t` opaque type is a handle for a position in a target buffer.

23 Format

24 `typedef uint64_t ompt_buffer_cursor_t;`
25 C / C++

19.4.4.9 `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++

Semantics

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`, see [Section 19.4.4.24](#)

19.4.4.10 `ompt_thread_t`

Summary

The `ompt_thread_t` enumeration type defines the valid thread type values.

Format

C / C++

```
typedef enum ompt_thread_t {  
    ompt_thread_initial           = 1,  
    ompt_thread_worker           = 2,  
    ompt_thread_other             = 3,  
    ompt_thread_unknown          = 4  
} ompt_thread_t;
```

C / C++

Semantics

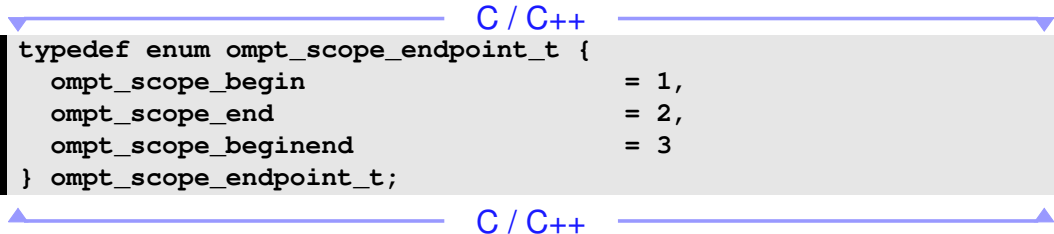
Any *initial thread* has thread type `ompt_thread_initial`. All *OpenMP threads* that are not initial threads have thread type `ompt_thread_worker`. A thread that an OpenMP implementation uses but that does not execute user code has thread type `ompt_thread_other`. Any thread that is created outside an OpenMP implementation and that is not an *initial thread* has thread type `ompt_thread_unknown`.

1 19.4.4.11 `ompt_scope_endpoint_t`

2 Summary

3 The `ompt_scope_endpoint_t` enumeration type defines valid scope endpoint values.

4 Format

5  C / C++

```
6 typedef enum ompt_scope_endpoint_t {  
7     ompt_scope_begin           = 1,  
8     ompt_scope_end            = 2,  
9     ompt_scope_beginend       = 3  
10 } ompt_scope_endpoint_t;
```

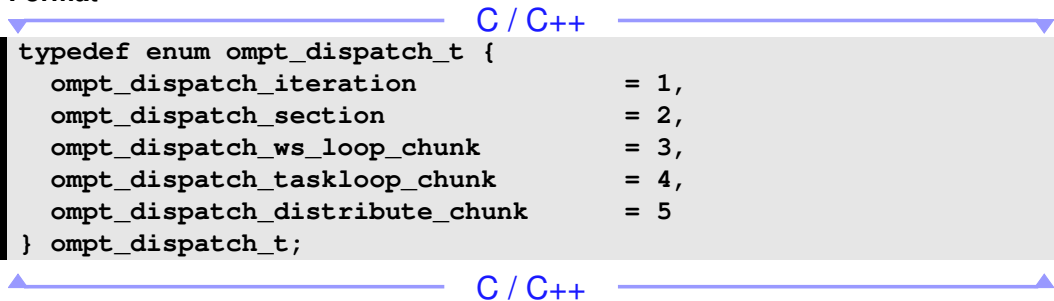
11 C / C++

10 19.4.4.12 `ompt_dispatch_t`

11 Summary

12 The `ompt_dispatch_t` enumeration type defines the valid dispatch kind values.

13 Format

14  C / C++

```
15 typedef enum ompt_dispatch_t {  
16     ompt_dispatch_iteration     = 1,  
17     ompt_dispatch_section       = 2,  
18     ompt_dispatch_ws_loop_chunk = 3,  
19     ompt_dispatch_taskloop_chunk = 4,  
20     ompt_dispatch_distribute_chunk = 5  
21 } ompt_dispatch_t;
```

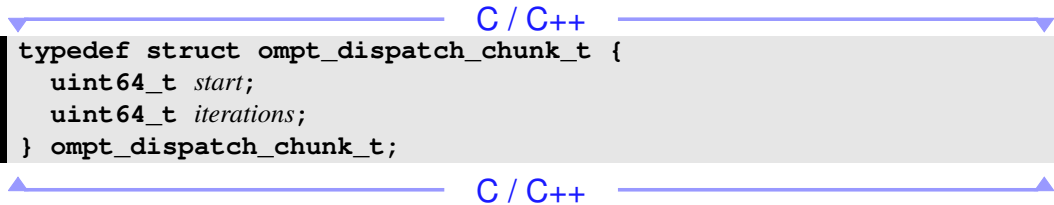
22 C / C++

21 19.4.4.13 `ompt_dispatch_chunk_t`

22 Summary

23 The `ompt_dispatch_chunk_t` type represents a the chunk information for a dispatched chunk.

24 Format

25  C / C++

```
26 typedef struct ompt_dispatch_chunk_t {  
27     uint64_t start;  
28     uint64_t iterations;  
29 } ompt_dispatch_chunk_t;
```

30 C / C++

Semantics

The `ompt_dispatch_chunk_t` type is a structure that holds information about a chunk of logical iterations of a loop nest. The `start` field specifies the first logical iteration of the chunk and the `iterations` field specifies the number of iterations in the chunk. Whether the chunk of a taskloop is contiguous is implementation defined.

19.4.4.14 `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_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_barrier_implicit_workshare = 8,
    ompt_sync_region_barrier_implicit_parallel = 9,
    ompt_sync_region_barrier_teams             = 10
} ompt_sync_region_t;
```

C / C++

19.4.4.15 `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_alloc_async              = 17,
    ompt_target_data_transfer_to_device_async = 18,
    ompt_target_data_transfer_from_device_async = 19,
```



```
1      ompt_target_data_delete_async          = 20
2  } ompt_target_data_op_t;
```

C / C++

19.4.4.16 ompt_work_t

Summary

The `ompt_work_t` enumeration type defines the valid work type values.

Format

C / C++

```
7  typedef enum ompt_work_t {
8      ompt_work_loop                = 1,
9      ompt_work_sections            = 2,
10     ompt_work_single_executor     = 3,
11     ompt_work_single_other        = 4,
12     ompt_work_workshare           = 5,
13     ompt_work_distribute          = 6,
14     ompt_work_taskloop            = 7,
15     ompt_work_scope               = 8,
16     ompt_work_loop_static         = 10,
17     ompt_work_loop_dynamic        = 11,
18     ompt_work_loop_guided         = 12,
19     ompt_work_loop_other          = 13
20 } ompt_work_t;
```

C / C++

19.4.4.17 ompt_mutex_t

Summary

The `ompt_mutex_t` enumeration type defines the valid mutex kind values.

Format

C / C++

```
25 typedef enum ompt_mutex_t {
26     ompt_mutex_lock                = 1,
27     ompt_mutex_test_lock           = 2,
28     ompt_mutex_nest_lock           = 3,
29     ompt_mutex_test_nest_lock      = 4,
30     ompt_mutex_critical            = 5,
31     ompt_mutex_atomic              = 6,
32     ompt_mutex_ordered             = 7
33 } ompt_mutex_t;
```

C / C++

19.4.4.18 ompt_native_mon_flag_t

Summary

The `ompt_native_mon_flag_t` enumeration type defines the valid native monitoring flag 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++

19.4.4.19 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_taskwait                   = 0x00000010,  
    ompt_task_underrferred                = 0x08000000,  
    ompt_task_untied                     = 0x10000000,  
    ompt_task_final                      = 0x20000000,  
    ompt_task_mergeable                  = 0x40000000,  
    ompt_task_merged                     = 0x80000000  
} ompt_task_flag_t;
```

C / C++

Semantics

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.

19.4.4.20 `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_taskwait_complete  = 8
} ompt_task_status_t;
```

C / C++

Semantics

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 for which the `detach` clause was specified 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 was 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 was fulfilled after the task completed execution of the associated structured block. The value `ompt_taskwait_complete` indicates completion of the dependent task that results from a `taskwait` construct with one or more `depend` clauses. The value `ompt_task_switch` is used for all other cases that a task was switched.

19.4.4.21 `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_nowait         = 9,
    ompt_target_enter_data_nowait = 10,
    ompt_target_exit_data_nowait = 11,
    ompt_target_update_nowait   = 12
} ompt_target_t;
```

C / C++

19.4.4.22 `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++

Semantics

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 primary 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 primary 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.

19.4.4.23 `ompt_target_map_flag_t`

Summary

The `ompt_target_map_flag_t` enumeration type defines the valid target map flag values.

Format

C / C++

```
typedef enum ompt_target_map_flag_t {  
    ompt_target_map_flag_to           = 0x01,  
    ompt_target_map_flag_from        = 0x02,  
    ompt_target_map_flag_alloc       = 0x04,  
    ompt_target_map_flag_release     = 0x08,  
    ompt_target_map_flag_delete      = 0x10,  
    ompt_target_map_flag_implicit    = 0x20,  
    ompt_target_map_flag_always      = 0x40,  
    ompt_target_map_flag_present     = 0x80,  
    ompt_target_map_flag_close       = 0x100,  
    ompt_target_map_flag_shared      = 0x200  
} ompt_target_map_flag_t;
```

C / C++

Semantics

The `ompt_target_map_flag_map-type` flag is set if the mapping operations have that *map-type*. If the *map-type* for the mapping operations is `tofrom`, both the `ompt_target_map_flag_to` and `ompt_target_map_flag_from` flags are set. The `ompt_target_map_flag_implicit` flag is set if the mapping operations result from implicit data-mapping rules. The `ompt_target_map_flag_map-type-modifier` flag is set if the mapping operations are specified with that *map-type-modifier*. The `ompt_target_map_flag_shared` flag is set if the original and corresponding storage are shared in the mapping operation.

19.4.4.24 `ompt_dependence_type_t`

Summary

The `ompt_dependence_type_t` enumeration type defines the valid task dependence type values.

Format

C / C++

```
typedef enum ompt_dependence_type_t {
    ompt_dependence_type_in           = 1,
    ompt_dependence_type_out          = 2,
    ompt_dependence_type_inout        = 3,
    ompt_dependence_type_mutexinoutset = 4,
    ompt_dependence_type_source       = 5,
    ompt_dependence_type_sink         = 6,
    ompt_dependence_type_inoutset     = 7
} ompt_dependence_type_t;
```

C / C++

19.4.4.25 ompt_severity_t

Summary

The `ompt_severity_t` enumeration type defines the valid severity values.

Format

C / C++

```
typedef enum ompt_severity_t {
    ompt_warning           = 1,
    ompt_fatal             = 2
} ompt_severity_t;
```

C / C++

19.4.4.26 ompt_cancel_flag_t

Summary

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++

19.4.4.27 `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++

Semantics

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

- Native Record Abstract Type, see [Section 19.4.3.3](#)

19.4.4.28 `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_implicit_parallel = 0x011,
    ompt_state_wait_barrier_implicit_workshare = 0x012,
    ompt_state_wait_barrier_explicit   = 0x014,
    ompt_state_wait_barrier_implementation = 0x015,
    ompt_state_wait_barrier_teams     = 0x016,

    ompt_state_wait_taskwait        = 0x020,
    ompt_state_wait_taskgroup       = 0x021,
```

```

1      ompt_state_wait_mutex           = 0x040,
2      ompt_state_wait_lock           = 0x041,
3      ompt_state_wait_critical       = 0x042,
4      ompt_state_wait_atomic        = 0x043,
5      ompt_state_wait_ordered       = 0x044,
6
7      ompt_state_wait_target         = 0x080,
8      ompt_state_wait_target_map     = 0x081,
9      ompt_state_wait_target_update  = 0x082,
10
11     ompt_state_idle                 = 0x100,
12     ompt_state_overhead             = 0x101,
13     ompt_state_undefined            = 0x102
14 } ompt_state_t;

```

C / C++

Semantics

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. The value **ompt_state_work_reduction** indicates that the thread is combining partial reduction results from threads in its team. An OpenMP implementation may never report a thread in this state; a thread that is combining partial reduction results may have its state reported as **ompt_state_work_parallel** or **ompt_state_overhead**. The value **ompt_state_wait_barrier_implicit_parallel** indicates that the thread is waiting at the implicit barrier at the end of a **parallel** region. The value **ompt_state_wait_barrier_implicit_workshare** indicates that the thread is waiting at an implicit barrier at the end of a worksharing construct. The value **ompt_state_wait_barrier_explicit** indicates that the thread is waiting in an explicit **barrier** region. The value **ompt_state_wait_barrier_implementation** indicates that the thread is waiting in a barrier not required by the OpenMP standard but introduced by an OpenMP implementation. The value **ompt_state_wait_barrier_teams** indicates that the thread is waiting at a barrier at the end of a **teams** region. The value **ompt_state_wait_taskwait** indicates that the thread is waiting at a **taskwait** construct. The value **ompt_state_wait_taskgroup** indicates that the thread is waiting at the end of a **taskgroup** construct. The value **ompt_state_wait_mutex** indicates that the thread is waiting for a mutex of an unspecified type. The value **ompt_state_wait_lock** indicates that the thread is waiting for a lock or nestable lock. The value **ompt_state_wait_critical** indicates that the thread is waiting to enter a **critical** region. The value **ompt_state_wait_atomic** indicates that the thread is waiting to enter an **atomic** region.

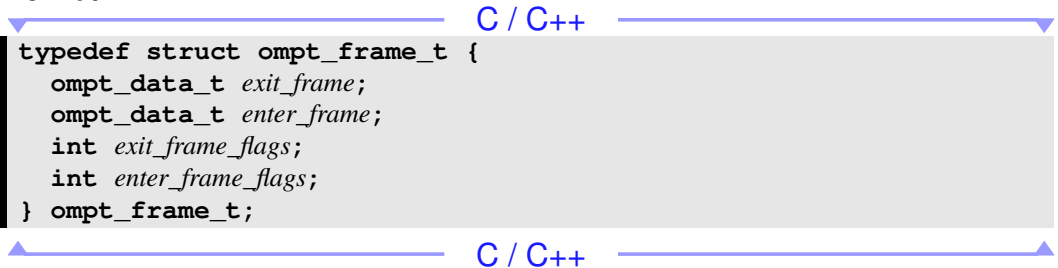
1 The value `ompt_state_wait_ordered` indicates that the thread is waiting to enter an
2 `ordered` region. The value `ompt_state_wait_target` indicates that the thread is waiting
3 for a `target` region to complete. The value `ompt_state_wait_target_map` indicates that
4 the thread is waiting for a target data mapping operation to complete. An implementation may
5 report `ompt_state_wait_target` for `target data` constructs. The value
6 `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. The value
9 `ompt_state_idle` indicates that the thread is idle, that is, it is not part of an OpenMP team.
10 The value `ompt_state_overhead` indicates that the thread is in the overhead state at any point
11 while executing within the OpenMP runtime, except while waiting at a synchronization point. The
12 value `ompt_state_undefined` indicates that the native thread is not created by the OpenMP
13 implementation.

14 19.4.4.29 `ompt_frame_t`

15 Summary

16 The `ompt_frame_t` type describes procedure frame information for an OpenMP task.

17 Format

18  C / C++
19 `typedef struct ompt_frame_t {`
20 `ompt_data_t exit_frame;`
21 `ompt_data_t enter_frame;`
22 `int exit_frame_flags;`
23 `int enter_frame_flags;`
24 `} ompt_frame_t;`

24 Semantics

25 Each `ompt_frame_t` object is associated with the task to which the procedure frames belong.
26 Each non-merged initial, implicit, explicit, or target task with one or more frames on the stack of a
27 native thread has an associated `ompt_frame_t` object.

28 The `exit_frame` field of an `ompt_frame_t` object contains information to identify the first
29 procedure frame executing the task region. The `exit_frame` for the `ompt_frame_t` object
30 associated with the *initial task* that is not nested inside any OpenMP construct is
31 `ompt_data_none`.

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
34 implementation or before executing a different task. If a task with frames on the stack is not
35 executing implementation code in the OpenMP runtime, the value of `enter_frame` for the
36 `ompt_frame_t` object associated with the task will be `ompt_data_none`.

1 For *exit_frame*, the *exit_frame_flags* and, for *enter_frame*, the *enter_frame_flags* field indicates that
2 the provided frame information points to a runtime or an application frame address. The same
3 fields also specify the kind of information that is provided to identify the frame, These fields are a
4 disjunction of values in the `ompt_frame_flag_t` enumeration type.

5 The lifetime of an `ompt_frame_t` object begins when a task is created and ends when the task is
6 destroyed. Tools should not assume that a frame structure remains at a constant location in memory
7 throughout the lifetime of the task. A pointer to an `ompt_frame_t` object is passed to some
8 callbacks; a pointer to the `ompt_frame_t` object of a task can also be retrieved by a tool at any
9 time, including in a signal handler, by invoking the `ompt_get_task_info` runtime entry point
10 (described in [Section 19.6.1.14](#)). A pointer to an `ompt_frame_t` object that a tool retrieved is
11 valid as long as the tool does not pass back control to the OpenMP implementation.

12
13 **Note** – A monitoring tool that uses asynchronous sampling can observe values of *exit_frame* and
14 *enter_frame* at inconvenient times. Tools must be prepared to handle `ompt_frame_t` objects
15 observed just prior to when their field values will be set or cleared.
16

17 19.4.4.30 `ompt_frame_flag_t`

18 Summary

19 The `ompt_frame_flag_t` enumeration type defines valid frame information flags.

20 Format

21 `C / C++`
22

```
typedef enum ompt_frame_flag_t {  
23     ompt_frame_runtime           = 0x00,  
24     ompt_frame_application      = 0x01,  
25     ompt_frame_cfa              = 0x10,  
26     ompt_frame_framepointer     = 0x20,  
27     ompt_frame_stackaddress     = 0x30  
} ompt_frame_flag_t;
```

28 Semantics

29 The value `ompt_frame_runtime` of the `ompt_frame_flag_t` type indicates that a frame
30 address is a procedure frame in the OpenMP runtime implementation. The value
31 `ompt_frame_application` of the `ompt_frame_flag_t` type indicates that a frame
32 address is a procedure frame in the OpenMP application.

33 Higher order bits indicate the kind of provided information that is unique for the particular frame
34 pointer. The value `ompt_frame_cfa` indicates that a frame address specifies a *canonical frame*
35 *address*. The value `ompt_frame_framepointer` indicates that a frame address provides the

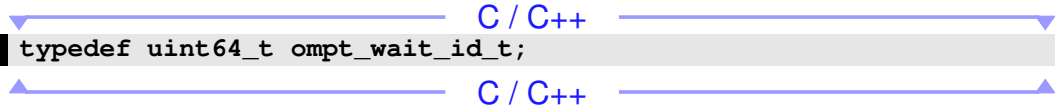
1 value of the frame pointer register. The value `ompt_frame_stackaddress` indicates that a
2 frame address specifies a pointer address that is contained in the current stack frame.

3 19.4.4.31 `ompt_wait_id_t`

4 Summary

5 The `ompt_wait_id_t` type describes wait identifiers for an OpenMP thread.

6 Format

7 
`typedef uint64_t ompt_wait_id_t;`
C / C++

8 Semantics

9 Each thread maintains a *wait identifier* of type `ompt_wait_id_t`. When a task that a thread
10 executes is waiting for mutual exclusion, the wait identifier of the thread indicates the reason that
11 the thread is waiting. A wait identifier may represent a critical section *name*, a lock, a program
12 variable accessed in an atomic region, or a synchronization object that is internal to an OpenMP
13 implementation. When a thread is not in a wait state then the value of the wait identifier of the
14 thread is undefined. `ompt_wait_id_none` is defined as an instance of type
15 `ompt_wait_id_t` with the value 0.

16 19.5 OMPT Tool Callback Signatures and Trace 17 Records

18 The C/C++ header file (`omp-tools.h`) provides the definitions of the types that are specified
19 throughout this subsection. Restrictions to the OpenMP tool callbacks are as follows:

20 Restrictions

- 21 • Tool callbacks may not use OpenMP directives or call any runtime library routines described
22 in [Chapter 18](#).
- 23 • Tool callbacks must exit by either returning to the caller or aborting.

24 19.5.1 Initialization and Finalization Callback Signature

25 19.5.1.1 `ompt_initialize_t`

26 Summary

27 A callback with type signature `ompt_initialize_t` initializes the 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++

Semantics

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 pointer to a tool initializer function 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 runtime library routine. The initializer returns a non-zero value if it succeeds; otherwise, the OMPT interface state changes to *inactive* as described in [Section 19.2.3](#).

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.

Cross References

- Tool Initialization and Finalization, see [Section 19.4.1](#)
- `omp_get_initial_device`, see [Section 18.7.8](#)
- `ompt_data_t`, see [Section 19.4.4.4](#)
- `ompt_start_tool`, see [Section 19.2.1](#)

19.5.1.2 ompt_finalize_t

Summary

A tool implements a finalizer with the type signature `ompt_finalize_t` to finalize its use of the OMPT interface.

Format

C / C++

```
typedef void (*ompt_finalize_t) (  
    ompt_data_t *tool_data  
);
```

C / C++

Semantics

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 must 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

- Tool Initialization and Finalization, see [Section 19.4.1](#)
- `ompt_data_t`, see [Section 19.4.4.4](#)
- `ompt_start_tool`, see [Section 19.2.1](#)

19.5.2 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 the 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_data_t`, see [Section 19.4.4.4](#)
- `ompt_id_t`, see [Section 19.4.4.3](#)

19.5.2.1 `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

```
C / C++  
typedef void (*ompt_callback_thread_begin_t) (  
    ompt_thread_t thread_type,  
    ompt_data_t *thread_data  
);  
C / C++
```

Trace Record

C / C++

```
typedef struct omp_t_record_thread_begin_t {  
    omp_thread_t thread_type;  
} omp_t_record_thread_begin_t;
```

C / C++

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` directive, see [Section 10.2](#)
- `teams` directive, see [Section 10.3](#)
- Initial Task, see [Section 12.8](#)
- `omp_data_t`, see [Section 19.4.4.4](#)
- `omp_thread_t`, see [Section 19.4.4.10](#)

19.5.2.2 `omp_callback_thread_end_t`

Summary

The `omp_callback_thread_end_t` type is used for callbacks that are dispatched when native threads are destroyed.

Format

C / C++

```
typedef void (*omp_callback_thread_end_t) (  
    omp_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` directive, see [Section 10.2](#)
- `teams` directive, see [Section 10.3](#)
- Initial Task, see [Section 12.8](#)
- Standard Trace Record Type, see [Section 19.4.3.4](#)
- `omp_data_t`, see [Section 19.4.4.4](#)

19.5.2.3 `ompt_callback_parallel_begin_t`

Summary

The `ompt_callback_parallel_begin_t` type is used for callbacks that are dispatched when a `parallel` or `teams` region starts.

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 behavior for accessing the frame object after the callback returned is unspecified.

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

1 **ompt_callback_parallel_begin_t** then *codeptr_ra* contains the return address of the call
2 to that runtime routine. If the implementation of the region is inlined then *codeptr_ra* contains the
3 return address of the callback invocation. If attribution to source code is impossible or
4 inappropriate, *codeptr_ra* may be *NULL*.

5 **Cross References**

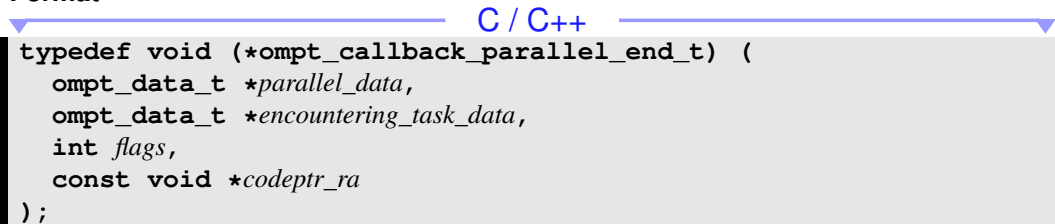
- 6 • **parallel** directive, see [Section 10.2](#)
- 7 • **teams** directive, see [Section 10.3](#)
- 8 • **ompt_data_t**, see [Section 19.4.4.4](#)
- 9 • **ompt_frame_t**, see [Section 19.4.4.29](#)
- 10 • **ompt_parallel_flag_t**, see [Section 19.4.4.22](#)

11 **19.5.2.4 ompt_callback_parallel_end_t**

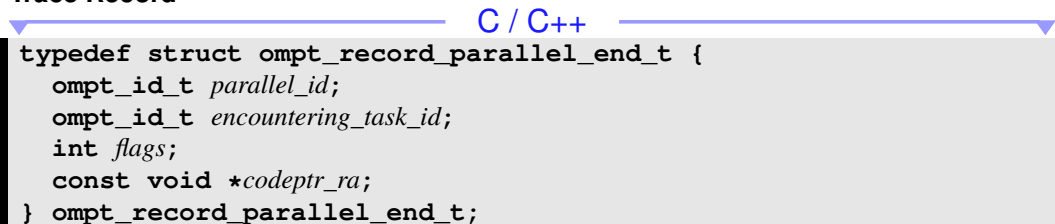
12 **Summary**

13 The **ompt_callback_parallel_end_t** type is used for callbacks that are dispatched when a
14 **parallel** or **teams** region ends.

15 **Format**

16  C / C++
17 `typedef void (*ompt_callback_parallel_end_t) (
18 ompt_data_t *parallel_data,
19 ompt_data_t *encountering_task_data,
20 int flags,
21 const void *codeptr_ra
);`

22 **Trace Record**

23  C / C++
24 `typedef struct ompt_record_parallel_end_t {
25 ompt_id_t parallel_id;
26 ompt_id_t encountering_task_id;
27 int flags;
28 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 callback invocation. If attribution to source code is impossible or inappropriate, *codeptr_ra* may be *NULL*.

Cross References

- **parallel** directive, see [Section 10.2](#)
- **teams** directive, see [Section 10.3](#)
- **ompt_data_t**, see [Section 19.4.4.4](#)
- **ompt_parallel_flag_t**, see [Section 19.4.4.22](#)

19.5.2.5 ompt_callback_work_t

Summary

The **ompt_callback_work_t** type is used for callbacks that are dispatched when worksharing regions and **taskloop** regions begin and end.

Format

```

C / C++
typedef void (*ompt_callback_work_t) (
    ompt_work_t work_type,
    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

C / C++

```
typedef struct ompt_record_work_t {
    ompt_work_t work_type;
    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 *work_type* 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 or **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** or **scope** 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 callback invocation. If attribution to source code is impossible or inappropriate, *codeptr_ra* may be *NULL*.

Cross References

- **taskloop** directive, see [Section 12.6](#)
- Work-Distribution Constructs, see [Chapter 11](#)
- **ompt_data_t**, see [Section 19.4.4.4](#)
- **ompt_scope_endpoint_t**, see [Section 19.4.4.11](#)
- **ompt_work_t**, see [Section 19.4.4.16](#)

19.5.2.6 `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.

If the `kind` argument is `ompt_dispatch_iteration`, the `value` field of the `instance` argument contains the logical iteration number. If the `kind` argument is `ompt_dispatch_section`, the `ptr` field of the `instance` argument contains a code address that identifies the structured block. In cases where a runtime routine implements the structured block associated with this callback, the `ptr` field of the `instance` argument contains the return address of the call to the runtime routine. In cases where the implementation of the structured block is inlined, the `ptr` field of the `instance` argument contains the return address of the invocation of this callback. If the `kind` argument is `ompt_dispatch_ws_loop_chunk`, `ompt_dispatch_taskloop_chunk` or `ompt_dispatch_distribute_chunk`, the `ptr` field of the `instance` argument points to a structure of type `ompt_dispatch_chunk_t` that contains the information for the chunk.

Cross References

- `sections` directive, see [Section 11.3](#)
- `taskloop` directive, see [Section 12.6](#)
- Worksharing-Loop Constructs, see [Section 11.5](#)
- `ompt_data_t`, see [Section 19.4.4.4](#)
- `ompt_dispatch_chunk_t`, see [Section 19.4.4.13](#)
- `ompt_dispatch_t`, see [Section 19.4.4.12](#)

19.5.2.7 `ompt_callback_task_create_t`

Summary

The `ompt_callback_task_create_t` type is used for callbacks that are dispatched when `task` regions 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  
);
```

C / C++

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;
```

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 associated with the encountering task. The behavior for accessing the frame object after the callback returned is unspecified.

The binding of the *new_task_data* argument is the generated task.

The *flags* argument indicates the kind of task (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 that runtime routine. If the implementation of the region is inlined then *codeptr_ra* contains the return address of the callback invocation. If attribution to source code is impossible or inappropriate, *codeptr_ra* may be *NULL*.

Cross References

- **task** directive, see [Section 12.5](#)
- Initial Task, see [Section 12.8](#)
- **ompt_data_t**, see [Section 19.4.4.4](#)
- **ompt_frame_t**, see [Section 19.4.4.29](#)
- **ompt_task_flag_t**, see [Section 19.4.4.19](#)

19.5.2.8 ompt_callback_dependences_t

Summary

The **ompt_callback_dependences_t** type is used for callbacks that are related to dependences and that are dispatched when new tasks are generated and when **ordered** constructs are encountered.

Format

```
typedef void (*ompt_callback_dependences_t) (  
    ompt_data_t *task_data,  
    const ompt_dependence_t *deps,  
    int ndeps  
);
```

Trace Record

C / C++

```
typedef struct ompt_record_dependencies_t {
    ompt_id_t task_id;
    ompt_dependence_t dep;
    int ndeps;
} ompt_record_dependencies_t;
```

C / C++

Description of Arguments

The binding of the *task_data* argument is the generated task for a depend clause on a task construct, the target task for a depend clause on a target construct respectively depend object in an asynchronous runtime routine, or the encountering implicit task for a depend clause of the ordered construct.

The *deps* argument lists dependences of the new task or the dependence vector of the ordered construct. Dependences denoted with depend objects are described in terms of their dependence semantics.

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

- **depend** clause, see [Section 15.9.5](#)
- **ordered** directive, see [Section 15.10.1](#)
- **ompt_data_t**, see [Section 19.4.4.4](#)
- **ompt_dependence_t**, see [Section 19.4.4.9](#)

19.5.2.9 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
);
```

C / C++

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;
```

C / C++

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 15.9.5](#)
- `ompt_data_t`, see [Section 19.4.4.4](#)

19.5.2.10 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

C / C++

```
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
);
```

C / C++

Trace Record

C / C++

```
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;
```

C / C++

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 or if the callback signals completion of a taskwait construct.

Cross References

- Task Scheduling, see [Section 12.9](#)
- `ompt_data_t`, see [Section 19.4.4.4](#)
- `ompt_task_status_t`, see [Section 19.4.4.20](#)

19.5.2.11 `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 or **teams** region. For the *implicit-task-end* and the *initial-task-end* events, this argument is *NULL*.

The binding of the *task_data* argument is the implicit task that executes the structured block of the parallel or **teams** 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 task (initial or implicit).

Cross References

- **parallel** directive, see [Section 10.2](#)
- **teams** directive, see [Section 10.3](#)
- **ompt_data_t**, see [Section 19.4.4.4](#)
- **ompt_scope_endpoint_t**, see [Section 19.4.4.11](#)

19.5.2.12 ompt_callback_masked_t

Summary

The **ompt_callback_masked_t** type is used for callbacks that are dispatched when **masked** regions start and end.

Format

```
C / C++
typedef void (*ompt_callback_masked_t) (
    ompt_scope_endpoint_t endpoint,
    ompt_data_t *parallel_data,
    ompt_data_t *task_data,
    const void *codeptr_ra
);
```

C / C++

Trace Record

C / C++

```
typedef struct ompt_record_masked_t {  
    ompt_scope_endpoint_t endpoint;  
    ompt_id_t parallel_id;  
    ompt_id_t task_id;  
    const void *codeptr_ra;  
} ompt_record_masked_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_masked_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 callback invocation. If attribution to source code is impossible or inappropriate, *codeptr_ra* may be *NULL*.

Cross References

- **masked** directive, see [Section 10.6](#)
- **ompt_data_t**, see [Section 19.4.4.4](#)
- **ompt_scope_endpoint_t**, see [Section 19.4.4.11](#)

19.5.2.13 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 *implicit-barrier-end* event at the end of a parallel region this argument is *NULL*. For the *implicit-barrier-wait-begin* and *implicit-barrier-wait-end* event at the end of a parallel region, whether this argument is *NULL* or points to the parallel data of the current parallel region is implementation defined.

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 callback invocation. If attribution to source code is impossible or inappropriate, *codeptr_ra* may be *NULL*.

Cross References

- `barrier` directive, see [Section 15.3.1](#)
- `taskgroup` directive, see [Section 15.4](#)
- `taskwait` directive, see [Section 15.5](#)
- Implicit Barriers, see [Section 15.3.2](#)
- Properties Common to All Reduction Clauses, see [Section 5.5.5](#)
- `ompt_data_t`, see [Section 19.4.4.4](#)
- `ompt_scope_endpoint_t`, see [Section 19.4.4.11](#)
- `ompt_sync_region_t`, see [Section 19.4.4.14](#)

19.5.2.14 `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

```
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  
);
```

Trace Record

```
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;
```

Description of Arguments

The *kind* argument indicates the kind of mutual exclusion event.

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 callback invocation. If attribution to source code is impossible or inappropriate, *codeptr_ra* may be `NULL`.

Cross References

- `atomic` directive, see [Section 15.8.5](#)
- `critical` directive, see [Section 15.2](#)
- `ompt_wait_id_t`, see [Section 19.4.4.31](#)
- `omp_init_lock` and `omp_init_nest_lock`, see [Section 18.9.1](#)
- `ompt_mutex_t`, see [Section 19.4.4.17](#)
- `ordered` Construct, see [Section 15.10](#)

19.5.2.15 `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  
);
```

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 callback invocation. If attribution to source code is impossible or inappropriate, *codeptr_ra* may be *NULL*.

Cross References

- **atomic** directive, see [Section 15.8.5](#)
- **critical** directive, see [Section 15.2](#)
- **omp_destroy_lock** and **omp_destroy_nest_lock**, see [Section 18.9.3](#)
- **ompt_wait_id_t**, see [Section 19.4.4.31](#)
- **omp_set_lock** and **omp_set_nest_lock**, see [Section 18.9.4](#)
- **omp_test_lock** and **omp_test_nest_lock**, see [Section 18.9.6](#)
- **omp_unset_lock** and **omp_unset_nest_lock**, see [Section 18.9.5](#)
- **ompt_mutex_t**, see [Section 19.4.4.17](#)
- **ordered** Construct, see [Section 15.10](#)

19.5.2.16 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.

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 callback invocation. If attribution to source code is impossible or inappropriate, *codeptr_ra* may be *NULL*.

Cross References

- **ompt_wait_id_t**, see [Section 19.4.4.31](#)
- **omp_set_lock** and **omp_set_nest_lock**, see [Section 18.9.4](#)
- **omp_test_lock** and **omp_test_nest_lock**, see [Section 18.9.6](#)
- **omp_unset_lock** and **omp_unset_nest_lock**, see [Section 18.9.5](#)
- **ompt_scope_endpoint_t**, see [Section 19.4.4.11](#)

19.5.2.17 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 callback invocation. If attribution to source code is impossible or inappropriate, *codeptr_ra* may be *NULL*.

Cross References

- **flush** directive, see [Section 15.8.6](#)
- **ompt_data_t**, see [Section 19.4.4.4](#)

19.5.2.18 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 callback invocation. If attribution to source code is impossible or inappropriate, *codeptr_ra* may be *NULL*.

Cross References

- **ompt_cancel_flag_t**, see [Section 19.4.4.26](#)

19.5.2.19 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++

Semantics

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 C 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 C 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 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 19.2.5](#).

Cross References

- Lookup Entry Points: `ompt_function_lookup_t`, see [Section 19.6.3](#)

19.5.2.20 `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  
);
```

Description of Arguments

The `device_num` argument identifies the logical device that is being finalized.

Semantics

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`, see [Section 19.5.2.24](#)

19.5.2.21 `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

```
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  
);
```

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 a 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 and Clauses, see [Chapter 13](#)

19.5.2.22 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 and Clauses, see [Chapter 13](#)

19.5.2.23 `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  
);
```

Semantics

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 a value less than the minimum requested buffer size in `*bytes` on entry to the callback, further recording of events for the device may be disabled until the next invocation of `ompt_start_trace`. This action causes the device to drop future trace records until recording is restarted. A first party tool may use the `ompt_get_buffer_limits` runtime entry point to determine the recommended number of bytes to provide when fulfilling the buffer request.

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 holds the minimum size of the buffer in bytes that is requested, which must not exceed the recommended buffer size returned by the `ompt_get_buffer_limits` runtime entry point for the same device. On return, it indicates size of the buffer to which `*buffer` points.

Cross References

- `ompt_buffer_t`, see [Section 19.4.4.7](#)
- `ompt_get_buffer_limits_t`, see [Section 19.6.2.6](#)

19.5.2.24 `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++

Semantics

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 primary 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 for 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_cursor_t`, see [Section 19.4.4.8](#)
- `ompt_buffer_t`, see [Section 19.4.4.7](#)

19.5.2.25 `ompt_callback_target_data_op_emi_t` and `ompt_callback_target_data_op_t`

Summary

The `ompt_callback_target_data_op_emi_t` and `ompt_callback_target_data_op_t` types are used for callbacks that are dispatched when a thread maps data to a device.

Format

C / C++

```
1 typedef void (*ompt_callback_target_data_op_emi_t) (  
2     ompt_scope_endpoint_t endpoint,  
3     ompt_data_t *target_task_data,  
4     ompt_data_t *target_data,  
5     ompt_id_t *host_op_id,  
6     ompt_target_data_op_t optype,  
7     void *src_addr,  
8     int src_device_num,  
9     void *dest_addr,  
10    int dest_device_num,  
11    size_t bytes,  
12    const void *codeptr_ra  
13 );  
14
```

```
15 typedef void (*ompt_callback_target_data_op_t) (  
16     ompt_id_t target_id,  
17     ompt_id_t host_op_id,  
18     ompt_target_data_op_t optype,  
19     void *src_addr,  
20     int src_device_num,  
21     void *dest_addr,  
22     int dest_device_num,  
23     size_t bytes,  
24     const void *codeptr_ra  
25 );
```

C / C++

Trace Record

C / C++

```
27 typedef struct ompt_record_target_data_op_t {  
28     ompt_id_t host_op_id;  
29     ompt_target_data_op_t optype;  
30     void *src_addr;  
31     int src_device_num;  
32     void *dest_addr;  
33     int dest_device_num;  
34     size_t bytes;  
35     ompt_device_time_t end_time;  
36     const void *codeptr_ra;  
37 } ompt_record_target_data_op_t;
```

C / C++

Semantics

A thread dispatches a registered `ompt_callback_target_data_op_emi` or `ompt_callback_target_data_op` callback when device memory is allocated or freed, as well as when data is copied to or from a device.

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 *endpoint* argument indicates that the callback signals the beginning or end of a scope.

The binding of the *target_task_data* argument is the target task region.

The binding of the *target_data* argument is the target region.

The *host_op_id* argument points to a tool-controlled integer value, which identifies a data operation on a target device.

The *optype* argument indicates the kind of data operation.

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.

Whether in some operations *src_addr* or *dest_addr* may point to an intermediate buffer is implementation defined.

The *bytes* argument indicates the size of data.

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_data_op_emi_t` or `ompt_callback_target_data_op_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 callback invocation. If attribution to source code is impossible or inappropriate, *codeptr_ra* may be *NULL*.

Restrictions

Restrictions to the `ompt_callback_target_data_op_emi` and `ompt_callback_target_data_op` callbacks are as follows:

- These callbacks must not be registered at the same time.

Cross References

- `map` clause, see [Section 5.8.3](#)
- `ompt_data_t`, see [Section 19.4.4.4](#)
- `ompt_id_t`, see [Section 19.4.4.3](#)
- `ompt_scope_endpoint_t`, see [Section 19.4.4.11](#)
- `ompt_target_data_op_t`, see [Section 19.4.4.15](#)

19.5.2.26 `ompt_callback_target_emi_t` and `ompt_callback_target_t`

Summary

The `ompt_callback_target_emi_t` and `ompt_callback_target_t` types are used for callbacks that are dispatched when a thread begins to execute a device construct.

Format

C / C++

```
typedef void (*ompt_callback_target_emi_t) (  
    ompt_target_t kind,  
    ompt_scope_endpoint_t endpoint,  
    int device_num,  
    ompt_data_t *task_data,  
    ompt_data_t *target_task_data,  
    ompt_data_t *target_data,  
    const void *codeptr_ra  
);
```

```
typedef void (*ompt_callback_target_t) (  
    ompt_target_t kind,  
    ompt_scope_endpoint_t endpoint,  
    int device_num,  
    ompt_data_t *task_data,  
    ompt_id_t target_id,  
    const void *codeptr_ra  
);
```

C / C++

Trace Record

C / C++

```
typedef struct ompt_record_target_t {
    ompt_target_t kind;
    ompt_scope_endpoint_t endpoint;
    int device_num;
    ompt_id_t task_id;
    ompt_id_t target_id;
    const void *codeptr_ra;
} 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 device number of the device that will execute the target region.

The binding of the *task_data* argument is the encountering task.

The binding of the *target_task_data* argument is the target task region. If a target region has no target task or if the target task is merged, this argument is **NULL**.

The binding of the *target_data* 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_emi_t** or **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 callback invocation. If attribution to source code is impossible or inappropriate, *codeptr_ra* may be *NULL*.

Restrictions

Restrictions to the **ompt_callback_target_emi** and **ompt_callback_target** callbacks are as follows:

- These callbacks must not be registered at the same time.

Cross References

- `target` directive, see [Section 13.8](#)
- `target data` directive, see [Section 13.5](#)
- `target enter data` directive, see [Section 13.6](#)
- `target exit data` directive, see [Section 13.7](#)
- `target update` directive, see [Section 13.9](#)
- `ompt_data_t`, see [Section 19.4.4.4](#)
- `ompt_id_t`, see [Section 19.4.4.3](#)
- `ompt_scope_endpoint_t`, see [Section 19.4.4.11](#)
- `ompt_target_t`, see [Section 19.4.4.21](#)

19.5.2.27 `ompt_callback_target_map_emi_t` and `ompt_callback_target_map_t`

Summary

The `ompt_callback_target_map_emi_t` and `ompt_callback_target_map_t` types are used for callbacks that are dispatched to indicate data mapping relationships.

Format

C / C++

```
typedef void (*ompt_callback_target_map_emi_t) (  
    ompt_data_t *target_data,  
    unsigned int nitems,  
    void **host_addr,  
    void **device_addr,  
    size_t *bytes,  
    unsigned int *mapping_flags,  
    const void *codeptr_ra  
);
```

```
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++

Semantics

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_emi** or **ompt_callback_target_map** callback to report the effect of all mappings or multiple **ompt_callback_target_map_emi** or **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_emi** or **ompt_callback_target_map** callbacks, these callbacks may be interleaved with **ompt_callback_target_data_op_emi** or **ompt_callback_target_data_op** callbacks used to report data operations associated with the mappings.

Description of Arguments

The binding of the *target_data* 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 sizes of data.

The *mapping_flags* argument indicates the kind of mapping operations, which may result from explicit **map** clauses or the implicit data-mapping rules defined in [Section 5.8](#). Flags for the mapping operations 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** or **ompt_callback_target_map_emi_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 callback invocation. If attribution to source code is impossible or inappropriate, *codeptr_ra* may be *NULL*.

Restrictions

Restrictions to the `ompt_callback_target_data_map_emi` and `ompt_callback_target_data_map` callbacks are as follows:

- These callbacks must not be registered at the same time.

Cross References

- `target` directive, see [Section 13.8](#)
- `target data` directive, see [Section 13.5](#)
- `target enter data` directive, see [Section 13.6](#)
- `target exit data` directive, see [Section 13.7](#)
- `ompt_callback_target_data_op_emi_t` and `ompt_callback_target_data_op_t`, see [Section 19.5.2.25](#)
- `ompt_data_t`, see [Section 19.4.4.4](#)
- `ompt_id_t`, see [Section 19.4.4.3](#)
- `ompt_target_map_flag_t`, see [Section 19.4.4.23](#)

19.5.2.28 `ompt_callback_target_submit_emi_t` and `ompt_callback_target_submit_t`

Summary

The `ompt_callback_target_submit_emi_t` and `ompt_callback_target_submit_t` types are used for callbacks that are dispatched before and after the host initiates creation of an initial task on a device.

Format

```

C / C++
typedef void (*ompt_callback_target_submit_emi_t) (
    ompt_scope_endpoint_t endpoint,
    ompt_data_t *target_data,
    ompt_id_t *host_op_id,
    unsigned int requested_num_teams
);

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++

Semantics

A thread dispatches a registered `ompt_callback_target_submit_emi` or `ompt_callback_target_submit` callback on the host before and after a target task initiates creation of an initial task on a device.

Description of Arguments

The *endpoint* argument indicates that the callback signals the beginning or end of a scope.

The binding of the *target_data* argument is the target region.

The *host_op_id* argument points to a tool-controlled integer value, which identifies an initial task on a 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 tool-controlled identifier that can be used to correlate a `ompt_record_target_kernel_t` record with its associated `ompt_callback_target_submit_emi` or `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.

Restrictions

Restrictions to the `ompt_callback_target_submit_emi` and `ompt_callback_target_submit` callbacks are as follows:

- These callbacks must not be registered at the same time.

Cross References

- `target` directive, see [Section 13.8](#)
- `ompt_data_t`, see [Section 19.4.4.4](#)
- `ompt_id_t`, see [Section 19.4.4.3](#)
- `ompt_scope_endpoint_t`, see [Section 19.4.4.11](#)

19.5.2.29 `ompt_callback_control_tool_t`

Summary

The `ompt_callback_control_tool_t` type is used for callbacks that dispatch *tool-control* events.

Format

```
typedef int (*ompt_callback_control_tool_t) (  
    uint64_t command,  
    uint64_t modifier,  
    void *arg,  
    const void *codeptr_ra  
);
```

Trace Record

```
typedef struct ompt_record_control_tool_t {  
    uint64_t command;  
    uint64_t modifier;  
    const void *codeptr_ra;  
} ompt_record_control_tool_t;
```

Semantics

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 `omp_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 18.14](#).

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 callback invocation. If attribution to source code is impossible or inappropriate, *codeptr_ra* may be *NULL*.

Constraints on Arguments

Tool-specific values for *command* must be ≥ 64 .

Cross References

- Tool Control Routine, see [Section 18.14](#)

19.5.2.30 `ompt_callback_error_t`

Summary

The `ompt_callback_error_t` type is used for callbacks that dispatch *runtime-error* events.

Format

```

C / C++
typedef void (*ompt_callback_error_t) (
    ompt_severity_t severity,
    const char *message,
    size_t length,
    const void *codeptr_ra
);
C / C++
```


Trace Record

C / C++

```
typedef struct ompt_record_error_t {  
    ompt_severity_t severity;  
    const char *message;  
    size_t length;  
    const void *codeptr_ra;  
} ompt_record_error_t;
```

C / C++

Semantics

A thread dispatches a registered `ompt_callback_error_t` callback when an `error` directive is encountered for which the `at (execution)` clause is specified.

Description of Arguments

The *severity* argument passes the specified severity level.

The *message* argument passes the C string from the `message` clause.

The *length* argument provides the length of the C string.

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_error_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 callback invocation. If attribution to source code is impossible or inappropriate, *codeptr_ra* may be `NULL`.

Cross References

- `error` directive, see [Section 8.5](#)
- `ompt_severity_t`, see [Section 19.4.4.25](#)

19.6 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

1 callback interface. The second class of lookup routines includes a unique lookup routine for each
2 kind of device that can return runtime entry points in a device's OMPT tracing interface.

3 The `omp-tools.h` C/C++ header file provides the definitions of the types that are specified
4 throughout this subsection.

5 **Binding**

6 The binding thread set for each of the entry points in this section is the encountering thread unless
7 otherwise specified. The binding task set is the task executing on the encountering thread.

8 **Restrictions**

9 Restrictions on OMPT runtime entry points are as follows:

- 10 • OMPT runtime entry points must not be called from a signal handler on a native thread
11 before a *native-thread-begin* or after a *native-thread-end* event.
- 12 • OMPT device runtime entry points must not be called after a *device-finalize* event for that
13 device.

14 **19.6.1 Entry Points in the OMPT Callback Interface**

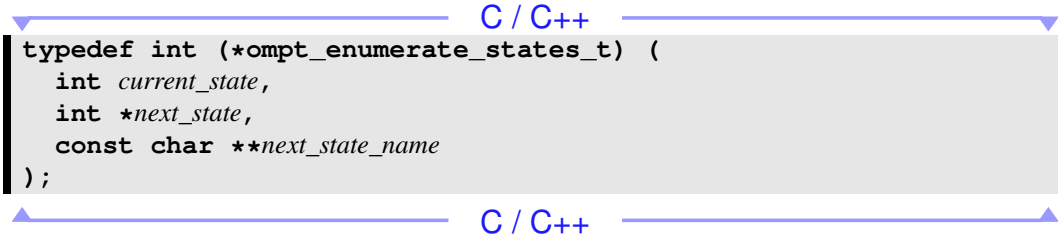
15 Entry points in the OMPT callback interface enable a tool to register callbacks for OpenMP events
16 and to inspect the state of an OpenMP thread while executing in a tool callback or a signal handler.
17 Pointers to these runtime entry points are obtained through the lookup function that is provided
18 through the OMPT initializer.

19 **19.6.1.1 `omp_enumerate_states_t`**

20 **Summary**

21 The `omp_enumerate_states_t` type is the type signature of the
22 `omp_enumerate_states` runtime entry point, which enumerates the thread states that an
23 OpenMP implementation supports.

24 **Format**

```
25  C / C++  
26 typedef int (*omp_enumerate_states_t) (  
27     int current_state,  
28     int *next_state,  
29     const char **next_state_name  
30 );
```

Semantics

An OpenMP implementation may support only a subset of the states that the `ompt_state_t` enumeration type defines. 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 that was 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 `0x102`.

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 program execution.

Cross References

- `ompt_state_t`, see [Section 19.4.4.28](#)

19.6.1.2 `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

C / C++

```
typedef int (*ompt_enumerate_mutex_impls_t) (  
    int current_impl,  
    int *next_impl,  
    const char **next_impl_name  
);
```

C / C++

Semantics

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 that was 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.

19.6.1.3 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

```
typedef ompt_set_result_t (*ompt_set_callback_t) (  
    ompt_callbacks_t event,  
    ompt_callback_t callback  
);
```

Semantics

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

Restrictions on the `ompt_set_callback` runtime entry point are as follows:

- The entry point must not return `ompt_set_impossible`.

Cross References

- Callbacks, see [Section 19.4.2](#)
- Monitoring Activity on the Host with OMPT, see [Section 19.2.4](#)
- `ompt_callback_t`, see [Section 19.4.4.1](#)
- `ompt_get_callback_t`, see [Section 19.6.1.4](#)
- `ompt_set_result_t`, see [Section 19.4.4.2](#)

19.6.1.4 `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

```
typedef int (*ompt_get_callback_t) (  
    ompt_callbacks_t event,  
    ompt_callback_t *callback  
);
```

Semantics

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 cannot be `NULL` and must point to valid storage.

Cross References

- Callbacks, see [Section 19.4.2](#)
- `ompt_callback_t`, see [Section 19.4.4.1](#)
- `ompt_set_callback_t`, see [Section 19.6.1.3](#)

19.6.1.5 `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 ompt_data_t *(*ompt_get_thread_data_t) (void);
```

C / C++

Semantics

Each OpenMP thread can have an associated thread data object of type `ompt_data_t`. The `ompt_get_thread_data` runtime entry point, which has type signature `ompt_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 `ompt_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 `ompt_data_none`. This runtime entry point is *async signal safe*.

Cross References

- `ompt_data_t`, see [Section 19.4.4.4](#)

19.6.1.6 `ompt_get_num_procs_t`

Summary

The `ompt_get_num_procs_t` type is the type signature of the `ompt_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 (*ompt_get_num_procs_t) (void);
```

C / C++

Binding

The binding thread set is all threads on the host device.

Semantics

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*.

19.6.1.7 `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 is all threads on a device.

Semantics

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

- `OMP_PLACES`, see [Section 21.1.5](#)
- *place-partition-var* ICV, see [Table 2.1](#)

19.6.1.8 `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 is all threads on a device.

Semantics

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.

1 **Description of Arguments**

2 The *place_num* argument specifies the place that is being queried.

3 The *ids* argument is an array in which the routine can return a vector of processor identifiers in the
4 specified place.

5 The *ids_size* argument indicates the size of the result array that is specified by *ids*.

6 **Effect**

7 If the *ids* array of size *ids_size* is large enough to contain all identifiers then they are returned in *ids*
8 and their order in the array is implementation defined. Otherwise, if the *ids* array is too small, the
9 values in *ids* when the function returns are unspecified. The routine always returns the number of
10 numerical identifiers of the processors that are available to the execution environment in the
11 specified place.

12 **19.6.1.9 ompt_get_place_num_t**

13 **Summary**

14 The **ompt_get_place_num_t** type is the type signature of the **ompt_get_place_num**
15 runtime entry point, which returns the place number of the place to which the current thread is
16 bound.

17 **Format**

```
18          | typedef int (*ompt_get_place_num_t) (void);
            | C / C++
            | C / C++
```

19 **Semantics**

20 When the current thread is bound to a place, **ompt_get_place_num** returns the place number
21 associated with the thread. The returned value is between 0 and one less than the value returned by
22 **ompt_get_num_places**, inclusive. When the current thread is not bound to a place, the routine
23 returns -1. This runtime entry point is *async signal safe*.

24 **19.6.1.10 ompt_get_partition_place_nums_t**

25 **Summary**

26 The **ompt_get_partition_place_nums_t** type is the type signature of the
27 **ompt_get_partition_place_nums** runtime entry point, which returns a list of place
28 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++

Semantics

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

- `OMP_PLACES`, see [Section 21.1.5](#)
- *place-partition-var* ICV, see [Table 2.1](#)

19.6.1.11 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

C / C++

```
typedef int (*ompt_get_proc_id_t) (void);
```

C / C++

Semantics

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*.

19.6.1.12 `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++

Semantics

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 an 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** region, **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 wait identifier of the thread. If `wait_id` is not `NULL` then the entry point assigns the value of the wait identifier of the thread to the object to which `wait_id` points. If the returned state is not one of the specified wait states then the value of the 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_wait_id_t`, see [Section 19.4.4.31](#)
- `ompt_enumerate_states_t`, see [Section 19.6.1.1](#)
- `ompt_state_t`, see [Section 19.4.4.28](#)

19.6.1.13 `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

```
typedef int (*ompt_get_parallel_info_t) (  
    int ancestor_level,  
    ompt_data_t **parallel_data,  
    int *team_size  
);
```

C / C++

C / C++

Semantics

During execution, an OpenMP program may employ nested parallel regions. The `ompt_get_parallel_info` runtime entry point, 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. Information about a parallel region may not be available if the ancestor level is 0; otherwise it must be available if the parallel region exists at the specified ancestor level. The entry point returns 2 if a parallel region exists at the specified ancestor level and the information is available, 1 if a parallel region exists at the specified ancestor level but the information is currently unavailable, and 0 otherwise.

A tool may use the pointer to the data object of a parallel region 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 or the new parallel team.

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

1 the primary thread for a parallel region to free storage for the region immediately after the region
2 ends, yet avoid having some other thread in the team that is executing the region potentially
3 reference the *parallel_data* object for the region after it has been freed.

4 **Description of Arguments**

5 The *ancestor_level* argument specifies the parallel region of interest by its ancestor level. Ancestor
6 level 0 refers to the innermost parallel region; information about enclosing parallel regions may be
7 obtained using larger values for *ancestor_level*.

8 The *parallel_data* argument returns the parallel data if the argument is not *NULL*.

9 The *team_size* argument returns the team size if the argument is not *NULL*.

10 **Effect**

11 If the runtime entry point returns 0 or 1, no argument is modified. Otherwise,
12 `ompt_get_parallel_info` has the following effects:

- 13 • If a non-null value was passed for *parallel_data*, the value returned in *parallel_data* is a
14 pointer to a data word that is associated with the parallel region at the specified level; and
- 15 • If a non-null value was passed for *team_size*, the value returned in the integer to which
16 *team_size* point is the number of threads in the team that is associated with the parallel region.

17 **Constraints on Arguments**

18 While argument *ancestor_level* is passed by value, all other arguments to the entry point must be
19 pointers to variables of the specified types or *NULL*.

20 **Cross References**

- 21 • `ompt_data_t`, see [Section 19.4.4.4](#)

22 **19.6.1.14 ompt_get_task_info_t**

23 **Summary**

24 The `ompt_get_task_info_t` type is the type signature of the `ompt_get_task_info`
25 runtime entry point, which returns information about the task, if any, at the specified ancestor level
26 in the current execution context.

27 **Format**

```
28 C / C++  
29 typedef int (*ompt_get_task_info_t) (  
30     int ancestor_level,  
31     int *flags,  
32     ompt_data_t **task_data,  
33     ompt_frame_t **task_frame,  
34     ompt_data_t **parallel_data,  
35     int *thread_num  
36 );
```

C / C++

Semantics

During execution, an OpenMP thread may be executing an OpenMP task. Additionally, the stack of the thread may contain procedure frames that are associated with suspended OpenMP tasks or OpenMP runtime system routines. To obtain information about any task on the stack of the current thread, 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.

Information about a task region may not be available if the ancestor level is 0; otherwise it must be available if the task region exists at the specified ancestor level. The entry point returns 2 if a task region exists at the specified ancestor level and the information is available, 1 if a task region exists 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;

- If a non-null value was passed for *task_data* then the value that is returned in the object to which it points is a pointer to a data word that is associated with the task at the specified level;
- If a non-null value was passed for *task_frame* then the value that is returned in the object to which *task_frame* points is a pointer to the **ompt_frame_t** structure that is associated with the task at the specified level;
- If a non-null value was passed for *parallel_data* then the value that is returned in the object to which *parallel_data* points is a pointer to a data word that is associated with the parallel region that contains the task at the specified level or, if the task at the specified level is an initial task, *NULL*; and
- If a non-null value was passed for *thread_num*, then the value that is returned in the object to which *thread_num* points indicates the number of the thread in the parallel region that is executing the task at the specified level.

Constraints on Arguments

While argument *ancestor_level* is passed by value, all other arguments to **ompt_get_task_info** must be pointers to variables of the specified types or *NULL*.

Cross References

- **ompt_data_t**, see [Section 19.4.4.4](#)
- **ompt_frame_t**, see [Section 19.4.4.29](#)
- **ompt_task_flag_t**, see [Section 19.4.4.19](#)

19.6.1.15 ompt_get_task_memory_t

Summary

The **ompt_get_task_memory_t** type is the type signature of the **ompt_get_task_memory** runtime entry point, which returns information about memory ranges that are associated with the task.

Format

C / C++

```
typedef int (*ompt_get_task_memory_t) (
    void **addr,
    size_t *size,
    int block
);
```

C / C++

Semantics

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 more memory ranges are 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.

19.6.1.16 `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  
);
```

Semantics

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`, see [Section 19.4.4.3](#)

19.6.1.17 `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++ _____
```

Semantics

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*.

19.6.1.18 `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

```
_____ C / C++ _____  
| typedef uint64_t (*ompt_get_unique_id_t) (void);  
_____ C / C++ _____
```

Semantics

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*.

19.6.1.19 `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);
```

Semantics

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.

19.6.2 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.

19.6.2.1 `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  
);
```

Semantics

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`, see [Section 19.4.4.5](#)

19.6.2.2 `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

```
typedef ompt_device_time_t (*ompt_get_device_time_t) (  
    ompt_device_t *device  
);
```

Semantics

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`, see [Section 19.4.4.5](#)
- `ompt_device_time_t`, see [Section 19.4.4.6](#)

19.6.2.3 `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

```
typedef double (*ompt_translate_time_t) (  
    ompt_device_t *device,  
    ompt_device_time_t time  
);
```

C / C++

C / C++

Semantics

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 `omp_get_wtime`.

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

- `omp_get_wtime`, see [Section 18.10.1](#)
- `ompt_device_t`, see [Section 19.4.4.5](#)
- `ompt_device_time_t`, see [Section 19.4.4.6](#)

19.6.2.4 `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

Restrictions on the `ompt_set_trace_ompt` runtime entry point are as follows:

- The entry point must not return `ompt_set_sometimes_paired`.

Cross References

- Callbacks, see [Section 19.4.2](#)
- Tracing Activity on Target Devices with OMPT, see [Section 19.2.5](#)
- `ompt_device_t`, see [Section 19.4.4.5](#)
- `ompt_set_result_t`, see [Section 19.4.4.2](#)

19.6.2.5 `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++

Semantics

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**.

Restrictions

Restrictions on the **ompt_set_trace_native** runtime entry point are as follows:

- The entry point must not return **ompt_set_sometimes_paired**.

Cross References

- Tracing Activity on Target Devices with OMPT, see [Section 19.2.5](#)
- **ompt_device_t**, see [Section 19.4.4.5](#)
- **ompt_native_mon_flag_t**, see [Section 19.4.4.18](#)
- **ompt_set_result_t**, see [Section 19.4.4.2](#)

19.6.2.6 ompt_get_buffer_limits_t

Summary

The **ompt_get_buffer_limits_t** type is the type signature of the **ompt_get_buffer_limits** runtime entry point, which returns the maximum number of concurrent buffer allocations and the recommended size of any buffer allocation that will be requested of the tool for a given device.

Format

C / C++

```
typedef void (*ompt_get_buffer_limits_t) (  
    ompt_device_t *device,  
    int *max_concurrent_allocs,  
    size_t *recommended_bytes  
);
```

C / C++

Semantics

The `ompt_get_buffer_limits` runtime entry point, which has type signature `ompt_get_buffer_limits_t`, returns the maximum number of concurrent buffer allocations and the recommended size of any buffer allocation that will be requested of the tool for a given device. A first party tool may use this entry point prior to a call to the `ompt_start_trace` entry point to determine the total size of the buffers that the implementation would need for tracing activity on the device at any given time.

The limits returned by this entry point remain the same on each successive call unless the `ompt_stop_trace` entry point is called for the same target device between the successive calls.

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 `*max_concurrent_allocs` argument indicates the maximum number of buffer allocations that may be requested by an implementation for tracing activity on the target device without the implementation dispatching a callback with the type signature `ompt_callback_buffer_complete_t` and the `buffer_owned` argument set to a non-zero value for any of the buffers.

The `*recommended_bytes` argument indicates the recommended buffer size of the buffer to be returned by the first party tool when the implementation dispatches a callback with the type signature `ompt_callback_buffer_request_t` for the target device.

Cross References

- `ompt_callback_buffer_complete_t`, see [Section 19.5.2.24](#)
- `ompt_callback_buffer_request_t`, see [Section 19.5.2.23](#)
- `ompt_device_t`, see [Section 19.4.4.5](#)
- `ompt_start_trace_t`, see [Section 19.6.2.7](#)
- `ompt_stop_trace_t`, see [Section 19.6.2.10](#)

19.6.2.7 `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  
);
```

Semantics

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 buffer in which a device can 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_callback_buffer_complete_t`, see [Section 19.5.2.24](#)
- `ompt_callback_buffer_request_t`, see [Section 19.5.2.23](#)
- `ompt_device_t`, see [Section 19.4.4.5](#)

19.6.2.8 `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

C / C++

```
typedef int (*ompt_pause_trace_t) (  
    ompt_device_t *device,  
    int begin_pause  
);
```

C / C++

Semantics

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`, see [Section 19.4.4.5](#)

19.6.2.9 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

C / C++

```
typedef int (*ompt_flush_trace_t) (  
    ompt_device_t *device  
);
```

C / C++

Semantics

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`, see [Section 19.4.4.5](#)

19.6.2.10 `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

```
typedef int (*ompt_stop_trace_t) (  
    ompt_device_t *device  
);
```

C / C++

Semantics

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 be 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`, see [Section 19.4.4.5](#)

19.6.2.11 `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++

Semantics

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_buffer_cursor_t`, see [Section 19.4.4.8](#)
- `ompt_device_t`, see [Section 19.4.4.5](#)

19.6.2.12 `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++

Semantics

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

- Record Type, see [Section 19.4.3.1](#)
- `ompt_buffer_cursor_t`, see [Section 19.4.4.8](#)
- `ompt_buffer_t`, see [Section 19.4.4.7](#)

19.6.2.13 `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

```
typedef ompt_record_ompt_t *(*ompt_get_record_ompt_t) (  
    ompt_buffer_t *buffer,  
    ompt_buffer_cursor_t current  
);
```

Semantics

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

- Standard Trace Record Type, see [Section 19.4.3.4](#)
- `ompt_buffer_cursor_t`, see [Section 19.4.4.8](#)
- `ompt_device_t`, see [Section 19.4.4.5](#)

19.6.2.14 `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

```
typedef void *(*ompt_get_record_native_t) (  
    ompt_buffer_t *buffer,  
    ompt_buffer_cursor_t current,  
    ompt_id_t *host_op_id  
);
```

Semantics

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_buffer_cursor_t`, see [Section 19.4.4.8](#)
- `ompt_buffer_t`, see [Section 19.4.4.7](#)
- `ompt_id_t`, see [Section 19.4.4.3](#)

19.6.2.15 `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++

Semantics

An OpenMP implementation may execute on a device that logs trace records in a native (device-specific) format that a tool cannot interpret directly. The `ompt_get_record_abstract` runtime entry point of a device, 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

- Native Record Abstract Type, see [Section 19.4.3.3](#)

19.6.3 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++

Semantics

An OpenMP implementation provides pointers to lookup routines that provide 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.

If the provided function name is unknown to the OpenMP implementation, the function returns *NULL*. In a compliant implementation, the lookup function provided by the tool initializer for the OMPT callback interface returns a valid function pointer for any OMPT runtime entry point name listed in Table 19.1.

A compliant implementation of a lookup function passed to a tool's `ompt_device_initialize` callback must provide non-*NULL* function pointers for all strings in Table 19.4, except for `ompt_set_trace_ompt` and `ompt_get_record_ompt`, as described in Section 19.2.5.

Description of Arguments

The *interface_function_name* argument is a C string that represents the name of a runtime entry point.

Cross References

- Entry Points in the OMPT Callback Interface, see [Section 19.6.1](#)
- Entry Points in the OMPT Device Tracing Interface, see [Section 19.6.2](#)
- Tracing Activity on Target Devices with OMPT, see [Section 19.2.5](#)
- `ompt_initialize_t`, see [Section 19.5.1.1](#)

20 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 that the third-party tool can load. 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 debuggers to inspect the OpenMP state of a live program or core file in an implementation-agnostic manner. That is, a third-party tool that uses OMPD should work with any conforming OpenMP implementation. An OpenMP implementer provides a library for OMPD that a third-party tool can dynamically load. The third-party tool can use the interface exported by the OMPD library to 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 the OpenMP program, or to find the addresses of symbols in it. 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 the third-party tool uses the API 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 directs the third-party tool to perform them by using the callback interface that the third-party tool exports.

The OMPD design insulates third-party 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 third-party tool are deployed, so that, for example, the third-party tool and the OpenMP program are not required to execute on the same machine.

Generally, the third-party tool does not interact directly with the OpenMP runtime but instead interacts with the runtime through the OMPD library. However, a few cases require the third-party tool to access the OpenMP runtime directly. These cases fall into two broad categories. The first is during initialization where the third-party 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 20.2.2](#) and [Section 20.2.3](#). The second category relates to arranging for the third-party 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 20.6](#). Each of these symbols corresponds to an event type. The OpenMP runtime must ensure that control passes through the appropriate named location when events occur. If the third-party tool requires notification of an event, it can plant a breakpoint at the matching

1 location. The location can, but may not, be a function. It can, for example, simply be a label.
2 However, the names of the locations must have external C linkage.

3 20.1 OMPD Interfaces Definitions

C / C++

4 A compliant implementation must supply a set of definitions for the OMPD runtime entry points,
5 OMPD third-party tool callback signatures, third-party tool interface functions and the special data
6 types of their parameters and return values. These definitions, which are listed throughout this
7 chapter, and their associated declarations shall be provided in a header file named `omp-tools.h`.
8 In addition, the set of definitions may specify other implementation-specific values.

9 The `ompd_dll_locations` variable, all OMPD third-party tool interface functions, and all
10 OMPD runtime entry points are external symbols with C linkage.

C / C++

11 20.2 Activating a Third-Party Tool

12 The third-party tool and the OpenMP program exist as separate processes. Thus, coordination is
13 required between the OpenMP runtime and the third-party tool for OMPD.

14 20.2.1 Enabling Runtime Support for OMPD

15 In order to support third-party tools, the OpenMP runtime may need to collect and to store
16 information that it may not otherwise maintain. The OpenMP runtime collects whatever
17 information is necessary to support OMPD if the environment variable `OMP_DEBUG` is set to
18 *enabled*.

19 Cross References

- 20 • `OMP_DEBUG`, see [Section 21.4.1](#)

21 20.2.2 `ompd_dll_locations`

22 Summary

23 The `ompd_dll_locations` global variable points to the locations of OMPD libraries that are
24 compatible with the OpenMP implementation.

25 Format

C

```
26 | extern const char **ompd_dll_locations;
```

C

Semantics

An OpenMP runtime may have more than one OMPD library. The third-party 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 pathname string pointers that provides the names of any compatible OMPD libraries. This variable must have `C` linkage. The third-party tool uses the name of the variable verbatim and, in particular, does not apply any name mangling before performing the look up.

The architecture on which the third-party tool and, thus, the OMPD library execute does not have to match the architecture on which the OpenMP program that is being examined executes. The third-party tool must interpret the contents of `ompd_dll_locations` to find a suitable OMPD library that matches its own architectural characteristics. On platforms that support different architectures (for example, 32-bit vs 64-bit), OpenMP implementations are encouraged to provide an OMPD library for each supported architecture that can handle OpenMP programs that run on any supported architecture. 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.

The `ompd_dll_locations` variable 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. Thus, if a third-party tool, such as a debugger, stops execution of the OpenMP program at any point at which `ompd_dll_locations` is non-null, the vector of strings to which it points shall be valid and complete.

Cross References

- `ompd_dll_locations_valid`, see [Section 20.2.3](#)

20.2.3 `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);
```

Semantics

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.

20.3 OMPD Data Types

This section defines OMPD data types.

20.3.1 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++ _____
```

20.3.2 Wait ID Type

Summary

A variable of `ompd_wait_id_t` type identifies the object on which a thread waits.

Format

```
_____ C / C++ _____  
| typedef uint64_t ompd_wait_id_t; |  
_____ C / C++ _____
```

Semantics

The values and meaning of `ompd_wait_id_t` are the same as those defined for the `ompt_wait_id_t` type.

Cross References

- `ompt_wait_id_t`, see [Section 19.4.4.31](#)

20.3.3 Basic Value Types

Summary

These definitions represent 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++

Semantics

The *ompd_addr_t* type represents an address in an OpenMP process with an unsigned integer type.

The *ompd_word_t* type represents a data word from the OpenMP runtime with a signed integer

type. The *ompd_seg_t* type represents a segment value with an unsigned integer type.

20.3.4 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++

Semantics

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.

Cross References

- Basic Value Types, see [Section 20.3.3](#)

20.3.5 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++

Semantics

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_flag_t` enumeration type.

Cross References

- Address Type, see [Section 20.3.4](#)
- Basic Value Types, see [Section 20.3.3](#)
- `ompt_frame_flag_t`, see [Section 19.4.4.30](#)

20.3.6 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++

Semantics

OpenMP runtimes may utilize different underlying devices, each represented by a device identifier. The device identifiers can vary in size and format and, thus, are not explicitly represented in the OMPD interface. Instead, a device identifier is passed across the interface via its `ompd_device_t` kind, its size in bytes and a pointer to where it is stored. The OMPD library and the third-party tool use the `ompd_device_t` kind to interpret the format of the device identifier that is referenced by the pointer argument. Each different device identifier kind is represented by a unique unsigned 64-bit integer value. Recommended values of `ompd_device_t` kinds are defined in the `ompd-types.h` header file, which is contained in the *Supplementary Source Code* package available via <https://www.openmp.org/specifications/>.

20.3.7 Native Thread Identifiers

Summary

The `ompd_thread_id_t` type provides information about native threads.

Format

```
typedef uint64_t ompd_thread_id_t;
```

Semantics

OpenMP runtimes may use different native thread implementations. Native thread identifiers for these implementations can vary in size and format and, thus, are not explicitly represented in the OMPD interface. Instead, a native thread identifier is passed across the interface via its `ompd_thread_id_t` kind, its size in bytes and a pointer to where it is stored. The OMPD library and the third-party tool use the `ompd_thread_id_t` kind to interpret the format of the native thread identifier that is referenced by the pointer argument. Each different native thread identifier kind is represented by a unique unsigned 64-bit integer value. Recommended values of `ompd_thread_id_t` kinds, and formats for some corresponding native thread identifiers, are defined in the `ompd-types.h` header file, which is contained in the *Supplementary Source Code* package available via <https://www.openmp.org/specifications/>.

20.3.8 OMPD Handle Types

Summary

The OMPD library defines handles for referring to address spaces, threads, parallel regions and tasks that are managed by the OpenMP runtime. The internal structures that these handles represent are opaque to the third-party tool.

Format

```
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;
```

Semantics

OMP uses handles for the following entities that are managed by the OpenMP runtime: 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. Handles are defined by the OMPD library and are opaque to the third-party tool. A handle remains constant and valid while the associated entity is managed by the OpenMP runtime or until it is released with the corresponding third-party tool interface routine for releasing handles of that type. If a tool receives notification of the end of the lifetime of a managed entity (see [Section 20.6](#)) or it releases the handle, the handle may no longer be referenced.

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 third-party 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.

20.3.9 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_teams          = 7,
    ompd_scope_target         = 8
} ompd_scope_t;
C / C++
```

Semantics

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 20.1](#).

TABLE 20.1: Mapping of Scope Type and OMPD Handles

Scope types	Handles
<i>ompd_scope_global</i>	Address space handle for the host device
<i>ompd_scope_address_space</i>	Any address space handle
<i>ompd_scope_thread</i>	Any thread handle
<i>ompd_scope_parallel</i>	Any parallel handle
<i>ompd_scope_implicit_task</i>	Task handle for an implicit task
<i>ompd_scope_teams</i>	Parallel handle for an implicit parallel region generated from a teams construct
<i>ompd_scope_target</i>	Parallel handle for an implicit parallel region generated from a target construct
<i>ompd_scope_task</i>	Any task handle

20.3.10 Team Generator Types

Summary

The `ompd_team_generator_t` type identifies the generator of a given team.

Format

```
C / C++
typedef enum ompd_team_generator_t {
    ompd_generator_program      = 0,
    ompd_generator_parallel     = 1,
    ompd_generator_teams        = 2,
    ompd_generator_target       = 3
} ompd_team_generator_t;
```

Semantics

The `ompd_team_generator_t` type represents the value of the *team-generator-var* ICV. The `ompd_generator_program` value indicates that the team is the initial team created at the start of the OpenMP program. The `ompd_generator_parallel`, `ompd_generator_teams`, and `ompd_generator_target` values indicate that the team was created by an encountered **parallel** construct, **teams** construct, or **target** construct, respectively.

20.3.11 ICV ID Type

Summary

The `ompd_icv_id_t` type identifies an OpenMP implementation ICV.

Format

C / C++

```
typedef uint64_t ompd_icv_id_t;
```

C / C++

Semantics

The `ompd_icv_id_t` type identifies OpenMP implementation ICVs. `ompd_icv_undefined` is an instance of this type with the value 0.

20.3.12 Tool Context Types

Summary

A third-party tool defines contexts to identify abstractions uniquely. The internal structures that these contexts represent are opaque to the OMPD library.

Format

C / C++

```
typedef struct _ompd_aspace_cont ompd_address_space_context_t;  
typedef struct _ompd_thread_cont ompd_thread_context_t;
```

C / C++

Semantics

A third-party tool uniquely defines an *address space context* to identify the address space for the process that it is monitoring. Similarly, it uniquely defines a *thread context* to identify a native thread of the process that it is monitoring. These contexts are opaque to the OMPD library.

20.3.13 Return Code Types

Summary

The `ompd_rc_t` type is the return code type of an OMPD operation.

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_incomplete = 11,
    ompd_rc_callback_error = 12,
    ompd_rc_incompatible_handle = 13
} ompd_rc_t;
```

C / C++

Semantics

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_incompatible_handle` is returned when the specified handle is incompatible with the query function;
- `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;

- 1 • `ompd_rc_device_read_error` is returned when a read operation failed on the device;
- 2 • `ompd_rc_device_write_error` is returned when a write operation failed on the
- 3 device;
- 4 • `ompd_rc_incompatible` is returned when this OMPD library is incompatible with the
- 5 OpenMP program or is not capable of handling it;
- 6 • `ompd_rc_nomem` is returned when a memory allocation fails;
- 7 • `ompd_rc_incomplete` is returned when the information provided on return is
- 8 incomplete, while the arguments are still set to valid values; and
- 9 • `ompd_rc_callback_error` is returned when the callback interface or any one of the
- 10 required callback routines provided by the third-party tool is invalid.

11 20.3.14 Primitive Type Sizes

12 Summary

13 The `ompd_device_type_sizes_t` type provides the size of primitive types in the OpenMP
14 architecture address space.

15 Format

C / C++

```
16       typedef struct ompd_device_type_sizes_t {  
17           uint8_t  sizeof_char;  
18           uint8_t  sizeof_short;  
19           uint8_t  sizeof_int;  
20           uint8_t  sizeof_long;  
21           uint8_t  sizeof_long_long;  
22           uint8_t  sizeof_pointer;  
23       } ompd_device_type_sizes_t;
```

C / C++

24 Semantics

25 The `ompd_device_type_sizes_t` type is used in operations through which the OMPD
26 library can interrogate the third-party tool about the size of primitive types for the target
27 architecture of the OpenMP runtime, as returned by the `sizeof` operator. The fields of
28 **`ompd_device_type_sizes_t`** give the sizes of the eponymous basic types used by the
29 OpenMP runtime. As the third-party tool and the OMPD library, by definition, execute on the same
30 architecture, the size of the fields can be given as `uint8_t`.

31 Cross References

- 32 • `ompd_callback_sizeof_fn_t`, see [Section 20.4.2.2](#)

20.4 OMPD Third-Party 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 third-party tool is examining. The OpenMP process on which the third-party 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 third-party 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 third-party tool to the OMPD library.

Description of Return Codes

All of the OMPD callback functions must return the following return codes or function-specific return codes:

- `ompd_rc_ok` on success; or
- `ompd_rc_stale_handle` if an invalid context argument is provided.

20.4.1 Memory Management of OMPD Library

`ompd_callback_memory_alloc_fn_t` (see [Section 20.4.1.1](#)) and `ompd_callback_memory_free_fn_t` (see [Section 20.4.1.2](#)) are provided by the third-party 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 third-party tool may use.

If the OMPD library is implemented in C++ then memory management operators, like `new` and `delete` and their variants, *must all* be overloaded and implemented in terms of the callbacks that the third-party tool provides. The OMPD library must be implemented in a manner such that any of its definitions of `new` or `delete` do not interfere with any that the third-party tool defines.

In some cases, the OMPD library must allocate memory to return results to the third-party tool. The third-party tool then owns this memory and has the responsibility to release it. Thus, the OMPD library and the third-party tool must use the same memory manager.

The OMPD library creates OMPD handles, which are opaque to the third-party tool and may have a complex internal structure. The third-party tool cannot determine if the handle pointers that the API returns correspond to discrete heap allocations. Thus, the third-party 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 OMPD API includes functions that the third-party tool must use when it no longer needs a handle.

1 A third-party tool creates contexts and passes them to the OMPD library. The OMPD library does
2 not release contexts; instead the third-party tool releases them after it releases any handles that may
3 reference the contexts.

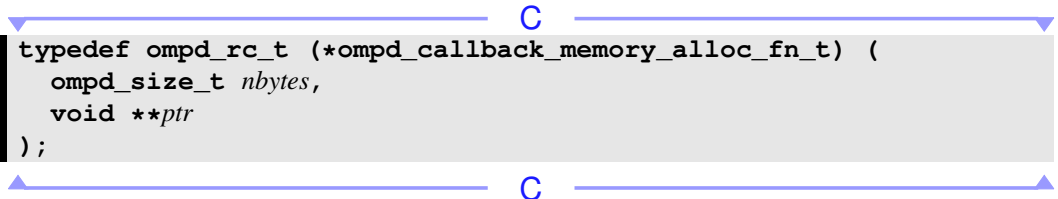
4 **20.4.1.1 ompd_callback_memory_alloc_fn_t**

5 **Summary**

6 The `ompd_callback_memory_alloc_fn_t` type is the type signature of the callback routine
7 that the third-party tool provides to the OMPD library to allocate memory.

8 **Format**

```
9 typedef ompd_rc_t (*ompd_callback_memory_alloc_fn_t) (  
10     ompd_size_t nbytes,  
11     void **ptr  
12 );
```



13 **Semantics**

14 The `ompd_callback_memory_alloc_fn_t` type is the type signature of the memory
15 allocation callback routine that the third-party tool provides. The OMPD library may call the
16 `ompd_callback_memory_alloc_fn_t` callback function to allocate memory.

17 **Description of Arguments**

18 The *nbytes* argument is the size in bytes of the block of memory to allocate.

19 The address of the newly allocated block of memory is returned in the location to which the *ptr*
20 argument points. The newly allocated block is suitably aligned for any type of variable and is not
21 guaranteed to be set to zero.

22 **Description of Return Codes**

23 Routines that use the `ompd_callback_memory_alloc_fn_t` type may return the general
24 return codes listed at the beginning of [Section 20.4](#).

25 **Cross References**

- 26 • Return Code Types, see [Section 20.3.13](#)
- 27 • Size Type, see [Section 20.3.1](#)
- 28 • The Callback Interface, see [Section 20.4.6](#)

29 **20.4.1.2 ompd_callback_memory_free_fn_t**

30 **Summary**

31 The `ompd_callback_memory_free_fn_t` type is the type signature of the callback routine
32 that the third-party tool provides to the OMPD library to deallocate memory.

Format

```
typedef ompd_rc_t (*ompd_callback_memory_free_fn_t) (  
    void *ptr  
);
```

Semantics

The `ompd_callback_memory_free_fn_t` type is the type signature of the memory deallocation callback routine that the third-party 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.

Description of Return Codes

Routines that use the `ompd_callback_memory_free_fn_t` type may return the general return codes listed at the beginning of [Section 20.4](#).

Cross References

- Return Code Types, see [Section 20.3.13](#)
- The Callback Interface, see [Section 20.4.6](#)
- `ompd_callback_memory_alloc_fn_t`, see [Section 20.4.1.1](#)

20.4.2 Context Management and Navigation

Summary

The third-party tool provides the OMPD library with callbacks to manage and to navigate context relationships.

20.4.2.1 `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 third-party tool provides to the OMPD library to map a native thread identifier to a third-party 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
);
```

Semantics

The `ompd_callback_get_thread_context_for_thread_id_fn_t` is the type signature of the context mapping callback routine that the third-party tool provides. This callback maps a native thread identifier to a third-party tool thread context. The native 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 third-party 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 third-party tool thread context.

Description of Return Codes

In addition to the general return codes listed at the beginning of [Section 20.4](#), routines that use the `ompd_callback_get_thread_context_for_thread_id_fn_t` type may also return the following return codes:

- `ompd_rc_bad_input` if a different value in `sizeof_thread_id` is expected for the native thread identifier kind given by `kind`; or
- `ompd_rc_unsupported` if the native thread identifier `kind` is not supported.

Restrictions

Restrictions on routines that use

`ompd_callback_get_thread_context_for_thread_id_fn_t` are as follows:

- The provided `thread_context` must be valid until the OMPD library returns from the OMPD third-party tool interface routine.

Cross References

- Native Thread Identifiers, see [Section 20.3.7](#)
- Return Code Types, see [Section 20.3.13](#)
- Size Type, see [Section 20.3.1](#)
- The Callback Interface, see [Section 20.4.6](#)
- Tool Context Types, see [Section 20.3.12](#)

20.4.2.2 ompd_callback_sizeof_fn_t

Summary

The `ompd_callback_sizeof_fn_t` type is the type signature of the callback routine that the third-party 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  
);
```

Semantics

The `ompd_callback_sizeof_fn_t` is the type signature of the type-size query callback routine that the third-party 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.

Description of Return Codes

Routines that use the `ompd_callback_sizeof_fn_t` type may return the general return codes listed at the beginning of [Section 20.4](#).

Cross References

- Primitive Type Sizes, see [Section 20.3.14](#)
- Return Code Types, see [Section 20.3.13](#)
- The Callback Interface, see [Section 20.4.6](#)
- Tool Context Types, see [Section 20.3.12](#)

20.4.3 Accessing Memory in the OpenMP Program or Runtime

The OMPD library cannot directly read from or write to memory of the OpenMP program. Instead the OMPD library must use callbacks that the third-party tool provides so that the third-party tool performs the operation.

20.4.3.1 `ompd_callback_symbol_addr_fn_t`

Summary

The `ompd_callback_symbol_addr_fn_t` type is the type signature of the callback that the third-party 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  
);
```

Semantics

The `ompd_callback_symbol_addr_fn_t` is the type signature of the symbol-address query callback routine that the third-party 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 third-party tool's representation of the address space of the process, core file, or device.

The *thread_context* argument is *NULL* for global memory accesses. 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. In this case, 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 third-party 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 address space. The symbol can correspond to any type of object, such as a variable,

1 thread local storage variable, function, or untyped label. The symbol can have local, global, or
2 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 it 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 name
6 may not be an exact match for the name seen by the third-party tool. If *file_name* is *NULL* then the
7 third-party tool first tries to find the symbol in the executable file, and, if the symbol is not found,
8 the third-party tool tries to find the symbol in the shared libraries in the order in which the shared
9 libraries are loaded into the address space. If *file_name* is non-null then the third-party tool first
10 tries to find the symbol in the libraries that match the name in the *file_name* argument, and, if the
11 symbol is not found, the third-party tool then uses the same procedure as when *file_name* is *NULL*.

12 The callback does not support finding either symbols that are dynamically allocated on the call
13 stack or statically allocated symbols that are defined within the scope of a function or subroutine.

14 The callback returns the address of the symbol in the location to which *symbol_addr* points.

15 Description of Return Codes

16 In addition to the general return codes listed at the beginning of [Section 20.4](#), routines that use the
17 `ompd_callback_symbol_addr_fn_t` type may also return the following return codes:

- 18 • `ompd_rc_error` if the requested symbol is not found; or
- 19 • `ompd_rc_bad_input` if no symbol name is provided.

20 Restrictions

21 Restrictions on routines that use the `ompd_callback_symbol_addr_fn_t` type are as
22 follows:

- 23 • The *address_space_context* argument must be non-null.
- 24 • The symbol that the *symbol_name* argument specifies must be defined.

25 Cross References

- 26 • Address Type, see [Section 20.3.4](#)
- 27 • Return Code Types, see [Section 20.3.13](#)
- 28 • The Callback Interface, see [Section 20.4.6](#)
- 29 • Tool Context Types, see [Section 20.3.12](#)

30 20.4.3.2 `ompd_callback_memory_read_fn_t`

31 Summary

32 The `ompd_callback_memory_read_fn_t` type is the type signature of the callback that the
33 third-party tool provides to read data (*read_memory*) or a string (*read_string*) from an OpenMP
34 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  
);
```

Semantics

The `ompd_callback_memory_read_fn_t` is the type signature of the read callback routines that the third-party tool provides.

The `read_memory` callback copies a block of data from `addr` within the address space given by `address_space_context` to the third-party tool `buffer`.

The `read_string` callback copies a string to which `addr` points, including the terminating null byte (`'\0'`), to the third-party 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 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 for global memory accesses 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 20.4.4](#)) to interpret the data.

Description of Return Codes

In addition to the general return codes listed at the beginning of [Section 20.4](#), routines that use the `ompd_callback_memory_read_fn_t` type may also return the following return codes:

- `ompd_rc_incomplete` if no terminating null byte is found while reading `nbytes` using the `read_string` callback; or
- `ompd_rc_error` if unallocated memory is reached while reading `nbytes` using either the `read_memory` or `read_string` callback.

Cross References

- Address Type, see [Section 20.3.4](#)
- Return Code Types, see [Section 20.3.13](#)
- Size Type, see [Section 20.3.1](#)
- The Callback Interface, see [Section 20.4.6](#)
- Tool Context Types, see [Section 20.3.12](#)
- Data Format Conversion: `ompd_callback_device_host_fn_t`, see [Section 20.4.4](#)

20.4.3.3 `ompd_callback_memory_write_fn_t`

Summary

The `ompd_callback_memory_write_fn_t` type is the type signature of the callback that the third-party 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  
);
```

Semantics

The `ompd_callback_memory_write_fn_t` is the type signature of the write callback routine that the third-party tool provides. The OMPD library may call this callback to have the third-party 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 for global memory accesses 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 20.4.4](#)) to render the data into a form that is compatible with the OpenMP runtime.

Description of Return Codes

Routines that use the `ompd_callback_memory_write_fn_t` type may return the general return codes listed at the beginning of [Section 20.4](#).

Cross References

- Address Type, see [Section 20.3.4](#)
- Return Code Types, see [Section 20.3.13](#)
- Size Type, see [Section 20.3.1](#)
- The Callback Interface, see [Section 20.4.6](#)
- Tool Context Types, see [Section 20.3.12](#)
- Data Format Conversion: `ompd_callback_device_host_fn_t`, see [Section 20.4.4](#)

20.4.4 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 third-party tool provides to convert data between the formats that the third-party 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  
);
```

Semantics

The architecture on which the third-party tool and the OMPD library execute may be different from the architecture on which the OpenMP program that is being examined executes. 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 third-party tool and OMPD library use and the ones that the OpenMP program use. The callback with the `ompd_callback_device_host_fn_t` type signature converts data between the 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.

Description of Return Codes

Routines that use the `ompd_callback_device_host_fn_t` type may return the general return codes listed at the beginning of [Section 20.4](#).

Cross References

- Return Code Types, see [Section 20.3.13](#)
- Size Type, see [Section 20.3.1](#)
- The Callback Interface, see [Section 20.4.6](#)
- Tool Context Types, see [Section 20.3.12](#)

20.4.5 `ompd_callback_print_string_fn_t`

Summary

The `ompd_callback_print_string_fn_t` type is the type signature of the callback that the third-party 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  
);
```

Semantics

The OMPD library may call the `ompd_callback_print_string_fn_t` callback function to emit output, such as logging or debug information. The third-party tool may set the `ompd_callback_print_string_fn_t` callback function to `NULL` to prevent the OMPD library from emitting output. The OMPD library 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.

Description of Return Codes

Routines that use the `ompd_callback_print_string_fn_t` type may return the general return codes listed at the beginning of [Section 20.4](#).

Cross References

- Return Code Types, see [Section 20.3.13](#)
- The Callback Interface, see [Section 20.4.6](#)

20.4.6 The Callback Interface

Summary

All OMPD library interactions with the OpenMP program must be through a set of callbacks that the third-party tool provides. These callbacks must also be used for allocating or releasing resources, such as memory, that the OMPD 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;
```

Semantics

The set of callbacks that the OMPD library must use is collected in the `ompd_callbacks_t` structure. An instance of this type is passed to the OMPD library as a parameter to `ompd_initialize` (see [Section 20.5.1.1](#)). Each field points to a function that the OMPD library must use either to interact with the OpenMP program or for memory operations.

1 The *alloc_memory* and *free_memory* fields are pointers to functions the OMPD library uses to
2 allocate and to release dynamic memory.

3 The *print_string* field points to a function that prints a string.

4 The architecture on which the OMPD library and third-party tool execute may be different from the
5 architecture on which the OpenMP program that is being examined executes. The *sizeof_type* field
6 points to a function that allows the OMPD library to determine the sizes of the basic integer and
7 pointer types that the OpenMP program uses. Because of the potential differences in the targeted
8 architectures, the conventions for representing data in the OMPD library and the OpenMP program
9 may be different. The *device_to_host* field points to a function that translates data from the
10 conventions that the OpenMP program uses to those that the third-party tool and OMPD library
11 use. The reverse operation is performed by the function to which the *host_to_device* field points.

12 The *symbol_addr_lookup* field points to a callback that the OMPD library can use to find the
13 address of a global or thread local storage symbol. The *read_memory*, *read_string* and
14 *write_memory* fields are pointers to functions for reading from and writing to global memory or
15 thread local storage in the OpenMP program.

16 The *get_thread_context_for_thread_id* field is a pointer to a function that the OMPD library can
17 use to obtain a thread context that corresponds to a native thread identifier.

18 Cross References

- 19 • Data Format Conversion: `ompd_callback_device_host_fn_t`, see [Section 20.4.4](#)
- 20 • `ompd_callback_get_thread_context_for_thread_id_fn_t`, see
21 [Section 20.4.2.1](#)
- 22 • `ompd_callback_memory_alloc_fn_t`, see [Section 20.4.1.1](#)
- 23 • `ompd_callback_memory_free_fn_t`, see [Section 20.4.1.2](#)
- 24 • `ompd_callback_memory_read_fn_t`, see [Section 20.4.3.2](#)
- 25 • `ompd_callback_memory_write_fn_t`, see [Section 20.4.3.3](#)
- 26 • `ompd_callback_print_string_fn_t`, see [Section 20.4.5](#)
- 27 • `ompd_callback_sizeof_fn_t`, see [Section 20.4.2.2](#)
- 28 • `ompd_callback_symbol_addr_fn_t`, see [Section 20.4.3.1](#)

29 20.5 OMPD Tool Interface Routines

30 This section defines the interface provided by the OMPD library to be used by the third-party tool.
31 Some interface routines require one or more specified threads to be *stopped* for the returned values
32 to be meaningful. In this context, a stopped thread is a thread that is not modifying the observable
33 OpenMP runtime state.

Description of Return Codes

All of the OMPD Tool Interface Routines must return function-specific return codes or any of the following return codes:

- `ompd_rc_stale_handle` if a provided handle is stale;
- `ompd_rc_bad_input` if an invalid value is provided for any input argument;
- `ompd_rc_callback` if a callback returned an unexpected error, which leads to a failure of the query;
- `ompd_rc_needs_state_tracking` if the information cannot be provided while the `debug-var` is disabled;
- `ompd_rc_ok` on success; or
- `ompd_rc_error` for any other error.

20.5.1 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 20.5.1.2](#)). If the tool supports the version that `ompd_get_api_version` returns, the tool starts the initialization by calling `ompd_initialize` (see [Section 20.5.1.1](#)) 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.

20.5.1.1 `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  
);
```

Semantics

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.

Description of Return Codes

This routine must return any of the general return codes listed at the beginning of [Section 20.5](#) or any of the following return codes:

- `ompd_rc_bad_input` if invalid callbacks are provided; or
- `ompd_rc_unsupported` if the requested API version cannot be provided.

Cross References

- Return Code Types, see [Section 20.3.13](#)
- The Callback Interface, see [Section 20.4.6](#)
- `ompd_get_api_version`, see [Section 20.5.1.2](#)

20.5.1.2 `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);
```

Semantics

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.

Description of Return Codes

This routine must return any of the general return codes listed at the beginning of [Section 20.5](#).

Cross References

- Return Code Types, see [Section 20.3.13](#)

20.5.1.3 `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);
```

Semantics

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 the *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 20.5.1.1](#)). Accordingly, the OMPD library must not use heap or stack memory for the string.

The signatures of `ompd_get_api_version` (see [Section 20.5.1.2](#)) 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.

Description of Return Codes

This routine must return any of the general return codes listed at the beginning of [Section 20.5](#).

Cross References

- Return Code Types, see [Section 20.3.13](#)

20.5.1.4 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);
```

Semantics

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 provided earlier during the call to `ompd_initialize`.

Description of Return Codes

This routine must return any of the general return codes listed at the beginning of [Section 20.5](#) or the following return code:

- `ompd_rc_unsupported` if the OMPD library is not initialized.

Cross References

- Return Code Types, see [Section 20.3.13](#)

20.5.2 Per OpenMP Process Initialization and Finalization

20.5.2.1 ompd_process_initialize

Summary

A tool calls `ompd_process_initialize` to obtain an address space handle for the host device 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 **host_handle  
);
```

Semantics

A tool calls `ompd_process_initialize` to obtain an address space handle for the host device 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 or core file. This call allows the OMPD library to confirm that it can handle the OpenMP process or core file that `context` identifies.

Description of Arguments

The `context` argument is an opaque handle that the tool provides to address an address space from the host device. On return, the `host_handle` argument provides an opaque handle to the tool for this address space, which the tool must release when it is no longer needed.

Description of Return Codes

This routine must return any of the general return codes listed at the beginning of [Section 20.5](#) or the following return code:

- `ompd_rc_incompatible` if the OMPD library is incompatible with the runtime library loaded in the process.

Cross References

- OMPD Handle Types, see [Section 20.3.8](#)
- Return Code Types, see [Section 20.3.13](#)
- Tool Context Types, see [Section 20.3.12](#)
- `ompd_rel_address_space_handle`, see [Section 20.5.2.3](#)

20.5.2.2 `ompd_device_initialize`

Summary

A tool calls `ompd_device_initialize` to obtain an address space handle for a non-host device that has at least one active target region.

Format

```
ompd_rc_t ompd_device_initialize(  
    ompd_address_space_handle_t *host_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  
);
```

Semantics

A tool calls `ompd_device_initialize` to obtain an address space handle for a non-host 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 *host_handle* argument is an opaque handle that the tool provides to reference the host device address space associated with an OpenMP process or core file. The *device_context* argument is an opaque handle that the tool provides to reference a non-host 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.

Description of Return Codes

This routine must return any of the general return codes listed at the beginning of [Section 20.5](#) or the following return code:

- `ompd_rc_unsupported` if the OMPD library has no support for the specific device.

Cross References

- OMPD Handle Types, see [Section 20.3.8](#)
- Return Code Types, see [Section 20.3.13](#)
- Size Type, see [Section 20.3.1](#)
- System Device Identifiers, see [Section 20.3.6](#)
- Tool Context Types, see [Section 20.3.12](#)

20.5.2.3 `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  
);
```

Semantics

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

Restrictions to the `ompd_rel_address_space_handle` routine are as follows:

- An address space context must not be used after the corresponding address space handle is released.

Description of Return Codes

This routine must return any of the general return codes listed at the beginning of [Section 20.5](#).

Cross References

- OMPD Handle Types, see [Section 20.3.8](#)
- Return Code Types, see [Section 20.3.13](#)

20.5.2.4 ompd_get_device_thread_id_kinds

Summary

The `ompd_get_device_thread_id_kinds` function returns a list of supported native thread identifier kinds and a corresponding list of their respective sizes.

Format

```
ompd_rc_t ompd_get_device_thread_id_kinds(  
    ompd_address_space_handle_t *device_handle,  
    ompd_thread_id_t **kinds,  
    ompd_size_t **thread_id_sizes,  
    int *count  
);
```

Semantics

The `ompd_get_device_thread_id_kinds` function returns an array of supported native thread identifier kinds and a corresponding array of their respective sizes for a given device. The OMPD library allocates storage for the arrays with the memory allocation callback that the tool provides. Each supported native thread identifier kind is guaranteed to be recognizable by the OMPD library and may be mapped to and from any OpenMP thread that executes on the device. The third-party tool owns the storage for the array of kinds and the array of sizes that is returned via the *kinds* and *thread_id_sizes* arguments, and it is responsible for freeing that storage.

Description of Arguments

The *device_handle* argument is a pointer to an opaque address space handle that represents a host device (returned by `ompd_process_initialize`) or a non-host device (returned by `ompd_device_initialize`). On return, the *kinds* argument is the address of a pointer to an array of native thread identifier kinds, the *thread_id_sizes* argument is the address of a pointer to an array of the corresponding native thread identifier sizes used by the OMPD library, and the *count* argument is the address of a variable that indicates the sizes of the returned arrays.

Description of Return Codes

This routine must return any of the general return codes listed at the beginning of [Section 20.5](#).

Cross References

- Native Thread Identifiers, see [Section 20.3.7](#)
- OMPD Handle Types, see [Section 20.3.8](#)
- Return Code Types, see [Section 20.3.13](#)
- Size Type, see [Section 20.3.1](#)

20.5.3 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.

20.5.4 Address Space Information

20.5.4.1 `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  
);
```

Semantics

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.

Description of Return Codes

This routine must return any of the general return codes listed at the beginning of [Section 20.5](#).

Cross References

- OMPD Handle Types, see [Section 20.3.8](#)
- Return Code Types, see [Section 20.3.13](#)

20.5.4.2 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  
);
```

Semantics

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.

Description of Return Codes

This routine must return any of the general return codes listed at the beginning of [Section 20.5](#).

Cross References

- OMPD Handle Types, see [Section 20.3.8](#)
- Return Code Types, see [Section 20.3.13](#)

20.5.5 Thread Handles

20.5.5.1 `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  
);
```

Semantics

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 team that is executing 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 represents the OpenMP thread number and selects the thread, the handle for which is to be returned. On return, the *thread_handle* argument is an opaque handle for the selected thread.

Description of Return Codes

This routine must return any of the general return codes listed at the beginning of [Section 20.5](#) or the following return code:

- `ompd_rc_bad_input` if the `thread_num` argument is greater than or equal to the `team-size-var` ICV or negative.

Restrictions

Restrictions on the `ompd_get_thread_in_parallel` function are as follows:

- The value of `thread_num` must be a non-negative integer smaller than the team size that was provided as the `team-size-var` ICV from `ompd_get_icv_from_scope`.

Cross References

- OMPD Handle Types, see [Section 20.3.8](#)
- Return Code Types, see [Section 20.3.13](#)
- `ompd_get_icv_from_scope`, see [Section 20.5.10.2](#)

20.5.5.2 `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  
);
```

Semantics

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.

Description of Return Codes

This routine must return any of the general return codes listed at the beginning of [Section 20.5](#) or any of the following return codes:

- `ompd_rc_bad_input` if a different value in `sizeof_thread_id` is expected for a thread kind of `kind`.
- `ompd_rc_unsupported` if the `kind` of thread is not supported.
- `ompd_rc_unavailable` if the thread is not an OpenMP thread.

Cross References

- Native Thread Identifiers, see [Section 20.3.7](#)
- OMPD Handle Types, see [Section 20.3.8](#)
- Return Code Types, see [Section 20.3.13](#)
- Size Type, see [Section 20.3.1](#)

20.5.5.3 `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  
);
```

Semantics

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.

Description of Return Codes

This routine must return any of the general return codes listed at the beginning of [Section 20.5](#).

Cross References

- OMPD Handle Types, see [Section 20.3.8](#)
- Return Code Types, see [Section 20.3.13](#)

20.5.5.4 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  
);
```

Semantics

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.

Description of Return Codes

This routine must return any of the general return codes listed at the beginning of [Section 20.5](#).

Cross References

- OMPD Handle Types, see [Section 20.3.8](#)
- Return Code Types, see [Section 20.3.13](#)

20.5.5.5 ompd_get_thread_id

Summary

The `ompd_get_thread_id` function 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  
);
```

Semantics

The `ompd_get_thread_id` function maps an OMPD thread handle to a native thread identifier. This call yields meaningful results only if the referenced OpenMP thread is stopped.

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.

Description of Return Codes

This routine must return any of the general return codes listed at the beginning of [Section 20.5](#) or any of the following return codes:

- `ompd_rc_bad_input` if a different value in `sizeof_thread_id` is expected for a thread kind of `kind`; or
- `ompd_rc_unsupported` if the `kind` of thread is not supported.

Cross References

- Native Thread Identifiers, see [Section 20.3.7](#)
- OMPD Handle Types, see [Section 20.3.8](#)
- Return Code Types, see [Section 20.3.13](#)
- Size Type, see [Section 20.3.1](#)

20.5.5.6 ompd_get_device_from_thread

Summary

The `ompd_get_device_from_thread` function obtains a pointer to the address space handle for a device on which an OpenMP thread is executing.

Format

```
ompd_rc_t ompd_get_device_from_thread(  
    ompd_thread_handle_t *thread_handle,  
    ompd_address_space_handle_t **device  
);
```

Semantics

The `ompd_get_device_from_thread` function obtains a pointer to the address space handle for a device on which an OpenMP thread is executing. The returned pointer will be the same as the address space handle pointer that was previously returned by a call to `ompd_process_initialize` (for a host device) or a call to `ompd_device_initialize` (for a non-host device). This call yields meaningful results only if the referenced OpenMP thread is stopped.

Description of Arguments

The `thread_handle` argument is a pointer to an opaque thread handle that represents an OpenMP thread. On return, the `device` argument is the address of a pointer to an OMPD address space handle.

Description of Return Codes

This routine must return any of the general return codes listed at the beginning of [Section 20.5](#).

Cross References

- OMPD Handle Types, see [Section 20.3.8](#)
- Return Code Types, see [Section 20.3.13](#)

20.5.6 Parallel Region Handles

20.5.6.1 `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 innermost parallel region.

Format

```
ompd_rc_t ompd_get_curr_parallel_handle(  
    ompd_thread_handle_t *thread_handle,  
    ompd_parallel_handle_t **parallel_handle  
);
```

Semantics

The `ompd_get_curr_parallel_handle` function enables the tool to obtain a pointer to the parallel handle for the innermost parallel region that is associated with an OpenMP thread. This call yields meaningful results only if the referenced OpenMP thread is stopped. The parallel handle is owned by the tool and it 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.

Description of Return Codes

This routine must return any of the general return codes listed at the beginning of [Section 20.5](#) or the following return code:

- `ompd_rc_unavailable` if the thread is not currently part of a team.

Cross References

- OMPD Handle Types, see [Section 20.3.8](#)
- Return Code Types, see [Section 20.3.13](#)
- `ompd_rel_parallel_handle`, see [Section 20.5.6.4](#)

20.5.6.2 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  
);
```

Semantics

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 team that is executing 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.

Description of Return Codes

This routine must return any of the general return codes listed at the beginning of [Section 20.5](#) or the following return code:

- `ompd_rc_unavailable` if no enclosing parallel region exists.

Cross References

- OMPD Handle Types, see [Section 20.3.8](#)
- Return Code Types, see [Section 20.3.13](#)
- `ompd_rel_parallel_handle`, see [Section 20.5.6.4](#)

20.5.6.3 `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  
);
```

Semantics

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 yields meaningful results only if at least one thread in the team that is executing the parallel region is stopped. A pointer to the parallel 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.

Description of Return Codes

This routine must return any of the general return codes listed at the beginning of [Section 20.5](#).

Cross References

- OMPD Handle Types, see [Section 20.3.8](#)
- Return Code Types, see [Section 20.3.13](#)
- `ompd_rel_parallel_handle`, see [Section 20.5.6.4](#)

20.5.6.4 `ompd_rel_parallel_handle`

Summary

The `ompd_rel_parallel_handle` function releases a parallel handle.

Format

```
ompd_rc_t ompd_rel_parallel_handle(  
    ompd_parallel_handle_t *parallel_handle  
);
```

Semantics

Parallel handles are opaque so tools cannot release them directly. Instead, a tool must pass a parallel 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.

Description of Return Codes

This routine must return any of the general return codes listed at the beginning of [Section 20.5](#).

Cross References

- OMPD Handle Types, see [Section 20.3.8](#)
- Return Code Types, see [Section 20.3.13](#)

20.5.6.5 `ompd_parallel_handle_compare`

Summary

The `ompd_parallel_handle_compare` function compares two parallel 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  
);
```

Semantics

The internal structure of parallel handles is opaque to tools. While tools can easily compare pointers to parallel 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 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.

Description of Return Codes

This routine must return any of the general return codes listed at the beginning of [Section 20.5](#).

Cross References

- OMPD Handle Types, see [Section 20.3.8](#)
- Return Code Types, see [Section 20.3.13](#)

20.5.7 Task Handles

20.5.7.1 `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  
);
```

Semantics

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 yields meaningful results 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.

Description of Return Codes

This routine must return any of the general return codes listed at the beginning of [Section 20.5](#) or the following return code:

- `ompd_rc_unavailable` if the thread is currently not executing a task.

Cross References

- OMPD Handle Types, see [Section 20.3.8](#)
- Return Code Types, see [Section 20.3.13](#)
- `ompd_rel_task_handle`, see [Section 20.5.7.5](#)

20.5.7.2 `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  
);
```

Semantics

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 yields meaningful results only if the thread that is executing the task that `task_handle` specifies is stopped while executing the task. 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.

Description of Return Codes

This routine must return any of the general return codes listed at the beginning of [Section 20.5](#) or the following return code:

- `ompd_rc_unavailable` if no generating task region exists.

Cross References

- OMPD Handle Types, see [Section 20.3.8](#)
- Return Code Types, see [Section 20.3.13](#)
- `ompd_rel_task_handle`, see [Section 20.5.7.5](#)

20.5.7.3 `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  
);
```

Semantics

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. An implicit task does not have a scheduling task. This call yields meaningful results only if the thread that is executing the task that *task_handle* specifies is stopped while executing the task. 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.

Description of Return Codes

This routine must return any of the general return codes listed at the beginning of [Section 20.5](#) or the following return code:

- `ompd_rc_unavailable` if no scheduling task exists.

Cross References

- OMPD Handle Types, see [Section 20.3.8](#)
- Return Code Types, see [Section 20.3.13](#)
- `ompd_rel_task_handle`, see [Section 20.5.7.5](#)

20.5.7.4 `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  
);
```

Semantics

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 to be returned. The `thread_num` argument is equal to the `thread-num-var` ICV value of the selected implicit task. On return, the `task_handle` argument points to a location that points to an opaque handle for the selected implicit task.

Description of Return Codes

This routine must return any of the general return codes listed at the beginning of [Section 20.5](#) or the following return code:

- `ompd_rc_bad_input` if the `thread_num` argument is greater than or equal to the `team-size-var` ICV or negative.

Restrictions

Restrictions on the `ompd_get_task_in_parallel` function are as follows:

- 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 `team-size-var` ICV that `ompd_get_icv_from_scope` returns.

Cross References

- OMPD Handle Types, see [Section 20.3.8](#)
- Return Code Types, see [Section 20.3.13](#)
- `ompd_get_icv_from_scope`, see [Section 20.5.10.2](#)

20.5.7.5 `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  
);
```

Semantics

Task handles are opaque to tools; thus 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.

Description of Return Codes

This routine must return any of the general return codes listed at the beginning of [Section 20.5](#).

Cross References

- OMPD Handle Types, see [Section 20.3.8](#)
- Return Code Types, see [Section 20.3.13](#)

20.5.7.6 `ompd_task_handle_compare`

Summary

The `ompd_task_handle_compare` function compares task handles.

1
2
3
4
5
6

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
);

```

7
8
9
10
11
12
13
14

Semantics

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.

15
16
17
18

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.

19
20

Description of Return Codes

This routine must return any of the general return codes listed at the beginning of [Section 20.5](#).

21
22
23

Cross References

- OMPD Handle Types, see [Section 20.3.8](#)
- Return Code Types, see [Section 20.3.13](#)

24

20.5.7.7 ompd_get_task_function

25
26
27

Summary

This `ompd_get_task_function` function returns the entry point of the code that corresponds to the body of a task.

28

Format

```

ompd_rc_t ompd_get_task_function (
    ompd_task_handle_t *task_handle,
    ompd_address_t *entry_point
);

```

29
30
31
32

Semantics

The `ompd_get_task_function` function returns the entry point of the code that corresponds to the body of code that the task executes. This call is meaningful only if the thread that is executing the task that `task_handle` specifies is stopped while executing the task.

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.

Description of Return Codes

This routine must return any of the general return codes listed at the beginning of [Section 20.5](#).

Cross References

- Address Type, see [Section 20.3.4](#)
- OMPD Handle Types, see [Section 20.3.8](#)
- Return Code Types, see [Section 20.3.13](#)

20.5.7.8 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  
);
```

Semantics

An OpenMP implementation maintains an `ompt_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 `ompt_frame_t` object of the task that `task_handle` identifies. This call yields meaningful results only if the thread that is executing the task that `task_handle` specifies is stopped while executing the task.

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 `ompt_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 `ompt_frame_t` object that is associated with the specified task.

Description of Return Codes

This routine must return any of the general return codes listed at the beginning of [Section 20.5](#).

Cross References

- Address Type, see [Section 20.3.4](#)
- Frame Information Type, see [Section 20.3.5](#)
- OMPD Handle Types, see [Section 20.3.8](#)
- Return Code Types, see [Section 20.3.13](#)
- `ompt_frame_t`, see [Section 19.4.4.29](#)

20.5.8 Querying Thread States

20.5.8.1 `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  
);
```

Semantics

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.

Description of Return Codes

This routine must return any of the general return codes listed at the beginning of [Section 20.5](#) or the following return code:

- `ompd_rc_bad_input` if an unknown value is provided in `current_state`.

Cross References

- OMPD Handle Types, see [Section 20.3.8](#)
- Return Code Types, see [Section 20.3.13](#)
- `ompt_state_t`, see [Section 19.4.4.28](#)

20.5.8.2 `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  
);
```

Semantics

The `ompd_get_state` function returns the state of an OpenMP thread. This call yields meaningful results only if the referenced OpenMP thread is stopped.

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.

Description of Return Codes

This routine must return any of the general return codes listed at the beginning of [Section 20.5](#).

Cross References

- OMPD Handle Types, see [Section 20.3.8](#)
- Return Code Types, see [Section 20.3.13](#)
- Wait ID Type, see [Section 20.3.2](#)
- `ompd_enumerate_states`, see [Section 20.5.8.1](#)

20.5.9 Display Control Variables

20.5.9.1 `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  
);
```

Semantics

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 "**icv-name=icv-value**".

On return, the third-party tool owns the vector and the strings. The OMPD 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 the `ompd_rel_display_control_vars` function 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.

Description of Return Codes

This routine must return any of the general return codes listed at the beginning of [Section 20.5](#).

Cross References

- OMPD Handle Types, see [Section 20.3.8](#)
- Return Code Types, see [Section 20.3.13](#)
- `ompd_initialize`, see [Section 20.5.1.1](#)
- `ompd_rel_display_control_vars`, see [Section 20.5.9.2](#)

20.5.9.2 `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  
);
```

Semantics

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.

Description of Return Codes

This routine must return any of the general return codes listed at the beginning of [Section 20.5](#).

Cross References

- Return Code Types, see [Section 20.3.13](#)
- `ompd_get_display_control_vars`, see [Section 20.5.9.1](#)

20.5.10 Accessing Scope-Specific Information

20.5.10.1 `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  
);
```

Semantics

An OpenMP implementation must support all ICVs listed in [Section 2.1](#). An OpenMP implementation may support additional implementation-specific variables. An implementation may store ICVs in a different scope than [Table 2.1](#) indicates. The `ompd_enumerate_icvs` function enables a tool to enumerate the ICVs that an OpenMP implementation supports and their related scopes.

1 When the *current* argument is set to the identifier of a supported ICV, `ompd_enumerate_icvs`
2 assigns the value, string name, and scope of the next ICV in the enumeration to the locations to
3 which the *next_id*, *next_icv_name*, and *next_scope* arguments point. On return, the third-party tool
4 owns the *next_icv_name* string. The OMPD library uses the memory allocation callback that the
5 tool provides to allocate the string storage; the tool is responsible for releasing the memory.

6 On return, the location to which the *more* argument points has the value of 1 whenever one or more
7 ICV are left in the enumeration. On return, that location has the value 0 when *current* is the last
8 ICV in the enumeration.

9 **Description of Arguments**

10 The *address_space_handle* argument identifies the address space. The *current* argument must be
11 an ICV that the OpenMP implementation supports. To begin enumerating the ICVs, a tool should
12 pass `ompd_icv_undefined` as the value of *current*. Subsequent calls to
13 `ompd_enumerate_icvs` should pass the value returned by the call in the *next_id* output
14 argument. On return, the *next_id* argument points to an integer with the value of the ID of the next
15 ICV in the enumeration. On return, the *next_icv_name* argument points to a character string with
16 the name of the next ICV. On return, the *next_scope* argument points to the scope enum value of the
17 scope of the next ICV. On return, the *more_enums* argument points to an integer with the value of 1
18 when more ICVs are left to enumerate and the value of 0 when no more ICVs are left.

19 **Description of Return Codes**

20 This routine must return any of the general return codes listed at the beginning of [Section 20.5](#) or
21 the following return code:

- 22 • `ompd_rc_bad_input` if an unknown value is provided in *current*.

23 **Cross References**

- 24 • ICV ID Type, see [Section 20.3.11](#)
- 25 • OMPD Handle Types, see [Section 20.3.8](#)
- 26 • OMPD Scope Types, see [Section 20.3.9](#)
- 27 • Return Code Types, see [Section 20.3.13](#)

28 **20.5.10.2 ompd_get_icv_from_scope**

29 **Summary**

30 The `ompd_get_icv_from_scope` function returns the value of an ICV.

31 **Format**

```
32 ompd_rc_t ompd_get_icv_from_scope (  
33     void *handle,  
34     ompd_scope_t scope,  
35     ompd_icv_id_t icv_id,  
36     ompd_word_t *icv_value  
37 );
```

Semantics

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

The provided `handle` must match the `scope` as defined in [Section 20.3.11](#).

The provided `scope` must match the scope for `icv_id` as requested by `ompd_enumerate_icvs`.

Description of Return Codes

This routine must return any of the general return codes listed at the beginning of [Section 20.5](#) or any of the following return codes:

- `ompd_rc_incompatible_handle` if the scope of the handle does not match the constraint;
- `ompd_rc_incompatible` if the ICV cannot be represented as an integer;
- `ompd_rc_incomplete` if only the first item of the ICV is returned in the integer (e.g., if `nthreads-var` is a list); or
- `ompd_rc_bad_input` if an unknown value is provided in `icv_id`.

Cross References

- ICV ID Type, see [Section 20.3.11](#)
- OMPD Handle Types, see [Section 20.3.8](#)
- OMPD Scope Types, see [Section 20.3.9](#)
- Return Code Types, see [Section 20.3.13](#)
- `ompd_enumerate_icvs`, see [Section 20.5.10.1](#)

20.5.10.3 `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  
);
```


Semantics

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 20.3.11](#).

The provided *scope* must match the scope for *icv_id* as requested by `ompd_enumerate_icvs`.

Description of Return Codes

This routine must return any of the general return codes listed at the beginning of [Section 20.5](#) or the following return code:

- `ompd_rc_incompatible_handle` if the scope of the handle does not match the constraint;
- `ompd_rc_bad_input` if an unknown value is provided in *icv_id*.

Cross References

- ICV ID Type, see [Section 20.3.11](#)
- OMPD Handle Types, see [Section 20.3.8](#)
- OMPD Scope Types, see [Section 20.3.9](#)
- Return Code Types, see [Section 20.3.13](#)
- `ompd_enumerate_icvs`, see [Section 20.5.10.1](#)

20.5.10.4 `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  
);
```

Semantics

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.

Description of Return Codes

This routine must return any of the general return codes listed at the beginning of [Section 20.5](#) or the following return code:

- `ompd_rc_unsupported` if the runtime library does not support OMPT.

Cross References

- OMPD Handle Types, see [Section 20.3.8](#)
- OMPD Scope Types, see [Section 20.3.9](#)
- Return Code Types, see [Section 20.3.13](#)
- `ompt_data_t`, see [Section 19.4.4.4](#)

20.6 Breakpoint Symbol Names 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.

20.6.1 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);
```

Semantics

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` directive, see [Section 10.2](#)
- `ompd_get_curr_parallel_handle`, see [Section 20.5.6.1](#)
- `ompd_get_curr_task_handle`, see [Section 20.5.7.1](#)

20.6.2 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);
```

Semantics

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` directive, see [Section 10.2](#)
- `ompd_get_curr_parallel_handle`, see [Section 20.5.6.1](#)
- `ompd_get_curr_task_handle`, see [Section 20.5.7.1](#)
- `ompd_rel_parallel_handle`, see [Section 20.5.6.4](#)

20.6.3 Beginning Teams Regions

Summary

Before starting the execution of an OpenMP **teams** region, the implementation executes `ompd_bp_teams_begin`.

Format

```
void ompd_bp_teams_begin(void);
```

Semantics

The OpenMP implementation must execute `ompd_bp_teams_begin` at every *teams-begin* event. At the point that the implementation reaches `ompd_bp_teams_begin`, the binding for `ompd_get_curr_parallel_handle` is the **teams** region that is beginning and the binding for `ompd_get_curr_task_handle` is the task that encountered the **teams** construct.

Cross References

- **teams** directive, see [Section 10.3](#)
- `ompd_get_curr_parallel_handle`, see [Section 20.5.6.1](#)
- `ompd_get_curr_task_handle`, see [Section 20.5.7.1](#)

20.6.4 Ending Teams Regions

Summary

After finishing the execution of an OpenMP **teams** region, the implementation executes `ompd_bp_teams_end`.

Format

```
void ompd_bp_teams_end(void);
```

Semantics

The OpenMP implementation must execute `ompd_bp_teams_end` at every *teams-end* event. At the point that the implementation reaches `ompd_bp_teams_end`, the binding for `ompd_get_curr_parallel_handle` is the **teams** region that is ending and the binding for `ompd_get_curr_task_handle` is the task that encountered the **teams** construct. After execution of `ompd_bp_teams_end`, any *parallel_handle* that was acquired for the **teams** region is invalid and should be released.

Cross References

- `teams` directive, see [Section 10.3](#)
- `ompd_get_curr_parallel_handle`, see [Section 20.5.6.1](#)
- `ompd_get_curr_task_handle`, see [Section 20.5.7.1](#)
- `ompd_rel_parallel_handle`, see [Section 20.5.6.4](#)

20.6.5 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);
```

Semantics

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 20.5.7.1](#)

20.6.6 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);
```

Semantics

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 20.5.7.1](#)
- `ompd_rel_task_handle`, see [Section 20.5.7.5](#)

20.6.7 Beginning OpenMP Threads

Summary

When starting an OpenMP thread, the implementation executes `ompd_bp_thread_begin`.

Format

```
void ompd_bp_thread_begin(void);
```

Semantics

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` directive, see [Section 10.2](#)
- Initial Task, see [Section 12.8](#)

20.6.8 Ending OpenMP Threads

Summary

When terminating an OpenMP thread, the implementation executes `ompd_bp_thread_end`.

Format

```
void ompd_bp_thread_end(void);
```

Semantics

The OpenMP implementation must execute `ompd_bp_thread_end` at every *native-thread-end* and *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` directive, see [Section 10.2](#)
- Initial Task, see [Section 12.8](#)
- `ompd_rel_thread_handle`, see [Section 20.5.5.3](#)

20.6.9 Beginning Target Regions

Summary

Before starting the execution of an OpenMP `target` region, the implementation executes `ompd_bp_target_begin`.

Format

```
void ompd_bp_target_begin(void);
```

Semantics

The OpenMP implementation must execute `ompd_bp_target_begin` at every *initial-task-begin* event that results from the execution of an initial task enclosing a `target` region. At the point that the implementation reaches `ompd_bp_target_begin`, the binding for `ompd_get_curr_parallel_handle` is the `target` region that is beginning and the binding for `ompd_get_curr_task_handle` is the initial task on the device.

Cross References

- `target` directive, see [Section 13.8](#)
- `ompd_get_curr_parallel_handle`, see [Section 20.5.6.1](#)
- `ompd_get_curr_task_handle`, see [Section 20.5.7.1](#)

20.6.10 Ending Target Regions

Summary

After finishing the execution of an OpenMP `target` region, the implementation executes `ompd_bp_target_end`.

Format

```
void ompd_bp_target_end(void);
```

Semantics

The OpenMP implementation must execute `ompd_bp_target_end` at every *initial-task-end* event that results from the execution of an initial task enclosing a **target** region. At the point that the implementation reaches `ompd_bp_target_end`, the binding for `ompd_get_curr_parallel_handle` is the **target** region that is ending and the binding for `ompd_get_curr_task_handle` is the initial task on the device. After execution of `ompd_bp_target_end`, any *parallel_handle* that was acquired for the **target** region is invalid and should be released.

Cross References

- **target** directive, see [Section 13.8](#)
- `ompd_get_curr_parallel_handle`, see [Section 20.5.6.1](#)
- `ompd_get_curr_task_handle`, see [Section 20.5.7.1](#)
- `ompd_rel_parallel_handle`, see [Section 20.5.6.4](#)

20.6.11 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);
```

Semantics

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 13.4](#)

20.6.12 Finalizing OpenMP Devices

Summary

When terminating an OpenMP thread, the implementation executes `ompd_bp_device_end`.

Format

```
void ompd_bp_device_end(void);
```


1 **Semantics**
2 The OpenMP implementation must execute **ompd_bp_device_end** at every *device-finalize*
3 event. This execution occurs after the thread executes all OpenMP regions. After execution of
4 **ompd_bp_device_end**, any *address_space_handle* that was acquired for this device is invalid
5 and should be released.

6 **Cross References**

- 7
 - Device Initialization, see [Section 13.4](#)

8
 - **ompd_rel_address_space_handle**, see [Section 20.5.2.3](#)

21 Environment Variables

This chapter describes the OpenMP environment variables that specify the settings of the ICVs that affect the execution of OpenMP programs (see [Chapter 2](#)). The names of the environment variables must be upper case. Unless otherwise specified, 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
```

As defined following [Table 2.2](#) in [Section 2.2](#), device-specific environment variables extend many of the environment variables defined in this chapter. If the corresponding environment variable for a specific device number is set, then the setting for that environment variable is used to set the value of the associated ICV of the device with the corresponding device number. If the corresponding environment variable that includes the `_DEV` suffix but no device number is set, then the setting of that environment variable is used to set the value of the associated ICV of any non-host device for which the device-number-specific corresponding environment variable is not set. The corresponding environment variable without a suffix sets the associated ICV of the host device. If the corresponding environment variable includes the `_ALL` suffix, the setting of that environment variable is used to set the value of the associated ICV of any host or non-host device for which corresponding environment variables that are device-number specific, have the `_DEV` suffix, or have no suffix are not set.

Restrictions

Restrictions to device-specific environment variables are as follows:

- Device-specific environment variables must not correspond to environment variables that initialize ICVs with global scope.
- Device-specific environment variables must not specify the initial device.

21.1 Parallel Region Environment Variables

This section defines environment variables that affect the operation of `parallel` regions.

21.1.1 OMP_DYNAMIC

The `OMP_DYNAMIC` environment variable controls dynamic adjustment of the number of threads to use for executing `parallel` regions by setting the initial value of the *dyn-var* ICV.

The value of this environment variable must be one of the following:

`true` | `false`

If the environment variable is set to `true`, the OpenMP implementation may adjust the number of threads to use for executing `parallel` regions in order to optimize the use of system resources. If the environment variable is set to `false`, the dynamic adjustment of the number of threads is disabled. The behavior of the program is implementation defined if the value of `OMP_DYNAMIC` is neither `true` nor `false`.

Example:

```
setenv OMP_DYNAMIC true
```

Cross References

- `parallel` directive, see [Section 10.2](#)
- *dyn-var* ICV, see [Table 2.1](#)
- `omp_get_dynamic`, see [Section 18.2.7](#)
- `omp_set_dynamic`, see [Section 18.2.6](#)

21.1.2 OMP_NUM_THREADS

The `OMP_NUM_THREADS` environment variable sets the number of threads to use for `parallel` regions by setting the initial value of the *nthreads-var* ICV. See [Chapter 2](#) for a comprehensive set of rules about the interaction between the `OMP_NUM_THREADS` environment variable, the `num_threads` clause, the `omp_set_num_threads` library routine and dynamic adjustment of threads, and [Section 10.2.1](#) for a complete algorithm that describes how the number of threads for a `parallel` region is determined.

1 The value of this environment variable must be a list of positive integer values. The values of the
2 list set the number of threads to use for **parallel** regions at the corresponding nested levels.

3 The behavior of the program is implementation defined if any value of the list specified in the
4 **OMP_NUM_THREADS** environment variable leads to a number of threads that is greater than an
5 implementation can support, or if any value is not a positive integer.

6 The **OMP_NUM_THREADS** environment variable sets the *max-active-levels-var* ICV to the number
7 of active levels of parallelism that the implementation supports if the **OMP_NUM_THREADS**
8 environment variable is set to a comma-separated list of more than one value. The value of the
9 *max-active-level-var* ICV may be overridden by setting **OMP_MAX_ACTIVE_LEVELS**. See
10 [Section 21.1.4](#) for details.

11 Example:

```
12 setenv OMP_NUM_THREADS 4,3,2
```

13 Cross References

- 14 • **OMP_MAX_ACTIVE_LEVELS**, see [Section 21.1.4](#)
- 15 • **num_threads** clause, see [Section 10.2.2](#)
- 16 • **parallel** directive, see [Section 10.2](#)
- 17 • *nthreads-var* ICV, see [Table 2.1](#)
- 18 • **omp_set_num_threads**, see [Section 18.2.1](#)

19 21.1.3 OMP_THREAD_LIMIT

20 The **OMP_THREAD_LIMIT** environment variable sets the maximum number of OpenMP threads
21 to use for a contention group by setting the *thread-limit-var* ICV. The value of this environment
22 variable must be a positive integer. The behavior of the program is implementation defined if the
23 requested value of **OMP_THREAD_LIMIT** is greater than the number of threads an implementation
24 can support, or if the value is not a positive integer.

25 Cross References

- 26 • *thread-limit-var* ICV, see [Table 2.1](#)

27 21.1.4 OMP_MAX_ACTIVE_LEVELS

28 The **OMP_MAX_ACTIVE_LEVELS** environment variable controls the maximum number of nested
29 active **parallel** regions by setting the initial value of the *max-active-levels-var* ICV. The value
30 of this environment variable must be a non-negative integer. The behavior of the program is
31 implementation defined if the requested value of **OMP_MAX_ACTIVE_LEVELS** is greater than the
32 maximum number of nested active parallel levels an implementation can support, or if the value is
33 not a non-negative integer.

Cross References

- *max-active-levels-var* ICV, see [Table 2.1](#)

21.1.5 OMP_PLACES

The **OMP_PLACES** environment variable sets the initial value of the *place-partition-var* ICV. A list of places can be specified in the **OMP_PLACES** environment variable. 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.

The **OMP_PLACES** environment variable can be defined using an explicit ordered list of comma-separated places. A place is defined by an unordered set of comma-separated non-negative numbers enclosed by braces, or a non-negative number. The meaning of the numbers and how the numbering is done are implementation defined. Generally, the numbers represent the smallest unit of execution exposed by the execution environment, typically a hardware thread.

Intervals may also be used to define places. Intervals can be specified using the *<lower-bound> : <length> : <stride>* notation to represent the following list of numbers: “*<lower-bound>*, *<lower-bound> + <stride>*, ..., *<lower-bound> + (<length> - 1)*<stride>*.” When *<stride>* is omitted, a unit stride is assumed. Intervals can specify numbers within a place as well as sequences of places.

An exclusion operator “!” can also be used to exclude the number or place immediately following the operator.

Alternatively, the abstract names listed in [Table 21.1](#) should be understood by the execution and runtime environment. The entities defined by the abstract names are implementation defined. An implementation may also add abstract names as appropriate for the target platform.

The abstract name may be appended with one or two positive numbers in parentheses, that is, *abstract_name(<num-places>)* or *abstract_name(<num-places> : <stride>)*, where *<num-places>* denotes the length of the place list and *<stride>* denotes the increment between consecutive places in the place list. When requesting fewer places than available on the system, the determination of which resources of type *abstract_name* are to be included in the place list is implementation defined. When requesting more resources than available, the length of the place list is implementation defined.

TABLE 21.1: Predefined Abstract Names for **OMP_PLACES**

Abstract Name	Meaning
threads	Each place corresponds to a single hardware thread on the device.
cores	Each place corresponds to a single core (having one or more hardware threads) on the device.
ll_caches	Each place corresponds to a set of cores that share the last level cache on the device.
numa_domains	Each place corresponds to a set of cores for which their closest memory on the device is: <ul style="list-style-type: none"> • the same memory; and • at a similar distance from the cores.
sockets	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.

$$\begin{aligned} \langle \text{list} \rangle & \models \langle \text{p-list} \rangle \mid \langle \text{aname} \rangle \\ \langle \text{p-list} \rangle & \models \langle \text{p-interval} \rangle \mid \langle \text{p-list} \rangle, \langle \text{p-interval} \rangle \\ \langle \text{p-interval} \rangle & \models \langle \text{place} \rangle : \langle \text{len} \rangle : \langle \text{stride} \rangle \mid \langle \text{place} \rangle : \langle \text{len} \rangle \mid \langle \text{place} \rangle \mid ! \langle \text{place} \rangle \\ \langle \text{place} \rangle & \models \{ \langle \text{res-list} \rangle \} \mid \langle \text{res} \rangle \\ \langle \text{res-list} \rangle & \models \langle \text{res-interval} \rangle \mid \langle \text{res-list} \rangle, \langle \text{res-interval} \rangle \\ \langle \text{res-interval} \rangle & \models \langle \text{res} \rangle : \langle \text{num-places} \rangle : \langle \text{stride} \rangle \mid \langle \text{res} \rangle : \langle \text{num-places} \rangle \mid \langle \text{res} \rangle \mid ! \langle \text{res} \rangle \\ \langle \text{aname} \rangle & \models \langle \text{word} \rangle (\langle \text{num-places} \rangle : \langle \text{stride} \rangle) \mid \langle \text{word} \rangle (\langle \text{num-places} \rangle) \mid \langle \text{word} \rangle \\ \langle \text{word} \rangle & \models \text{sockets} \mid \text{cores} \mid \text{ll_caches} \mid \text{numa_domains} \\ & \quad \mid \text{threads} \mid \langle \text{implementation-defined abstract name} \rangle \\ \langle \text{res} \rangle & \models \textit{non-negative integer} \\ \langle \text{num-places} \rangle & \models \textit{positive integer} \\ \langle \text{stride} \rangle & \models \textit{integer} \\ \langle \text{len} \rangle & \models \textit{positive integer} \end{aligned}$$

1 Examples:

```
2 setenv OMP_PLACES threads
3 setenv OMP_PLACES "threads (4) "
4 setenv OMP_PLACES "threads (8:2) "
5 setenv OMP_PLACES
6     "{0,1,2,3},{4,5,6,7},{8,9,10,11},{12,13,14,15}"
7 setenv OMP_PLACES "{0:4},{4:4},{8:4},{12:4}"
8 setenv OMP_PLACES "{0:4}:4:4"
```

9 where each of the last three definitions corresponds to the same 4 places including the smallest
10 units of execution exposed by the execution environment numbered, in turn, 0 to 3, 4 to 7, 8 to 11,
11 and 12 to 15.

12 Cross References

- 13 • *place-partition-var* ICV, see [Table 2.1](#)

14 21.1.6 OMP_PROC_BIND

15 The `OMP_PROC_BIND` environment variable sets the initial value of the *bind-var* ICV. The value
16 of this environment variable is either **true**, **false**, or a comma separated list of **primary**,
17 **close**, or **spread**. The values of the list set the thread affinity policy to be used for parallel
18 regions at the corresponding nested level.

19 If the environment variable is set to **false**, the execution environment may move OpenMP threads
20 between OpenMP places, thread affinity is disabled, and **proc_bind** clauses on **parallel**
21 constructs are ignored.

22 Otherwise, the execution environment should not move OpenMP threads between OpenMP places,
23 thread affinity is enabled, and the initial thread is bound to the first place in the *place-partition-var*
24 ICV prior to the first active parallel region. An initial thread that is created by a **teams** construct is
25 bound to the first place in its *place-partition-var* ICV before it begins execution of the associated
26 structured block.

27 If the environment variable is set to **true**, the thread affinity policy is implementation defined but
28 must conform to the previous paragraph. The behavior of the program is implementation defined if
29 the value in the `OMP_PROC_BIND` environment variable is not **true**, **false**, or a comma
30 separated list of **primary**, **close**, or **spread**. The behavior is also implementation defined if
31 an initial thread cannot be bound to the first place in the *place-partition-var* ICV.

32 The `OMP_PROC_BIND` environment variable sets the *max-active-levels-var* ICV to the number of
33 active levels of parallelism that the implementation supports if the `OMP_PROC_BIND` environment
34 variable is set to a comma-separated list of more than one element. The value of the
35 *max-active-level-var* ICV may be overridden by setting `OMP_MAX_ACTIVE_LEVELS`. See
36 [Section 21.1.4](#) for details.

1 Examples:

```
2 setenv OMP_PROC_BIND false  
3 setenv OMP_PROC_BIND "spread, spread, close"
```

4 Cross References

- 5 • `OMP_MAX_ACTIVE_LEVELS`, see [Section 21.1.4](#)
- 6 • `proc_bind` clause, see [Section 10.2.4](#)
- 7 • `parallel` directive, see [Section 10.2](#)
- 8 • `teams` directive, see [Section 10.3](#)
- 9 • Controlling OpenMP Thread Affinity, see [Section 10.2.3](#)
- 10 • `bind-var` ICV, see [Table 2.1](#)
- 11 • `max-active-levels-var` ICV, see [Table 2.1](#)
- 12 • `place-partition-var` ICV, see [Table 2.1](#)
- 13 • `omp_get_proc_bind`, see [Section 18.3.1](#)

14 21.2 Program Execution Environment Variables

15 This section defines environment variables that affect program execution.

16 21.2.1 OMP_SCHEDULE

17 The `OMP_SCHEDULE` environment variable controls the schedule kind and chunk size of all
18 worksharing-loop directives that have the schedule kind `runtime`, by setting the value of the
19 `run-sched-var` ICV. The value of this environment variable takes the form `[modifier:]kind[, chunk]`,
20 where:

- 21 • `modifier` is one of `monotonic` or `nonmonotonic`;
- 22 • `kind` is one of `static`, `dynamic`, `guided`, or `auto`;
- 23 • `chunk` is an optional positive integer that specifies the chunk size.

24 If the `modifier` is not present, the `modifier` is set to `monotonic` if `kind` is `static`; for any other
25 `kind` it is set to `nonmonotonic`.

26 If `chunk` is present, white space may be on either side of the “,”. See [Section 11.5.3](#) for a detailed
27 description of the schedule kinds.

28 The behavior of the program is implementation defined if the value of `OMP_SCHEDULE` does not
29 conform to the above format.

1 Examples:

```
2 setenv OMP_SCHEDULE "guided, 4"  
3 setenv OMP_SCHEDULE "dynamic"  
4 setenv OMP_SCHEDULE "nonmonotonic:dynamic, 4"
```

5 Cross References

- 6 • `schedule` clause, see [Section 11.5.3](#)
- 7 • `run-sched-var` ICV, see [Table 2.1](#)

8 21.2.2 OMP_STACKSIZE

9 The `OMP_STACKSIZE` environment variable controls the size of the stack for threads created by
10 the OpenMP implementation, by setting the value of the `stacksize-var` ICV. The environment
11 variable does not control the size of the stack for an initial thread. The value of this environment
12 variable takes the form `size[unit]`, where:

- 13 • `size` is a positive integer that specifies the size of the stack for threads that are created by the
14 OpenMP implementation.
- 15 • `unit` is **B**, **K**, **M**, or **G** and specifies whether the given size is in Bytes, Kilobytes (1024 Bytes),
16 Megabytes (1024 Kilobytes), or Gigabytes (1024 Megabytes), respectively. If `unit` is present,
17 white space may occur between `size` and it, whereas if `unit` is not present then **K** is assumed.

18 The behavior of the program is implementation defined if `OMP_STACKSIZE` does not conform to
19 the above format, or if the implementation cannot provide a stack with the requested size.

20 Examples:

```
21 setenv OMP_STACKSIZE 2000500B  
22 setenv OMP_STACKSIZE "3000 k "  
23 setenv OMP_STACKSIZE 10M  
24 setenv OMP_STACKSIZE " 10 M "  
25 setenv OMP_STACKSIZE "20 m "  
26 setenv OMP_STACKSIZE " 1G"  
27 setenv OMP_STACKSIZE 20000
```

28 Cross References

- 29 • `stacksize-var` ICV, see [Table 2.1](#)

30 21.2.3 OMP_WAIT_POLICY

31 The `OMP_WAIT_POLICY` environment variable provides a hint to an OpenMP implementation
32 about the desired behavior of waiting threads by setting the `wait-policy-var` ICV. A compliant
33 OpenMP implementation may or may not abide by the setting of the environment variable. The
34 value of this environment variable must be one of the following:

1 **active** | **passive**

2 The **active** value specifies that waiting threads should mostly be active, consuming processor
3 cycles, while waiting. An OpenMP implementation may, for example, make waiting threads spin.
4 The **passive** value specifies that waiting threads should mostly be passive, not consuming
5 processor cycles, while waiting. For example, an OpenMP implementation may make waiting
6 threads yield the processor to other threads or go to sleep. The details of the **active** and
7 **passive** behaviors are implementation defined. The behavior of the program is implementation
8 defined if the value of **OMP_WAIT_POLICY** is neither **active** nor **passive**.

9 Examples:

```
10       setenv OMP_WAIT_POLICY ACTIVE  
11       setenv OMP_WAIT_POLICY active  
12       setenv OMP_WAIT_POLICY PASSIVE  
13       setenv OMP_WAIT_POLICY passive
```

14 **Cross References**

- 15 • *wait-policy-var* ICV, see [Table 2.1](#)

16 **21.2.4 OMP_DISPLAY_AFFINITY**

17 The **OMP_DISPLAY_AFFINITY** environment variable sets the *display-affinity-var* ICV so that
18 the runtime displays formatted affinity information for the initial device. Affinity information is
19 printed for all OpenMP threads in each parallel region upon first entering it. Also, if the information
20 accessible by the format specifiers listed in [Table 21.2](#) changes for any thread in the parallel region
21 changes then thread affinity information for all threads in that region is again displayed. If the
22 thread affinity for each respective parallel region at each nesting level has already been displayed
23 and the thread affinity has not changed, then the information is not displayed again. Thread affinity
24 information for threads in the same parallel region may be displayed in any order. The value of the
25 **OMP_DISPLAY_AFFINITY** environment variable may be set to one of these values:

26 **true** | **false**

27 The **true** value instructs the runtime to display the OpenMP thread affinity information, and uses
28 the format setting defined in the *affinity-format-var* ICV. The runtime does not display the OpenMP
29 thread affinity information when the value of the **OMP_DISPLAY_AFFINITY** environment
30 variable is **false** or undefined. For all values of the environment variable other than **true** or
31 **false**, the display action is implementation defined.

32 Example:

```
33       setenv OMP_DISPLAY_AFFINITY TRUE
```

34 For this example, an OpenMP implementation displays thread affinity information during program
35 execution, in a format given by the *affinity-format-var* ICV. The following is a sample output:

```
1 nesting_level= 1, thread_num= 0, thread_affinity= 0,1
2 nesting_level= 1, thread_num= 1, thread_affinity= 2,3
```

Cross References

- `OMP_AFFINITY_FORMAT`, see [Section 21.2.5](#)
- Controlling OpenMP Thread Affinity, see [Section 10.2.3](#)
- `affinity-format-var` ICV, see [Table 2.1](#)
- `display-affinity-var` ICV, see [Table 2.1](#)

21.2.5 OMP_AFFINITY_FORMAT

The `OMP_AFFINITY_FORMAT` environment variable sets the initial value of the `affinity-format-var` ICV which defines the format when displaying OpenMP thread affinity information. The value of this environment variable is case sensitive and leading and trailing whitespace is significant. Its value is a character string that may contain as substrings one or more field specifiers (as well as other characters). The format of each field specifier is

```
%[[[0].] size ] type
```

where each specifier must contain the percent symbol (%) and a type, that must be either a single character short name or its corresponding long name delimited with curly braces, such as `%n` or `%{thread_num}`. A literal percent is specified as `%%`. Field specifiers can be provided in any order. The behavior is implementation defined for field specifiers that do not conform to this format.

The `0` modifier indicates whether or not to add leading zeros to the output, following any indication of sign or base. The `.` modifier indicates the output should be right justified when `size` is specified. By default, output is left justified. The minimum field length is `size`, which is a decimal digit string with a non-zero first digit. If no `size` is specified, the actual length needed to print the field will be used. If the `0` modifier is used with `type` of `A`, `{thread_affinity}`, `H`, `{host}`, or a type that is not printed as a number, the result is unspecified. Any other characters in the format string that are not part of a field specifier will be included literally in the output.

TABLE 21.2: 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 <i>level</i> 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.

1 Implementations may define additional field types. If an implementation does not have information
 2 for a field type or an unknown field type is part of a field specifier, "undefined" is printed for this
 3 field when displaying the OpenMP thread affinity information.

4 Example:

```
5 setenv OMP_AFFINITY_FORMAT  
6 "Thread Affinity: %0.3L %.8n %.15{thread_affinity} %.12H"
```

7 The above example causes an OpenMP implementation to display OpenMP thread affinity
 8 information in the following form:

```
9 Thread Affinity: 001 0 0-1,16-17 nid003  
10 Thread Affinity: 001 1 2-3,18-19 nid003
```

11 **Cross References**

- 12 • Controlling OpenMP Thread Affinity, see [Section 10.2.3](#)
- 13 • *affinity-format-var* ICV, see [Table 2.1](#)

- 1 • `omp_get_ancestor_thread_num`, see [Section 18.2.16](#)
- 2 • `omp_get_level`, see [Section 18.2.15](#)
- 3 • `omp_get_num_teams`, see [Section 18.4.1](#)
- 4 • `omp_get_num_threads`, see [Section 18.2.2](#)
- 5 • `omp_get_thread_num`, see [Section 18.2.4](#)
- 6 • `omp_get_thread_num`, see [Section 18.2.4](#)

7 **21.2.6 OMP_CANCELLATION**

8 The `OMP_CANCELLATION` environment variable sets the initial value of the *cancel-var* ICV. The
9 value of this environment variable must be one of the following:

10 **true|false**

11 If the environment variable is set to **true**, the effects of the **cancel** construct and of cancellation
12 points are enabled (i.e., cancellation is enabled). If the environment variable is set to **false**,
13 cancellation is disabled and the **cancel** construct and cancellation points are effectively ignored.
14 The behavior of the program is implementation defined if `OMP_CANCELLATION` is set to neither
15 **true** nor **false**.

16 **Cross References**

- 17 • **cancel** directive, see [Section 16.1](#)
- 18 • *cancel-var* ICV, see [Table 2.1](#)

19 **21.2.7 OMP_AVAILABLE_DEVICES**

20 The `OMP_AVAILABLE_DEVICES` environment variable sets the *available-devices-var* ICV and
21 determines the available non-host devices and their device numbers by permitting selection of
22 devices from the set of supported accessible devices and by ordering them. This ICV is initialized
23 before any other ICV that uses a device number, depends on the number of available devices, or
24 permits device-specific environment variables. After the *available-devices-var* ICV is initialized,
25 only those devices that the ICV identifies are available and the `omp_get_num_devices` routine
26 returns the number of devices stored in the ICV.

27 The value of this environment variable must be a comma-separated list. Each item is either a trait
28 specification as specified in the following or *****. A ***** expands to all accessible and supported devices
29 while a trait specification expands to a possibly empty set of accessible and supported devices for
30 which the specification is fulfilled. After expansion, further selection via an optional array
31 subscript syntax and removal of devices that appear in previous items, each item contains an
32 unordered set of devices. A consecutive unique device number is then assigned to each device in

1 the sets, starting with device number zero, where the device number of the first device in an item is
2 the total number of devices in all previous items.

3 Traits are specified by the case-insensitive trait name followed by the argument in parentheses. The
4 permitted traits are **kind**(*kind-name*), **isa**(*isa-name*), **arch**(*arch-name*), and
5 **vendor**(*vendor-name*), where the names are as specified in [Section 7.1](#) and the *OpenMP*
6 *Additional Definitions* document; the *kind-name* **host** is not permitted. Multiple traits can be
7 combined using the binary operators **&&** and **||** to require both or either trait, respectively.
8 Parentheses can be used for grouping, but are optional except that **&&** and **||** may not appear in the
9 same grouping level. The unary **!** operator inverts the meaning of the immediately following trait
10 or parenthesized group.

11 Each trait specification or ***** yields a (possibly zero-sized) array of non-host devices with the lowest
12 array element, if it exists, having index zero. The C/C++ syntax [*index*] can be used to select an
13 element and the array section syntax for C/C++ as specified in [Section 3.2.5](#) can be used to specify
14 a subset of elements. Any array element specified by the subscript that is outside the bounds of the
15 array resulting from the trait specification or ***** is silently excluded.

16 Cross References

- 17 • Device Directives and Clauses, see [Chapter 13](#)
- 18 • *available-devices-var* ICV, see [Table 2.1](#)

19 21.2.8 OMP_DEFAULT_DEVICE

20 The **OMP_DEFAULT_DEVICE** environment variable sets the initial value of the *default-device-var*
21 ICV. The value of this environment variable must be a comma-separated list, each item being either
22 a non-negative integer value that denotes the device number, a trait specification with an optional
23 subscript selector, or one of the following case-insensitive string literals: **initial** to specify the
24 host device, **invalid** to specify the device number **omp_invalid_device**, or **default** to
25 set the ICV as if this environment variable was not specified (see [Section 1.3](#)).

26 The trait specification is as described for **OMP_AVAILABLE_DEVICES** (see [Section 21.2.7](#)),
27 except that in addition the trait *device_num*(*device number*) may be specified, **host** is permitted as
28 *kind-name*. The device numbers yielded by the trait specification are sorted in ascending order by
29 device number; the array-element syntax as described in **OMP_AVAILABLE_DEVICES** can be
30 used to select an element from the set. If an item is an empty set, non-existing element, or does not
31 evaluate to an available device, the next item is evaluated; otherwise, the *default-device-var* ICV is
32 set to the first value of the set. However, **initial**, **invalid**, and **default** always match. If
33 none of the list items match, the *default-device-var* ICV is set to **omp_invalid_device**.

34 Cross References

- 35 • Device Directives and Clauses, see [Chapter 13](#)
- 36 • *default-device-var* ICV, see [Table 2.1](#)

21.2.9 OMP_TARGET_OFFLOAD

The **OMP_TARGET_OFFLOAD** environment variable sets the initial value of the *target-offload-var* ICV. Its value must be one of the following:

mandatory | **disabled** | **default**

The **mandatory** value specifies that the effect of any device construct or device memory routine that uses a device that is unavailable or not supported by the implementation, or uses a non-conforming device number, is as if the **omp_invalid_device** device number was used. 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](#).

Example:

```
% setenv OMP_TARGET_OFFLOAD mandatory
```

Cross References

- Device Directives and Clauses, see [Chapter 13](#)
- Device Memory Routines, see [Section 18.8](#)
- *target-offload-var* ICV, see [Table 2.1](#)

21.2.10 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 [Table 2.1](#)

21.3 OMPT Environment Variables

This section defines environment variables that affect operation of the OMPT tool interface.

21.3.1 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

- OMPT Interface, see [Chapter 19](#)
- *tool-var* ICV, see [Table 2.1](#)

21.3.2 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. Whether the value of this environment variable is case sensitive is implementation defined.

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

- OMPT Interface, see [Chapter 19](#)
- *tool-libraries-var* ICV, see [Table 2.1](#)
- `ompt_start_tool`, see [Section 19.2.1](#)

21.3.3 OMP_TOOL_VERBOSE_INIT

The `OMP_TOOL_VERBOSE_INIT` environment variable sets the *tool-verbose-init-var* ICV, which controls whether an OpenMP implementation will verbosely log the registration of a tool. The value of this environment variable must be one of the following:

`disabled` | `stdout` | `stderr` | `<filename>`

If `OMP_TOOL_VERBOSE_INIT` is set to any value other than case insensitive `disabled`, `stdout`, or `stderr`, the value is interpreted as a filename and the OpenMP runtime will try to log to a file with prefix *filename*. If the value is interpreted as a filename, whether it is case sensitive is implementation defined. If opening the logfile fails, the output will be redirected to `stderr`. If `OMP_TOOL_VERBOSE_INIT` is not defined, the default value for *tool-verbose-init-var* is `disabled`. Support for logging to `stdout` or `stderr` is implementation defined. Unless *tool-verbose-init-var* is `disabled`, the OpenMP runtime will log the steps of the tool activation process defined in [Section 19.2.2](#) to a file with a name that is constructed using the provided filename prefix. The format and detail of the log is implementation defined. At a minimum, the log will contain one of the following:

- That the *tool-var* ICV is disabled;
- An indication that a tool was available in the address space at program launch; or
- The path name of each tool in `OMP_TOOL_LIBRARIES` that is considered for dynamic loading, whether dynamic loading was successful, and whether the `ompt_start_tool` function is found in the loaded library.

In addition, if an `ompt_start_tool` function is called the log will indicate whether or not the tool will use the OMPT interface.

Example:

```
% setenv OMP_TOOL_VERBOSE_INIT disabled
% setenv OMP_TOOL_VERBOSE_INIT STDERR
% setenv OMP_TOOL_VERBOSE_INIT ompt_load.log
```

Cross References

- OMPT Interface, see [Chapter 19](#)
- *tool-verbose-init-var* ICV, see [Table 2.1](#)

21.4 OMPD Environment Variables

This section defines environment variables that affect operation of the OMPD tool interface.

21.4.1 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

- Enabling Runtime Support for OMPD, see [Section 20.2.1](#)
- OMPD Interface, see [Chapter 20](#)
- *debug-var* ICV, see [Table 2.1](#)

21.5 Memory Allocation Environment Variables

This section defines environment variables that affect memory allocations.

21.5.1 OMP_ALLOCATOR

The `OMP_ALLOCATOR` environment variable sets the initial value of the *def-allocator-var* ICV that specifies the default allocator for allocation calls, directives and clauses that do not specify an allocator. The following grammar describes the values accepted for the `OMP_ALLOCATOR` environment variable.

$\langle \text{allocator} \rangle \models \langle \text{predef-allocator} \rangle \mid \langle \text{predef-mem-space} \rangle \mid \langle \text{predef-mem-space} \rangle : \langle \text{traits} \rangle$
 $\langle \text{traits} \rangle \models \langle \text{trait} \rangle = \langle \text{value} \rangle \mid \langle \text{trait} \rangle = \langle \text{value} \rangle, \langle \text{traits} \rangle$
 $\langle \text{predef-allocator} \rangle \models \text{one of the predefined allocators from Table 6.3}$
 $\langle \text{predef-mem-space} \rangle \models \text{one of the predefined memory spaces from Table 6.1}$
 $\langle \text{trait} \rangle \models \text{one of the allocator trait names from Table 6.2}$
 $\langle \text{value} \rangle \models \text{one of the allowed values from Table 6.2} \mid \text{non-negative integer}$
 $\mid \langle \text{predef-allocator} \rangle$

1 The *value* can be an integer only if the *trait* accepts a numerical value, for the **fb_data** *trait* the
 2 *value* can only be *predef-allocator*. If the value of this environment variable is not a predefined
 3 allocator, then a new allocator with the given predefined memory space and optional traits is
 4 created and set as the *def-allocator-var* ICV. If the new allocator cannot be created, the
 5 *def-allocator-var* ICV will be set to **omp_default_mem_alloc**.

6 Example:

```

7 setenv OMP_ALLOCATOR omp_high_bw_mem_alloc
8 setenv OMP_ALLOCATOR omp_large_cap_mem_space:alignment=16, \
9 pinned=true
10 setenv OMP_ALLOCATOR omp_high_bw_mem_space:pool_size=1048576, \
11 fallback=allocator_fb,fb_data=omp_low_lat_mem_alloc

```

12 Cross References

- 13 • Memory Allocators, see [Section 6.2](#)
- 14 • *def-allocator-var* ICV, see [Table 2.1](#)

15 21.6 Teams Environment Variables

16 This section defines environment variables that affect the operation of **teams** regions.

17 21.6.1 OMP_NUM_TEAMS

18 The **OMP_NUM_TEAMS** environment variable sets the maximum number of teams created by a
 19 **teams** construct by setting the *nteam*s-*var* ICV. The value of this environment variable must be a
 20 positive integer. The behavior of the program is implementation defined if the requested value of
 21 **OMP_NUM_TEAMS** is greater than the number of teams that an implementation can support, or if
 22 the value is not a positive integer.

Cross References

- **teams** directive, see [Section 10.3](#)
- *ntteams-var* ICV, see [Table 2.1](#)

21.6.2 OMP_TEAMS_THREAD_LIMIT

The **OMP_TEAMS_THREAD_LIMIT** environment variable sets the maximum number of OpenMP threads that can execute tasks 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** directive, see [Section 10.3](#)
- *teams-thread-limit-var* ICV, see [Table 2.1](#)

21.7 OMP_DISPLAY_ENV

The **OMP_DISPLAY_ENV** environment variable instructs the runtime to display the information as described in the **omp_display_env** routine section ([Section 18.15](#)). The value of the **OMP_DISPLAY_ENV** environment variable may be set to one of these values:

true | **false** | **verbose**

If the environment variable is set to **true**, the effect is as if the **omp_display_env** routine is called with the *verbose* argument set to *false* at the beginning of the program. If the environment variable is set to **verbose**, the effect is as if the **omp_display_env** routine is called with the *verbose* argument set to *true* at the beginning of the program. If the environment variable is undefined or set to **false**, the runtime does not display any information. For all values of the environment variable other than **true**, **false**, and **verbose**, the displayed information is unspecified.

Example:

```
% setenv OMP_DISPLAY_ENV true
```

For the output of the above example, see [Section 18.15](#).

Cross References

- Environment Display Routine, see [Section 18.15](#)

A OpenMP Implementation-Defined Behaviors

This appendix summarizes the behaviors that are described as implementation defined in the OpenMP 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.

Chapter 1:

- **Processor:** A hardware unit that is implementation defined (see [Section 1.2.1](#)).
- **Device:** An implementation-defined logical execution engine (see [Section 1.2.1](#)).
- **Device pointer:** An *implementation-defined handle* that refers to a device address (see [Section 1.2.6](#)).
- **Supported active levels of parallelism:** The maximum number of active parallel regions that may enclose any region of code in the program is implementation defined (see [Section 1.2.7](#)).
- **Deprecated features:** For any *deprecated* feature, whether any modifications provided by its replacement feature (if any) apply to the deprecated feature is implementation defined (see [Section 1.2.7](#)).
- **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 elements or structure elements) is implementation defined but is no larger than the base language requires. The manner in which a program can obtain the referenced device address from a device pointer, outside the mechanisms specified by OpenMP, is implementation defined (see [Section 1.4.1](#)).

Chapter 2:

- **Internal control variables:** The initial values of *dyn-var*, *nthreads-var*, *run-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*, *num-procs-var* and *def-allocator-var* are implementation defined (see [Section 2.2](#)).

Chapter 3:

- C / C++
- A pragma directive that uses **omp_x** as the first processing token is implementation defined (see [Section 3.1](#)).
- C / C++

C++

- The attribute namespace of an attribute specifier or the optional namespace qualifier within a **sequence** attribute that uses **omp_x** is implementation defined (see [Section 3.1](#)).
- Whether a **throw** executed inside a region that arises from an exception-aborting directive results in runtime error termination is implementation defined (see [Section 3.1](#)).

C++

Fortran

- Any directive that uses **omx** or **omp_x** in the sentinel is implementation defined (see [Section 3.1](#)).

Fortran

Chapter 4:

- **Loop-iteration spaces and vectors:** The particular integer type used to compute the iteration count for the collapsed loop is implementation defined (see [Section 4.4.2](#)).

Chapter 5:

Fortran

- **Data-sharing attributes:** The data-sharing attributes of dummy arguments that do not have the **VALUE** attribute are implementation defined if the associated actual argument is shared unless the actual argument is a scalar variable, structure, an array that is not a pointer or assumed-shape array, or a simply contiguous array section (see [Section 5.1.2](#)).
- **threadprivate directive:** If the conditions for values of data in the threadprivate objects of threads (other than an initial thread) to persist between two consecutive active parallel regions do not all hold, the allocation status of an allocatable variable in the second region is implementation defined (see [Section 5.2](#)).

Fortran

- **is_device_ptr clause:** Support for pointers created outside of the OpenMP device data management routines is implementation defined (see [Section 5.4.7](#)).

Fortran

- **has_device_addr and use_device_addr clauses:** The result of inquiring about list item properties other than the **CONTIGUOUS** attribute, storage location, storage size, array bounds, character length, association status and allocation status is implementation defined (see [Section 5.4.9](#) and [Section 5.4.10](#)).

Fortran

- **aligned clause:** If the *alignment* modifier is not specified, the default alignments for SIMD instructions on the target platforms are implementation defined (see [Section 5.11](#)).

Chapter 6:

- **Memory spaces:** The actual storage resources that each memory space defined in Table 6.1 represents are implementation defined. The mechanism that provides the constant value of the variables allocated in the `omp_const_mem_space` memory space is implementation defined (see Section 6.1).
- **Memory allocators:** The minimum size for partitioning allocated memory over storage resources is implementation defined. The default value for the `pool_size` allocator trait (see Table 6.2) is implementation defined. The memory spaces associated with the predefined `omp_cgroup_mem_alloc`, `omp_pteam_mem_alloc` and `omp_thread_mem_alloc` allocators (see Table 6.3) are implementation defined (see Section 6.2).

Chapter 7:

- **OpenMP context:** The accepted *isa-name* values for the *isa* trait, the accepted *arch-name* values for the *arch* trait, the accepted *extension-name* values for the *extension* trait and whether the `dispatch` construct is added to the *construct* set are implementation defined (see Section 7.1).
- **Metadirectives:** The number of times that each expression of the context selector of a `when` clause is evaluated is implementation defined (see Section 7.4.1).
- **Declare variant directives:** If two replacement candidates have the same score then their order is implementation defined. The number of times each expression of the context selector of a `match` clause is evaluated is implementation defined. For calls to `constexpr` base functions that are evaluated in constant expressions, whether any variant replacement occurs is implementation defined. Any differences that the specific OpenMP context requires in the prototype of the variant from the base function prototype are implementation defined (see Section 7.5).
- **declare simd directive:** If a SIMD version is created and the `simdlen` clause is not specified, the number of concurrent arguments for the function is implementation defined (see Section 7.7).
- **Declare target directives:** Whether the same version is generated for different devices, or whether a version that is called in a `target` region differs from the version that is called outside a `target` region, is implementation defined (see Section 7.8).

Chapter 8:

- **requires directive:** Support for any feature specified by a requirement clause on a `requires` directive is implementation defined (see Section 8.2).

Chapter 9:

- **unroll construct:** If no clauses are specified, if and how the loop is unrolled is implementation defined. If the `partial` clause is specified without an *unroll-factor* argument then the unroll factor is a positive integer that is implementation defined (see Section 9.2).

Chapter 10:

- **Dynamic adjustment of threads:** Providing the ability to adjust the number of threads dynamically is implementation defined (see [Section 10.2.1](#)).
- **Thread affinity:** For the **close** thread affinity policy, if $T > P$ and P does not divide T evenly, the exact number of threads in a particular place is implementation defined. For the **spread** thread affinity, if $T > P$ and P does not divide T evenly, the exact number of threads in a particular subpartition is implementation defined. The determination of whether the affinity request can be fulfilled is implementation defined. If the affinity request cannot be fulfilled, then the affinity of threads in the team is implementation defined (see [Section 10.2.3](#)).
- **teams construct:** The number of teams that are created is implementation defined, but it is greater than or equal to the lower bound and less than or equal to the upper bound values of the **num_teams** clause if specified. If the **num_teams** clause is not specified, the number of teams is less than or equal to the value of the *ntteams-var* ICV if its value is greater than zero. Otherwise it is an implementation defined value greater than or equal to 1 (see [Section 10.3](#)).
- **simd construct:** The number of iterations that are executed concurrently at any given time is implementation defined (see [Section 10.5](#)).

Chapter 11:

- **single construct:** The method of choosing a thread to execute the structured block each time the team encounters the construct is implementation defined (see [Section 11.1](#)).
- **sections construct:** The method of scheduling the structured block sequences among threads in the team is implementation defined (see [Section 11.3](#)).
- **Worksharing-loop directive:** The schedule that is used is implementation defined if the **schedule** clause is not specified or if the specified schedule has the kind **auto**. The value of *simd_width* for the **simd** schedule modifier is implementation defined (see [Section 11.5](#)).
- **distribute construct:** If no **dist_schedule** clause is specified then the schedule for the **distribute** construct is implementation defined (see [Section 11.6](#)).

Chapter 12:

- **taskloop construct:** The number of loop iterations assigned to a task created from a **taskloop** construct is implementation defined, unless the **grainsize** or **num_tasks** clause is specified (see [Section 12.6](#)).

C++

- **taskloop construct:** For **firstprivate** variables of class type, the number of invocations of copy constructors to perform the initialization is implementation defined (see [Section 12.6](#)).

C++

1 **Chapter 13:**

- 2 • **thread_limit clause:** The maximum number of threads that participate in executing
3 tasks in the contention group that each team initiates is implementation defined if no
4 **thread_limit** clause is specified on the construct. Otherwise, it has the implementation
5 defined upper bound of the *teams-thread-limit-var* ICV, if the value of this ICV is greater
6 than zero (see [Section 13.3](#)).

7 **Chapter 14:**

- 8 • **interop Construct:** The *foreign-runtime-id* values for the **prefer_type** clause that the
9 implementation supports, including non-standard names compatible with this clause, and the
10 default choice when the implementation supports multiple values are implementation defined
11 (see [Section 14.1](#)).

12 **Chapter 15:**

- 13 • **atomic construct:** A compliant implementation may enforce exclusive access between
14 **atomic** regions that update different storage locations. The circumstances under which this
15 occurs are implementation defined. If the storage location designated by *x* is not size-aligned
16 (that is, if the byte alignment of *x* is not a multiple of the size of *x*), then the behavior of the
17 atomic region is implementation defined (see [Section 15.8.5](#)).

18 **Chapter 16:**

- 19 • None.

20 **Chapter 17:**

- 21 • None.

22 **Chapter 18:**

- 23 • Runtime Routine names that begin with the **omp_x** prefix are implementation-defined
24 extensions to the OpenMP Runtime API (see [Chapter 18](#)).

25 **C / C++**

- 26 • **Runtime library definitions:** The enum types for **omp_allocator_handle_t**,
27 **omp_event_handle_t**, **omp_interop_fr_t** and **omp_memspace_handle_t** are
28 implementation defined. The integral or pointer type for **omp_interop_t** is
29 implementation defined. The value of the **omp_invalid_device** enumerator is
implementation defined (see [Section 18.1](#)).

30 **C / C++**

31 **Fortran**

- 32 • **Runtime library definitions:** Whether the include file **omp_lib.h** or the module
33 **omp_lib** (or both) is provided is implementation defined. Whether the **omp_lib.h** file
34 provides derived-type definitions or those routines that require an explicit interface is
35 implementation defined. Whether any of the OpenMP runtime library routines that take an
36 argument are extended with a generic interface so arguments of different **KIND** type can be
accommodated is implementation defined. The value of the **omp_invalid_device**
named constant is implementation defined (see [Section 18.1](#)).

Fortran

- 1 • **omp_set_num_threads routine:** If the argument is not a positive integer, the behavior is
2 implementation defined (see [Section 18.2.1](#)).
- 3 • **omp_set_schedule routine:** For implementation-specific schedule kinds, the values and
4 associated meanings of the second argument are implementation defined (see [Section 18.2.9](#)).
- 5 • **omp_get_schedule routine:** The value returned by the second argument is
6 implementation defined for any schedule kinds other than **static**, **dynamic** and **guided**
7 (see [Section 18.2.10](#)).
- 8 • **omp_get_supported_active_levels routine:** The number of active levels of
9 parallelism supported by the implementation is implementation defined, but must be positive
10 (see [Section 18.2.12](#)).
- 11 • **omp_set_max_active_levels routine:** If the argument is a negative integer then the
12 behavior is implementation defined. If the argument is less than the *active-levels-var* ICV, the
13 *max-active-levels-var* ICV is set to an implementation-defined value between the value of the
14 argument and the value of *active-levels-var*, inclusive (see [Section 18.2.13](#)).
- 15 • **omp_get_place_proc_ids routine:** The meaning of the non-negative numerical
16 identifiers returned by the **omp_get_place_proc_ids** routine is implementation
17 defined. The order of the numerical identifiers returned in the array *ids* is implementation
18 defined (see [Section 18.3.4](#)).
- 19 • **omp_set_affinity_format routine:** When called from within any **parallel** or
20 **teams** region, the binding thread set (and binding region, if required) for the
21 **omp_set_affinity_format** region and the effect of this routine are implementation
22 defined (see [Section 18.3.8](#)).
- 23 • **omp_get_affinity_format routine:** When called from within any **parallel** or
24 **teams** region, the binding thread set (and binding region, if required) for the
25 **omp_get_affinity_format** region is implementation defined (see [Section 18.3.9](#)).
- 26 • **omp_display_affinity routine:** If the *format* argument does not conform to the
27 specified format then the result is implementation defined (see [Section 18.3.10](#)).
- 28 • **omp_capture_affinity routine:** If the *format* argument does not conform to the
29 specified format then the result is implementation defined (see [Section 18.3.11](#)).
- 30 • **omp_set_num_teams routine:** If the argument does not evaluate to a positive integer, the
31 behavior of this routine is implementation defined (see [Section 18.4.3](#)).
- 32 • **omp_set_teams_thread_limit routine:** If the argument is not a positive integer, the
33 behavior is implementation defined (see [Section 18.4.5](#)).
- 34 • **omp_pause_resource_all routine:** The behavior of this routine is implementation
35 defined if the argument kind is not listed in [Section 18.6.1](#) (see [Section 18.6.2](#)).
- 36 • **omp_target_memcpy_rect** and **omp_target_memcpy_rect_async routines:**
37 The maximum number of dimensions supported is implementation defined, but must be at

1 least three (see [Section 18.8.6](#) and [Section 18.8.8](#)).

- 2 • **Lock routines:** If a lock contains a synchronization hint, the effect of the hint is
3 implementation defined (see [Section 18.9](#)).
- 4 • **Interoperability routines:** Implementation-defined properties may use zero and positive
5 values for properties associated with an `omp_interop_t` object (see [Section 18.12](#)).

6 **Chapter 19:**

- 7 • **Tool callbacks:** If a tool attempts to register a callback listed in [Table 19.3](#), whether the
8 registered callback may never, sometimes or always invoke this callback for the associated
9 events is implementation defined (see [Section 19.2.4](#)).
- 10 • **Device tracing:** Whether a target device supports tracing or not is implementation defined; if
11 a target device does not support tracing, a `NULL` may be supplied for the `lookup` function to
12 the device initializer of a tool (see [Section 19.2.5](#)).
- 13 • **ompt_set_trace_ompt and ompt_get_record_ompt runtime entry points:**
14 Whether a device-specific tracing interface defines this runtime entry point, indicating that it
15 can collect traces in OMPT format, is implementation defined. The kinds of trace records
16 available for a device is implementation defined (see [Section 19.2.5](#)).
- 17 • **Native record abstract type:** The meaning of a `hwid` value for a device is implementation
18 defined (see [Section 19.4.3.3](#)).
- 19 • **ompt_dispatch_chunk_t type:** Whether the chunk of a taskloop is contiguous is
20 implementation defined (see [Section 19.4.4.13](#)).
- 21 • **ompt_record_abstract_t type:** The set of OMPT thread states supported is
22 implementation defined (see [Section 19.4.4.28](#)).
- 23 • **ompt_callback_sync_region_t callback type:** For the *implicit-barrier-wait-begin*
24 and *implicit-barrier-wait-end* events at the end of a parallel region, whether the
25 `parallel_data` argument is `NULL` or points to the parallel data of the current parallel
26 region is implementation defined (see [Section 19.5.2.13](#)).
- 27 • **ompt_callback_target_data_op_emi_t and**
28 **ompt_callback_target_data_op_t callback types:** Whether in some operations
29 `src_addr` or `dest_addr` might point to an intermediate buffer is implementation defined (see
30 [Section 19.5.2.25](#)).
- 31 • **ompt_get_place_proc_ids_t entry point type:** The meaning of the numerical
32 identifiers returned is implementation defined. The order of `ids` returned in the array is
33 implementation defined (see [Section 19.6.1.8](#)).
- 34 • **ompt_get_partition_place_nums_t entry point type:** The order of the identifiers
35 returned in the array `place_nums` is implementation defined (see [Section 19.6.1.10](#)).
- 36 • **ompt_get_proc_id_t entry point type:** The meaning of the numerical identifier
37 returned is implementation defined (see [Section 19.6.1.11](#)).

Chapter 20:

- **ompd_callback_print_string_fn_t callback type:** The value of *category* is implementation defined (see [Section 20.4.5](#)).
- **ompd_parallel_handle_compare operation:** The means by which parallel region handles are ordered is implementation defined (see [Section 20.5.6.5](#)).
- **ompd_task_handle_compare operation:** The means by which task handles are ordered is implementation defined (see [Section 20.5.7.6](#)).

Chapter 21:

- **OMP_DYNAMIC environment variable:** If the value is neither **true** nor **false**, the behavior of the program is implementation defined (see [Section 21.1.1](#)).
- **OMP_NUM_THREADS environment variable:** If any value of the list specified leads to a number of threads that is greater than the implementation can support, or if any value is not a positive integer, then the behavior of the program is implementation defined (see [Section 21.1.2](#)).
- **OMP_THREAD_LIMIT environment variable:** If the requested value is greater than the number of threads an implementation can support, or if the value is not a positive integer, the behavior of the program is implementation defined (see [Section 21.1.3](#)).
- **OMP_MAX_ACTIVE_LEVELS environment variable:** If the value is a negative integer or is greater than the maximum number of nested active parallel levels that an implementation can support then the behavior of the program is implementation defined (see [Section 21.1.4](#)).
- **OMP_PLACES environment variable:** The meaning of the numbers specified in the environment variable and how the numbering is done are implementation defined. The precise definitions of the abstract names are implementation defined. An implementation may add implementation-defined abstract names as appropriate for the target platform. When creating a place list of *n* elements by appending the number *n* to an abstract name, the determination of which resources to include in the place list is implementation defined. When requesting more resources than available, the length of the place list is also implementation defined. The behavior of the program is implementation defined when the execution environment cannot map a numerical value (either explicitly defined or implicitly derived from an interval) within the **OMP_PLACES** list to a processor on the target platform, or if it maps to an unavailable processor. The behavior is also implementation defined when the **OMP_PLACES** environment variable is defined using an abstract name (see [Section 21.1.5](#)).
- **OMP_PROC_BIND environment variable:** If the value is not **true**, **false**, or a comma separated list of **primary**, **close**, or **spread**, the behavior is implementation defined. The behavior is also implementation defined if an initial thread cannot be bound to the first place in the OpenMP place list. The thread affinity policy is implementation defined if the value is **true** (see [Section 21.1.6](#)).
- **OMP_SCHEDULE environment variable:** If the value does not conform to the specified format then the behavior of the program is implementation defined (see [Section 21.2.1](#)).

- 1 ● **OMP_STACKSIZE environment variable:** If the value does not conform to the specified
2 format or the implementation cannot provide a stack of the specified size then the behavior is
3 implementation defined (see [Section 21.2.2](#)).
- 4 ● **OMP_WAIT_POLICY environment variable:** The details of the **active** and **passive**
5 behaviors are implementation defined (see [Section 21.2.3](#)).
- 6 ● **OMP_DISPLAY_AFFINITY environment variable:** For all values of the environment
7 variables other than **true** or **false**, the display action is implementation defined (see
8 [Section 21.2.4](#)).
- 9 ● **OMP_AFFINITY_FORMAT environment variable:** Additional implementation-defined
10 field types can be added (see [Section 21.2.5](#)).
- 11 ● **OMP_CANCELLATION environment variable:** If the value is set to neither **true** nor
12 **false**, the behavior of the program is implementation defined (see [Section 21.2.6](#)).
- 13 ● **OMP_TARGET_OFFLOAD environment variable:** The support of **disabled** is
14 implementation defined (see [Section 21.2.9](#)).
- 15 ● **OMP_TOOL_LIBRARIES environment variable:** Whether the value of the environment
16 variable is case sensitive is implementation defined (see [Section 21.3.2](#)).
- 17 ● **OMP_TOOL_VERBOSE_INIT environment variable:** Support for logging to **stdout** or
18 **stderr** is implementation defined. Whether the value of the environment variable is case
19 sensitive when it is treated as a filename is implementation defined. The format and detail of
20 the log is implementation defined (see [Section 21.3.3](#)).
- 21 ● **OMP_DEBUG environment variable:** If the value is neither **disabled** nor **enabled**, the
22 behavior is implementation defined (see [Section 21.4.1](#)).
- 23 ● **OMP_NUM_TEAMS environment variable:** If the value is not a positive integer or is greater
24 than the number of teams that an implementation can support, the behavior of the program is
25 implementation defined (see [Section 21.6.1](#)).
- 26 ● **OMP_TEAMS_THREAD_LIMIT environment variable:** If the value is not a positive integer
27 or is greater than the number of threads that an implementation can support, the behavior of
28 the program is implementation defined (see [Section 21.6.2](#)).

B Features History

This appendix summarizes the major changes between OpenMP API versions since version 2.5.

B.1 Deprecated Features

The following features were deprecated in Version 6.0:

- None.

B.2 Version 5.2 to 6.0 Differences

- All features deprecated in versions 5.2, 5.1 and 5.0 were removed.
- The environment variable syntax was extended to support initializing ICVs for host and non-host devices with a single environment variable (see [Section 2.2](#) and [Chapter 21](#)).

C++

- The **decl** attribute was added to improve the attribute syntax for declarative directives (see [Section 3.1](#)).

C++

Fortran

- The definitions of locator list items and assignable OpenMP types were extended to include function references that have data pointer results (see [Section 3.2.1](#)).

Fortran

C / C++

- Array section definition was extended to permit, where explicitly allowed, omission of length when the size of the array dimension is not known (see [Section 3.2.5](#)).

C / C++

- The semantics of the **use_device_ptr** and **use_device_addr** clauses on a **target data** construct were altered to imply a reference count update on entry and exit from the region for the corresponding objects that they reference in the device data environment (see [Section 5.4.8](#) and [Section 5.4.10](#)).

C++

- The circumstances under which implicitly declared reduction identifiers are supported for variables of class type were clarified (see [Section 5.5.3](#) and [Section 5.5.5](#)).

C++

- The property of the *map-type* modifier was changed to “default” such that it can be freely placed and omitted even if other modifiers are used (see [Section 5.8.3](#)).
- The **map** clause was extended to permit mapping of assumed-size arrays (see [Section 5.8.3](#)).
- The **groupprivate** directive was added to specify that variables should be privatized with respect to a contention group (see [Section 5.12](#)).
- The **local** clause was added to the **declare target** directive to specify that variables should be replicated locally for each device (see [Section 5.13](#)).
- The allocator trait **part_size** was added to specify the size of the **interleaved** allocator partitions (see [Section 6.2](#)).
- The **pin_device**, **preferred_device** and **target_access** memory allocator traits were defined to provide greater control of memory allocations that may be accessible from multiple devices (see [Section 6.2](#)).
- The **device** value of the **access** allocator trait was defined as the default **access** allocator trait and to provide the semantics that an allocator with the trait corresponds to memory that all threads on a specific device can access. The semantics of an allocator with the **all** value were updated to correspond to memory that all threads in the system can access (see [Section 6.2](#)).
- The **interop** operation of the **append_args** clause was extended to allow specification of all modifiers of the **init** clause (see [Section 7.5.3](#) and [Section 14.1.2](#)).
- The **dispatch** construct was extended with the **interop** clause to support appending arguments specific to a call site (see [Section 7.6](#) and [Section 7.6.1](#)).
- The **reverse** construct was added to reverse the iteration order of a loop (see [Section 9.3](#)).
- The **interchange** construct was added to permute the order of loops in a loop nest (see [Section 9.4](#)).
- The **apply** clause was added to enable more flexible composition of loop-transforming constructs (see [Section 9.5](#)).
- The **omp_curr_progress_width** identifier (see [Section 10.1](#)), **safesync** clause on the **parallel** construct (see [Section 10.2.5](#)) and the **omp_get_max_progress_width** runtime routine (see [Section 18.7.2](#)) were added to control which synchronizing threads are guaranteed to make progress eventually.
- The *prescriptiveness* modifier was added to the **num_threads** clause and **strict** semantics were defined for the clause (see [Section 10.2.2](#)).

- The **memscope** clause was added to the **atomic** and **flush** constructs to allow the binding thread set to span multiple devices (see [Section 15.8.4](#)).
- New routines were added to obtain memory spaces and memory allocators to allocate remote and shared memory (see [Section 18.13](#)).
- The **omp_get_memspace_num_resources** routine was added to be able to query the number of available resources of a memory space (see [Section 18.13.12](#)).
- The **omp_get_submemspace** routine was added to obtain a memory space with a subset of the original memory space resources (see [Section 18.13.13](#)).
- The **ompt_get_buffer_limits** runtime entry point was added to the OMPT device tracing interface so that a first party tool can obtain an upper limit on the sizes of the trace buffers that it should make available to the implementation (see [Section 19.5.2.23](#) and [Section 19.6.2.6](#)).
- The environment variable **OMP_AVAILABLE_DEVICES** was added and the environment variable **OMP_DEFAULT_DEVICE** was extended to support device selection by traits (see [Chapter 21](#)).
- The environment variable **OMP_PLACES** was extended to support an increment between consecutive places when creating a place list from an abstract name (see [Section 21.1.5](#)).

B.3 Version 5.1 to 5.2 Differences

- The *explicit-task-var* ICV has replaced the *implicit-task-var* ICV and has the opposite meaning and semantics (see [Chapter 2](#)). The **omp_in_explicit_task** routine was added to query if a code region is executed from an explicit task region (see [Section 18.5.2](#)).
- Major reorganization and numerous changes were made to improve the quality of the specification of OpenMP syntax and to increase consistency of restrictions and their wording. These changes frequently result in the possible perception of differences to preceding versions of the OpenMP specification. However, those differences almost always resolve ambiguities, which may nonetheless have implications for existing implementations and programs.
- For OpenMP directives, reserved the **omp** sentinel (see [Section 3.1](#), [Section 3.1.1](#) and [Section 3.1.2](#)) and, for implementation-defined directives that extend the OpenMP directives reserved the **omp_x** sentinel for C/C++ and free source form Fortran (see [Section 3.1](#) and [Section 3.1.2](#)) and the **om_x** sentinel for fixed source form Fortran to accommodate character position requirements (see [Section 3.1.1](#)). Reserved clause names that begin with the **omp_x** prefix for implementation-defined clauses on OpenMP directives (see [Section 3.2](#)). Reserved names in the base language that start with the **omp_x** and **omp_x** prefix and reserved the **omp** and **omp_x** namespaces (see [Chapter 4](#)) for the OpenMP runtime API and for implementation-defined extensions to that API (see [Chapter 18](#)).

- 1 • Allowed any clause that can be specified on a paired **end** directive to be specified on the
2 directive (see [Section 3.1](#)), including the **copyprivate** clause (see [Section 5.7.2](#)) and the
3 **nowait** clause in Fortran (see [Section 15.6](#)).
- 4 • Allowed **if** clause on **teams** construct (see [Section 3.4](#) and [Section 10.3](#)).
- 5 • For consistency with the syntax of other definitions of the clause, the syntax of the **destroy**
6 clause on the **depobj** construct with no argument was deprecated (see [Section 3.5](#)).
- 7 • For consistency with the syntax of other clauses, the syntax of the **linear** clause that
8 specifies its argument and *linear-modifier* as *linear-modifier (list)* was deprecated and the
9 **step** modifier was added for specifying the linear step (see [Section 5.4.6](#)).
- 10 • The *minus* (–) operator for reductions was deprecated (see [Section 5.5.5](#)).
- 11 • The syntax of modifiers without comma separators in the **map** clause was deprecated (see
12 [Section 5.8.3](#)).
- 13 • To support the complete range of user-defined mappers and to improve consistency of **map**
14 clause usage, the **declare mapper** directive was extended to accept *iterator-modifier* and
15 the **present** *map-type-modifier* (see [Section 5.8.3](#) and [Section 5.8.7](#)).
- 16 • If a matching mapped list item is not found in the data environment, the pointer retains its
17 original value as per the firstprivate semantics (see ??).
- 18 • The **enter** clause was added as a synonym for the **to** clause on the declare target directive,
19 and the corresponding **to** clause was deprecated to reduce parsing ambiguity (see
20 [Section 5.8.4](#) and [Section 7.8](#)).

Fortran

- 21 • Metadirectives (see [Section 7.4](#)), assumption directives (see [Section 8.3](#)), **nothing**
22 directives (see [Section 8.4](#)), **error** directives (see [Section 8.5](#)) and loop transformation
23 constructs (see [Chapter 9](#)) were added to the list of directives that are allowed in a pure
24 procedure (see [Chapter 3](#)).
- 25 • The **allocators** construct was added to support the use of OpenMP allocators for
26 variables that are allocated by a Fortran **ALLOCATE** statement, and the application of
27 **allocate** directives to an **ALLOCATE** statement was deprecated (see [Section 6.7](#)).
- 28 • For consistency with other constructs with associated base language code, the **dispatch**
29 construct was extended to allow an optional paired **end** directive to be specified (see
30 [Section 7.6](#)).

Fortran

- 31 • To support the full range of allocators and to improve consistency with the syntax of other
32 clauses, the argument that specified the arguments of the **uses_allocators** as a
33 comma-separated list in which each list item is a *clause-argument-specification* of the form
34 *allocator[(traits)]* was deprecated (see [Section 6.8](#)).

- 1 • To improve code clarity and to reduce ambiguity in this specification, the **otherwise**
2 clause was added as a synonym for the **default** clause on metadirectives and the
3 corresponding **default** clause syntax was deprecated (see [Section 7.4.2](#)).

C / C++

- 4 • To improve overall syntax consistency and to reduce redundancy, the delimited form of the
5 **declare target** directive was deprecated (see [Section 7.8.2](#)).

C / C++

- 6 • The behavior of the **order** clause with the **concurrent** parameter was changed so that it
7 only affects whether a loop schedule is reproducible if a modifier is explicitly specified (see
8 [Section 10.4](#)).
- 9 • Support for the **allocate** and **firstprivate** clauses on the **scope** directive was
10 added (see [Section 11.2](#)).
- 11 • The **ompt_callback_work** callback work types for worksharing loop were added (see
12 [Section 11.5](#)).
- 13 • To simplify usage, the **map** clause on a **target enter data** or **target exit data**
14 construct now has a default map type that provides the same behavior as the **to** or **from** map
15 types, respectively (see [Section 13.6](#) and [Section 13.7](#)).
- 16 • The **interop** construct was updated to allow the **init** clause to accept an *interop_type* in
17 any position of the modifier list (see [Section 14.1](#)).
- 18 • The **doacross** clause was added as a synonym for the **depend** clause with the keywords
19 **source** and **sink** as *dependence-type* modifiers and the corresponding **depend** clause
20 syntax was deprecated to improve code clarity and to reduce parsing ambiguity. Also, the
21 **omp_cur_iteration** keyword was added to represent an iteration vector that refers to
22 the current logical iteration (see [Section 15.9.6](#)).

23 B.4 Version 5.0 to 5.1 Differences

- 24 • Full support of C11, C++11, C++14, C++17, C++20 and Fortran 2008 was completed (see
25 [Section 1.7](#)).
- 26 • Various changes throughout the specification were made to provide initial support of Fortran
27 2018 (see [Section 1.7](#)).
- 28 • To support device-specific ICV settings the environment variable syntax was extended to
29 support device-specific variables (see [Section 2.2](#) and [Chapter 21](#)).
- 30 • The OpenMP directive syntax was extended to include C++ attribute specifiers (see
31 [Section 3.1](#)).
- 32 • The **omp_all_memory** reserved locator was added (see [Section 3.1](#)), and the **depend**
33 clause was extended to allow its use (see [Section 15.9.5](#)).

- 1 • Support for **private** and **firstprivate** as an argument to the **default** clause in C
2 and C++ was added (see [Section 5.4.1](#)).
- 3 • Support was added so that iterators may be defined and used in a **map** clause (see
4 [Section 5.8.3](#)) or in data-motion clause on a **target update** directive (see [Section 13.9](#)).
- 5 • The **present** argument was added to the **defaultmap** clause (see [Section 5.8.6](#)).
- 6 • Support for the **align** clause on the **allocate** directive and **allocator** and **align**
7 modifiers on the **allocate** clause was added (see [Chapter 6](#)).
- 8 • The *target_device* trait set was added to the OpenMP context (see [Section 7.1](#)), and the
9 **target_device** selector set was added to context selectors (see [Section 7.2](#)).
- 10 • For C/C++, the declare variant directive was extended to support elision of preprocessed
11 code and to allow enclosed function definitions to be interpreted as variant functions (see
12 [Section 7.5](#)).
- 13 • The **declare variant** directive was extended with new clauses (**adjust_args** and
14 **append_args**) that support adjustment of the interface between the original function and
15 its variants (see [Section 7.5](#)).
- 16 • The **dispatch** construct was added to allow users to control when variant substitution
17 happens and to define additional information that can be passed as arguments to the function
18 variants (see [Section 7.6](#)).
- 19 • Support was added for indirect calls to the device version of a procedure or function in
20 **target** regions (see [Section 7.8](#)).
- 21 • Assumption directives were added to allow users to specify invariants (see [Section 8.3](#)).
- 22 • To support clarity in metadirectives, the **nothing** directive was added (see [Section 8.4](#)).
- 23 • To allow users to control the compilation process and runtime error actions, the **error**
24 directive was added (see [Section 8.5](#)).
- 25 • Loop transformation constructs were added (see [Chapter 7](#)).
- 26 • The **masked** construct was added to support restricting execution to a specific thread to
27 replace the deprecated **master** construct (see [Section 10.6](#)).
- 28 • The **scope** directive was added to support reductions without requiring a **parallel** or
29 worksharing region (see [Section 11.2](#)).
- 30 • The **grainsize** and **num_tasks** clauses for the **taskloop** construct were extended
31 with a **strict** modifier to ensure a deterministic distribution of logical iterations to tasks
32 (see [Section 12.6](#)).
- 33 • The **thread_limit** clause was added to the **target** construct to control the upper bound
34 on the number of threads in the created contention group (see [Section 13.8](#)).

- 1 • The **has_device_addr** clause was added to the **target** construct to allow access to
2 variables or array sections that already have a device address (see [Section 13.8](#)).
- 3 • The **interop** directive was added to enable portable interoperability with foreign execution
4 contexts used to implement OpenMP (see [Section 14.1](#)). Runtime routines that facilitate use
5 of **omp_interop_t** objects were also added (see [Section 18.12](#)).
- 6 • The **nowait** clause was added to the **taskwait** directive to support insertion of
7 non-blocking join operations in a task dependence graph (see [Section 15.5](#)).
- 8 • Support was added for compare-and-swap and (for C and C++) minimum and maximum
9 atomic operations through the **compare** clause. Support was also added for the specification
10 of the memory order to apply to a failed comparing atomic operation with the **fail** clause
11 (see [Section 15.8.5](#)).
- 12 • Specification of the **seq_cst** clause on a **flush** construct was allowed, with the same
13 meaning as a **flush** construct without a list and without a clause (see [Section 15.8.6](#)).
- 14 • To support inout sets, the **inoutset** argument was added to the **depend** clause (see
15 [Section 15.9.5](#)).
- 16 • The **omp_set_num_teams** and **omp_set_teams_thread_limit** runtime routines
17 were added to control the number of teams and the size of those teams on the **teams**
18 construct (see [Section 18.4.3](#) and [Section 18.4.5](#)). Additionally, the **omp_get_max_teams**
19 and **omp_get_teams_thread_limit** runtime routines were added to retrieve the
20 values that will be used in the next **teams** construct (see [Section 18.4.4](#) and [Section 18.4.6](#)).
- 21 • The **omp_target_is_accessible** runtime routine was added to test whether host
22 memory is accessible from a given device (see [Section 18.8.4](#)).
- 23 • To support asynchronous device memory management, **omp_target_memcpy_async**
24 and **omp_target_memcpy_rect_async** runtime routines were added (see
25 [Section 18.8.7](#) and [Section 18.8.8](#)).
- 26 • The **omp_get_mapped_ptr** runtime routine was added to support obtaining the device
27 pointer that is associated with a host pointer for a given device (see [Section 18.8.11](#)).
- 28 • The **omp_calloc**, **omp_realloc**, **omp_aligned_alloc** and
29 **omp_aligned_calloc** API routines were added (see [Section 18.13](#)).
- 30 • For the **omp_alloctrait_key_t** enum, the **omp_atv_serialized** value was added
31 and the **omp_atv_default** value was changed (see [Section 18.13.1](#)).
- 32 • The **omp_display_env** runtime routine was added to provide information about ICVs
33 and settings of environment variables (see [Section 18.15](#)).
- 34 • The **ompt_scope_beginend** value was added to the **ompt_scope_endpoint_t**
35 enum to indicate the coincident beginning and end of a scope (see [Section 19.4.4.11](#)).

- 1 • The `ompt_sync_region_barrier_implicit_workshare`,
2 `ompt_sync_region_barrier_implicit_parallel`, and
3 `ompt_sync_region_barrier_teams` values were added to the
4 `ompt_sync_region_t` enum (see [Section 19.4.4.14](#)).
- 5 • Values for asynchronous data transfers were added to the `ompt_target_data_op_t`
6 enum (see [Section 19.4.4.15](#)).
- 7 • The `ompt_state_wait_barrier_implementation` and
8 `ompt_state_wait_barrier_teams` values were added to the `ompt_state_t`
9 enum (see [Section 19.4.4.28](#)).
- 10 • The `ompt_callback_target_data_op_emi_t`,
11 `ompt_callback_target_emi_t`, `ompt_callback_target_map_emi_t`, and
12 `ompt_callback_target_submit_emi_t` callbacks were added to support external
13 monitoring interfaces (see [Section 19.5.2.25](#), [Section 19.5.2.26](#), [Section 19.5.2.27](#) and
14 [Section 19.5.2.28](#)).
- 15 • The `ompt_callback_error_t` type was added (see [Section 19.5.2.30](#)).
- 16 • The `OMP_PLACES` syntax was extended (see [Section 21.1.5](#)).
- 17 • The `OMP_NUM_TEAMS` and `OMP_TEAMS_THREAD_LIMIT` environment variables were
18 added to control the number and size of teams on the `teams` construct (see [Section 21.6.1](#)
19 and [Section 21.6.2](#)).

20 B.5 Version 4.5 to 5.0 Differences

- 21 • The memory model was extended to distinguish different types of flush operations according
22 to specified flush properties (see [Section 1.4.4](#)) and to define a happens before order based on
23 synchronizing flush operations (see [Section 1.4.5](#)).
- 24 • Various changes throughout the specification were made to provide initial support of C11,
25 C++11, C++14, C++17 and Fortran 2008 (see [Section 1.7](#)).
- 26 • Full support of Fortran 2003 was completed (see [Section 1.7](#)).
- 27 • The `target-offload-var` internal control variable (see [Chapter 2](#)) and the
28 `OMP_TARGET_OFFLOAD` environment variable (see [Section 21.2.9](#)) were added to support
29 runtime control of the execution of device constructs.
- 30 • Control over whether nested parallelism is enabled or disabled was integrated into the
31 `max-active-levels-var` internal control variable (see [Section 2.2](#)), the default value of which is
32 now implementation defined, unless determined according to the values of the
33 `OMP_NUM_THREADS` (see [Section 21.1.2](#)) or `OMP_PROC_BIND` (see [Section 21.1.6](#))
34 environment variables.

- 1 • Support for array shaping (see [Section 3.2.4](#)) and for array sections with non-unit strides in C
2 and C++ (see [Section 3.2.5](#)) was added to facilitate specification of discontinuous storage,
3 and the **target update** construct (see [Section 13.9](#)) and the **depend** clause (see
4 [Section 15.9.5](#)) were extended to allow the use of shape-operators (see [Section 3.2.4](#)).
- 5 • Iterators (see [Section 3.2.6](#)) were added to support expressions in a list that expand to
6 multiple expressions.
- 7 • The canonical loop form was defined for Fortran and, for all base languages, extended to
8 permit non-rectangular loop nests (see [Section 4.4.1](#)).
- 9 • The *relational-op* in the *canonical loop form* for C/C++ was extended to include **!=** (see
10 [Section 4.4.1](#)).
- 11 • To support conditional assignment to lastprivate variables, the **conditional** modifier was
12 added to the **lastprivate** clause (see [Section 5.4.5](#)).
- 13 • The **inscan** modifier for the **reduction** clause (see [Section 5.5.8](#)) and the **scan**
14 directive (see [Section 5.6](#)) were added to support inclusive and exclusive scan computations.
- 15 • To support task reductions, the **task** modifier was added to the **reduction** clause (see
16 [Section 5.5.8](#)), the **task_reduction** clause (see [Section 5.5.9](#)) was added to the
17 **taskgroup** construct (see [Section 15.4](#)), and the **in_reduction** clause (see
18 [Section 5.5.10](#)) was added to the **task** (see [Section 12.5](#)) and **target** (see [Section 13.8](#))
19 constructs.
- 20 • To support taskloop reductions, the **reduction** (see [Section 5.5.8](#)) and **in_reduction**
21 (see [Section 5.5.10](#)) clauses were added to the **taskloop** construct (see [Section 12.6](#)).
- 22 • The description of the **map** clause was modified to clarify the mapping order when multiple
23 *map-types* are specified for a variable or structure members of a variable on the same
24 construct. The **close map-type-modifier** was added as a hint for the runtime to allocate
25 memory close to the target device (see [Section 5.8.3](#)).
- 26 • The capability to map C/C++ pointer variables and to assign the address of device memory
27 that is mapped by an array section to them was added. Support for mapping of Fortran
28 pointer and allocatable variables, including pointer and allocatable components of variables,
29 was added (see [Section 5.8.3](#)).
- 30 • The **defaultmap** clause (see [Section 5.8.6](#)) was extended to allow selecting the
31 data-mapping or data-sharing attributes for any of the scalar, aggregate, pointer, or
32 allocatable classes on a per-region basis. Additionally it accepts the **none** parameter to
33 support the requirement that all variables referenced in the construct must be explicitly
34 mapped or privatized.
- 35 • The **declare mapper** directive was added to support mapping of data types with direct
36 and indirect members (see [Section 5.8.7](#)).

- 1 • Predefined memory spaces (see [Section 6.1](#)), predefined memory allocators and allocator
2 traits (see [Section 6.2](#)) and directives, clauses and API routines (see [Chapter 6](#) and
3 [Section 18.13](#)) to use them were added to support different kinds of memories.
- 4 • Metadirectives (see [Section 7.4](#)) and declare variant directives (see [Section 7.5](#)) were added
5 to support selection of directive variants and declared function variants at a call site,
6 respectively, based on compile-time traits of the enclosing context.
- 7 • Support for nested **declare target** directives was added (see [Section 7.8](#)).
- 8 • The **requires** directive (see [Section 8.2](#)) was added to support applications that require
9 implementation-specific features.
- 10 • The **teams** construct (see [Section 10.3](#)) was extended to support execution on the host
11 device without an enclosing **target** construct (see [Section 13.8](#)).
- 12 • The **loop** construct and the **order (concurrent)** clause were added to support
13 compiler optimization and parallelization of loops for which iterations may execute in any
14 order, including concurrently (see [Section 10.4](#) and [Section 11.7](#)).
- 15 • The collapse of associated loops that are imperfectly nested loops was defined for the **simd**
16 (see [Section 10.5](#)), worksharing-loop (see [Section 11.5](#)), **distribute** (see [Section 11.6](#))
17 and **taskloop** (see [Section 12.6](#)) constructs.
- 18 • The **simd** construct (see [Section 10.5](#)) was extended to accept the **if**, **nontemporal**, and
19 **order (concurrent)** clauses and to allow the use of **atomic** constructs within it.
- 20 • The default loop schedule modifier for worksharing-loop constructs without the **static**
21 schedule and the **ordered** clause was changed to **nonmonotonic** (see [Section 11.5](#)).
- 22 • The **affinity** clause was added to the **task** construct (see [Section 12.5](#)) to support hints
23 that indicate data affinity of explicit tasks.
- 24 • The **detach** clause for the **task** construct (see [Section 12.5](#)) and the
25 **omp_fulfill_event** runtime routine (see [Section 18.11.1](#)) were added to support
26 execution of detachable tasks.
- 27 • The **taskloop** construct (see [Section 12.6](#)) was added to the list of constructs that can be
28 canceled by the **cancel** construct (see [Section 16.1](#)).
- 29 • To support mutually exclusive inout sets, a **mutexinoutset** *dependence-type* was added
30 to the **depend** clause (see [Section 12.9](#) and [Section 15.9.5](#)).
- 31 • The semantics of the **use_device_ptr** clause for pointer variables was clarified and the
32 **use_device_addr** clause for using the device address of non-pointer variables inside the
33 **target data** construct was added (see [Section 13.5](#)).
- 34 • To support reverse offload, the **ancestor** modifier was added to the **device** clause for the
35 **target** construct (see [Section 13.8](#)).

- 1 • To reduce programmer effort, implicit declare target directives for some functions (C, C++,
2 Fortran) and subroutines (Fortran) were added (see [Section 13.8](#) and [Section 7.8](#)).
- 3 • The **target update** construct (see [Section 13.9](#)) was modified to allow array sections that
4 specify discontinuous storage.
- 5 • The **to** and **from** clauses on the **target update** construct (see [Section 13.9](#)), the
6 **depend** clause on task generating constructs (see [Section 15.9.5](#)), and the **map** clause (see
7 [Section 5.8.3](#)) were extended to allow any lvalue expression as a list item for C/C++.
- 8 • Lock hints were renamed to synchronization hints, and the old names were deprecated (see
9 [Section 15.1](#)).
- 10 • The **depend** clause was added to the **taskwait** construct (see [Section 15.5](#)).
- 11 • To support acquire and release semantics with weak memory ordering, the **acq_rel**,
12 **acquire**, and **release** clauses were added to the **atomic** construct (see [Section 15.8.5](#))
13 and **flush** construct (see [Section 15.8.6](#)), and the memory ordering semantics of implicit
14 flushes on various constructs and runtime routines were clarified (see [Section 15.8.7](#)).
- 15 • The **atomic** construct was extended with the **hint** clause (see [Section 15.8.5](#)).
- 16 • The **depend** clause (see [Section 15.9.5](#)) was extended to support iterators and to support
17 depend objects that can be created with the new **depobj** construct.
- 18 • New combined constructs **master taskloop**, **parallel master**,
19 **parallel master taskloop**, **master taskloop simd**
20 **parallel master taskloop simd** (see [Section 17.3](#)) were added.
- 21 • The **omp_set_nested** and **omp_get_nested** routines and the **OMP_NESTED**
22 environment variable were deprecated.
- 23 • The **omp_get_supported_active_levels** routine was added to query the number of
24 active levels of parallelism supported by the implementation (see [Section 18.2.12](#)).
- 25 • Runtime routines **omp_set_affinity_format** (see [Section 18.3.8](#)),
26 **omp_get_affinity_format** (see [Section 18.3.9](#)), **omp_set_affinity** (see
27 [Section 18.3.10](#)), and **omp_capture_affinity** (see [Section 18.3.11](#)) and environment
28 variables **OMP_DISPLAY_AFFINITY** (see [Section 21.2.4](#)) and **OMP_AFFINITY_FORMAT**
29 (see [Section 21.2.5](#)) were added to provide OpenMP runtime thread affinity information.
- 30 • The **omp_pause_resource** and **omp_pause_resource_all** runtime routines were
31 added to allow the runtime to relinquish resources used by OpenMP (see [Section 18.6.1](#) and
32 [Section 18.6.2](#)).
- 33 • The **omp_get_device_num** runtime routine (see [Section 18.7.6](#)) was added to support
34 determination of the device on which a thread is executing.
- 35 • Support for a first-party tool interface (see [Chapter 19](#)) was added.
- 36 • Support for a third-party tool interface (see [Chapter 20](#)) was added.

- Support for controlling offloading behavior with the `OMP_TARGET_OFFLOAD` environment variable was added (see [Section 21.2.9](#)).
- Stubs for Runtime Library Routines (previously Appendix A) were moved to a separate document.
- Interface Declarations (previously Appendix B) were moved to a separate document.

B.6 Version 4.0 to 4.5 Differences

- Support for several features of Fortran 2003 was added (see [Section 1.7](#)).
- The `if` clause was extended to take a *directive-name-modifier* that allows it to apply to combined constructs (see [Section 3.4](#)).
- The implicit data-sharing attribute for scalar variables in `target` regions was changed to `firstprivate` (see [Section 5.1.1](#)).
- Use of some C++ reference types was allowed in some data sharing attribute clauses (see [Section 5.4](#)).
- The `ref`, `val`, and `uval` modifiers were added to the `linear` clause (see [Section 5.4.6](#)).
- Semantics for reductions on C/C++ array sections were added and restrictions on the use of arrays and pointers in reductions were removed (see [Section 5.5.8](#)).
- Support was added to the map clauses to handle structure elements (see [Section 5.8.3](#)).
- To support unstructured data mapping for devices, the `map` clause (see [Section 5.8.3](#)) was updated and the `target enter data` (see [Section 13.6](#)) and `target exit data` (see [Section 13.7](#)) constructs were added.
- 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 7.8](#)).
- The `simdlen` clause was added to the `simd` construct (see [Section 10.5](#)) to support specification of the exact number of iterations desired per SIMD chunk.
- A parameter was added to the `ordered` clause of the worksharing-loop construct (see [Section 11.5](#)) and clauses were added to the `ordered` construct (see [Section 15.10](#)) 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 11.5](#)).
- The `priority` clause was added to the `task` construct (see [Section 12.5](#)) 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

1 priority value (see [Section 18.5.1](#)) and the `OMP_MAX_TASK_PRIORITY` environment
2 variable was added to control the maximum priority value allowed (see [Section 21.2.10](#)).

- 3 • The `taskloop` construct (see [Section 12.6](#)) was added to support nestable parallel loops
4 that create OpenMP tasks.
- 5 • To support interaction with native device implementations, the `use_device_ptr` clause
6 was added to the `target data` construct (see [Section 13.5](#)) and the `is_device_ptr`
7 clause was added to the `target` construct (see [Section 13.8](#)).
- 8 • The `nowait` and `depend` clauses were added to the `target` construct (see [Section 13.8](#))
9 to improve support for asynchronous execution of `target` regions.
- 10 • The `private`, `firstprivate` and `defaultmap` clauses were added to the `target`
11 construct (see [Section 13.8](#)).
- 12 • The `hint` clause was added to the `critical` construct (see [Section 15.2](#)).
- 13 • The `source` and `sink` dependence types were added to the `depend` clause (see
14 [Section 15.9.5](#)) to support doacross loop nests.
- 15 • To support a more complete set of device construct shortcuts, the `target parallel`,
16 `target parallel worksharing-loop`, `target parallel worksharing-loop SIMD`, and
17 `target simd` (see [Section 17.3](#)) combined constructs were added.
- 18 • Query functions for OpenMP thread affinity were added (see [Section 18.3.2](#) to
19 [Section 18.3.7](#)).
- 20 • Device memory routines were added to allow explicit allocation, deallocation, memory
21 transfers, and memory associations (see [Section 18.8](#)).
- 22 • The lock API was extended with lock routines that support storing a hint with a lock to select
23 a desired lock implementation for a lock's intended usage by the application code (see
24 [Section 18.9.2](#)).
- 25 • C/C++ Grammar (previously Appendix B) was moved to a separate document.

26 B.7 Version 3.1 to 4.0 Differences

- 27 • Various changes throughout the specification were made to provide initial support of Fortran
28 2003 (see [Section 1.7](#)).
- 29 • C/C++ array syntax was extended to support array sections (see [Section 3.2.5](#)).
- 30 • The `reduction` clause (see [Section 5.5.8](#)) was extended and the `declare reduction`
31 construct (see [Section 5.5.11](#)) was added to support user defined reductions.

- 1 • The `proc_bind` clause (see [Section 10.2.3](#)), the `OMP_PLACES` environment variable (see
2 [Section 21.1.5](#)), and the `omp_get_proc_bind` runtime routine (see [Section 18.3.1](#)) were
3 added to support thread affinity policies.
- 4 • SIMD directives were added to support SIMD parallelism (see [Section 10.5](#)).
- 5 • Implementation defined task scheduling points for untied tasks were removed (see
6 [Section 12.9](#)).
- 7 • Device directives (see [Chapter 13](#)), the `OMP_DEFAULT_DEVICE` environment variable (see
8 [Section 21.2.8](#)), and the `omp_set_default_device`, `omp_get_default_device`,
9 `omp_get_num_devices`, `omp_get_num_teams`, `omp_get_team_num`, and
10 `omp_is_initial_device` routines were added to support execution on devices.
- 11 • The `taskgroup` construct (see [Section 15.4](#)) was added to support deep task
12 synchronization.
- 13 • The `atomic` construct (see [Section 15.8.5](#)) was extended to support atomic swap with the
14 `capture` clause, to allow new atomic update and capture forms, and to support sequentially
15 consistent atomic operations with a new `seq_cst` clause.
- 16 • The `depend` clause (see [Section 15.9.5](#)) was added to support task dependences.
- 17 • The `cancel` construct (see [Section 16.1](#)), the `cancellation_point` construct (see
18 [Section 16.2](#)), the `omp_get_cancellation` runtime routine (see [Section 18.2.8](#)), and
19 the `OMP_CANCELLATION` environment variable (see [Section 21.2.6](#)) were added to support
20 the concept of cancellation.
- 21 • The `OMP_DISPLAY_ENV` environment variable (see [Section 21.7](#)) was added to display the
22 value of ICVs associated with the OpenMP environment variables.
- 23 • Examples (previously Appendix A) were moved to a separate document.

24 B.8 Version 3.0 to 3.1 Differences

- 25 • The `bind-var` ICV (see [Section 2.1](#)) and the `OMP_PROC_BIND` environment variable (see
26 [Section 21.1.6](#)) were added to support control of whether threads are bound to processors.
- 27 • Data environment restrictions were changed to allow `intent(in)` and `const`-qualified
28 types for the `firstprivate` clause (see [Section 5.4.4](#)).
- 29 • Data environment restrictions were changed to allow Fortran pointers in `firstprivate`
30 (see [Section 5.4.4](#)) and `lastprivate` (see [Section 5.4.5](#)) clauses.
- 31 • New reduction operators `min` and `max` were added for C and C++ (see [Section 5.5](#)).
- 32 • The `nthreads-var` ICV was modified to be a list of the number of threads to use at each nested
33 parallel region level, and the algorithm for determining the number of threads used in a
34 parallel region was modified to handle a list (see [Section 10.2.1](#)).

- 1 • The **final** and **mergeable** clauses (see [Section 12.5](#)) were added to the **task** construct
2 to support optimization of task data environments.
- 3 • The **taskyield** construct (see [Section 12.7](#)) was added to allow user-defined task
4 scheduling points.
- 5 • The **atomic** construct (see [Section 15.8.5](#)) was extended to include **read**, **write**, and
6 **capture** forms, and an **update** clause was added to apply the already existing form of the
7 **atomic** construct.
- 8 • The nesting restrictions in [Section 17.1](#) were clarified to disallow closely-nested OpenMP
9 regions within an **atomic** region so that an **atomic** region can be consistently defined with
10 other OpenMP regions to include all code in the **atomic** construct.
- 11 • The **omp_in_final** runtime library routine (see [Section 18.5.3](#)) was added to support
12 specialization of final task regions.
- 13 • Descriptions of examples (previously Appendix A) were expanded and clarified.
- 14 • Incorrect use of **omp_integer_kind** in Fortran interfaces was replaced with
15 **selected_int_kind(8)**.

16 B.9 Version 2.5 to 3.0 Differences

- 17 • The definition of active **parallel** region was changed so that a **parallel** region is
18 active if it is executed by a team that consists of more than one thread (see [Section 1.2.2](#)).
- 19 • The concept of tasks was added to the execution model (see [Section 1.2.5](#) and [Section 1.3](#)).
- 20 • The OpenMP memory model was extended to cover atomicity of memory accesses (see
21 [Section 1.4.1](#)). The description of the behavior of **volatile** in terms of **flush** was
22 removed.
- 23 • The definition of the *nest-var*, *dyn-var*, *nthreads-var* and *run-sched-var* internal control
24 variables (ICVs) were modified to provide one copy of these ICVs per task instead of one
25 copy for the whole program (see [Chapter 2](#)). The **omp_set_num_threads** and
26 **omp_set_dynamic** runtime library routines were specified to support their use from
27 inside a **parallel** region (see [Section 18.2.1](#) and [Section 18.2.6](#)).
- 28 • The *thread-limit-var* ICV, the **omp_get_thread_limit** runtime library routine and the
29 **OMP_THREAD_LIMIT** environment variable were added to support control of the maximum
30 number of threads (see [Section 2.1](#), [Section 18.2.11](#) and [Section 21.1.3](#)).
- 31 • The *max-active-levels-var* ICV, **omp_set_max_active_levels** and
32 **omp_get_max_active_levels** runtime library routines, and
33 **OMP_MAX_ACTIVE_LEVELS** environment variable were added to support control of the
34 number of nested active **parallel** regions (see [Section 2.1](#), [Section 18.2.13](#),
35 [Section 18.2.14](#) and [Section 21.1.4](#)).

- 1 • The *stacksize-var* ICV and the **OMP_STACKSIZE** environment variable were added to
2 support control of thread stack sizes (see [Section 2.1](#) and [Section 21.2.2](#)).
- 3 • The *wait-policy-var* ICV and the **OMP_WAIT_POLICY** environment variable were added to
4 control the desired behavior of waiting threads (see [Section 2.1](#) and [Section 21.2.3](#)).
- 5 • Predetermined data-sharing attributes were defined for Fortran assumed-size arrays (see
6 [Section 5.1.1](#)).
- 7 • Static class members variables were allowed in **threadprivate** directives (see
8 [Section 5.2](#)).
- 9 • Invocations of constructors and destructors for private and threadprivate class type variables
10 was clarified (see [Section 5.2](#), [Section 5.4.3](#), [Section 5.4.4](#), [Section 5.7.1](#) and [Section 5.7.2](#)).
- 11 • The use of Fortran allocatable arrays was allowed in **private**, **firstprivate**,
12 **lastprivate**, **reduction**, **copyin** and **copyprivate** clauses (see [Section 5.2](#),
13 [Section 5.4.3](#), [Section 5.4.4](#), [Section 5.4.5](#), [Section 5.5.8](#), [Section 5.7.1](#) and [Section 5.7.2](#)).
- 14 • Support for **firstprivate** was added to the **default** clause in Fortran (see
15 [Section 5.4.1](#)).
- 16 • Implementations were precluded from using the storage of the original list item to hold the
17 new list item on the primary thread for list items in the **private** clause, and the value was
18 made well defined on exit from the **parallel** region if no attempt is made to reference the
19 original list item inside the **parallel** region (see [Section 5.4.3](#)).
- 20 • Data environment restrictions were changed to allow **intent (in)** and **const**-qualified
21 types for the **firstprivate** clause (see [Section 5.4.4](#)).
- 22 • Data environment restrictions were changed to allow Fortran pointers in **firstprivate**
23 (see [Section 5.4.4](#)) and **lastprivate** (see [Section 5.4.5](#)).
- 24 • New reduction operators **min** and **max** were added for C and C++ (see [Section 5.5](#)).
- 25 • Determination of the number of threads in **parallel** regions was updated (see
26 [Section 10.2.1](#)).
- 27 • The assignment of iterations to threads in a loop construct with a **static** schedule kind was
28 made deterministic (see [Section 11.5](#)).
- 29 • The worksharing-loop construct was extended to support association with more than one
30 perfectly nested loop through the **collapse** clause (see [Section 11.5](#)).
- 31 • Iteration variables for worksharing-loops were allowed to be random access iterators or of
32 unsigned integer type (see [Section 11.5](#)).
- 33 • The schedule kind **auto** was added to allow the implementation to choose any possible
34 mapping of iterations in a loop construct to threads in the team (see [Section 11.5](#)).
- 35 • The **task** construct (see [Chapter 12](#)) was added to support explicit tasks.

- 1 • The **taskwait** construct (see [Section 15.5](#)) was added to support task synchronization.
- 2 • The runtime library routines **omp_set_schedule** and **omp_get_schedule** were
- 3 added to set and to retrieve the value of the *run-sched-var* ICV (see [Section 18.2.9](#) and
- 4 [Section 18.2.10](#)).
- 5 • The **omp_get_level** runtime library routine was added to return the number of nested
- 6 **parallel** regions that enclose the task that contains the call (see [Section 18.2.15](#)).
- 7 • The **omp_get_ancestor_thread_num** runtime library routine was added to return the
- 8 thread number of the ancestor of the current thread (see [Section 18.2.16](#)).
- 9 • The **omp_get_team_size** runtime library routine was added to return the size of the
- 10 thread team to which the ancestor of the current thread belongs (see [Section 18.2.17](#)).
- 11 • The **omp_get_active_level** runtime library routine was added to return the number of
- 12 active **parallel** regions that enclose the task that contains the call (see [Section 18.2.18](#)).
- 13 • Lock ownership was defined in terms of tasks instead of threads (see [Section 18.9](#)).

Index

Symbols

`_OPENMP` macro, [73](#), [636](#), [637](#), [646](#)

A

acquire flush, [32](#)
adjust_args, [199](#)
affinity, [237](#)
affinity, [274](#)
align, [179](#)
aligned, [171](#)
allocate, [180](#), [182](#)
allocator, [180](#)
allocators, [184](#)
append_args, [200](#)
apply Clause, [230](#)
array sections, [67](#)
array shaping, [66](#)
assumes, [221](#)
assumption clauses, [220](#)
assumption directives, [219](#)
at, [216](#)
atomic, [321](#)
atomic, [319](#)
atomic construct, [651](#)
attribute clauses, [111](#)
attributes, data-mapping, [150](#)
attributes, data-sharing, [99](#)
auto, [263](#)

B

barrier, [310](#)
barrier, implicit, [312](#)
base language format, [77](#)
begin declare target, [213](#)
begin declare variant, [202](#)
begin metadirective, [196](#)

begin assumes, [222](#)

bind, [268](#)

branch, [209](#)

C

cancel, [343](#)
cancellation constructs, [343](#)
 cancel, [343](#)
 cancellation point, [347](#)
cancellation point, [347](#)
canonical loop nest form, [88](#)
capture, atomic, [321](#)
clause format, [59](#)
clauses
 adjust_args, [199](#)
 affinity, [274](#)
 align, [179](#)
 aligned, [171](#)
 allocate, [182](#)
 allocator, [180](#)
 append_args, [200](#)
 apply Clause, [230](#)
 assumption, [220](#)
 at, [216](#)
 atomic, [319](#)
 attribute data-sharing, [111](#)
 bind, [268](#)
 branch, [209](#)
 collapse, [96](#)
 copyin, [146](#)
 copyprivate, [148](#)
 data copying, [146](#)
 data-sharing, [111](#)
 default, [112](#)
 defaultmap, [162](#)
 depend, [334](#)

- `destroy`, 76
- `detach`, 275
- `device`, 286
- `device_type`, 285
- `dist_schedule`, 266
- `doacross`, 337
- `enter`, 161
- `exclusive`, 146
- extended-atomic*, 319
- `filter`, 249
- `final`, 271
- `firstprivate`, 114
- `from`, 169
- full*, 227
- `grainsize`, 279
- `has_device_addr`, 124
- `hint`, 306, 308
- if Clause, 75
- `in_reduction`, 140
- `inclusive`, 145
- `indirect`, 214
- `init`, 303
- `initializer`, 132
- `interop`, 205
- `is_device_ptr`, 123
- `lastprivate`, 117
- `linear`, 120
- `link`, 162
- `local`, 174
- `map`, 152
- `match`, 198
- memory-order*, 319
- `memscope`, 320
- `mergeable`, 270
- `message`, 224
- `nocontext`, 206
- `nogroup`, 318
- `nontemporal`, 246
- `novariants`, 206
- `nowait`, 317
- `num_tasks`, 279
- `num_teams`, 243
- `num_threads`, 236
- `order`, 243
- `ordered`, 97
- `otherwise`, 195
- parallelization-level*, 342
- partial*, 228
- `permutation`, 230
- `priority`, 271
- `private`, 113
- `proc_bind`, 239
- `reduction`, 136
- requirement*, 218
- `safelen`, 247
- `safesync`, 239
- `schedule`, 262
- `severity`, 223
- `shared`, 113
- `simdlen`, 247
- `sizes`, 226
- `task_reduction`, 139
- `thread_limit`, 287
- `to`, 169
- `uniform`, 170
- `untied`, 270
- `update`, 332
- `use`, 304
- `use_device_addr`, 125
- `use_device_ptr`, 123
- `uses_allocators`, 185
- `when`, 194
- `collapse`, 96
- combined and composite directive
 - names, 353
- combined construct semantics, 354
- `compare`, `atomic`, 321
- compilation sentinels, 73, 74
- compliance, 36
- composite constructs, 355
- composition of constructs, 349
- conditional compilation, 73
- consistent loop schedules, 98
- construct syntax, 50
- constructs
 - `allocators`, 184

- atomic**, 321
- barrier**, 310
- cancel**, 343
- cancellation constructs, 343
- cancellation point**, 347
- combined constructs, 354
- composite constructs, 355
- critical**, 309
- depobj**, 333
- device constructs, 285
- dispatch**, 204
- distribute**, 264
- do**, 261
- flush**, 325
- for**, 260
- interop**, 301
- loop**, 267
- masked**, 248
- ordered**, 339–341
- parallel**, 232
- reverse**, 228
- scope**, 252
- sections**, 253
- simd**, 245
- single**, 250
- target**, 293
- target data**, 289
- target enter data**, 290
- target exit data**, 292
- target update**, 298
- task**, 272
- taskgroup**, 314
- tasking constructs, 270
- taskloop**, 276
- taskwait**, 315
- taskyield**, 280
- teams**, 240
- interchange**, 229
- tile**, 225
- unroll**, 227
- work-distribution, 250
- workshare**, 255
- worksharing, 250

- worksharing-loop construct, 258
- controlling OpenMP thread affinity, 237
- copyin**, 146
- copyprivate**, 148
- critical**, 309

D

- data copying clauses, 146
- data environment, 99
- data terminology, 14
- data-mapping control, 150
- data-motion clauses, 167
- data-sharing attribute clauses, 111
- data-sharing attribute rules, 99
- declare mapper**, 164
- declare reduction**, 141
- declare simd**, 207
- Declare Target, 209
- declare target**, 211
- declare variant**, 201
- declare variant, 197
- default**, 112
- defaultmap**, 162
- depend**, 334
- depend object, 332
- dependences, 331
- depobj**, 333
- deprecated features, 656
- destroy**, 76
- detach**, 275
- device**, 286
- device constructs
 - device constructs, 285
 - target**, 293
 - target update**, 298
- device data environments, 29, 290, 292
- device directives, 285
- device information routines, 391
- device memory routines, 396
- device_type**, 285
- directive format, 51
- directive syntax, 50
- directives
 - allocate**, 180

- assumes**, 221
- assumptions, 219
- begin assumes**, 222
- begin declare target**, 213
- begin declare variant**, 202
- begin metadirective**, 196
- declare mapper**, 164
- declare reduction**, 141
- declare simd**, 207
- Declare Target, 209
- declare target**, 211
- declare variant**, 201
- declare variant, 197
- error**, 222
- groupprivate**, 172
- memory management directives, 175
- metadirective**, 193, 195
- nothing**, 222
- requires**, 216
- scan Directive**, 143
- section**, 254
- threadprivate**, 104
- variant directives, 187
- dispatch**, 204
- dist_schedule**, 266
- distribute**, 264
- do**, 261
- doacross**, 337
- dynamic**, 262
- dynamic thread adjustment, 650

E

- enter**, 161
- environment display routine, 458
- environment variables, 628
 - OMP_AFFINITY_FORMAT**, 637
 - OMP_ALLOCATOR**, 644
 - OMP_AVAILABLE_DEVICES**, 639
 - OMP_CANCELLATION**, 639
 - OMP_DEBUG**, 644
 - OMP_DEFAULT_DEVICE**, 640
 - OMP_DISPLAY_AFFINITY**, 636
 - OMP_DISPLAY_ENV**, 646
 - OMP_DYNAMIC**, 629

- OMP_MAX_ACTIVE_LEVELS**, 630
- OMP_MAX_TASK_PRIORITY**, 641
- OMP_NUM_TEAMS**, 645
- OMP_NUM_THREADS**, 629
- OMP_PLACES**, 631
- OMP_PROC_BIND**, 633
- OMP_SCHEDULE**, 634
- OMP_STACKSIZE**, 635
- OMP_TARGET_OFFLOAD**, 641
- OMP_TEAMS_THREAD_LIMIT**, 646
- OMP_THREAD_LIMIT**, 630
- OMP_TOOL**, 642
- OMP_TOOL_LIBRARIES**, 642
- OMP_TOOL_VERBOSE_INIT**, 643
- OMP_WAIT_POLICY**, 635

- event, 426
- event callback registration, 466
- event callback signatures, 495
- event routines, 426
- exclusive**, 146
- execution model, 24
- extended-atomic*, 319

F

- features history, 656
- filter**, 249
- final**, 271
- firstprivate**, 114
- fixed source form conditional compilation
 - sentinels, 73
- fixed source form directives, 57
- flush**, 325
- flush operation, 30
- flush synchronization, 32
- flush-set, 30
- for**, 260
- frames, 491
- free source form conditional compilation
 - sentinel, 74
- free source form directives, 58
- from**, 169
- full*, 227

G

glossary, 2
grainsize, 279
groupprivate, 172
guided, 262

H

happens before, 32
has_device_addr, 124
header files, 356
hint, 308
history of features, 656

I

ICVs (internal control variables), 40
if Clause, 75
implementation, 647
implementation terminology, 21
implicit barrier, 312
implicit data-mapping attribute rules, 150
implicit flushes, 327
in_reduction, 140
include files, 356
inclusive, 145
indirect, 214
informational and utility directives, 216
init, 303
internal control variables, 647
internal control variables (ICVs), 40
interop, 205
interoperability, 301
Interoperability routines, 427
introduction, 1
is_device_ptr, 123
iterators, 70

L

lastprivate, 117
linear, 120
link, 162
list item privatization, 109
local, 174
lock routines, 415
loop, 267

loop concepts, 88
loop iteration spaces, 94
loop iteration vectors, 94
loop terminology, 9
loop-transforming constructs, 225

M

map, 152
mapper, 152
mapper identifiers, 152
masked, 248
match, 198
memory allocators, 176
memory management, 175
memory management directives
 memory management directives, 175
memory management routines, 433
memory model, 28
memory spaces, 175
memory-order, 319
memscope, 320
mergeable, 270
message, 224
metadirective, 193
metadirective, 195
modifier
 task-dependence-*typetask-dependence-*
 type, 331
modifying and retrieving ICV values, 45
modifying ICVs, 42

N

nesting of regions, 349
nocontext, 206
nogroup, 318
nontemporal, 246
normative references, 36
nothing, 222
novariants, 206
nowait, 317
num_tasks, 279
num_teams, 243
num_threads, 236

O

`OMP_AFFINITY_FORMAT`, 637
`omp_aligned_alloc`, 446
`omp_aligned_calloc`, 449
`omp_alloc`, 446
`OMP_ALLOCATOR`, 644
`OMP_AVAILABLE_DEVICES`, 639
`omp_calloc`, 449
`OMP_CANCELLATION`, 639
`omp_capture_affinity`, 381
`omp_curr_progress_width`, 232
`OMP_DEBUG`, 644
`OMP_DEFAULT_DEVICE`, 640
`omp_destroy_allocator`, 443
`omp_destroy_lock`, 419
`omp_destroy_nest_lock`, 419
`OMP_DISPLAY_AFFINITY`, 636
`omp_display_affinity`, 381
`OMP_DISPLAY_ENV`, 646
`omp_display_env`, 458
`OMP_DYNAMIC`, 629
`omp_free`, 448
`omp_fulfill_event`, 426
`omp_get_active_level`, 372
`omp_get_affinity_format`, 379
`omp_get_ancestor_thread_num`, 370
`omp_get_cancellation`, 364
`omp_get_default_allocator`, 445
`omp_get_default_device`, 393
`omp_get_device_allocator`, 441
`omp_get_device_and_host_allocator`, 441
`omp_get_device_and_host_memspace`, 437
`omp_get_device_memspace`, 437
`omp_get_device_num`, 395
`omp_get_devices_all_allocator`, 441
`omp_get_devices_all_memspace`, 437
`omp_get_devices_allocator`, 441
`omp_get_devices_and_host_allocator`, 441
`omp_get_devices_and_host_memspace`, 437
`omp_get_devices_memspace`, 437
`omp_get_dynamic`, 364
`omp_get_initial_device`, 396
`omp_get_interop_int`, 429
`omp_get_interop_name`, 431
`omp_get_interop_ptr`, 430
`omp_get_interop_rc_desc`, 433
`omp_get_interop_str`, 431
`omp_get_interop_type_desc`, 432
`omp_get_level`, 370
`omp_get_mapped_ptr`, 414
`omp_get_max_active_levels`, 369
`omp_get_max_progress_width`, 392
`omp_get_max_task_priority`, 387
`omp_get_max_teams`, 385
`omp_get_max_threads`, 361
`omp_get_memspaces_num_resources`, 453
`omp_get_num_devices`, 394
`omp_get_num_interop_properties`, 429
`omp_get_num_places`, 374
`omp_get_num_procs`, 392
`omp_get_num_teams`, 383
`omp_get_num_threads`, 360
`omp_get_partition_num_places`, 377
`omp_get_partition_place_nums`, 378
`omp_get_place_num`, 376
`omp_get_place_num_procs`, 375
`omp_get_place_proc_ids`, 376
`omp_get_proc_bind`, 373
`omp_get_schedule`, 367
`omp_get_submemspace`, 454
`omp_get_supported_active_levels`, 368
`omp_get_team_num`, 383
`omp_get_team_size`, 371
`omp_get_teams_thread_limit`, 386
`omp_get_thread_limit`, 367

omp_get_thread_num, 362
 omp_get_wtick, 425
 omp_get_wtime, 425
 omp_in_explicit_task, 388
 omp_in_final, 388
 omp_in_parallel, 362
 omp_init_allocator, 439
 omp_init_lock, 417, 418
 omp_init_nest_lock, 417, 418
 omp_is_initial_device, 395
 OMP_MAX_ACTIVE_LEVELS, 630
 OMP_MAX_TASK_PRIORITY, 641
 OMP_NUM_TEAMS, 645
 OMP_NUM_THREADS, 629
 omp_pause_resource, 389
 omp_pause_resource_all, 391
 OMP_PLACES, 631
 OMP_PROC_BIND, 633
 omp_realloc, 451
 OMP_SCHEDULE, 634
 omp_set_affinity_format, 378
 omp_set_default_allocator, 444
 omp_set_default_device, 393
 omp_set_dynamic, 363
 omp_set_lock, 420
 omp_set_max_active_levels, 368
 omp_set_nest_lock, 420
 omp_set_num_teams, 384
 omp_set_num_threads, 359
 omp_set_schedule, 365
 omp_set_teams_thread_limit, 385
 OMP_STACKSIZE, 635
 omp_target_alloc, 396
 omp_target_associate_ptr, 410
 omp_target_disassociate_ptr, 412
 omp_target_free, 398
 omp_target_is_accessible, 401
 omp_target_is_present, 400
 omp_target_memcpy, 402
 omp_target_memcpy_async, 405
 omp_target_memcpy_rect, 403
 omp_target_memcpy_rect_async, 407
 OMP_TARGET_OFFLOAD, 641
 OMP_TEAMS_THREAD_LIMIT, 646
 omp_test_lock, 423
 omp_test_nest_lock, 423
 OMP_THREAD_LIMIT, 630
 OMP_TOOL, 642
 OMP_TOOL_LIBRARIES, 642
 OMP_TOOL_VERBOSE_INIT, 643
 omp_unset_lock, 422
 omp_unset_nest_lock, 422
 OMP_WAIT_POLICY, 635
 ompd_bp_device_begin, 626
 ompd_bp_device_end, 626
 ompd_bp_parallel_begin, 620
 ompd_bp_parallel_end, 621
 ompd_bp_target_begin, 625
 ompd_bp_target_end, 625
 ompd_bp_task_begin, 623
 ompd_bp_task_end, 623
 ompd_bp_teams_begin, 622
 ompd_bp_teams_end, 622
 ompd_bp_thread_begin, 624
 ompd_bp_thread_end, 624
 ompd_callback_device_host
 _fn_t, 583
 ompd_callback_get_thread
 _context_for_thread_id
 _fn_t, 576
 ompd_callback_memory_alloc
 _fn_t, 575
 ompd_callback_memory_free
 _fn_t, 575
 ompd_callback_memory_read
 _fn_t, 580
 ompd_callback_memory_write
 _fn_t, 582
 ompd_callback_print_string
 _fn_t, 584
 ompd_callback_sizeof_fn_t, 578
 ompd_callback_symbol_addr
 _fn_t, 579
 ompd_callbacks_t, 585
 ompd_dll_locations_valid, 564

- `ompd_dll_locations`, 563
- `ompt_callback_buffer`
 - `_complete_t`, 519
- `ompt_callback_buffer`
 - `_request_t`, 519
- `ompt_callback_cancel_t`, 514
- `ompt_callback_control`
 - `_tool_t`, 529
- `ompt_callback_dependences_t`, 503
- `ompt_callback_dispatch_t`, 501
- `ompt_callback_error_t`, 530
- `ompt_callback_device`
 - `_finalize_t`, 517
- `ompt_callback_device`
 - `_initialize_t`, 515
- `ompt_callback_flush_t`, 513
- `ompt_callback_implicit`
 - `_task_t`, 506
- `ompt_callback_masked_t`, 507
- `ompt_callback_mutex`
 - `_acquire_t`, 510
- `ompt_callback_mutex_t`, 511
- `ompt_callback_nest_lock_t`, 512
- `ompt_callback_parallel`
 - `_begin_t`, 497
- `ompt_callback_parallel`
 - `_end_t`, 498
- `ompt_callback_sync`
 - `_region_t`, 508
- `ompt_callback_device`
 - `_load_t`, 517
- `ompt_callback_device`
 - `_unload_t`, 518
- `ompt_callback_target_data`
 - `_emi_op_t`, 520
- `ompt_callback_target_data`
 - `_op_t`, 520
- `ompt_callback_target_emi_t`, 523
- `ompt_callback_target`
 - `_map_emi_t`, 525
- `ompt_callback_target_map_t`, 525
- `ompt_callback_target`
 - `_submit_emi_t`, 527
- `ompt_callback_target`
 - `_submit_t`, 527
- `ompt_callback_target_t`, 523
- `ompt_callback_task`
 - `_create_t`, 502
- `ompt_callback_task`
 - `_dependence_t`, 504
- `ompt_callback_task`
 - `_schedule_t`, 505
- `ompt_callback_thread`
 - `_begin_t`, 495
- `ompt_callback_thread_end_t`, 496
- `ompt_callback_work_t`, 499
- OpenMP allocator structured blocks, 80
- OpenMP argument lists, 63
- OpenMP atomic structured blocks, 82
- OpenMP compliance, 36
- OpenMP context-specific structured blocks, 80
- OpenMP function dispatch structured blocks, 81
- OpenMP operations, 65
- OpenMP stylized expressions, 79
- OpenMP types, 77
- order**, 243
- ordered**, 97, 339–341
- otherwise**, 195

P

- parallel**, 232
- parallelism generating constructs, 232
- parallelization-level*, 342
- partial*, 228
- permutation**, 230
- priority**, 271
- private**, 113
- proc_bind**, 239

R

- read, atomic**, 321
- initializer**, 132
- reduction**, 136
- reduction clauses, 126
- release flush, 32

requirement, 218
requires, 216
reserved locators, 65
resource relinquishing routines, 389
reverse, 228
runtime, 263
runtime library definitions, 356
runtime library routines, 356

S

safelen, 247
safesync, 239
scan Directive, 143
schedule, 262
scheduling, 282
scope, 252
section, 254
sections, 253
severity, 223
shared, 113
simd, 245
simdlen, 247
Simple Lock Routines, 415
single, 250
sizes, 226
stand-alone directives, 56
static, 262
strong flush, 30
structured blocks, 79
synchronization constructs, 306
synchronization constructs and clauses, 306
synchronization hint type, 306
synchronization hints, 306
synchronization terminology, 10

T

target, 293
target data, 289
target memory routines, 396
target update, 298
task, 272
task scheduling, 282
task-dependence-type, 331
task_reduction, 139

taskgroup, 314
tasking constructs, 270
tasking routines, 387
tasking terminology, 12
taskloop, 276
taskwait, 315
taskyield, 280
teams, 240
teams region routines, 382
thread affinity, 237
thread affinity routines, 373
thread team routines, 359
thread_limit, 287
threadprivate, 104
interchange, 229
tile, 225
timer, 425
timing routines, 425
to, 169
tool control, 455
tool initialization, 463
tool interfaces definitions, 460, 563
tools header files, 460, 563
tracing device activity, 467
types
 sync_hint, 306

U

uniform, 170
unroll, 227
untied, 270
update, 332
update, atomic, 321
use, 304
use_device_addr, 125
use_device_ptr, 123
uses_allocators, 185

V

variables, environment, 628
variant directives, 187

W

wait identifier, 493

wall clock timer, [425](#)
error, [222](#)
when, [194](#)
work-distribution
 constructs, [250](#)
work-distribution constructs, [250](#)
workshare, [255](#)
worksharing
 constructs, [250](#)
worksharing constructs, [250](#)
worksharing-loop construct, [258](#)
write, atomic, [321](#)