



OpenMP Technical Report 12: Version 6.0 Preview 2

This Technical Report is the second 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 features of preview 1 with full support for C23, including C attribute syntax, and C++23. It introduces new C/C++ attributes, extensions to data mapping clauses, and new loop transformations. Support for free-agent threads, to extend support for OpenMP tasks, and the coexecute directive, to enhance device support for Fortran, were added. This preview also contains several clarifications, corrections, and refinements of the OpenMP API. See Appendix B.2 for the complete list of changes relative to version 5.2.

EDITORS

Bronis R. de Supinski

Michael Klemm

November 9, 2023

Expires November 6, 2024

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 previously.

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 2 November 2023

Copyright ©1997-2023 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 draft version includes the following internal GitHub issues (corresponding Trac ticket numbers in parentheses when they exist) applied to the 5.2 LaTeX sources: 1121 (189), 1479 (547), 1495 (563), 1584 (652), 1585 (653), 1843-1844 (911-912), 1935, 1946, 2019-2020, 2038, 2158, 2186-2187, 2653, 2691, 2721, 2734, 2736-2737, 2740, 2757, 2784, 2902, 3006-3007, 3027, 3058, 3151, 3161, 3164, 3168, 3171, 3180, 3184-3185, 3189-3190, 3193, 3199, 3201, 3206-3210, 3212, 3216-3217, 3220, 3222, 3229-3232, 3234, 3237-3238, 3240-3241, 3245, 3255, 3258-3259, 3267-3268, 3278-3279, 3281, 3285, 3290, 3293, 3296, 3301, 3303-3304, 3326-3327, 3331, 3335-3337, 3341, 3345, 3347, 3351, 3353, 3367, 3379, 3384, 3396, 3406, 3419, 3425, 3437-3441, 3449, 3452-3453, 3455, 3459-3460, 3467, 3475, 3488, 3490-3491, 3493, 3503, 3506-3509, 3512, 3514, 3516-3517, 3530, 3543, 3547, 3549-3550, 3555, 3558, 3560, 3574-3575, 3577, 3582, 3585, 3590, 3594-3595, 3601, 3609, 3612, 3615, 3640, 3645, 3647, 3654, 3657, 3662-3663, 3668, 3678, 3680, 3705, 3709

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

Contents

I	Definitions	1
1	Overview of the OpenMP API	2
1.1	Scope	2
1.2	Glossary	2
1.3	Execution Model	42
1.4	Memory Model	46
1.4.1	Structure of the OpenMP Memory Model	46
1.4.2	Device Data Environments	47
1.4.3	Memory Management	48
1.4.4	The Flush Operation	48
1.4.5	Flush Synchronization and Happens-Before Order	50
1.4.6	OpenMP Memory Consistency	52
1.5	Tool Interfaces	53
1.5.1	OMPT	53
1.5.2	OMPD	53
1.6	OpenMP Compliance	54
1.7	Normative References	54
1.8	Organization of this Document	56
2	Internal Control Variables	58
2.1	ICV Descriptions	58
2.2	ICV Initialization	60
2.3	Modifying and Retrieving ICV Values	64
2.4	How the Per-Data Environment ICVs Work	66
2.5	ICV Override Relationships	68

3	Environment Variables	69
3.1	Parallel Region Environment Variables	70
3.1.1	OMP_DYNAMIC	70
3.1.2	OMP_NUM_THREADS	70
3.1.3	OMP_THREAD_LIMIT	71
3.1.4	OMP_MAX_ACTIVE_LEVELS	71
3.1.5	OMP_PLACES	72
3.1.6	OMP_PROC_BIND	74
3.2	Program Execution Environment Variables	75
3.2.1	OMP_SCHEDULE	75
3.2.2	OMP_STACKSIZE	76
3.2.3	OMP_WAIT_POLICY	76
3.2.4	OMP_DISPLAY_AFFINITY	77
3.2.5	OMP_AFFINITY_FORMAT	78
3.2.6	OMP_CANCELLATION	80
3.2.7	OMP_AVAILABLE_DEVICES	80
3.2.8	OMP_DEFAULT_DEVICE	81
3.2.9	OMP_TARGET_OFFLOAD	81
3.2.10	OMP_THREADS_RESERVE	82
3.2.11	OMP_MAX_TASK_PRIORITY	84
3.3	OMPT Environment Variables	84
3.3.1	OMP_TOOL	84
3.3.2	OMP_TOOL_LIBRARIES	84
3.3.3	OMP_TOOL_VERBOSE_INIT	85
3.4	OMPD Environment Variables	86
3.4.1	OMP_DEBUG	86
3.5	Memory Allocation Environment Variables	87
3.5.1	OMP_ALLOCATOR	87
3.6	Teams Environment Variables	88
3.6.1	OMP_NUM_TEAMS	88
3.6.2	OMP_TEAMS_THREAD_LIMIT	88
3.7	OMP_DISPLAY_ENV	88

4	Directive and Construct Syntax	90
4.1	Directive Format	91
4.1.1	Fixed Source Form Directives	97
4.1.2	Free Source Form Directives	98
4.2	Clause Format	99
4.2.1	OpenMP Argument Lists	103
4.2.2	Reserved Locators	105
4.2.3	OpenMP Operations	106
4.2.4	Array Shaping	106
4.2.5	Array Sections	107
4.2.6	iterator Modifier	110
4.3	Conditional Compilation	112
4.3.1	Fixed Source Form Conditional Compilation Sentinels	113
4.3.2	Free Source Form Conditional Compilation Sentinel	114
4.4	<i>directive-name-modifier</i> Modifier	114
4.5	if Clause	119
4.6	destroy Clause	120
5	Base Language Formats and Restrictions	122
5.1	OpenMP Types and Identifiers	122
5.2	OpenMP Stylized Expressions	124
5.3	Structured Blocks	124
5.3.1	OpenMP Allocator Structured Blocks	125
5.3.2	OpenMP Function Dispatch Structured Blocks	126
5.3.3	OpenMP Atomic Structured Blocks	127
5.4	Loop Concepts	134
5.4.1	Canonical Loop Nest Form	134
5.4.2	OpenMP Loop-Iteration Spaces and Vectors	140
5.4.3	collapse Clause	142
5.4.4	ordered Clause	143
5.4.5	Consistent Loop Schedules	144
5.4.6	Canonical Loop Sequence Form	145
5.4.7	looprange Clause	146

II	Directives and Clauses	147
6	Data Environment	148
6.1	Data-Sharing Attribute Rules	148
6.1.1	Variables Referenced in a Construct	148
6.1.2	Variables Referenced in a Region but not in a Construct	152
6.2	threadprivate Directive	153
6.3	List Item Privatization	158
6.4	Data-Sharing Attribute Clauses	161
6.4.1	default Clause	161
6.4.2	shared Clause	162
6.4.3	private Clause	163
6.4.4	firstprivate Clause	164
6.4.5	lastprivate Clause	167
6.4.6	linear Clause	170
6.4.7	is_device_ptr Clause	173
6.4.8	use_device_ptr Clause	174
6.4.9	has_device_addr Clause	175
6.4.10	use_device_addr Clause	176
6.5	Reduction and Induction Clauses and Directives	177
6.5.1	OpenMP Reduction and Induction Identifiers	177
6.5.2	OpenMP Reduction and Induction Expressions	177
6.5.3	Implicitly Declared OpenMP Reduction Identifiers	182
6.5.4	Implicitly Declared OpenMP Induction Identifiers	183
6.5.5	Properties Common to Reduction and induction Clauses	184
6.5.6	Properties Common to All Reduction Clauses	186
6.5.7	Reduction Scoping Clauses	188
6.5.8	Reduction Participating Clauses	189
6.5.9	reduction Clause	189
6.5.10	task_reduction Clause	192
6.5.11	in_reduction Clause	193
6.5.12	induction Clause	194
6.5.13	declare reduction Directive	196
6.5.14	combiner Clause	198

6.5.15	initializer Clause	198
6.5.16	declare induction Directive	199
6.5.17	inductor Clause	201
6.5.18	collector Clause	201
6.6	scan Directive	202
6.6.1	inclusive Clause	204
6.6.2	exclusive Clause	205
6.7	Data Copying Clauses	205
6.7.1	copyin Clause	205
6.7.2	copyprivate Clause	207
6.8	Data-Mapping Control	209
6.8.1	Implicit Data-Mapping Attribute Rules	209
6.8.2	Mapper Identifiers and mapper Modifiers	211
6.8.3	map Clause	212
6.8.4	enter Clause	221
6.8.5	link Clause	222
6.8.6	defaultmap Clause	222
6.8.7	declare mapper Directive	224
6.9	Data-Motion Clauses	227
6.9.1	to Clause	229
6.9.2	from Clause	230
6.10	uniform Clause	231
6.11	aligned Clause	231
6.12	groupprivate Directive	232
6.13	local Clause	235
7	Memory Management	236
7.1	Memory Spaces	236
7.2	Memory Allocators	237
7.3	align Clause	240
7.4	allocator Clause	241
7.5	allocate Directive	242
7.6	allocate Clause	244
7.7	allocators Construct	246

7.8	uses_allocators Clause	246
8	Variant Directives	249
8.1	OpenMP Contexts	249
8.2	Context Selectors	251
8.3	Matching and Scoring Context Selectors	254
8.4	Metadirectives	255
8.4.1	when Clause	256
8.4.2	otherwise Clause	257
8.4.3	metadirective	258
8.4.4	begin metadirective	258
8.5	Declare Variant Directives	259
8.5.1	match Clause	260
8.5.2	adjust_args Clause	261
8.5.3	append_args Clause	262
8.5.4	declare variant Directive	264
8.5.5	begin declare variant Directive	265
8.6	dispatch Construct	267
8.6.1	interop Clause	268
8.6.2	novariants Clause	269
8.6.3	nocontext Clause	269
8.7	declare simd Directive	270
8.7.1	<i>branch</i> Clauses	272
8.8	Declare Target Directives	273
8.8.1	declare target Directive	275
8.8.2	begin declare target Directive	278
8.8.3	indirect Clause	279
9	Informational and Utility Directives	281
9.1	error Directive	281
9.2	at Clause	282
9.3	message Clause	283
9.4	severity Clause	283

9.5	requires Directive	284
9.5.1	<i>requirement</i> Clauses	285
9.6	Assumption Directives	291
9.6.1	<i>assumption</i> Clauses	292
9.6.2	assumes Directive	297
9.6.3	assume Directive	298
9.6.4	begin assumes Directive	298
9.7	nothing Directive	298
10	Loop-Transforming Constructs	300
10.1	tile Construct	301
10.1.1	sizes Clause	302
10.2	unroll Construct	302
10.2.1	full Clause	303
10.2.2	partial Clause	304
10.3	reverse Construct	304
10.4	interchange Construct	305
10.4.1	permutation Clause	306
10.5	fuse Construct	306
10.6	apply Clause	307
11	Parallelism Generation and Control	309
11.1	omp_curr_progress_width Identifier	309
11.2	parallel Construct	309
11.2.1	Determining the Number of Threads for a parallel Region	312
11.2.2	num_threads Clause	314
11.2.3	Controlling OpenMP Thread Affinity	315
11.2.4	proc_bind Clause	317
11.2.5	safesync Clause	318
11.3	teams Construct	319
11.3.1	num_teams Clause	322
11.4	order Clause	322
11.5	simd Construct	324
11.5.1	nontemporal Clause	325

11.5.2	safelen Clause	326
11.5.3	simdlen Clause	326
11.6	masked Construct	327
11.6.1	filter Clause	328
12	Work-Distribution Constructs	329
12.1	single Construct	330
12.2	scope Construct	331
12.3	sections Construct	332
12.3.1	section Directive	334
12.4	workshare Construct	334
12.5	coexecute Construct	337
12.6	Worksharing-Loop Constructs	339
12.6.1	for Construct	341
12.6.2	do Construct	342
12.6.3	schedule Clause	343
12.7	distribute Construct	345
12.7.1	dist_schedule Clause	347
12.8	loop Construct	348
12.8.1	bind Clause	350
13	Tasking Constructs	352
13.1	untied Clause	352
13.2	mergeable Clause	353
13.3	final Clause	353
13.4	threadset Clause	354
13.5	priority Clause	355
13.6	task Construct	355
13.6.1	affinity Clause	358
13.6.2	detach Clause	359
13.7	taskloop Construct	360
13.7.1	grainsize Clause	363
13.7.2	num_tasks Clause	364
13.8	taskyield Construct	364

13.9	Initial Task	365
13.10	Task Scheduling	366
14	Device Directives and Clauses	369
14.1	device_type Clause	369
14.2	device Clause	370
14.3	thread_limit Clause	371
14.4	Device Initialization	372
14.5	target data Construct	373
14.6	target enter data Construct	374
14.7	target exit data Construct	376
14.8	target Construct	378
14.9	target update Construct	383
15	Interoperability	386
15.1	interop Construct	386
15.1.1	OpenMP Foreign Runtime Identifiers	388
15.1.2	init Clause	388
15.1.3	use Clause	389
15.2	Interoperability Requirement Set	390
16	Synchronization Constructs and Clauses	391
16.1	Synchronization Hints	391
16.1.1	Synchronization Hint Type	391
16.1.2	hint Clause	393
16.2	critical Construct	394
16.3	Barriers	396
16.3.1	barrier Construct	396
16.3.2	Implicit Barriers	397
16.3.3	Implementation-Specific Barriers	399
16.4	taskgroup Construct	399
16.5	taskwait Construct	401
16.6	nowait Clause	403
16.7	nogroup Clause	404

16.8	OpenMP Memory Ordering	405
16.8.1	<i>memory-order</i> Clauses	405
16.8.2	<i>atomic</i> Clauses	409
16.8.3	<i>extended-atomic</i> Clauses	411
16.8.4	memscope Clause	414
16.8.5	atomic Construct	415
16.8.6	flush Construct	419
16.8.7	Implicit Flushes	421
16.9	OpenMP Dependences	425
16.9.1	<i>task-dependence-type</i> Modifier	425
16.9.2	Depend Objects	426
16.9.3	update Clause	426
16.9.4	depobj Construct	427
16.9.5	depend Clause	428
16.9.6	doacross Clause	431
16.10	ordered Construct	433
16.10.1	Stand-alone ordered Construct	434
16.10.2	Block-associated ordered Construct	435
16.10.3	<i>parallelization-level</i> Clauses	437
17	Cancellation Constructs	439
17.1	<i>cancel-directive-name</i> Clauses	439
17.2	cancel Construct	440
17.3	cancellation point Construct	444
18	Composition of Constructs	445
18.1	Nesting of Regions	445
18.2	Clauses on Combined and Composite Constructs	446
18.3	Combined and Composite Directive Names	449
18.4	Combined Construct Semantics	450
18.5	Composite Construct Semantics	451

III Runtime Library Routines 452

19 Runtime Library Routines 453

19.1	Runtime Library Definitions	454
19.2	Thread Team Routines	457
19.2.1	<code>omp_set_num_threads</code>	457
19.2.2	<code>omp_get_num_threads</code>	457
19.2.3	<code>omp_get_max_threads</code>	458
19.2.4	<code>omp_get_thread_num</code>	459
19.2.5	<code>omp_in_parallel</code>	459
19.2.6	<code>omp_set_dynamic</code>	460
19.2.7	<code>omp_get_dynamic</code>	461
19.2.8	<code>omp_get_cancellation</code>	462
19.2.9	<code>omp_set_schedule</code>	462
19.2.10	<code>omp_get_schedule</code>	464
19.2.11	<code>omp_get_thread_limit</code>	465
19.2.12	<code>omp_get_supported_active_levels</code>	465
19.2.13	<code>omp_set_max_active_levels</code>	466
19.2.14	<code>omp_get_max_active_levels</code>	467
19.2.15	<code>omp_get_level</code>	467
19.2.16	<code>omp_get_ancestor_thread_num</code>	468
19.2.17	<code>omp_get_team_size</code>	469
19.2.18	<code>omp_get_active_level</code>	470
19.3	Thread Affinity Routines	470
19.3.1	<code>omp_get_proc_bind</code>	470
19.3.2	<code>omp_get_num_places</code>	472
19.3.3	<code>omp_get_place_num_procs</code>	472
19.3.4	<code>omp_get_place_proc_ids</code>	473
19.3.5	<code>omp_get_place_num</code>	474
19.3.6	<code>omp_get_partition_num_places</code>	474
19.3.7	<code>omp_get_partition_place_nums</code>	475
19.3.8	<code>omp_set_affinity_format</code>	476
19.3.9	<code>omp_get_affinity_format</code>	477
19.3.10	<code>omp_display_affinity</code>	478

19.3.11	<code>omp_capture_affinity</code>	478
19.4	Teams Region Routines	480
19.4.1	<code>omp_get_num_teams</code>	480
19.4.2	<code>omp_get_team_num</code>	480
19.4.3	<code>omp_set_num_teams</code>	481
19.4.4	<code>omp_get_max_teams</code>	482
19.4.5	<code>omp_set_teams_thread_limit</code>	483
19.4.6	<code>omp_get_teams_thread_limit</code>	484
19.5	Tasking Routines	484
19.5.1	<code>omp_get_max_task_priority</code>	484
19.5.2	<code>omp_in_explicit_task</code>	485
19.5.3	<code>omp_in_final</code>	485
19.5.4	<code>omp_is_free_agent</code>	486
19.5.5	<code>omp_ancestor_is_free_agent</code>	487
19.6	Resource Relinquishing Routines	488
19.6.1	<code>omp_pause_resource</code>	488
19.6.2	<code>omp_pause_resource_all</code>	490
19.7	Device Information Routines	491
19.7.1	<code>omp_get_num_procs</code>	491
19.7.2	<code>omp_get_max_progress_width</code>	492
19.7.3	<code>omp_set_default_device</code>	492
19.7.4	<code>omp_get_default_device</code>	493
19.7.5	<code>omp_get_num_devices</code>	493
19.7.6	<code>omp_get_device_num</code>	494
19.7.7	<code>omp_is_initial_device</code>	495
19.7.8	<code>omp_get_initial_device</code>	495
19.8	Device Memory Routines	496
19.8.1	<code>omp_target_alloc</code>	496
19.8.2	<code>omp_target_free</code>	498
19.8.3	<code>omp_target_is_present</code>	499
19.8.4	<code>omp_target_is_accessible</code>	500
19.8.5	<code>omp_target_memcpy</code>	501
19.8.6	<code>omp_target_memcpy_rect</code>	503

19.8.7	<code>omp_target_memcpy_async</code>	505
19.8.8	<code>omp_target_memcpy_rect_async</code>	507
19.8.9	<code>omp_target_memset</code>	510
19.8.10	<code>omp_target_memset_async</code>	512
19.8.11	<code>omp_target_associate_ptr</code>	514
19.8.12	<code>omp_target_disassociate_ptr</code>	516
19.8.13	<code>omp_get_mapped_ptr</code>	518
19.9	Lock Routines	519
19.9.1	<code>omp_init_lock</code> and <code>omp_init_nest_lock</code>	521
19.9.2	<code>omp_init_lock_with_hint</code> and <code>omp_init_nest_lock_with_hint</code>	522
19.9.3	<code>omp_destroy_lock</code> and <code>omp_destroy_nest_lock</code>	523
19.9.4	<code>omp_set_lock</code> and <code>omp_set_nest_lock</code>	524
19.9.5	<code>omp_unset_lock</code> and <code>omp_unset_nest_lock</code>	526
19.9.6	<code>omp_test_lock</code> and <code>omp_test_nest_lock</code>	527
19.10	Timing Routines	529
19.10.1	<code>omp_get_wtime</code>	529
19.10.2	<code>omp_get_wtick</code>	530
19.11	Event Routine	530
19.11.1	<code>omp_fulfill_event</code>	530
19.12	Interoperability Routines	531
19.12.1	<code>omp_get_num_interop_properties</code>	532
19.12.2	<code>omp_get_interop_int</code>	533
19.12.3	<code>omp_get_interop_ptr</code>	534
19.12.4	<code>omp_get_interop_str</code>	535
19.12.5	<code>omp_get_interop_name</code>	535
19.12.6	<code>omp_get_interop_type_desc</code>	536
19.12.7	<code>omp_get_interop_rc_desc</code>	537
19.13	Memory Management Routines	538
19.13.1	Memory Management Types	538
19.13.2	Memory Space Routines	541
19.13.3	<code>omp_init_allocator</code>	544
19.13.4	Memory Allocator Routines	545

19.13.5	<code>omp_destroy_allocator</code>	548
19.13.6	<code>omp_set_default_allocator</code>	549
19.13.7	<code>omp_get_default_allocator</code>	550
19.13.8	<code>omp_alloc</code> and <code>omp_aligned_alloc</code>	550
19.13.9	<code>omp_free</code>	552
19.13.10	<code>omp_calloc</code> and <code>omp_aligned_calloc</code>	553
19.13.11	<code>omp_realloc</code>	555
19.13.12	<code>omp_get_memspace_num_resources</code>	557
19.13.13	<code>omp_get_submemspace</code>	558
19.14	Tool Control Routine	559
19.15	Environment Display Routine	562

IV Tool Interfaces 564

20 OMPT Interface 565

20.1	OMPT Interfaces Definitions	565
20.2	Activating a First-Party Tool	565
20.2.1	<code>omp_start_tool</code>	565
20.2.2	Determining Whether a First-Party Tool Should be Initialized	567
20.2.3	Initializing a First-Party Tool	568
20.2.4	Monitoring Activity on the Host with OMPT	571
20.2.5	Tracing Activity on Target Devices with OMPT	572
20.3	Finalizing a First-Party Tool	576
20.4	OMPT Data Types	576
20.4.1	Tool Initialization and Finalization	576
20.4.2	Callbacks	577
20.4.3	Tracing	578
20.4.4	Miscellaneous Type Definitions	580
20.5	OMPT Tool Callback Signatures and Trace Records	598
20.5.1	Initialization and Finalization Callback Signature	598
20.5.2	Event Callback Signatures and Trace Records	600
20.6	OMPT Runtime Entry Points for Tools	637
20.6.1	Entry Points in the OMPT Callback Interface	637

20.6.2	Entry Points in the OMPT Device Tracing Interface	654
20.6.3	Lookup Entry Points: <code>ompt_function_lookup_t</code>	665
21	OMPD Interface	667
21.1	OMPD Interfaces Definitions	668
21.2	Activating a Third-Party Tool	668
21.2.1	Enabling Runtime Support for OMPD	668
21.2.2	<code>ompd_dll_locations</code>	668
21.2.3	<code>ompd_dll_locations_valid</code>	669
21.3	OMPD Data Types	670
21.3.1	Size Type	670
21.3.2	Wait ID Type	670
21.3.3	Basic Value Types	671
21.3.4	Address Type	671
21.3.5	Frame Information Type	671
21.3.6	System Device Identifiers	672
21.3.7	Native Thread Identifiers	673
21.3.8	OMPD Handle Types	673
21.3.9	OMPD Scope Types	674
21.3.10	Team Generator Types	675
21.3.11	ICV ID Type	676
21.3.12	Tool Context Types	676
21.3.13	Return Code Types	676
21.3.14	Primitive Type Sizes	678
21.4	OMPD Third-Party Tool Callback Interface	678
21.4.1	Memory Management of OMPD Library	679
21.4.2	Context Management and Navigation	681
21.4.3	Accessing Memory in the OpenMP Program or Runtime	683
21.4.4	Data Format Conversion: <code>ompd_callback_device_host_fn_t</code>	687
21.4.5	<code>ompd_callback_print_string_fn_t</code>	688
21.4.6	The Callback Interface	689
21.5	OMPD Tool Interface Routines	691
21.5.1	Per OMPD Library Initialization and Finalization	691
21.5.2	Per OpenMP Process Initialization and Finalization	695

21.5.3	Thread and Signal Safety	698
21.5.4	Address Space Information	698
21.5.5	Thread Handles	700
21.5.6	Parallel Region Handles	705
21.5.7	Task Handles	709
21.5.8	Querying Thread States	716
21.5.9	Display Control Variables	718
21.5.10	Accessing Scope-Specific Information	720
21.6	Breakpoint Symbol Names for OMPD	724
21.6.1	Beginning Parallel Regions	724
21.6.2	Ending Parallel Regions	725
21.6.3	Beginning Teams Regions	726
21.6.4	Ending Teams Regions	726
21.6.5	Beginning Task Regions	727
21.6.6	Ending Task Regions	727
21.6.7	Beginning OpenMP Threads	728
21.6.8	Ending OpenMP Threads	728
21.6.9	Beginning Target Regions	729
21.6.10	Ending Target Regions	729
21.6.11	Initializing OpenMP Devices	730
21.6.12	Finalizing OpenMP Devices	730

V Appendices 732

A OpenMP Implementation-Defined Behaviors 733

B Features History 743

B.1	Deprecated Features	743
B.2	Version 5.2 to 6.0 Differences	743
B.3	Version 5.1 to 5.2 Differences	747
B.4	Version 5.0 to 5.1 Differences	750
B.5	Version 4.5 to 5.0 Differences	752
B.6	Version 4.0 to 4.5 Differences	756
B.7	Version 3.1 to 4.0 Differences	758

B.8	Version 3.0 to 3.1 Differences	758
B.9	Version 2.5 to 3.0 Differences	759

Index		762
--------------	--	------------

List of Figures

20.1 First-Party Tool Activation Flow Chart	567
---	-----

List of Tables

2.1	ICV Scopes and Descriptions	58
2.2	ICV Initial Values	61
2.3	Ways to Modify and to Retrieve ICV Values	64
2.4	ICV Override Relationships	68
3.1	Predefined Abstract Names for OMP_PLACES	72
3.2	Available Field Types for Formatting OpenMP Thread Affinity Information	78
3.3	Reservation Types for OMP_THREADS_RESERVE	82
4.1	Syntactic Properties for Clauses , Arguments and Modifiers	101
6.1	Implicitly Declared C/C++ Reduction Identifiers	182
6.2	Implicitly Declared Fortran Reduction Identifiers	183
6.3	Implicitly Declared C/C++ Induction Identifiers	184
6.4	Implicitly Declared Fortran Induction Identifiers	184
6.5	Map-Type Decay of Map Type Combinations	225
7.1	Predefined Memory Spaces	236
7.2	Allocator Traits	237
7.3	Predefined Allocators	239
12.1	ompt_callback_work Callback Work Types for Worksharing-Loop	340
13.1	ompt_callback_task_create Callback Flags Evaluation	357
19.1	Required Values of the ompt_interop_property_t enum Type	532
19.2	Required Values for the ompt_interop_rc_t enum Type	533
19.3	Standard Tool Control Commands	560
20.1	OMPT Callback Interface Runtime Entry Point Names and Their Type Signatures	570
20.2	Callbacks for which ompt_set_callback Must Return ompt_set_always	572
20.3	OMPT Tracing Interface Runtime Entry Point Names and Their Type Signatures	573
20.4	Association of dev1 and dev2 arguments for target data operations	628
21.1	Mapping of Scope Type and OMPD Handles	675

1

Part I

2

Definitions

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) 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

construct selector set	A selector sets that may match the construct trait set . 249 , 252–254 , 260
device selector set	A selector sets that may match the device trait set . 252–254
implementation selector set	A selector sets that may match the implementation trait set . 252–254

target_device selector set	A selector sets that may match the target device trait set . 252–254
user selector set	A selector sets that may match traits in the dynamic trait set . 252, 254
accessible device	The host device or any non-host device accessible for execution. 62, 80, 290
acquire flush	A flush that has the acquire flush property . 32, 36, 49–51, 417, 420, 422–425
acquire flush property	A flush with the acquire flush property orders memory operations that follow the flush after memory operations performed by a different thread that synchronizes with it. 3, 18, 420
active level	An active parallel region that encloses a given region at some point in the execution of an OpenMP program . The number of active levels is the number of active parallel regions that encloses the given region . 3, 36, 465, 466, 734
active parallel region	A parallel region comprised of implicit tasks that are being executed by a team to which multiple threads are assigned. 3, 38, 58, 59, 74, 154, 155, 460, 466, 469, 733
active target region	A target region that is executed on a device other than the device that encountered the target construct . 67
address range	The addresses of a contiguous set of storage locations . 13, 18, 25, 29, 35, 501
address space	A collection of logical, virtual, or physical memory address ranges that contain code, stack, and/or data. Address ranges within an address space need not be contiguous. An address space consists of one or more segments . 3, 18, 28, 33, 40, 289, 501, 567, 568, 676, 681, 682, 684, 702
address space context	A tool context that refers to an address space within an OpenMP process . 676
address space handle	A handle that refers to an address space within an OpenMP process . 675, 705
affected loop nest	The subset of canonical loop nests of an associated loop sequence that are selected by the looprange clause . 146, 300, 307
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 . 3, 15, 19, 30, 34, 39, 41, 46, 105, 155, 223, 359, 733
all tasks	All tasks participating in the OpenMP program . 8, 189, 233, 238
all threads	All OpenMP threads participating in the OpenMP program . A specific usage of the term may be explicitly limited to all threads on a given device or OpenMP thread pool . 3, 8, 47, 52, 169, 415
allocator	A memory allocator . 3, 237–243, 245–247, 287, 381
allocator trait	A trait of an allocator . 237–239
ancestor thread	For a given thread , its parent thread or one of the ancestor threads of its parent thread . 3, 468, 469, 487, 747

array element	A single member of an array as defined by the base language . 4, 184, 204, 205
array item	An array, an array section , or an array element . 448
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. 4, 6, 7, 12, 26, 34, 81, 104, 107–109, 174–176, 178, 179, 181, 184, 185, 190, 191, 195, 204, 205, 213, 214, 217, 218, 220, 225, 227, 429, 430
assigned list item	A list item to which assignment is performed as the result of a data-motion clause . 228–230
assigned thread	A thread that has been assigned an implicit task of a parallel region . 30, 37, 38, 42, 43, 459
associated device	The associated device of a memory allocator is the device that is specified when the memory allocator is created; If the associated memory space is a predefined memory space , the associated device is the current device . 4, 46
associated iteration	A logical iteration of the associated loops of a loop-nest-associated directive . 33, 303, 339
associated iteration space	The logical iteration space of the associated loops of a loop-nest-associated directive . 340, 347
associated loop	A loop from a canonical loop nest or a DO CONCURRENT loop in Fortran that is controlled by a given loop-nest-associated directive . 4, 10, 22–24, 33, 41, 96, 140–144, 149–151, 163, 168, 171, 190, 203, 299–301, 303–305, 349, 360, 363, 364, 434
associated loop sequence	The associated canonical loop sequence of a loop-sequence-associated directive . 3, 146, 300
associated memory space	The associated memory space of a memory allocator is the memory space that is specified when the memory allocator is created. 4, 26, 237, 239
assumed-size array	For C/C++, an array section for which the number of array elements is assumed. For Fortran, an assumed-size array in the base language . 4, 42, 107, 109, 150, 151, 160, 174, 176, 212, 213, 218, 219
assumption directive	A directive that provides invariants that specify additional information about the expected properties of the program that can optionally be used for optimization. An implementation may ignore this information without altering the behavior of the program. 4, 291, 294
assumption scope	The scope for which the invariants specified by an assumption directive must hold. 291–298
async signal safe	The guarantee that interruption by signal delivery will not interfere with a set of operations. An async signal safe runtime entry point is safe to call from a signal handler . 4, 600, 624, 642, 643, 645, 646, 649, 651–653
atomic captured update	An atomic update operation that is specified by an atomic construct on which the capture clause is present. 131, 412, 416

atomic conditional update	An atomic update operation that is specified by an atomic construct on which the compare clause is present. 129, 412, 413, 416–419
atomic operation	An operation that is specified by an atomic construct or is implicitly performed by the OpenMP implementation and that atomically accesses and/or modifies a specific storage location . 5, 31–33, 47, 49–52, 215, 216, 239, 391, 417–419, 423
atomic read	An atomic operation that is specified by an atomic construct on which the read clause is present. 128, 410, 416
atomic scope	The set of threads that may concurrently access or modify a given storage location with atomic operations , where at least one of the operations modifies the storage location . 47, 51, 239, 415
atomic update	An atomic operation that is specified by an atomic construct on which the update clause is present. 4, 5, 129, 410, 412, 416, 417, 419
atomic write	An atomic operation that is specified by an atomic construct on which the write clause is present. 129, 411, 416
attach-ineligible	A pointer variable for which pointer attachment may not be performed. 214
attached pointer	A pointer variable in a device data environment that, as a result of a mapping operation , becomes the base pointer of a given data entity that also exists in the device data environment . 30, 216, 220, 227, 228, 381
barrier	A point in the execution of a program encountered by a team , beyond which no thread in the team may execute until all threads in the team have reached the barrier and all explicit tasks generated for execution by the team have executed to completion. If cancellation has been requested, threads may proceed to the end of the canceled region even if some threads in the team have not reached the barrier . 5, 18, 20, 43–45, 207, 310, 327, 329–335, 339, 346, 366, 367, 396, 398, 399, 403, 417, 421–423, 441, 595
base address	If a data entity has a base pointer , the address of the first storage location of the implicit array of its base pointer ; otherwise, if the data entity has a base variable , the address of the first storage location of its base variable ; otherwise, the address of the first storage location of the data entity. 18, 174, 176, 213

base array	<p>For C/C++, a containing array of a given lvalue expression or array section that does not appear in the expression of any of its other containing arrays.</p> <p>For Fortran, a containing array of a given variable or array section that does not appear in the designator of any of its other containing arrays.</p> <p>COMMENT: For the array section $(*p0).x0[k1].p1->p2[k2].x1[k3].x2[4][0:n]$, where identifiers pi have a pointer type declaration and identifiers xi have an array type declaration, the base array is: $(*p0).x0[k1].p1->p2[k2].x1[k3].x2$.</p> <p>6, 448</p>
base expression	<p>The base array of a given array section or array element, if it exists; otherwise, the base pointer of the array section or array element.</p> <p>COMMENT: For the array section $(*p0).x0[k1].p1->p2[k2].x1[k3].x2[4][0:n]$, where identifiers pi have a pointer type declaration and identifiers xi have an array type declaration, the base expression is: $(*p0).x0[k1].p1->p2[k2].x1[k3].x2$.</p> <p>More examples for C/C++:</p> <ul style="list-style-type: none"> • The base expression for $x[i]$ and for $x[i:n]$ is x, if x is an array or pointer. • The base expression for $x[5][i]$ and for $x[5][i:n]$ is x, if x is a pointer to an array or x is 2-dimensional array. • The base expression for $y[5][i]$ and for $y[5][i:n]$ is $y[5]$, if y is an array of pointers or y is a pointer to a pointer. <p>Examples for Fortran:</p> <ul style="list-style-type: none"> • The base expression for $x(i)$ and for $x(i:j)$ is x. <p>6, 108, 109, 175, 176, 185, 210, 213, 214</p>
base function	<p>A function that is declared and defined in the base language. 14, 32, 41, 252, 253, 259–266</p>
base language	<p>A programming language that serves as the foundation of the OpenMP specification.</p> <p>Section 1.7 lists the current base languages for the OpenMP API.</p> <p>2, 4, 6–8, 16, 19, 28, 30, 31, 33, 35, 36, 42, 45, 46, 51, 54–56, 90, 93, 94, 97, 98, 105, 107, 108, 110, 122–124, 128, 134, 139, 140, 153, 159, 176, 177, 185, 186, 195, 197, 200, 211, 214, 225, 226, 240–242, 246, 247, 261, 264, 266, 291, 336, 388, 416, 436, 733</p>
base language thread	<p>A thread of execution that defines a single flow of control within the program and that may execute concurrently with other base language threads, as specified by the base language. 6, 45</p>

base pointer

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

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

COMMENT: For the [array section](#)

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

`(*p0).x0[k1].p1->p2.`

[5–7](#), [13](#), [26](#), [150](#), [176](#), [191](#), [195](#), [214–216](#), [218](#), [219](#), [379](#), [447](#), [448](#)

base program

A program written in a [base language](#). [28](#), [122](#)

base variable

For a given data entity that is a [variable](#) or [array section](#), a [variable](#) denoted by a [base language](#) identifier that is either the data entity or is a [containing array](#) or [containing structure](#) of the data entity.

COMMENT:

Examples for C/C++:

- The data entities x , $x[i]$, $x[:n]$, $x[i].y[j]$ and $x[i].y[:n]$, where x and y have array type declarations, all have the [base variable](#) x .
- The lvalue expressions and [array sections](#) $p[i]$, $p[:n]$, $p[i].y[j]$ and $p[i].y[:n]$, where p has a pointer type and $p[i].y$ has an array type, has a [base pointer](#) p but does not have a [base variable](#).

Examples for Fortran:

- The data objects x , $x(i)$, $x(:n)$, $x(i)\%y(j)$ and $x(i)\%y(:n)$, where x and y have array type declarations, all have the [base variable](#) x .
- The data objects $p(i)$, $p(:n)$, $p(i)\%y(j)$ and $p(i)\%y(:n)$, where p has a pointer type and $p(i)\%y$ has an array type, has a [base pointer](#) p but does not have a [base variable](#).
- For the associated pointer p , p is both its [base variable](#) and [base pointer](#).

[5](#), [7](#), [155](#), [176](#), [209](#), [210](#), [219](#), [380](#), [447](#), [448](#)

binding implicit task

The [implicit task](#) of the current [team](#) assigned to the [encountering thread](#). [8](#), [20](#), [66](#), [315](#)

binding region	The enclosing region that determines the execution context and limits the scope of the effects of the bound region is called the binding region . The binding region is not defined for regions for which the binding thread set is all threads or the encountering thread , nor is it defined for regions for which the binding task set is all tasks . 8, 29, 44, 144, 337, 348–350, 396, 433, 436, 440, 444, 468, 476, 477
binding task set	The set of tasks that are affected by, or provide the context for, the execution of a region . The binding task set for a given region can be all tasks , the current team tasks , all tasks in the contention group , all tasks of the current team that are generated in the region , the binding implicit task , or the generating task . 8, 64, 267, 373, 374, 376, 378, 383, 387, 399, 404, 466, 486, 487, 511, 513
binding thread set	The set of threads that are affected by, or provide the context for, the execution of a region . The binding thread set for a given region can be all threads on a specified set of devices, all threads that are executing tasks in a contention group , all primary threads that are executing the initial tasks of an enclosing teams region , the current team , or the encountering thread . 8, 29, 41, 44, 166, 169, 309, 319, 323, 324, 327, 329–332, 334, 337–339, 345, 348–350, 352, 356, 360, 361, 394, 396, 401, 404, 415–417, 420, 427, 434, 435, 440, 441, 444, 468, 469, 476, 477, 746
bounds-independent loop	For a structured block sequence , an enclosed canonical loop nest where none of its loops have loop bounds that depend on the execution of a preceding executable statement in the sequence. 145
C pointer	For C/C++, a base language pointer variable . For Fortran, a variable of type <code>C_PTR</code> . 16, 174
callback	A tool callback . 8, 32, 53, 54, 187, 218, 275, 281, 311, 320, 328, 330, 332–334, 336, 338, 340, 346, 357, 361, 362, 372, 373, 375, 377, 380, 384, 395, 397–400, 402, 418, 421, 430, 433, 435, 442, 512, 561, 562, 566, 571, 573, 576, 580, 581, 667, 681
callback dispatch	Callback dispatch processes a registered callback when an associated event occurs in a manner consistent with the return code provided when a first-party tool registered the callback . 8, 581, 659
callback registration	Callback registration provides a tool callback to an OpenMP implementation to enable callback dispatch . 8, 32, 569, 571
cancellable construct	A construct that has the cancellable property . 8, 439, 440, 444
cancellable property	The property that a construct is a cancellable construct . 8, 309, 332, 341, 342, 399, 439
cancellation	An action that cancels (that is, aborts) a region and causes executing implicit tasks or explicit tasks to proceed to the end of the canceled region . 5, 9, 45, 329, 396–398, 422, 425, 439–444

cancellation point	A point at which implicit tasks and explicit tasks check if cancellation has been requested. If cancellation has been observed, they perform the cancellation . 40 , 45 , 59 , 396 , 398 , 422 , 425 , 440–444
candidate	A replacement candidate . 255 , 259
canonical frame address	An address associated with a procedure frame on a call stack that was the value of the stack pointer immediately prior to calling the procedure for which the frame represents the invocation. 597
canonical loop nest	A loop nest that complies with the rules and restrictions defined in Section 5.4.1 . 3 , 4 , 8 , 9 , 17 , 19 , 22–24 , 95 , 134–136 , 139 , 140 , 142 , 145 , 146 , 168 , 202 , 299 , 300 , 303 , 307 , 344
canonical loop sequence	A sequence of canonical loop nests that complies with the rules and restrictions defined in Section 5.4.6 . 4 , 19 , 23 , 24 , 95 , 135 , 145 , 146 , 300 , 744 , 746
child task	A task is a <i>child-task</i> of its generating task region . The region of a child task is not part of its generating task region . 9 , 15 , 18 , 34 , 37 , 401 , 423
chunk	A contiguous non-empty subset of the collapsed iterations of a loop-collapsing construct . 339 , 343–346 , 348 , 360 , 451
class type	For C++, variables declared with one of the class , struct , or union keywords. 155 , 159 , 160 , 165 , 166 , 168 , 169 , 182 , 186 , 191 , 206–208 , 217 , 219 , 381
clause	A mechanism to specify customized directive behavior. xix , 3–5 , 9 , 10 , 12–15 , 17 , 24 , 26 , 27 , 30–33 , 43 , 45 , 46 , 48 , 59 , 62 , 65 , 67–69 , 90 , 91 , 93 , 94 , 99–106 , 109–112 , 120–122 , 140–144 , 146 , 148–152 , 154 , 155 , 158–166 , 168 , 169 , 171–177 , 181 , 184–186 , 188–233 , 235 , 236 , 241–247 , 250 , 252 , 253 , 255–296 , 299–309 , 312–315 , 318 , 319 , 321–324 , 326–335 , 339 , 343–356 , 359–361 , 363 , 364 , 369–371 , 373–384 , 387–391 , 393 , 395 , 401–423 , 425–430 , 432–439 , 441–443 , 446–448 , 451 , 470 , 514 , 744–746 , 748 , 749 , 757
clause group	A clause set for which restrictions or properties related to their use on all directives are specified. 272 , 285 , 292 , 405 , 409 , 411 , 437 , 439 , 746
clause set	A set of clauses for which restrictions on their use or other properties of their use on a given directive are specified. 9 , 148 , 285 , 292 , 361
clause-list trait	A trait that is defined with properties that match the clauses that may be specified for a given directive . 249 , 250 , 252
closely nested construct	A construct nested inside another construct with no other construct nested between them. 336 , 338 , 350 , 442–444
closely nested region	A region nested inside another region with no parallel region nested between them. 29 , 194 , 329 , 351 , 442 , 444
code block	A contiguous region of memory that contains code of an OpenMP program to be executed on a device . 372
collapsed iteration	A logical iteration of the collapsed loops of a loop-collapsing construct . 9 , 10 , 22 , 33 , 41 , 158 , 171 , 172 , 182 , 195 , 202–204 , 323 , 324 , 327 , 339 , 340 , 343–346 , 348 , 349 , 360 , 423 , 436 , 451

collapsed iteration space	The logical iteration space of the collapsed loops of a loop-collapsing construct . 142, 203, 326, 343, 348
collapsed logical iteration	A collapsed iteration . 142, 158
collapsed loop	For a loop-collapsing construct , the outermost associated loop or one that is controlled by the collapse clause. 9, 10, 23, 142, 158, 171, 324, 325, 339, 344–346, 348, 349, 361
collective step expression	An expression in terms of a step expression and a collector that eliminates recursive calculation in an induction operation . 10, 22, 182
collector	A binary operator used to eliminate recursion in an induction operation . 10, 22, 202
collector expression	A OpenMP stylized expression that evaluates to the value of the collective step expression of a collapsed iteration . 21, 182–184, 200, 202
combined construct	A construct that corresponds to a combined directive . 10, 11, 22, 34, 120, 190, 249, 292, 319, 321, 323, 436, 446–448
combined directive	A directive that is a shortcut for specifying one directive immediately nested inside another directive . A combined directive is semantically identical to explicitly specifying the first directive containing one instance of the second directive and no other statements. 10, 11, 101, 292, 447, 449
combined target construct	A combined construct that is composed of a target construct along with another construct . 209, 210, 448
combiner expression	An OpenMP stylized expression that specifies how a reduction combines partial results into a single value. 31, 178, 179, 185, 186, 198, 203
compatible context selector	The context selector that matches the OpenMP context in which a directive is encountered. 254–256, 259
compatible map type	A map type that is consistent with data-motion attribute of a given data-motion clause . 227, 229, 230
compilation unit	For C/C++, a translation unit. For Fortran, a program unit. 15, 48, 95, 156, 157, 221, 234, 242, 243, 245, 284–286, 291, 297, 381
compile-time error termination	Error termination preformed during compilation. 45, 285, 314
compliant implementation	An implementation of the OpenMP specification that compiles and executes any conforming program as defined by the specification. A compliant implementation may exhibit unspecified behavior when compiling or executing a non-conforming program . 2, 10, 14, 20, 40, 44, 54, 56, 76, 77, 90, 344, 417, 667
composite construct	A construct that corresponds to a composite directive . 11, 22, 34, 120, 190, 202, 249, 292, 319, 321, 436, 446, 447, 451

composite directive	A directive that is composed of two (or more) directives but does not have identical semantics to specifying one of the directives immediately nested inside the other. A composite directive either adds semantics not included in the directives from which it is composed or provides an effective nesting of the one directives inside the other that would otherwise be non-conforming. 10, 11, 101, 292, 447, 449
conforming device number	A device number that may be used in a conforming program . 46, 237, 370
conforming program	An OpenMP program that follows all rules and restrictions of the OpenMP specification. 2, 10, 11, 27, 28, 40, 42, 54, 255, 300, 344
constituent construct	For a given combined construct or composite construct , a construct from which it, or any one of its constituent constructs , is composed. 11, 22, 34, 120, 190, 191, 447
constituent directive	For a given combined directive or composite directive , a construct from which it, or any one of its constituent directives , is composed. 11, 101
construct	An executable directive and its paired end directive (if any) and the associated structured block (if any) not including the code in any called procedures . That is, the lexical extent of an executable directive . 2–5, 8–12, 14–30, 32–46, 54, 59, 60, 63, 65–68, 74, 91, 94, 96, 103–105, 111, 120–122, 130, 131, 141, 143, 144, 148–152, 154, 155, 158, 159, 161–163, 165, 166, 168, 169, 171, 173–177, 186, 188–191, 193–195, 202, 203, 207, 209, 210, 212–219, 223–225, 227, 241, 245–247, 249, 262, 263, 267–271, 286–288, 292, 293, 295, 296, 301, 303, 305, 307–312, 319–324, 327, 330, 331, 333–339, 341–343, 345–356, 359–362, 364, 369–384, 386, 387, 390, 391, 393–397, 399–418, 420–430, 432–437, 439–448, 451, 514, 566, 595, 598, 675, 710, 746, 748, 751, 757
construct trait set	The trait set that consists of all enclosing constructs at a given point in an OpenMP program up to a target construct . 2, 13, 249, 250, 252, 254, 270

containing array	<p>For C/C++, a non-subscripted array (a containing array) to which a series of zero or more array subscript operators and/or . (dot) operators are applied to yield a given lvalue expression or array section for which storage is contained by the array.</p> <p>For Fortran, an array (a containing array) without the POINTER attribute and without a subscript list to which a series of zero or more array subscript operators and/or component selectors are applied to yield a given variable or array section for which storage is contained by the array.</p> <p>COMMENT: An array is a containing array of itself. For the array section $(*p0).x0[k1].p1->p2[k2].x1[k3].x2[4][0:n]$, where identifiers p_i have a pointer type declaration and identifiers x_i have an array type declaration, the containing arrays are: $(*p0).x0[k1].p1->p2[k2].x1$ and $(*p0).x0[k1].p1->p2[k2].x1[k3].x2$.</p> <p>6, 7, 12, 106, 215, 218, 219</p>
containing structure	<p>For C/C++, a structure to which a series of zero or more . (dot) operators and/or array subscript operators are applied to yield a given lvalue expression or array section for which storage is contained by the structure.</p> <p>For Fortran, a structure to which a series of zero or more component selectors and/or array subscript selectors are applied to yield a given variable or array section for which storage is contained by the structure.</p> <p>COMMENT: A structure is a containing structure of itself.</p> <p>For C/C++, a structure pointer p to which the $->$ operator applies is equivalent to the application of a . (dot) operator to $(*p)$ for the purposes of determining containing structures.</p> <p>For the array section $(*p0).x0[k1].p1->p2[k2].x1[k3].x2[4][0:n]$, where identifiers p_i have a pointer type declaration and identifiers x_i have an array type declaration, the containing structures are: $(*p0).x0[k1].p1$, $(*p0).x0[k1].p1.p2[k2]$ and $(*p0).x0[k1].p1.p2[k2].x1[k3]$</p> <p>7, 12, 215, 218, 219</p>
contention group	<p>All implicit tasks and their descendent tasks that are generated in an implicit parallel region, R, and in all nested regions for which R is the innermost enclosing implicit parallel region. 8, 23, 28, 33, 35, 42–45, 59, 60, 71, 82, 233, 238, 289, 309, 313, 318, 371, 394, 415</p>
context selector	<p>The specification of an OpenMP context in which a construct is encountered for use in clauses and modifiers. 10, 17, 35, 251–256, 259–261, 265, 266, 284</p>
context-matching construct	<p>A construct that has the context-matching property. 252</p>

context-matching property	The property that a directive adds a trait of the same name to the construct trait set of the current OpenMP context . 12, 267, 309, 319, 324, 341, 342, 378
corresponding base pointer initialization	For a given data entity that has a base pointer , an assignment to the base pointer such that any lexical reference to the data entity or a subobject of the data entity in a target region refers to its corresponding data entity or subobject in the device data environment . 216, 379
corresponding list item	A list item in a device data environment that corresponds to an original list item . 13, 24, 176, 212, 215–217, 219–221, 227–229, 274, 291, 378, 383, 745
corresponding pointer	A corresponding list item for which the an original list item may be used as a base pointer . 29, 215, 220
corresponding storage	An address range in a device data environment that corresponds to, but may be distinct from, an address range in the device data environments of the encountering device . 13, 25, 30, 33, 174, 213, 214, 216, 217, 219, 228
corresponding storage block	A storage block that is used as corresponding storage . 47, 48, 215, 216
current device	The device on which the current task is executing. 20, 47, 49, 58, 370
current task	For a given thread , the task corresponding to the task region that it is executing. 13, 17, 20, 212, 262, 399, 401, 460, 466, 487
current task region	The region that corresponds to the current task . 44, 324, 396, 401, 440, 441
current team	All threads in the team executing the innermost enclosing parallel region . 8, 29, 33, 38, 60, 152, 324, 327, 328, 330–332, 334, 339, 354, 396, 399, 401, 434, 435, 440, 444, 469, 595
current team tasks	All tasks encountered by the corresponding team . The implicit tasks constituting the parallel region and any descendent tasks encountered during the execution of these implicit tasks are included in this set of tasks . 8, 238
data environment	The variables associated with the execution of a given region . 13–15, 20, 25–27, 29, 37, 43, 45, 47, 48, 58, 64, 66, 67, 148, 193, 207, 208, 212, 227, 326, 356, 359, 360, 373, 374, 376, 378, 383
data-environment attribute	A data-sharing attribute or a data-mapping attribute . 13, 148
data-environment attribute clause	A clause that explicitly determines the data-environment attributes of the list items in its list argument. 148, 224
data-mapping attribute	The relationship of an entity in a given device data environment to the version of that entity in the data environment of the enclosing context . 13, 18, 21, 148, 151, 209, 210, 223
data-mapping attribute clause	A clause that explicitly determines the data-mapping attributes of the list items in its list argument. 14, 18, 27, 47, 148, 209, 221, 373, 374, 376, 378

data-mapping construct	A construct that has the data-mapping property . 150
data-mapping property	The property of a construct on which a data-mapping attribute clause may be specified. 14, 373, 374, 376, 378
data-motion attribute	The data-movement relationship between a given device data environment and the version of that entity in the data environment of the enclosing context . 10, 227
data-motion clause	A clause that specifies data movement between a device set that is specified by the construct on which it appears. 4, 10, 211, 225, 227–230, 383
data-sharing attribute	The relationship of an entity in a given data environment to the version of that entity in the enclosing context . 13, 14, 18, 21, 30, 148, 150–153, 161, 210, 223, 374, 376, 378, 383
data-sharing attribute clause	A clause that explicitly determines the data-sharing attributes of the list items in its list argument. 18, 148, 150, 151, 158–161, 163, 177, 349, 360, 378, 380
declarative directive	A directive that may only be placed in a declarative context and results in one or more declarations only; it is not associated with the immediate execution of any user code or implementation code . 14, 93, 94, 97, 103, 153, 196, 199, 224, 232, 242, 256, 264, 265, 270, 275, 278, 292
declare target directive	A declarative directive that ensures that procedures and/or variables can be executed or accessed on a device . 25, 27, 47, 178, 233, 249, 273–276, 278, 279, 285, 290, 291
declare variant directive	A declarative directive that declare a function variant for a given base function . 249, 259, 260, 265, 266
declare-target property	The property that a directive applies to procedures and/or variables to ensure that they can be executed or accessed on a device . 275, 278
defined	For variables , the property of having a valid value. For C, for the contents of variables , the property of having a valid value. For C++, for the contents of variables of POD (plain old data) type, the property of having a valid value. For variables of non-POD class type, the property of having been constructed but not subsequently destructed. For Fortran, for the contents of variables , the property of having a valid value. For the allocation or association status of variables , the property of having a valid status. COMMENT: Programs that rely upon variables that are not defined are non-conforming programs . 14, 40, 72, 73, 226
dependence	An ordering relation between two instances of executable code that must be enforced by a compliant implementation . 16, 17, 37, 425–428, 430, 432, 434, 513

dependent task	A task that because of a task dependence cannot be executed until its predecessor tasks have completed. 30, 37, 367, 401, 402, 423–425, 428–430, 513
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 fully specified in the version of the specification immediately prior to the one in which it is first deprecated . In most cases, a new feature replaces the deprecated feature. Unless otherwise specified, whether any modifications provided by the replacement feature apply to the deprecated feature is implementation defined . 15, 196, 733, 743, 747–749, 751, 755
descendent task	A task that is the child task of a task region or of a region that corresponds to one of its descendent tasks . 12, 13, 15, 361, 367, 423, 441
detachable task	An explicit task that only completes after an associated event variable that represents an <i>allow-completion event</i> is fulfilled and execution of the associated structured block has completed. 356, 359, 423, 424
device	An implementation-defined logical execution engine. COMMENT: A device could have one or more processors . 3, 4, 9, 13–18, 20, 21, 26–30, 36, 40, 42, 43, 46–48, 58, 59, 67, 80, 81, 175, 209, 212, 217, 221, 227, 235, 238–240, 249, 250, 252, 254, 262, 273, 274, 289, 290, 369, 372, 375, 377–379, 381, 384, 388, 389, 415, 460, 488, 496, 498, 501, 511, 513, 543, 547, 573, 627, 637, 684, 733, 743, 747
device address	An address of an object that may be referenced on a target device . 16, 47, 173–175, 289, 290, 733
device construct	A construct that has the device property . 2, 15, 16, 36, 217, 285, 288–291, 370
device data environment	The initial data environment associated with a device . 5, 13, 14, 16, 24, 25, 30, 47, 48, 67, 148, 173–176, 193, 209, 212–217, 219–221, 227, 228, 274, 290, 373, 374, 376, 378, 381–383, 510–513, 747
device global requirement property	The property that a requirement clause indicates requirements for the behavior of device constructs that a program requires the implementation to support across all compilation units . 285
device local variable	A variable with static storage duration that is replicated for each device by the OpenMP implementation. Its name provides access to a different block of storage for each device . A variable that is part of an aggregate variable cannot be made a device local variable independently of the other components, except for static data members of C++ classes. If a variable is made a device local variable , its components are also device local variables . 15, 47, 149, 218, 235, 273, 274, 290, 733

device number	A number that the OpenMP implementation assigns to a device or otherwise may be used in an OpenMP program to refer to a device . 11, 46, 58, 59, 62, 63, 240, 370, 378, 511, 513, 627
device pointer	An implementation defined handle that refers to a device address and is represented by a C pointer . 47, 173, 174, 262, 289, 390, 733
device procedure	A function (for C/C++ and Fortran) or subroutine (for Fortran) that can be executed on a target device , as part of a target region . 36, 222, 274, 285, 288–291
device property	The property of a construct that accepts the device clause. 15, 275, 278, 373, 374, 376, 378, 383, 386
device trait set	The trait set that consists of traits that define the characteristics of the device being targeted by the compiler at that point in the OpenMP program . 2, 249, 250
device-affecting construct	A construct that has the device-affecting property . 380
device-affecting property	The property that a device construct can modify the state of the device data environment of a specified target device . 16, 373, 374, 376, 378, 383
device-specific environment variable	An alternative OpenMP environment variable that controls of the behavior of the program only with respect to a particular device or set of devices . 62, 63
directive	A base language mechanism to specify OpenMP program behavior. 2, 4, 9–11, 13, 14, 16–18, 22, 24–26, 28, 31–33, 36, 40, 42, 45–48, 52, 54, 56, 59, 69, 90–103, 105–107, 109, 122, 125–130, 136, 139–144, 146, 148, 149, 151–153, 155–158, 161, 162, 168, 171, 172, 178, 185, 186, 190, 196–200, 202–205, 209, 211, 213, 215, 216, 221–226, 233, 234, 236, 238, 242–245, 247, 249, 250, 252, 253, 255–258, 264–268, 270, 271, 274, 276–286, 288–292, 297–300, 303, 308, 310, 312, 314, 315, 321, 323, 324, 334, 336, 349, 352, 359, 360, 370, 375, 377–381, 383, 387, 389–391, 395, 402, 403, 405, 409, 417, 421–424, 439, 443, 449, 745, 746, 748–751
directive variant	A directive specification that can be used in a metadirective . 32, 255–258
divergent threads	Two threads that have reached different points in user code or otherwise have reached a common point via calls from different points in user code. 31, 45
doacross dependence	A dependence between executable code corresponding to stand-alone ordered regions from two doacross iterations : the sink iteration and the source iteration , where the source iteration precedes the sink iteration in the doacross iteration space . The doacross dependence is fulfilled when the executable code from the source iteration has completed. 16, 34, 425, 432, 434
doacross iteration	A logical iteration of a doacross loop nest . 16, 17, 34, 424, 425, 432, 434
doacross iteration space	The logical iteration space of a doacross loop nest . 16, 432

doacross logical iteration	A doacross iteration . 432
doacross loop nest	A canonical loop nest that has cross-iteration dependences between its logical iterations as specified by the use of stand-alone ordered constructs , such that executable code from a logical iteration is dependent on the executable code of one or more earlier logical iterations . COMMENT: The argument of the ordered clause on a worksharing-loop construct identifies the loops associated with the doacross loop nest . 16, 17, 432, 434, 757
dynamic context selector	Any context selector that is not a static context selector . 266
dynamic replacement candidate	A replacement candidate that may be selected at run time to replace a given metadirective . 255, 256, 259
dynamic trait set	The trait set that consists of traits that define the dynamic properties of an OpenMP program at a given point in its execution. 3, 249, 250, 252
enclosing context	For C/C++, the innermost scope enclosing a directive . For Fortran, the innermost scoping unit enclosing a directive . 13, 14, 29, 151, 152, 195, 197, 200, 208, 255, 269, 270, 333, 336, 338, 346, 347
encountering device	For a given construct , the device on which the encountering task of the construct executes. 13, 25, 29, 229, 230
encountering task	For a given region , the current task of the encountering thread . 17, 37, 45, 227, 263, 281, 310, 319, 320, 340, 354, 359, 361, 373, 387, 397, 398, 402, 403, 440–442, 469
encountering thread	For a given region , the thread that encounters the corresponding construct . 7, 8, 17, 21, 32, 43, 44, 309, 310, 315, 316, 318, 319, 349, 350, 356, 378, 387, 420, 427, 468, 469, 474, 476, 477, 486, 487, 747
ending address	The address of the last storage location of a list item or, for a mapped variable of its original list item . 18, 25, 213
environment variable	Unless specifically stated otherwise, an OpenMP environment variable . 62
error termination event	A fatal action performed in response to an error. 10, 33, 45, 314, 745 A point of interest in the execution of a thread . 8, 15, 37, 39, 53, 54, 187, 217, 218, 274, 275, 281, 310, 311, 320, 327, 328, 330–334, 336, 338, 340, 346, 356, 359, 361, 372, 373, 375, 377, 379, 380, 384, 395–402, 417, 418, 421, 423, 424, 430, 433–435, 442, 511, 512, 514, 561, 565, 568, 569, 571, 581, 641, 667, 668
exception-aborting directive	A directive that has the exception-aborting property . 295, 735
exception-aborting property	For C++, the property of a directive to be implementation defined whether an exceptions is caught or results in a runtime error termination . 17, 90, 378

exclusive scan computation	A scan computation for which the value read does not include the updates performed in the same logical iteration . 203
executable directive	A directive that appears in an executable context and results in implementation code and/or prescribes the manner in which associated user code must execute. 11, 24, 36, 42, 90, 93, 94, 125, 136, 246, 255, 267, 281, 282, 301, 302, 304–306, 309, 319, 324, 327, 330–332, 334, 337, 341, 342, 345, 348, 355, 360, 364, 373, 374, 376, 378, 383, 386, 394, 396, 399, 401, 415, 419, 427, 434, 435, 440, 444
explicit barrier	A barrier that is specified by a barrier construct . 396
explicit region	A region that corresponds to either a construct of the same name or a library routine call that explicitly appears in the program. 35, 42, 90, 338, 653
explicit task	A task that is not an implicit task . 5, 8, 9, 15, 18, 19, 29, 33, 37, 44–46, 59, 190, 191, 310, 315, 352, 356, 360–362, 366, 396, 424, 444
explicit task region	A region that corresponds to an explicit task . 32, 47, 163, 356
explicitly determined data-mapping attribute	A data-mapping attribute that is determined due to the presence of a list item on a data-mapping attribute clause . 209
explicitly determined data-sharing attribute	A data-sharing attribute that is determined due to the presence of a list item on a data-sharing attribute clause . 148, 151, 162
extended address range	The address range that starts from the minimum of the starting address and the base address and ends with maximum of the ending address and the base address of an original list item . 25, 213
extension trait	A trait that is implementation defined . 249, 250
final task	A task that forces all of its child tasks to become final tasks and included tasks . 18, 59, 352, 354, 357, 359
first-party tool	A tool that executes in the address space of the program that it is monitoring. 8, 27, 28, 53, 562, 565, 567
flush	An operation that a thread performs to enforce consistency between its view and the view of any other threads of memory . 3, 18, 20, 32, 35, 39, 45, 48–52, 329, 391, 415, 420–422
flush property	A property that determines the manner in which a flush enforces memory consistency . Any flush has one or more of the following: the strong flush property , the release flush property , and the acquire flush property . 50
flush-set	The set of variables upon which a strong flush operates. 49
foreign execution context	A context that is instantiated from a foreign runtime environment in order to facilitate execution on a given device . 18, 387, 388, 751
foreign runtime environment	A runtime environment that exists outside the OpenMP runtime with which the OpenMP implementation may interoperate. 18, 386
foreign task	An instance of executable code that is executed in a foreign execution context . 387, 388

frame	A storage area on the stack of a thread that is associated with a procedure invocation. A frame includes space for one or more saved registers and often also includes space for saved arguments, local variables, and padding for alignment. 9 , 19 , 596 , 597 , 649
free-agent thread	An unassigned thread on which an explicit task is scheduled for execution or a primary thread for an explicit parallel region that was a free-agent thread when it encountered the parallel construct. 19 , 32 , 36 , 59 , 82 , 83 , 315 , 367 , 486 , 487 , 737 , 744 , 747
function variant	A definition of a function that may be used as an alternative to the base language definition. 14 , 32 , 41 , 249 , 259–265 , 267–269
generated loop	A loop that is generated by a loop-transforming construct and is one of the resulting loops that replace the construct . 136 , 140 , 143 , 300 , 301 , 303 , 307 , 308
generated loop nest	A canonical loop nest that is generated by a loop-transforming construct . 300
generated loop sequence	A canonical loop sequence that is generated by a loop-transforming construct . 300
generating task	For a given region , the task for which execution by a thread generated the region . 8 , 19 , 66 , 67 , 267 , 356 , 373 , 374 , 376 , 378 , 383 , 387 , 424 , 466 , 486 , 487 , 511 , 513 , 710
generating task region	For a given region , the region that corresponds to its generating task . 9 , 21 , 26 , 40 , 710 , 711
global	A program aspect such as a scope that covers the whole OpenMP program . 20 , 58 , 62 , 243
groupprivate variable	A variable that is replicated, one instance per a specified group of tasks , by the OpenMP implementation. Its name provides access to a different block of storage for each specified group. A variable that is part of an aggregate variable cannot be made a groupprivate variable independently of the other components, except for static data members of C++ classes. If a variable is made a groupprivate variable , its components are also groupprivate variables with respect to the same group. 19 , 149 , 218 , 233 , 234 , 274 , 276 , 278 , 339 , 379
handle	An opaque reference that uniquely identifies an abstraction. 3 , 16 , 26 , 29 , 37 , 41 , 219 , 237 , 388 , 389 , 646 , 700 , 702 , 703
happens before	For an event <i>A</i> to happen before an event <i>B</i> , <i>A</i> must precede <i>B</i> in happens-before order . 51
happens-before order	An asymmetric relation that is consistent with simply happens-before order and, for C/C++, the “happens before” order defined by the base language . 19 , 239 , 290
hardware thread	An indivisible hardware execution unit on which only one OpenMP thread can execute at a time. 31 , 72 , 73 , 309
host address	An address of an object that may be referenced on the host device . 20 , 290

host device	The device on which the OpenMP program begins execution. 3, 19, 21, 27, 36, 43, 44, 46, 48, 63, 81, 216, 239, 250, 289, 319, 369, 373, 375–377, 380, 381, 384
host pointer	A pointer that refers to a host address . 289, 290
ICV	Acronym form for internal control variable . 20, 28, 33, 58, 60, 62–69, 71, 74, 76, 78, 80–83, 154, 241, 252, 267, 287, 312, 315, 316, 318, 319, 322, 340, 344, 355, 356, 360, 370, 371, 374, 376, 378, 383, 422, 425, 440, 441, 459, 466, 475, 567, 568
ICV scope	A context that contains one copy of a given ICV and defines the extent in which the ICV controls program behavior; the ICV scope may be the OpenMP program (i.e., global), the current device , the binding implicit task , or the data environment of the current task . 20, 58, 62, 64, 66, 67, 374, 376, 378, 383
idle thread	An unassigned thread that is not currently executing any task . 366, 595
implementation code	Implicit code that is introduced by the OpenMP implementation. 14, 18, 32, 34, 596
implementation defined	Behavior that must be documented by the implementation, and is allowed to vary among different compliant implementations . An implementation is allowed to define it as unspecified behavior . 15–18, 36, 40, 45–47, 54, 62, 67, 71–73, 76, 77, 83, 90, 91, 97–99, 142, 153, 155, 173, 175, 232, 236, 237, 239, 240, 250, 253, 254, 256, 259, 260, 264, 270, 273, 281, 283, 284, 303, 304, 313–317, 319, 322, 324, 330, 333, 340, 344, 346, 361, 371, 386, 388, 389, 391, 393, 417, 466, 476, 477, 561, 571, 573, 627, 733–738
implementation trait set	The trait set that consists of traits that describe the functionality supported by the OpenMP implementation at that point in the OpenMP program . 2, 249, 250
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> . For Fortran, the set of array elements for a given array pointer. COMMENT: For C/C++, the implicit array for pointer p with type <i>T</i> (*)[10] consists of all accessible elements p[i][j], for all <i>i</i> and <i>j</i> =0,1,...,9. 5, 219
implicit barrier	A barrier that is specified as part of the semantics of a construct other than the barrier construct . 337, 397–399, 403, 441
implicit flush	A flush that is specified as part of the semantics of a construct other than the flush construct . 423
implicit parallel region	An inactive parallel region that is not generated from a parallel construct . Implicit parallel regions surround the whole OpenMP program , all target regions , and all teams regions . 12, 21, 22, 33, 42–44, 233, 315, 321, 350, 675

implicit task	A task generated by an implicit parallel region or generated when a parallel construct is encountered during execution. 3, 4, 7–9, 12, 13, 18, 21, 22, 28–30, 35, 37, 38, 42, 43, 47, 58, 60, 66, 67, 152, 165, 189, 190, 205, 207, 208, 310, 311, 315, 316, 318, 329–340, 346, 421, 422, 424, 444, 475
implicit task region	A region that corresponds to an implicit task . 42, 67
implicitly determined data-mapping attribute	A data-mapping attribute that applies to an entity for which no data-mapping attribute is otherwise determined. 209, 216, 223
implicitly determined data-sharing attribute	A data-sharing attribute that applies to an entity for which no data-sharing attribute is otherwise determined. 148, 151, 160, 161, 209, 211, 223
inactive parallel region	A parallel region comprised of one implicit task and, thus, is being executed by a team comprised of only its primary thread . 21, 469
inactive target region	A target region that is executed on the same device that encountered the target construct . 67, 216
included task	A task for which execution is sequentially included in the generating task region . That is, an included task is an underrferred task and executed by the encountering thread . 18, 21, 26, 32, 46, 352, 356, 374, 376, 378, 383, 387, 401, 403, 511
inclusive scan computation	A scan computation for which the value read includes the updates performed in the same logical iteration . 202
indirect device invocation	An indirect call to the device version of a procedure on a device other than the host device , through a function pointer (C/C++), a pointer to a member function (C++) or a procedure pointer (Fortran) that refers to the host version of the procedure . 279
induction expression	A collector expression or a inductor expression . 177, 178
induction operation	A recurrence operation that expresses the value of a variable as a function, the inductor , applied to its previous value and a step expression . For an induction operation performed on a loop on the induction variable x and a loop-invariant step expression s , $x_i = x_{i-1} \oplus s$, $i > 0$, where x_i is the value of x at the start of collapsed iteration i , x_0 is the value of x before any tasks enter the loop, and the binary operator \oplus is the inductor . For some inductors , the induction operation can be expressed in a non-recursive closed form as $x_i = x_0 \oplus s_i = x_0 \oplus (s \otimes i)$ where $s_i = s \otimes i$. The expression s_i is the collective step expression of iteration i and the binary operator \otimes is the collector . 10, 22, 35, 40, 177, 181, 195, 202
induction variable	A variable for which an induction operation determines its values. 22, 181, 199, 200
inductor	A binary operator used by an induction operation . 22, 181

inductor expression	An OpenMP stylized expression that specifies how an induction operation determines a new value of an induction variable from its previous value and a step expression . 21, 181, 183–186, 195, 200, 201
informational directive	A <i>directive</i> that is neither declarative nor executable, but otherwise conveys user code properties to the compiler. 93, 281, 284, 292, 297, 298
initial task	An implicit task associated with an implicit parallel region . 8, 22, 33, 43, 44, 67, 190, 315, 320, 338, 346, 371, 379, 424
initial task region	A region that corresponds to an initial task . 42, 43, 58, 59, 422, 424, 460
initial team	The team that comprises an initial thread executing an implicit parallel region . 37, 43, 59, 319, 346, 348
initial thread	The thread that executes an implicit parallel region . 22, 29, 30, 39, 42, 43, 74, 76, 154, 319, 320, 337, 345, 346, 350, 422, 424, 585, 734
initializer expression	An OpenMP stylized expression that determines the initializer for the private copies of reduction list items . 31, 179–182, 185, 186, 199, 203
input phase	The portion of a logical iteration that contains all computations that update a list item for which a scan computation is performed. 40, 202, 203
internal control variable	A conceptual variable that specifies runtime behavior of a set of threads or tasks in an OpenMP program . 20, 58
interoperability requirement set	A logical set of properties of each task to which directives add or remove and that other constructs that have interoperability semantics can query. 262, 263, 267, 403, 404
intervening code	For two consecutive associated loops in a canonical loop nest , user code that appears inside the loop body of the outer associated loop but outside the loop body of the inner associated loop . 30, 136, 142
iteration count	The number of times that the loop body of a given loop is executed. 140–142, 360
leaf construct	For a given combined construct or composite construct , a constituent construct that is not itself a combined construct or composite construct . 292, 436, 446–448
league	The set of teams formed by a teams construct , each of which is associated with a different contention group . 37, 43, 59, 190, 319, 320, 347, 348
lexicographic order	The total order of two logical iteration vectors $\omega_a = (i_1, \dots, i_n)$ 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 $k \in \{1, \dots, m - 1\}$. 301
list	A comma-separated set. 13, 14, 23, 30, 148, 156, 186, 196, 199, 227, 278
list item	A member of a list . 4, 13, 14, 17, 18, 22, 25, 27, 29, 34, 148–150, 155–163, 165, 166, 168, 169, 171–176, 178, 179, 181–195, 202–210, 212–222, 225–228, 232–235, 262, 263, 267, 268, 274–278, 360, 373, 374, 376, 378–383, 420, 421, 426, 427, 441, 442

logical iteration	An instance of the executed loop body of a canonical loop nest , denoted by a number in the logical iteration space of the loops that indicates the order in which the logical iteration would be executed relative to the other logical iterations in a sequential execution. 4 , 9 , 16–18 , 21–23 , 40 , 142 , 144 , 190 , 299 , 300 , 303 , 305 , 307 , 360–364 , 749
logical iteration space	For a canonical loop nest , the sequence $0, \dots, N - 1$ where N is the number of distinct logical iterations . 4 , 10 , 16 , 23 , 142
logical iteration vector	An n -tuple (i_1, \dots, i_n) that identifies a logical iteration of a canonical loop nest , where n is the loop nest depth and i_k is the logical iteration number of the k^{th} loop, from outermost to innermost. 23 , 31 , 301
logical iteration vector space	The set of logical iteration vectors that each correspond to a logical iteration of a canonical loop nest . 144 , 301
loop body	A structured block that encompasses the executable statements that are iteratively executed by a loop statement. 22 , 23 , 136
loop iteration variable	A variable that determines the iteration space of a loop. 23 , 140 , 141 , 149–151 , 168 , 171 , 300 , 361 , 432
loop nest depth	For a canonical loop nest , the maximal number of loops, including the outermost loop, that can be associated with a loop-nest-associated directive . 23 , 140
loop sequence length	For a canonical loop sequence , the number of consecutive canonical loop nests regardless of their nesting into blocks. 145 , 146
loop-collapsing construct	A loop-nest-associated construct for which some number of outer associated loops may be collapsed loops . 9 , 10 , 158 , 171 , 323
loop-iteration vector	An n -tuple (i_1, \dots, i_n) that identifies a logical iteration of the associated loops of a loop-nest-associated directive , where n is the number of associated loops and i_k is the value of the loop iteration variable of the k^{th} associated loop , from outermost to innermost. 23 , 140 , 141 , 432
loop-iteration vector space	The set of loop-iteration vectors that each correspond to a logical iteration of the associated loops of a loop-nest-associated directive . 140 , 141
loop-nest-associated construct	A loop-nest-associated directive and its associated loops . 23 , 41 , 96 , 144 , 432
loop-nest-associated directive	An executable directive for which the associated user code must be a canonical loop nest . 4 , 23 , 24 , 33 , 94–96 , 136 , 140 , 150 , 171 , 195 , 300 , 301 , 436
loop-sequence-associated construct	A loop-sequence-associated directive and its associated loops . 24 , 146
loop-sequence-associated directive	An executable directive for which the associated user code must be a canonical loop sequence . 4 , 24 , 94 , 95 , 300

loop-sequence-transforming construct	A loop-sequence-associated construct with the loop-transforming property . 300
loop-transforming construct	A loop-transforming directive and its associated loops . 19, 135, 136, 140, 145, 299, 300, 308
loop-transforming directive	A directive with the loop-transforming property . 24, 300
loop-transforming property	The property that a construct is replaced by the loops that result from applying the transformation as defined by its directive to its associated loops . 24, 298, 301, 302, 304–306
loosely structured block	A block of zero or more executable constructs (including OpenMP constructs), where the first executable construct (if any) is not a Fortran BLOCK construct, with a single entry at the top and a single exit at the bottom. 35, 95
map-entering clause	A map clause that, if it appears on a map-entering construct , specifies that the reference count of corresponding list items is increased and, as a result, may enter the device data environment . 24, 213, 215, 217, 291, 375
map-entering construct	A construct that has the map-entering property . 24, 213, 215, 217, 219
map-entering property	A property of a construct that a map-entering clause may appear on it. 24, 213, 373, 374, 378
map-exiting clause	A map clause that, if it appears on a map-exiting construct , specifies that the reference count of corresponding list items is decreased and, as a result, may exit the device data environment . 24, 213, 377
map-exiting construct	A construct that has the map-exiting property . 24, 216
map-exiting property	A property of a construct that a map-exiting clause may appear on it. 24, 213, 373, 376, 378
map-type decay	The process that determines the final map-type of each mapping operation that results from mapping a variable with a user-defined mapper . 214, 225
map-type modifier	A modifier that has the map-type-modifying property . 214
map-type-modifying property	A modifier with the map-type-modifying property modifies the behavior of the map-type of a mapping operation . 24, 25, 214
mappable storage block	A contiguous address range in memory that contains a set of mapped list items . 215, 216, 219, 228

mappable type	<p>A type that is valid for a mapped variable. If a type is composed from other types (such as the type of an array element or a structure element) and any of the other types are not mappable types then the type is not a mappable type.</p> <p>For C, the type must be a complete type.</p> <p>For C++, the type must be a complete type; in addition, for class types:</p> <ul style="list-style-type: none"> • All member functions accessed in any target region must appear in a declare target directive. <p>For Fortran, no restrictions on the type except that for derived types:</p> <ul style="list-style-type: none"> • All type-bound procedures accessed in any target region must appear in a declare target directive. <p>COMMENT: Pointer types are mappable types but the memory block to which the pointer refers is not mapped.</p> <p>25, 219, 221, 222, 228</p>
mapped address range	The address range that starts from the starting address and ends with the ending address of an original list item . 25, 213
mapped variable	An original variable in a data environment with a corresponding variable in a device data environment . The original and corresponding variables may share storage. 17, 25, 34, 381, 382
mapper	An operation that defines how variables of given type are to be mapped or updated with respect to a device data environment . 40, 122, 175, 209, 211, 214, 219, 220, 224–230
mapping operation	An operation that establishes or removes a correspondence between a variable in one data environment and another variable in a device data environment . 5, 24, 25, 33, 47, 215–217, 291, 745
mapping-only construct	A construct that establishes correspondences between the data environment of the encountering device but otherwise does not affect the associated structured block (if any). 25, 216
mapping-only property	The property that a construct is a mapping-only construct . 373, 374, 376
matchable candidate	A mapped variable for which corresponding storage was created in a device data environment . 25, 213
matched candidate	A matchable candidate for which its mapped address range or its extended address range corresponds to the address range of the original list item . 174, 213, 219
memory	A storage resource to store and to retrieve variable accessible by threads . 3, 9, 18, 25, 26, 32, 35, 36, 38, 39, 46–49, 52, 59, 105, 106, 169, 235–240, 289, 290, 405–409, 415, 420, 429, 496, 510–513, 556, 686
memory allocator	An OpenMP object that fulfills requests to allocate and to deallocate memory for program variables from the storage resources of its associated memory space . 3, 4, 48, 59, 219, 237–246, 287, 381, 556, 747

memory space	A representation of storage resources from which memory can be allocated or deallocated. More than one memory space may exist. 4, 26, 36, 48, 219, 236, 239, 248, 543, 747
mergeable task	A task that may be a merged task if it is an undelayed task or an included task . 36, 353, 357, 387, 401
merged task	A task for which the data environment , inclusive of ICVs, is the same as that of its generating task region . 26, 357
metadirective	A directive that conditionally resolves to another directive . 16, 17, 32, 93, 255–258, 292, 749
modifier	A mechanism to specify customized clause behavior. <i>xix</i> , 12, 24, 25, 68, 100–103, 105, 110, 112, 168, 169, 171, 186, 203, 211, 213, 214, 218, 219, 227, 228, 232, 247, 248, 262, 308, 339, 344, 346, 387, 389, 426, 433, 448, 748, 749
mutually exclusive tasks	Tasks that may be executed in any order, but not at the same time. 367, 429
name-list trait	A trait that is defined with properties that match the names that identify a particular instances of the trait that are effective at a given point in an OpenMP program . 249–251, 253
named pointer	For C/C++, the base pointer of a given lvalue expression or array section , or the base pointer of one of its named pointers . For Fortran, the base pointer of a given variable or array section , or the base pointer of one of its named pointers . COMMENT: For the array section $(*p0).x0[k1].p1 \rightarrow p2[k2].x1[k3].x2[4][0:n]$, where identifiers p_i have a pointer type declaration and identifiers x_i have an array type declaration, the named pointers are: $p0$, $(*p0).x0[k1].p1$, and $(*p0).x0[k1].p1 \rightarrow p2$. 26, 106
native thread	An execution entity upon which an OpenMP thread may be implemented. 26, 28, 31, 39, 42, 44, 60, 76, 77, 310, 311, 320, 323, 585, 595, 596, 600, 601, 637, 673, 676, 684, 698, 701–705
native thread context	A tool context that refers to a native thread . 676, 681, 682, 684, 686, 687, 690
native thread handle	A handle that refers to a native thread . 675, 700–705
native thread identifier	An identifier for a native thread defined by a native thread implementation. 79, 673, 681, 682, 690, 697, 698, 701, 702, 704
native trace record	A trace record for an OpenMP device that is in a device-specific format. 574
nested construct	A construct (lexically) enclosed by another construct . 449
nested region	A region (dynamically) enclosed by another region . That is, a region generated from the execution of another region or one of its nested regions . 12, 27, 29, 42, 329

new list item	An instance of a list item created for the data environment of the construct on which a privatization clause or a data-mapping attribute clause specified. 30, 40, 158, 159, 163, 165, 166, 168, 171, 173, 174, 195, 203, 214, 215, 217
non-conforming program	An OpenMP program that is not a conforming program . 10, 14, 40, 426, 498
non-host declare target directive	A declare target directive that does not specify a device_type clause with host . 274
non-host device	A device that is not the host device . 3, 36, 46, 59, 62, 63, 289, 351, 369, 381
non-null pointer	A pointer that is not NULL . 498, 533–535, 566, 568, 573, 598, 599
non-null value	A value that is not NULL . 556, 576, 648, 650, 664, 669, 684–687, 717
non-property trait	A trait that is specified without additional properties . 249, 250, 253
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. 139, 140, 143, 144, 301, 345, 348, 363, 364
non-sequentially consistent atomic construct	An atomic construct for which the seq_cst clause is not specified 52
NULL	A null pointer. For C and C++, the value NULL or the value nullptr . For Fortran, the value C_NULL_PTR . 27, 85, 262, 477–479, 496, 498, 500, 504, 509, 511, 513, 517, 518, 531, 534–536, 551, 553, 555, 556, 561, 566, 568, 573, 602, 604, 605, 608, 611–616, 618–622, 627, 629, 632, 635, 636, 640, 641, 646–650, 666, 669, 670, 684, 686, 687, 689, 718, 741
OMPD	An interface that helps a third-party tool inspect the OpenMP state of a program that has begun execution. 27, 39, 42, 53, 54, 59, 667, 676, 681, 682, 684, 690, 702
OMPD library	A dynamically loadable library that implements the OMPD interface. 667, 698
OMPT	An interface that helps a first-party tool monitor the execution of an OpenMP program . 42, 53, 397, 565–568, 571, 598, 599
OMPT active	An OMPT interface state in which the OpenMP implementation is prepared to accept runtime calls from a first-party tool and will dispatch any registered callbacks and in which a first-party tool can invoke runtime entry points if not otherwise restricted. 561, 568
OMPT inactive	An OMPT interface state in which the OpenMP implementation will not make any callbacks and in which a first-party tool cannot invoke runtime entry points . 561, 567, 568, 598
OMPT interface state	A state that indicates the permitted interactions between a first-party tool and the OpenMP implementation. 27, 28, 561, 567, 568, 598
OMPT pending	An OMPT interface state in which the OpenMP implementation can only call functions to initialize a first-party tool and in which a first-party tool cannot invoke runtime entry points . 567, 568

OpenMP Additional Definitions document	A document that exists outside of the OpenMP specification and defines additional values that may be used in a conforming program . The OpenMP Additional Definitions document is available via https://www.openmp.org/specifications/ . 28, 80, 250, 386, 388, 454, 456
OpenMP API routine	A runtime library routine that is defined by the OpenMP implementation and that can be called from user code via the OpenMP API. 32, 58, 289, 290, 296
OpenMP architecture	The architecture on which a region executes. 28, 567
OpenMP context	The execution context of an OpenMP program , including the active constructs , the execution devices , OpenMP functionality supported by the implementation and any available dynamic values as represented by a set of traits . 10, 12, 13, 35, 249, 251, 252, 254–256, 259–261, 264, 266, 270, 284
OpenMP environment variable	A variable that is part of the runtime environment in which an OpenMP program executes and that a user may set to control the behavior of the program, typically through the initialization of an ICV . 16, 17, 58, 63
OpenMP process	A collection of one or more native threads and address spaces . An OpenMP process may contain native threads and address spaces for multiple OpenMP architectures . At least one native thread in an OpenMP process is mapped to an OpenMP thread . An OpenMP process may be live or a core file. 3, 28, 671, 676, 684
OpenMP program	A program that consists of a base program that is annotated with OpenMP directives or that calls OpenMP API runtime library routines. 3, 9, 11, 16, 17, 19–22, 26–28, 32, 39, 40, 42, 44–48, 52, 53, 58, 60, 69, 122, 152, 155, 161, 171, 188, 221, 225, 226, 236, 237, 249, 250, 256, 290, 299, 321, 339, 346, 355, 381, 382, 391, 394, 418–420, 426, 498, 561, 562, 565, 567, 568, 596, 597, 667, 669, 733
OpenMP stylized expression	A base language expression that is subject to restrictions that enable its use within an OpenMP implementation. 10, 22, 177
OpenMP thread	A logical execution entity with a stack and associated thread-specific memory subject to the semantics and constraints of this specification and may be implemented upon a native thread . 3, 19, 26, 28, 29, 31, 38, 44–46, 315, 700–702, 704, 705, 737
OpenMP thread pool	The set of all threads that may execute a task of a contention group and, thus, are ever available to be assigned to a team that executes implicit tasks of the contention group . 3, 33, 39, 42, 44, 354, 367
original list item	The instance of a list item in the data environment of the enclosing context . 13, 17, 18, 25, 29, 34, 158, 159, 162, 165, 166, 168, 169, 171, 173–176, 179, 185, 186, 188–190, 192, 193, 195, 203, 206, 212, 215–217, 220, 221, 227, 228, 230, 274, 291, 346, 348, 383, 745
original pointer	An original list item that corresponds to a corresponding pointer . 216

original storage	An address range in a data environment of a encountering device . 29, 33, 47, 216–219
original storage block	A storage block that is used as original storage . 47, 48, 215
orphaned construct	A construct that gives rise to a region for which the binding thread set is the current team , but is not nested within another construct that gives rise to the binding region . 435
parallel handle	A handle that refers to a parallel region . 675
parallel region	A region that has a set of associated implicit tasks and an associated team of threads that execute those tasks . 3, 19, 21, 29, 30, 35, 38, 41, 43, 44, 59, 67, 315, 329–332, 334, 339, 349, 350, 356, 360, 396–399, 423, 459
parallelism-generating construct	A construct that has the parallelism-generating property . 169, 300
parallelism-generating property	The property that a construct enables parallel execution by generating one or more teams , explicit tasks , or SIMD instructions . 29, 309, 319, 324, 355, 360, 374, 376, 378, 383
parent device	For a given target region , the device on which the corresponding target construct was encountered. 193, 288, 370, 378
parent thread	The thread that encountered the parallel construct and generated a parallel region is the parent thread of each thread that executes a task region that binds to that parallel region . The primary thread of a parallel region is the same thread as its parent thread with respect to any resources associated with an OpenMP thread . The thread that encounters a target or teams construct is not the parent thread of the initial thread of the corresponding target or teams region . 3, 29, 43
partitioned construct	A construct that has the partitioned property . 29, 329
partitioned property	The property of a construct that is a work-distribution construct for which any encountered user code in the corresponding region , excluding code from nested regions that are not closely nested regions , is executed by only one thread from its binding thread set . 29, 330, 332, 334, 337, 341, 342, 345, 348
partitioned work-sharing construct	A construct that is both a partitioned construct and a worksharing construct . 29, 43
partitioned work-sharing region	A region that corresponds to a partitioned worksharing construct . 445
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. 136, 143, 203, 301, 305
persistent self map	A self map for which the corresponding storage remains present in the device data environment , as if it has an infinite reference count. 47, 290, 733

place	An unordered set of processors on a device . 30, 38, 43, 59, 60, 72–74, 315–318, 474, 475, 734, 737, 743
place list	The ordered list that describes all OpenMP places available to the execution environment. 30, 72, 319, 734, 743
place number	A number that uniquely identifies a place in the place list , with zero identifying the first place in the place list , and each consecutive whole number identifying the next place in the place list . 474, 475
place partition	An ordered list that corresponds to a contiguous interval in the place list . It describes the places currently available to the execution environment for a given parallel region . 38, 60, 315–318
pointer attachment	The process of making a pointer variable an attached pointer . 5, 215, 217
predecessor task	A task that must complete before its dependent tasks can be executed. 15, 37, 375, 377, 379, 384, 401, 424, 429, 430
predetermined data-sharing attribute	A data-sharing attribute that applies regardless of the clauses that are specified on a given construct . 148–151, 160, 162, 209, 224
preprocessed code	For C/C++, a sequence of preprocessing tokens that result from the first six phases of translation, as defined by the base language . 266, 750
primary thread	An assigned thread that has thread number 0. A primary thread may be an initial thread or the thread that encounters a parallel construct , forms a team , generates a set of implicit tasks , and then executes one of those tasks as thread number 0. 8, 19, 21, 29, 30, 38, 43, 44, 154, 206, 309, 310, 316, 317, 328, 330, 424, 459
private variable	With respect to a given set of task regions or SIMD lanes that bind to the same parallel region , a variable for which the name provides access to a different block of storage for each task region or SIMD lane . A variable that is part of an aggregate variable cannot be made a private variable independently of other components. If a variable is privatized, its components are also private variables . 30, 46, 47, 159, 160, 205, 207, 343, 347, 348
privatization clause	The clause that may result in private variables that are new list items . 27, 148, 160
procedure	A function (for C/C++ and Fortran) or subroutine (for Fortran). 9, 11, 14, 19, 21, 33, 39, 54, 90, 123, 171, 172, 178, 226, 249, 253, 260, 264, 265, 270–274, 276–279, 323, 327, 337–339, 369, 379, 381, 446, 596, 597, 649, 684, 750
processor	An implementation-defined hardware unit on which one or more threads can execute. 15, 30, 59, 73, 77
product order	The partial order of two logical iteration vectors $\omega_a = (i_1, \dots, i_n)$ 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 $k \in \{1, \dots, n\}$. 301

program order	An ordering of operations performed by the same thread as determined by the execution sequence of operations specified by the base language . COMMENT: For versions of C and C++ that include base language support for threading, program order corresponds to the <i>sequenced-before</i> relation between operations performed by the same thread . 31, 34, 50–52
progress unit	An implementation-defined set of consecutive hardware threads on which native threads may execute a common stream of instructions. If any two OpenMP threads that execute on those native threads serially execute diverging user code then they become divergent threads . 45, 309, 318
property	A characteristic of an OpenMP feature. 8, 9, 13–17, 22, 24–27, 29, 31, 34, 38, 39, 41, 101, 250–252, 254, 257, 262, 263, 267, 403, 404
pure property	The property that a directive has no observable side effects or state, yielding the same result every time it is encountered. 90, 153, 196, 199, 202, 224, 232, 242, 258, 264, 270, 275, 281, 297, 298, 301, 302, 304–306, 324
read-modify-write	An atomic operation that reads and writes to a given storage location . COMMENT: Any <i>atomic-update</i> is a read-modify-write operation. 31, 50
reduction clause	A reduction scoping clause or a reduction participating clause . 158, 161, 177–179, 184–186, 188–190, 192, 194, 196, 197
reduction expression	A combiner expression or a initializer expression . 177, 178
reduction participating clause	A clause that defines the participants in a reduction. 31, 177, 189, 193
reduction scoping clause	A clause that defines the region in which a reduction is computed. 31, 177, 188–190, 192, 193, 361, 442

region	<p>All code encountered during a specific instance of the execution of a given construct, structured block sequence or OpenMP library routine. A region includes any code in called routines as well as any implementation code. The generation of a task at the point where a task-generating construct is encountered is a part of the region of the encountering thread. However, an explicit task region that corresponds to a task-generating construct is not part of the region of the encountering thread unless it is an included task region. The point where a target or teams directive is encountered is a part of the region of the encountering thread, but the region that corresponds to the target or teams directive is not. A region may also be thought of as the dynamic or runtime extent of a construct or of an OpenMP library routine.</p> <p>During the execution of an OpenMP program, a construct may give rise to many regions. 3–5, 8, 9, 13, 15–19, 21, 22, 25, 27–47, 50–52, 58–60, 65, 67, 74, 90, 96, 132, 133, 144, 148, 152–155, 158, 159, 166, 169, 175, 177, 178, 186, 188–190, 192–194, 206–208, 213, 215–218, 220, 227, 228, 238, 239, 242, 245, 247, 267, 269, 273, 287–289, 295, 298, 309–312, 314, 316, 319–321, 323–325, 327, 329–339, 345–347, 349–352, 356, 359–361, 364, 366, 367, 370, 373, 374, 376, 378–384, 387, 391, 393–401, 415–418, 420–425, 427, 433–436, 439–445, 459, 460, 466, 468–470, 476, 477, 486, 487, 511–514, 561, 595, 598, 646, 648, 649, 700, 733, 735, 746</p>
registered callback	A callback for which callback registration has been performed. 8 , 53 , 569 , 571
release flush	A flush that has the release flush property . 32 , 36 , 49–51 , 417 , 420 , 422–425
release flush property	A flush with the release flush property orders memory operations that precede the flush before memory operations performed by a different thread with which it synchronizes. 18 , 32 , 420
release sequence	A set of modifying atomic operations that are associated with a release flush that may establish a synchronizes-with relation between the release flush and an acquire flush . 50 , 51 , 423
replacement candidate	A directive variant or function variant that may be selected to replace a metadirective or base function . 9 , 17 , 255 , 256 , 259 , 261 , 264
reservation type	A thread-reservation type . 82
reserved thread	A thread that is restricted in the type of thread as which it can be used. A thread can be a structured thread or free-agent thread . 39 , 82
reverse-offload region	A region that is associated with a target construct that specifies a device clause with the ancestor device-modifier . 274
routine	Unless specifically stated otherwise, an OpenMP API routine . 58 , 63–65 , 366 , 380 , 381 , 459 , 469 , 486 , 487 , 510–513 , 747
runtime entry point	A function interface provided by an OpenMP runtime for use by a tool . A runtime entry point is typically not associated with a global function symbol. 4 , 27 , 28 , 32 , 571 , 573 , 574 , 580 , 596 , 637 , 641 , 646 , 647 , 649

runtime error termination	Error termination performed during execution. 17, 45, 90, 215, 217, 227, 314, 370, 488, 496, 735
scalar variable	For C/C++, a scalar-variable, as defined by the base language . For Fortran, a scalar variable with intrinsic type, as defined by the base language , excluding character type. 138, 150, 153, 169, 210, 211, 736
scan computation	The last generalized prefix sum, as defined in Section 6.6 . 18, 21, 22, 33, 40, 190, 191, 202, 203
scan phase	The portion of an associated iteration that includes all statements that read the result of a scan computation . 202–204
schedulable task set	If the thread is a structured thread , the set of tasks bound to the current team . If the thread is an unassigned thread , any explicit task in the contention group associated with the current OpenMP thread pool . 366, 367
schedule kind	The manner in which the collapsed iterations of associated loops are to be distributed among a set of threads that cooperatively execute the associated loops , as specified by a loop-nest-associated directive or the run-sched-var ICV. 60, 67, 339, 340, 344
segment	A portion of an address space associated with a set of address ranges. 3, 671
selector set	Unless specifically stated otherwise, a trait selector set . 2, 3, 253
self map	A mapping operation for which the corresponding storage is the same as its original storage . 30, 215–217, 291, 745
separated construct	A construct for which its associated structured block is split into multiple structured block sequences by a separating directive . 33, 96, 202, 203
separating directive	A directive that splits a structured block that is associated with a construct , the separated construct into multiple structured block sequences . 33, 96, 203–205
sequential part	All code encountered during the execution of an initial task region that is not part of a parallel region that corresponds to a parallel construct or a task region corresponding to a task construct . Instead, it is enclosed by an implicit parallel region . COMMENT: Executable statements in called procedures may be in both a sequential part and any number of explicit parallel regions at different points in the program execution. 33, 154, 476, 477
sequentially consistent atomic operation	An atomic operation that is specified by An atomic construct for which the seq_cst clause is specified. 52
shape-operator	For C/C++, an array shaping operator that reinterprets a pointer expression as an array with one or more specified dimensions. 227

shared variable	With respect to a given set of task regions that bind to the same parallel region , a variable for which the name provides access to the same block of storage for each task region . A variable that is part of an aggregate variable cannot be made a shared variable independently of the other components, except for static datamembers of C++ classes. 34 , 46 , 49–52 , 410–412
sibling task	Two tasks are each a sibling task of the other if they are child tasks of the same task regions . 34 , 37 , 425 , 428–430
signal	A software interrupt delivered to a thread . 4 , 34 , 698
signal handler	A function called asynchronously when a signal is delivered to a thread . 4 , 596 , 637 , 698
SIMD	Single Instruction, Multiple Data, a lock-step parallelization paradigm. 171 , 249 , 270 , 271 , 327
SIMD chunk	A set of iterations executed concurrently, each by a SIMD lane , by a single thread by means of SIMD instructions . 34 , 271 , 324 , 326 , 757
SIMD construct	A simd construct or a combined construct or composite construct for which the simd construct is a constituent construct . 344
SIMD instruction	A single machine instruction that can operate on multiple data elements. 29 , 34 , 42 , 232 , 324
SIMD lane	A software or hardware mechanism capable of processing one data element from a SIMD instruction . 30 , 34 , 44 , 46 , 158 , 159 , 163 , 171 , 172 , 188 , 189 , 195 , 324
SIMD loop	A loop that includes at least one SIMD chunk . 231 , 270 , 271
simdizable construct	A construct that has the simdizable property . 324 , 436
simdizable property	The property that a construct may be encountered during execution of a simd region . 34 , 301 , 302 , 304–306 , 324 , 348 , 415 , 435
simply contiguous array section	An array section that statically can be determined to have contiguous storage or that, in Fortran, has the CONTIGUOUS attribute. 153 , 736
simply happens before	For an event <i>A</i> to simply happen before an event <i>B</i> , <i>A</i> must precede <i>B</i> in simply happens-before order . 51
simply happens-before order	An ordering relation that is consistent with program order and the synchronizes-with relation . 19 , 34 , 51
sink iteration	A doacross iteration for which executable code, because of a doacross dependence , cannot execute until executable code from the source iteration has completed. 16 , 432
source iteration	A doacross iteration for which executable code must complete execution before executable code from another doacross iteration can execute due to a doacross dependence . 16 , 34 , 432
stand-alone directive	A construct in which no user code is associated, but may produce implementation code . 97
starting address	The address of the first storage location of a list item or, for a mapped variable of its original list item . 18 , 25 , 213

static context selector	The context selector for which the OpenMP context can be fully determined at compile time. 17 , 255 , 257 , 259
static storage duration	For C/C++, the lifetime of an object with static storage duration, as defined by the base language . For Fortran, the lifetime of a variable with a SAVE attribute, implicit or explicit, a common block object or a variable declared in a module. 15 , 47 , 150 , 152 , 156 , 162 , 180 , 221 , 222 , 229 , 234 , 243 , 274 , 290 , 379 , 733
step expression	A loop-invariant expression used by an induction operation . 10 , 22 , 181 , 182 , 185 , 199 , 200
storage block	The physical storage that corresponds to an address range in memory . 13 , 29 , 35 , 47 , 48
storage location	A storage block in memory . 3 , 5 , 17 , 31 , 34 , 46–48 , 127 , 132 , 133 , 171 , 174 , 175 , 190 , 193 , 195 , 213 , 326 , 415–421 , 428–430
strictly nested region	A region nested inside another region with no other explicit region nested between them. 347 , 351
strictly structured block	A single Fortran BLOCK construct, with a single entry at the top and a single exit at the bottom. 35 , 95 , 336
string literal	For C/C++, a string literal. For Fortran, a character literal constant. 388
strong flush	A flush that has the strong flush property . 18 , 49 , 50 , 52 , 417 , 420
strong flush property	A flush with the strong flush property flushes a set of variables from the temporary view of the memory of the current thread to the memory . 18 , 35 , 420
structure	A structure is a variable that contains one or more variables . For C/C++, implemented using struct types. For C++, implemented using class types. For Fortran, implemented using derived types. 12 , 35 , 153 , 213 , 214 , 219 , 220 , 229 , 230 , 380 , 566 , 568 , 576 , 598 , 599 , 736
structured block	For C/C++, an executable statement, possibly compound, with a single entry at the top and a single exit at the bottom, or an OpenMP construct . For Fortran, a strictly structured block or a loosely structured block . 11 , 15 , 23 , 25 , 33 , 40 , 42 , 46 , 74 , 96 , 124–132 , 136 , 140 , 168 , 174–176 , 204–208 , 249 , 267 , 271 , 300 , 303 , 310 , 320 , 327 , 330 , 331 , 333 , 335–338 , 340 , 346 , 356 , 366 , 395 , 401 , 415–418 , 423 , 424 , 433
structured block sequence	For C/C++, a sequence of zero or more executable statements (including constructs) that together have a single entry at the top and a single exit at the bottom. For Fortran, a block of zero or more executable constructs (including OpenMP constructs) with a single entry at the top and a single exit at the bottom. 8 , 32 , 33 , 96 , 125 , 136 , 145 , 168 , 169 , 202–205 , 332–334
structured parallelism	Parallel execution through the implicit tasks of (possibly nested) parallel regions by the set of structured threads in a contention group . 82 , 83

structured thread	A thread that is assigned to a team and is not a free-agent thread . 32 , 33 , 35 , 60 , 82 , 83 , 313 , 744
subsidiary directive	A directive that is not an executable directive and that appears only as part of a construct . 93 , 202 , 333 , 334
subtask	A portion of a task region between two consecutive task scheduling points in which a thread cannot switch from executing one task to executing another task . 44
supported active levels	An implementation defined maximum number of active levels of parallelism. 733
supported device	The host device or any non-host device supported by the implementation for execution of target code for which the device-related requirements of the requires directive are fulfilled. 62 , 80
synchronization construct	A construct that orders the completion of code executed by different threads . 391
synchronization hint	An indicator of the expected dynamic behavior or suggested implementation of a synchronization mechanism. 391–393
synchronizes with	For an event <i>A</i> to synchronize with an event <i>B</i> , a synchronizes-with relation must exist from <i>A</i> to <i>B</i> . 3 , 50 , 51 , 423–425
synchronizes-with relation	An asymmetric relation that relates a release flush to an acquire flush , or, for C/C++, any pair of events <i>A</i> and <i>B</i> such that <i>A</i> “synchronizes with” <i>B</i> according to the base language , and establishes memory consistency between their respective executing threads . 32 , 34 , 36 , 50
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. 15 , 16 , 36 , 42 , 43 , 53 , 58 , 59 , 174–176 , 193 , 215 , 217 , 218 , 227 , 229 , 230 , 250 , 290 , 370 , 372 , 373 , 375 , 376 , 379 , 384 , 565 , 659
target device trait set	The trait set that consists of traits that define the characteristics of a device that the implementation supports. 3 , 249 , 250 , 252 , 254
target memory space	A memory space that is associated with at least one device that is not the current device when it is created. 239 , 543 , 545 , 547
target task	A mergeable task and untied task that is generated by a device construct or a call to a device memory routine and that coordinates activity between the current device and the target device . 43 , 67 , 193 , 218 , 374–380 , 383 , 384 , 422 , 424 , 511–514
target variant	A version of a device procedure that can only be executed as part of a target region . 249

task	A specific instance of executable code and its data environment that the OpenMP implementation can schedule for execution by a team . 3 , 8 , 9 , 13 , 15 , 18–22 , 26 , 28–30 , 32–34 , 36 , 37 , 39 , 40 , 42–48 , 58–60 , 66 , 67 , 154 , 158 , 159 , 162 , 163 , 165 , 187–190 , 192–195 , 213 , 214 , 216 , 217 , 219 , 233 , 238 , 296 , 309–311 , 313 , 315 , 318 , 320 , 328 , 330 , 331 , 333 , 334 , 336 , 338 , 340 , 346 , 352–357 , 359–361 , 363 , 364 , 366 , 367 , 371 , 387 , 388 , 390 , 394–404 , 415 , 417 , 418 , 423–425 , 428–430 , 433 , 434 , 441 , 442 , 444 , 451 , 595 , 596 , 598 , 649 , 710
task completion	A condition that is satisfied when a thread reaches the end of the executable code that is associated with the task and any <i>allow-completion event</i> that is created for the task has been fulfilled. 37 , 356
task dependence	A dependence between two sibling tasks : the dependent task and a previously generated predecessor task . The task dependence is fulfilled when the predecessor task has completed. 15 , 37 , 367 , 425 , 426 , 428 , 429 , 513 , 514
task handle	A handle that refers to a task region . 675 , 710
task region	A region consisting of all code encountered during the execution of a task . 13 , 15 , 29 , 30 , 34 , 36 , 37 , 39 , 40 , 43 , 44 , 47 , 154 , 165 , 310 , 319 , 367 , 374 , 376 , 378 , 383 , 421 , 422 , 441 , 486 , 596 , 649
task scheduling point	A point during the execution of the current task region at which it can be suspended to be resumed later; or the point of task completion , after which the executing thread may switch to a different task region . 36 , 44 , 154 , 187 , 310 , 356 , 366 , 396 , 397 , 399 , 401 , 416 , 421 , 422 , 511 , 513
task synchronization construct	A taskwait , taskgroup , or a barrier construct. 44 , 356
task-generating construct	A construct that has the task-generating property . 32 , 44 , 150–152 , 429 , 430 , 445 , 746
task-generating property	The property that a construct generates one or more explicit tasks that are child tasks of the encountering task . 37 , 355 , 360 , 374 , 376 , 378 , 383
taskgroup set	A set of tasks that are logically grouped by a taskgroup region , such that a task is a member of the taskgroup set if and only if its task region is nested in the taskgroup region and it binds to the same parallel region as the taskgroup region . 37 , 399 , 441
team	A set of one or more assigned threads assigned to execute the set of implicit tasks of a parallel region . 3 , 5 , 7 , 13 , 21–23 , 28–30 , 36–45 , 59 , 67 , 154 , 172 , 190 , 191 , 195 , 205–207 , 309 , 310 , 315–320 , 322 , 327 , 329–335 , 339 , 340 , 343–348 , 350 , 371 , 394 , 396 , 397 , 416 , 424 , 436 , 443 , 459 , 648 , 700 , 735 , 738
team number	A number that the OpenMP implementation assigns to an initial team . If the initial team is not part of a league formed by a teams construct then the team number is zero; otherwise, the team number is a non-negative integer less than the number of initial teams in the league . 37 , 60 , 348

team-executed construct	A construct that has the team-executed property . 44
team-executed property	The property that a construct gives rise to a team-executed region . 38, 330–332, 334, 341, 342, 348, 396
team-executed region	A region that is executed by all or none of the threads in the current team . 38, 44, 445
team-generating construct	A construct that has the team-generating property . 445
team-generating property	The property that a construct generates a parallel region . 38, 309, 319, 378
team-worker thread	A thread that is assigned to a team but is not the primary thread . It executes one of the implicit tasks that is generated when the team is formed for an active parallel region . 41, 43
temporary view	The state of memory that is accessible to a particular thread . 420
third-party tool	A tool that executes as a separate process from the process that it is monitoring and potentially controlling. 27, 53, 667, 668, 681, 682, 684
thread	Unless specifically stated otherwise, an OpenMP thread . 3–5, 8, 13, 16–19, 22, 25, 28, 29, 31–54, 58–60, 62, 71, 74, 76, 82, 83, 90, 153–155, 165, 166, 172, 187, 188, 190, 191, 195, 205–208, 217, 218, 238, 239, 275, 281, 282, 289, 290, 295, 309–320, 327–340, 343–346, 349, 350, 352, 354, 356, 360, 361, 366, 367, 371, 372, 375, 377, 380, 384, 391, 392, 394–402, 404, 415–418, 420–425, 430, 433–436, 440–444, 459, 469, 474, 511, 512, 561, 565, 585, 595, 597, 598, 646, 648, 649, 653, 700, 710, 734, 735, 738, 751
thread affinity	A binding of threads to places within the current place partition . 58, 59, 74, 78, 154, 315, 316, 470, 734, 737, 738
thread number	For an assigned thread , a non-negative number assigned by the OpenMP implementation. For threads within the same team , zero identifies the primary thread and subsequent consecutive numbers identify any worker threads of the team . For an unassigned thread , the value <code>omp_unassigned_thread</code> . 30, 60, 154, 309, 310, 315, 318, 328, 343, 459, 468, 700
thread state	The state associated with a thread . Also, an enumeration type that describes the current OpenMP activity of a thread . Only one of the enumeration values can apply to a thread at any time. 44, 53, 565, 646
thread-exclusive construct	A construct that has the thread-exclusive property . 445
thread-exclusive property	The property that a construct when encountered by multiple threads in the current team is executed by only one thread at a time. 38, 394, 435
thread-limiting construct	A construct that has the thread-limiting property . 90

thread-limiting property	For C++, the property that a construct limits the thread that can catch an exception thrown in the corresponding region to the thread that threw the exception. 38 , 309 , 319 , 327 , 330–332 , 355 , 378 , 394 , 435
thread-pool-worker thread	A thread in an OpenMP thread pool that is not the initial thread . 585
thread-reservation type	The type specified for a reserved thread . 32 , 82
thread-safe procedure	A procedure that performs the intended function even when executed concurrently (by multiple native threads). 54
thread-set	The set of threads for which a flush may enforce memory consistency . 48 , 49 , 51 , 52 , 415 , 420 , 422
threadprivate memory	The set of threadprivate variables associated with each thread . 46
threadprivate variable	A variable that is replicated, one instance per thread , by the OpenMP implementation. Its name then provides access to a different block of storage for each thread . A variable that is part of an aggregate variable cannot be made a threadprivate variable independently of the other components, except for static data members of C++ classes. If a variable is made a threadprivate variable , its components are also threadprivate variables . 39 , 153–157 , 205 , 206 , 323 , 339 , 380
tied task	A task that, when its task region is suspended, can be resumed only by the same thread that was executing it before suspension. That is, the task is tied to that thread . 44 , 352 , 367
tool	Code that can observe and/or modify the execution of an application. 2 , 18 , 32 , 38 , 39 , 42 , 53 , 54 , 59 , 60 , 372 , 373 , 561 , 562 , 565–568 , 573 , 580 , 598 , 599 , 649 , 698
tool callback	A function that a tool provides to an OpenMP implementation to invoke when an associated event occurs. 8 , 53 , 397 , 433 , 451 , 641
tool context	An opaque reference provided by a tool to an OMPD library . A tool context uniquely identifies an abstraction. 3 , 26 , 39 , 676 , 681
trace record	A data structure in which to store information associated with an occurrence of an event . 26 , 573 , 574 , 637
trait	An aspect of an OpenMP implementation or the execution of an OpenMP program . 3 , 9 , 13 , 16–18 , 20 , 26–28 , 36 , 39 , 236–242 , 245 , 247 , 249 , 250 , 252–254 , 266 , 284 , 737 , 743
trait selector	A member of a trait selector set . 249 , 251–255 , 257 , 260 , 266
trait selector set	A set of traits that are specified to match the trait set at a given point in an OpenMP program . 33 , 39 , 251
trait set	A grouping of related traits . 11 , 16 , 17 , 20 , 36 , 39 , 249 , 252 , 254
unassigned thread	A thread that is not currently assigned to any team . 19 , 20 , 33 , 38 , 42 , 43 , 354 , 367 , 459 , 595

underrred task	A task for which execution is not deferred with respect to its generating task region . That is, its generating task region is suspended until execution of the structured block associated with the underrred task is completed. 21, 26, 40, 357, 361, 424
undefined	For variables , the property of not being defined , that is, of not having a valid value. 48, 442, 641
unified address space	An address space that is used by all devices . 289
unit of work	In constructs that use units of work , a single or multiple executable statements that will be executed by a single thread and are part of the same structured block . A structured block can consist of one or more units of work ; the number of units of work into which a structured block is split is allowed to vary among different compliant implementations . 40, 334, 335, 337, 338, 605
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 OpenMP program . Such unspecified behavior may result from: <ul style="list-style-type: none"> • Issues that this specification documents as having unspecified behavior. • A non-conforming program. • A conforming program exhibiting an implementation defined behavior. 10, 20, 40, 46–48, 55, 90, 175, 185, 238, 245, 289, 355, 379, 381, 398
untied task	A task that, when its task region is suspended, can be resumed by any thread in the team . That is, the task is not tied to any thread . 36, 44, 155, 352, 357, 367
update value	The update value of a new list item used for a scan computation is, for a given logical iteration , the value of the new list item on completion of its input phase. 40, 203
user-defined cancellation point	A cancellation point that is specified by a cancellation point construct . 444
user-defined induction	An induction operation that is defined by a declare induction directive . 201, 202
user-defined mapper	A mapper that is defined by a declare mapper directive . 24, 122, 214, 224, 225, 227
user-defined reduction	An reduction operation that is defined by a declare reduction directive . 196, 198, 443
utility directive	A directive that facilitates interactions with the compiler and/or supports code readability; it may be either informational or executable. 93, 281, 282, 298

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. COMMENT: An array element or structure element is a variable that is part of an aggregate variable . 3, 6–9, 12–15, 18, 19, 22–26, 30, 34, 35, 39–41, 46–52, 54, 58, 94, 96, 103–105, 111, 121, 122, 126, 134, 137–141, 148–162, 165, 166, 168, 169, 172, 175, 178–182, 186, 191, 195–197, 200, 205–211, 214, 218–226, 233, 234, 236, 240–244, 246, 247, 253, 256, 259, 261, 263, 268–270, 273–278, 290, 300, 312, 319, 324, 328, 329, 335, 336, 338, 339, 343, 346, 348, 349, 354, 359, 361, 369, 373–376, 378–383, 420, 421, 432, 573, 598, 641, 647, 649, 733, 748
variant substitution	The replacement of a call to a base function by a call to a function variant . 259, 267, 268
wait identifier	A unique opaque handle associated with each data object (for example, a lock) that the OpenMP runtime uses to enforce mutual exclusion and potentially to cause a thread to wait actively or passively. 597, 598, 646
white space	A non-empty sequence of space and/or horizontal tab characters. 69, 76, 78, 91, 92, 97–100, 113, 114
work distribution	The manner in which execution of a region that corresponds to a work-distribution construct is assigned to threads . 142
work-distribution construct	A construct that has the work-distribution property . 2, 29, 41, 165, 166, 169, 329, 349
work-distribution property	The property that a construct is cooperatively executed by threads in the binding thread set of the corresponding region . 41, 330–332, 334, 337, 341, 342, 345, 348
work-distribution region	A region that corresponds to a work-distribution construct . 166, 169, 329
worker thread	Unless specifically stated otherwise, a team-worker thread . 38, 311
worksharing construct	A construct that has the worksharing property . 29, 41, 43, 44, 166, 172, 189–191, 195, 329, 333, 339, 349, 398, 443, 447
worksharing property	The property of a construct that is a work-distribution construct that is executed by the team of the innermost enclosing parallel region and includes, by default, an implicit barrier. 41, 330–332, 334, 341, 342, 348
worksharing region	A region that corresponds to a worksharing construct . 44, 166, 190, 329, 397, 422
worksharing-loop construct	A construct that has the worksharing-loop property . 17, 41, 190, 195, 339–344, 434, 436, 441, 443
worksharing-loop property	The property of a worksharing construct that is a loop-nest-associated construct that distributes the collapsed iterations of the associated loops among the threads in the team . 41, 341, 342
worksharing-loop region	A region that corresponds to a worksharing-loop construct . 339, 340, 434–436

**zero-offset
assumed-size ar-
ray**

An **assumed-size array** for which the lower bound is zero. 174, 210, 214

1.3 Execution Model

A compliant implementation must follow the abstract execution model that the supported **base language** and OpenMP specification define, as observable from the results of user code in a **conforming program**. These results do not include output from external monitoring **tools** or **tools** that use the OpenMP **tool** interfaces (i.e., **OMPT** and **OMPD**), which may reflect deviations from the execution model such as the unprescribed use of additional **native threads**, **SIMD instruction**, alternate loop transformations, or other **target devices** to facilitate parallel execution of the program.

The OpenMP API consists of several **directives**, runtime routines and two tool interfaces. Some **directives** allow customization of **base language** declarations while other **directives** specify details of program execution. Such **executable directives** may be lexically associated with **base language** code. Each **executable directive** and any such associated **base language** code forms a **construct**. An **OpenMP program** executes **regions**, which consist of all code encountered by **native threads**.

Some **regions** are implicit but many are **explicit regions**, which correspond to a specific instance of a **construct** or runtime routine. Execution is composed of **nested regions** since a given **region** may encounter additional **constructs** and runtime routines. References to **regions**, particularly **explicit regions** or **nested regions**, that correspond to a specific type of **construct** or runtime routine usually include the name of that **construct** or runtime routine to identify the type of **region** that results.

With the OpenMP API, multiple **threads** execute **tasks** defined implicitly or explicitly by OpenMP **directives** and their associated user code, if any. An implementation may use of multiple **devices** for a given execution of an **OpenMP program**. Using different numbers of **threads** may result in different numeric results because of changes in the association of numeric operations.

Each device executes a set of one or more **contention groups**. Each **contention group** consists of a set of **tasks** that an associated set of **threads**, an **OpenMP thread pool**, executes. The lifetime of the **OpenMP thread pool** is the same as that of the **contention group**. The **threads** that are associated with each **contention group** are distinct from **threads** associated with any other **contention group**. **Threads** cannot migrate to executed tasks of a different **contention group**.

Each **OpenMP thread pool** has an **initial thread**, which may be the **thread** that starts execution of a **region** that is not nested within any other **region**, or which may be the **thread** that starts execution of the **structured block** associated with a **target** or **teams** construct. Each **initial thread** executes sequentially; the code that it encounters is part of an **implicit task region**, called an **initial task region**, that is generated by the **implicit parallel region** that surrounds all code executed by the **initial thread**. The other **threads** in the **OpenMP thread pool** associated with a **contention group** are **unassigned threads**. An **implicit task** is assigned to each of those **threads**. When a **task** encounters a **parallel construct**, some of the **unassigned threads** become **assigned threads** that are assigned to the **team** of that **parallel region**.

1 The **thread** that executes the **implicit parallel region** that surrounds the whole program executes on
2 the **host device**. An implementation may support other **devices** besides the **host device**. If
3 supported, these **devices** are available to the **host device** for *offloading* code and data. Each **device**
4 has its own **contention groups**.

5 A **task** that encounters a **target** construct generates a new **target task**; its **region** encloses the
6 **target region**. The **target task** is complete after the **target region** completes execution. When
7 a **target task** executes, an **initial thread** executes the enclosed **target region**. The **initial thread**
8 executes sequentially, as if the **target region** is part of an **initial task region** that an **implicit**
9 **parallel region** generates. The **initial thread** may execute on the requested **target device**, if it is
10 available. If the **target device** does not exist or the implementation does not support it, all **target**
11 **regions** associated with that **device** execute on the **host device**. Otherwise, the implementation
12 ensures that the **target region** executes as if it were executed in the **data environment** of the **target**
13 **device** unless an **if clause** is present and the **if clause** expression evaluates to *false*.

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

20 Any **thread** that encounters a **parallel** construct becomes the **primary thread** of the new **team**
21 that consists of itself and zero or more additional **unassigned threads** that are then assigned to that
22 **team** as **team-worker threads**. Those **threads** remain **assigned threads** for the lifetime of that **team**.
23 A set of **implicit tasks**, one per **thread**, is generated. The code inside the **parallel** construct
24 defines the code for each **implicit task**. A different **thread** in the **team** is assigned to each **implicit**
25 **task**, which is tied, that is, only that assigned **thread** ever executes it. The **task region** of the **task**
26 being executed by the **encountering thread** is suspended, and each member of the new **team**
27 executes its **implicit task**. The **primary thread** is the **parent thread** of any **thread** that executes a **task**
28 that is bound to the **parallel region**. An implicit **barrier** occurs at the end of the **parallel region**.
29 Only the **primary thread** resumes execution beyond the end of that **region**, resuming the suspended
30 **task region**. The other **threads** again become **unassigned threads**. A single program can specify any
31 number of **parallel** constructs.

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

38 When any **team** encounters a **partitioned worksharing construct**, the work inside the **construct** is
39 divided into work partitions, each of which is executed by one member of the **team**, instead of the
40 work being executed redundantly by each **thread**. An implicit **barrier** occurs at the end of any **region**
41 that corresponds to a **worksharing construct** for which the **nowait** clause is not specified.

1 Redundant execution of code by every **thread** in the **team** resumes after the end of the **worksharing**
2 **construct**. **Regions** that correspond to **team-executed constructs**, including all **worksharing regions**
3 and **barrier regions**, are executed by the current **team** such that all **threads** in the **team** execute the
4 **team-executed regions** in the same order.

5 When a **loop construct** is encountered, the iterations of the loop associated with the **construct** are
6 executed in the context of its **encountering threads**, as determined according to its **binding region**. If
7 the **loop region** binds to a **teams region**, the **region** is encountered by the set of **primary thread**
8 that execute the **teams region**. If the **loop region** binds to a **parallel region**, the **region** is
9 encountered by the **team** that execute the **parallel region**. Otherwise, the **region** is encountered
10 by a single **thread**. If the **loop region** binds to a **teams region**, the **encountering threads** may
11 continue execution after the **loop** region without waiting for all iterations to complete; the
12 iterations are guaranteed to complete before the end of the **teams region**. Otherwise, all iterations
13 must complete before the **encountering thread** continue execution after the **loop region**. All
14 **threads** that encounter the **loop construct** may participate in the execution of the iterations. Only
15 one **thread** may execute any given iteration.

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

18 When any **thread** encounters a **task-generating construct**, one or more **explicit tasks** are generated.
19 Explicitly generated **tasks** are scheduled onto **threads** of the **task binding thread set**, subject to the
20 availability of the **threads** to execute work. Thus, execution of the new **task** could be immediate, or
21 deferred until later according to task scheduling constraints and **thread** availability. Completion of
22 all **explicit tasks** bound to a given **parallel region** is guaranteed before the **primary thread** leaves the
23 implicit **barrier** at the end of the **region**. Completion of a subset of all **explicit tasks** bound to a
24 given **parallel region** may be specified through the use of **task synchronization constructs**.
25 Completion of all **explicit tasks** bound to an **implicit parallel region** is guaranteed when the
26 associated **initial task** completes. The **initial task** on the **host device** that begins a typical **OpenMP**
27 **program** is guaranteed to end by the time that the program exits.

28 **Threads** are allowed to suspend the **current task region** at a **task scheduling point** in order to execute
29 a different **task**. Thus, each **task** consists of a set of one or more **subtasks** that each correspond to
30 the portion of the **task region** between any two consecutive **task scheduling points** that the **task**
31 encounters. If the **task region** of a **tied task** is suspended, the initially assigned **thread** later resumes
32 execution of the next **subtask** of the suspended **task region**. If the **task region** of an **untied task** is
33 suspended, any **thread** in the **binding thread set** of the **task** may resume execution of its next **subtask**.

34 **OpenMP threads** are logical execution entities that are mapped to **native threads** for actual
35 execution. OpenMP does not dictate the details of the implementation of **native threads** and, instead,
36 specifies requirements on the **thread state** of **OpenMP threads**. As long as those requirements are
37 met, a **compliant implementation** may map the same **OpenMP thread** differently (i.e., to different
38 **native threads**) for different portions of its execution (e.g., for the execution of different **subtasks**).
39 Similarly, while the lifetime of an **OpenMP thread** and its **OpenMP thread pool** is identical to that
40 of the associated **contention group**, OpenMP does not specify the lifetime of any **native threads** to
41 which it is mapped. **Native threads** may be created at any time and may be terminated at any time.

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

10 If a **task** encounters a **cancel construct**, any other *cancel-directive-name clauses*, 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**. **Tasks** check if **cancellation** has been activated for their **region** at
13 **cancellation points** and, if so, also resume execution at the end of the canceled **region**.

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

18 OpenMP specifies circumstances that cause **error termination**. If **compile-time error termination** is
19 specified, the effect is as if an **error directive** for which *sev-level* is **fatal** and *action-time* is
20 **compilation** is encountered. If **runtime error termination** is specified, the effect is as if an
21 **error directive** for which *sev-level* is **fatal** and *action-time* is **execution** is encountered.

22 A **construct** that creates a **data environment** creates it at the time that the **construct** is encountered.
23 The description of a **construct** defines whether it creates a **data environment**. Synchronization
24 **constructs** and library routines are available in the OpenMP API to coordinate **tasks** and their data
25 accesses. In addition, library routines and environment variables are available to control or to query
26 the runtime environment of **OpenMP programs**. The scope of OpenMP synchronization
27 mechanisms may be limited to the **contention group** of the **encountering task**. Except where
28 explicitly specified, any effect of the mechanisms between **contention groups** is **implementation**
29 **defined**. Section 1.4 details the OpenMP memory model, including the effect of these features.

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

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

1 threads from making progress if executing the [explicit tasks](#) as [included tasks](#) would ensure that
2 they make progress.

3 Each [device](#) is identified by a [device number](#). The [device number](#) for the [host device](#) is the value of
4 the total number of [non-host devices](#), while each [non-host device](#) has a unique [device number](#) that
5 is greater than or equal to zero and less than the [device number](#) for the [host device](#). Additionally,
6 the constant `omp_initial_device` can be used as an alias for the [host device](#) and the constant
7 `omp_invalid_device` can be used to specify an invalid [device number](#). A [conforming device](#)
8 [number](#) is either a non-negative integer that is less than or equal to `omp_get_num_devices()`
9 or equal to `omp_initial_device` or `omp_invalid_device`.

10 1.4 Memory Model

11 1.4.1 Structure of the OpenMP Memory Model

12 The OpenMP API provides a relaxed-consistency, shared-memory model. All [OpenMP threads](#)
13 have access to a place to store and to retrieve [variables](#), called the [memory](#). A given [storage](#)
14 [location](#) in the [memory](#) may be associated with one or more [devices](#), such that only [threads](#) on
15 [associated devices](#) have access to it. In addition, each [thread](#) is allowed to have its own *temporary*
16 *view* of the [memory](#). The temporary view of [memory](#) for each [thread](#) is not a required part of the
17 OpenMP [memory](#) model, but can represent any kind of intervening structure, such as machine
18 registers, cache, or other local storage, between the [thread](#) and the [memory](#). The temporary view of
19 memory allows the [thread](#) to cache [variables](#) and thereby to avoid going to [memory](#) for every
20 reference to a [variable](#). Each [thread](#) also has access to another type of [memory](#) that must not be
21 accessed by other [threads](#), called [threadprivate memory](#).

22 A [directive](#) that accepts data-sharing attribute [clauses](#) determines two kinds of access to [variables](#)
23 used in the associated [structured block](#) of the [directive](#): [shared variables](#) and [private variable](#). Each
24 [variable](#) referenced in the [structured block](#) has an original [variable](#), which is the [variable](#) by the
25 same name that exists in the [OpenMP program](#) immediately outside the [construct](#). Each reference
26 to a [shared variable](#) in the [structured block](#) becomes a reference to the original variable. For each
27 [private variable](#) referenced in the [structured block](#), a new version of the original [variable](#) (of the
28 same type and size) is created in [memory](#) for each [task](#) or [SIMD lane](#) that contains code associated
29 with the [directive](#). Creation of the new version does not alter the value of the original [variable](#).
30 However, attempts to access the original [variable](#) from within the [region](#) that corresponds to the
31 [directive](#) result in [unspecified behavior](#); see [Section 6.4.3](#) for additional details. References to a
32 [private variable](#) in the [structured block](#) refer to the private version of the original [variable](#) for the
33 current [task](#) or [SIMD lane](#). The relationship between the value of the original [variable](#) and the
34 initial or final value of the private version depends on the exact clause that specifies it. Details of
35 this issue, as well as other issues with privatization, are provided in [Chapter 6](#).

36 The minimum size at which a [memory](#) update may also read and write back adjacent [variables](#) that
37 are part of an [aggregate variable](#) is [implementation defined](#) but is no larger than the [base language](#)
38 requires.

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

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

13 A [private variable](#) in a [task region](#) that subsequently generates an inner nested [parallel region](#) is
14 permitted to be made shared for [implicit tasks](#) in the inner [parallel region](#). A [private variable](#) in
15 a [task region](#) can also be shared by an [explicit task region](#) generated during its execution. However,
16 the programmer must use synchronization that ensures that the lifetime of the [variable](#) does not end
17 before completion of the [explicit task region](#) sharing it. Any other access by one [task](#) to the [private](#)
18 [variables](#) of another [task](#) results in [unspecified behavior](#).

19 A [storage location](#) in [memory](#) that is associated with a given [device](#) has a [device address](#) that may
20 be dereferenced by a [thread](#) executing on that [device](#), but it may not be generally accessible from
21 other [devices](#). A different [device](#) may obtain a [device pointer](#) that refers to this [device address](#). The
22 manner in which an [OpenMP program](#) can obtain the referenced [device address](#) from a [device](#)
23 [pointer](#), outside of mechanisms specified by OpenMP, is [implementation defined](#). Unless otherwise
24 specified, the [atomic scope](#) of a [storage location](#) is [all threads](#) on the [current device](#).

25 1.4.2 Device Data Environments

26 When an [OpenMP program](#) begins, an implicit [target data region](#) for each [device](#) surrounds
27 the whole program. Each [device](#) has a [device data environment](#) that is defined by its implicit
28 [target data region](#). Any [declare target directives](#) and [directives](#) that accept [data-mapping](#)
29 [attribute clauses](#) determine how an [original storage block](#) in a [data environment](#) is mapped to a
30 [corresponding storage block](#) in a [device data environment](#). Additionally, if a [variable](#) with [static](#)
31 [storage duration](#) has [original storage](#) that is accessible on a [device](#), and the [variable](#) is not a [device](#)
32 [local variable](#), it may be treated as if its storage is mapped with a [persistent self map](#) in the implicit
33 [target data region](#) of the [device](#); whether this happens is [implementation defined](#).

34 When an [original storage block](#) is mapped to a [device data environment](#) and a [corresponding](#)
35 [storage block](#) is not present in the [device data environment](#), a new [corresponding storage block](#) (of
36 the same type and size as the [original storage block](#)) is created in the [device data environment](#).
37 Conversely, the [original storage block](#) becomes the [corresponding storage block](#) of the new [storage](#)
38 [block](#) in the [device data environment](#) of the [device](#) that performs a [mapping operation](#).

39 The [corresponding storage block](#) in the [device data environment](#) may share storage with the [original](#)

1 storage block. Writes to the corresponding storage block may alter the value of the original storage
2 block. Section 1.4.6 discusses the impact of this possibility on memory consistency. When a task
3 executes in the context of a device data environment, references to the original storage block refer
4 to the corresponding storage block in the device data environment. If an original storage block is
5 not currently mapped and a corresponding storage block does not exist in the device data
6 environment then accesses to the original storage block result in unspecified behavior unless the
7 unified_shared_memory clause is specified on a requires directive for the compilation
8 unit.

9 The relationship between the value of the original storage block and the initial or final value of the
10 corresponding storage block depends on the map-type. Details of this issue, as well as other issues
11 with mapping a variable, are provided in Section 6.8.3.

12 The original storage block in a data environment and a corresponding storage block in a device data
13 environment may share storage. Without intervening synchronization data races can occur.

14 If a storage block has a corresponding storage block with which it does not share storage, a write to
15 a storage location designated by the storage block causes the value at the corresponding storage
16 block to become undefined.

17 1.4.3 Memory Management

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

24 An OpenMP program can use a memory allocator to allocate memory in which to store variables.
25 This memory will be allocated from the storage resources of the memory space associated with the
26 memory allocator. Memory allocators are also used to deallocate previously allocated memory.
27 When a memory allocator is not used to allocate memory, OpenMP does not prescribe the storage
28 resource for the allocation; the memory for the variables may be allocated in any storage resource.

29 1.4.4 The Flush Operation

30 The memory model has relaxed-consistency because the temporary view of memory of a thread is
31 not required to be consistent with memory at all times. A value written to a variable can remain in
32 that temporary view until it is forced to memory at a later time. Likewise, a read from a variable
33 may retrieve the value from that temporary view, unless it is forced to read from memory. OpenMP
34 flush operations are used to enforce consistency between the temporary view of memory of a thread
35 and memory, or between the temporary views of multiple threads.

36 A flush has an associated thread-set that constrains the threads for which it enforces memory

1 consistency. Consistency is only guaranteed to be enforced between the view of **memory** of these
2 **threads**. Unless otherwise stated, the **thread-set** of a **flush** only includes all **threads** on the **current**
3 **device**.

4 If a **flush** is a **strong flush**, it enforces consistency between the temporary view of a **thread** and
5 **memory**. A **strong flush** is applied to a set of **variable** called the **flush-set**. A **strong flush** restricts
6 how an implementation may reorder **memory** operations. Implementations must not reorder the
7 code for a **memory** operation for a given **variable**, or the code for a **flush** for the **variable**, with
8 respect to a **strong flush** that refers to the same **variable**.

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

22 A **strong flush** provides a guarantee of consistency between the temporary view of a **thread** and
23 **memory**. Therefore, a **strong flush** can be used to guarantee that a value written to a **variable** by one
24 **thread** may be read by a second **thread**. To accomplish this, the programmer must ensure that the
25 second **thread** has not written to the **variable** since its last **strong flush** of the **variable**, and that the
26 following sequence of events are completed in this specific order:

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

31 If a **flush** is a **release flush** or **acquire flush**, it can enforce consistency between the views of **memory**
32 of two synchronizing **threads**. A **release flush** guarantees that any prior operation that writes or
33 reads a **shared variable** will appear to be completed before any operation that writes or reads the
34 same **shared variable** and follows an **acquire flush** with which the **release flush** synchronizes (see
35 [Section 1.4.5](#) for more details on **flush** synchronization). A **release flush** will propagate the values
36 of all **shared variables** in its temporary view to **memory** prior to the **thread** performing any
37 subsequent **atomic operation** that may establish a synchronization. An **acquire flush** will discard
38 any value of a **shared variable** in its temporary view to which the **thread** has not written since last
39 performing a **release flush**, and it will load any value of a **shared variable** propagated by a **release**

1 flush that synchronizes with it (according to the synchronizes-with relation) into its temporary view
2 so that it may be subsequently read. Therefore, release flushes and acquire flushes may also be used
3 to guarantee that a value written to a variable by one thread may be read by a second thread. To
4 accomplish this, the programmer must ensure that the second thread has not written to the variable
5 since its last acquire flush, and that the following sequence of events happen in this specific order:

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

Note – OpenMP synchronization operations, described in Chapter 16 and in Section 19.9, are recommended for enforcing this order. Synchronization through variables is possible but is not recommended because the proper timing of flushes is difficult.

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

1.4.5 Flush Synchronization and Happens-Before Order

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

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

- If a release flush is performed on entry to an atomic operation, that atomic operation starts its release sequence.
- If a release flush is performed in an implicit flush region, an atomic operation that is provided by the implementation and that modifies an internal synchronization variable starts its release sequence.
- If a release flush is performed by an explicit flush region, any atomic operation that modifies a shared variable and follows the flush region in the program order of its thread starts an associated release sequence.

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

- 4 • If an **acquire flush** is performed on exit from an **atomic operation**, that **atomic operation** is its
5 associated **atomic operation**.
- 6 • If an **acquire flush** is performed in an implicit **flush region**, an **atomic operation** that is
7 provided by the implementation and that reads an internal synchronization **variable** is its
8 associated **atomic operation**.
- 9 • If an **acquire flush** is performed by an explicit **flush region**, any **atomic operation** that reads
10 a **shared variable** and precedes the **flush region** in the **program order** of its **thread** is an
11 associated **atomic operation**.

12 The **atomic scope** of the internal synchronization **variable** that is used in implicit **flush regions** is
13 the intersection of the **thread-sets** of the synchronizing **flushes**.

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

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

18 An operation *X* **simply happens before** an operation *Y*, that is, *X* precedes *Y* in **simply**
19 **happens-before order**, if any of the following conditions are satisfied:

- 20 1. *X* and *Y* are performed by the same **thread**, and *X* precedes *Y* in the **program order** of the
21 **thread**;
- 22 2. *X* **synchronizes with** *Y* according to the **flush** synchronization conditions explained above or
23 according to the definition of the “synchronizes with” relation in the **base language**, if such a
24 definition exists; or
- 25 3. Another operation, *Z*, exists such that *X* **simply happens before** *Z* and *Z* **simply happens**
26 **before** *Y*.

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

- 28 1. *X* “happens before” *Y*, as defined in the **base language** if such a definition exists; or
- 29 2. *X* **simply happens before** *Y*.

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

1.4.6 OpenMP Memory Consistency

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

- If two operations performed by different [threads](#) are [sequentially consistent atomic operations](#) or they are [strong flushes](#) that flush the same [variable](#), then they must be completed as if in some sequential order, seen by all affected [threads](#).
- If two operations performed by the same [thread](#) are [sequentially consistent atomic operations](#) or they access, modify, or, with a [strong flush](#), flush the same [variable](#), then they must be completed as if in the [program order](#) of that [thread](#), as seen by all affected [threads](#).
- If two operations are performed by different [threads](#) and one *happens before* the other, then they must be completed as if in that *happens before* order, as seen by all affected [threads](#), if:
 - both operations access or modify the same [variable](#);
 - both operations are [strong flushes](#) that flush the same [variable](#); or
 - both operations are [sequentially consistent atomic operations](#).
- Any two [atomic operations](#) from different [atomic regions](#) must be completed as if in the same order as the [strong flushes](#) implied in their [regions](#), as seen by all affected [threads](#).

The [flush](#) operation can be specified using the [flush directive](#), and is also implied at various locations in an [OpenMP program](#); see [Section 16.8.6](#) for details.

Note – Since [flushes](#) by themselves cannot prevent data races, explicit [flushes](#) are only useful in combination with [non-sequentially consistent atomic constructs](#).

[OpenMP programs](#) that:

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

behave as though operations on [shared variables](#) were simply interleaved in an order consistent with the order in which they are performed by each [thread](#). The relaxed consistency model is invisible for such programs, and any explicit [flushes](#) in such programs are redundant.

1.5 Tool Interfaces

The OpenMP API includes two [tool](#) interfaces, [OMPT](#) and [OMPD](#), to enable development of high-quality, portable, [tools](#) that support monitoring, performance, or correctness analysis and debugging of [OpenMP programs](#) developed using any implementation of the OpenMP API. An implementation of the OpenMP API may differ from the abstract execution model described by its specification. The ability of [tools](#) that use [OMPT](#) or [OMPD](#) to observe such differences does not constrain implementations of the OpenMP API in any way.

1.5.1 OMPT

The [OMPT](#) interface, which is intended for [first-party tools](#), provides the following:

- A mechanism to initialize a [first-party tool](#);
- 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 [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 [events](#) for which [tool callbacks](#) are invoked. For some [events](#), OpenMP implementations must guarantee that a [registered callback](#) will be invoked for each occurrence of the [event](#). For other [events](#), OpenMP implementations are permitted to invoke a [registered callback](#) for some or no occurrences of the [event](#); for such [events](#), however, OpenMP implementations are encouraged to invoke [tool callbacks](#) on as many occurrences of the [event](#) as is practical. [Section 20.2.4](#) 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, [OMPT](#) includes a Fortran binding only for `omp_control_tool`; all other [OMPT](#) functionality is supported 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 21 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 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 **procedures** provided by the **base language** must be **thread-safe procedures** 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 **procedures** by different **threads** must produce correct results (although not necessarily the same as serial execution results, as in the case of random number generation **procedures**).

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.

- 1 ● ISO/IEC 9899:2011, *Information Technology - Programming Languages - C*.
2 This OpenMP API specification refers to ISO/IEC 9899:2011 as C11.
- 3 ● ISO/IEC 9899:2018, *Information Technology - Programming Languages - C*.
4 This OpenMP API specification refers to ISO/IEC 9899:2018 as C18.
- 5 ● ISO/IEC 9899:2023, *Information Technology - Programming Languages - C*.
6 This OpenMP API specification refers to ISO/IEC 9899:2023 as C23.
- 7 ● ISO/IEC 14882:1998, *Information Technology - Programming Languages - C++*.
8 This OpenMP API specification refers to ISO/IEC 14882:1998 as C++98.
- 9 ● ISO/IEC 14882:2011, *Information Technology - Programming Languages - C++*.
10 This OpenMP API specification refers to ISO/IEC 14882:2011 as C++11.
- 11 ● ISO/IEC 14882:2014, *Information Technology - Programming Languages - C++*.
12 This OpenMP API specification refers to ISO/IEC 14882:2014 as C++14.
- 13 ● ISO/IEC 14882:2017, *Information Technology - Programming Languages - C++*.
14 This OpenMP API specification refers to ISO/IEC 14882:2017 as C++17.
- 15 ● ISO/IEC 14882:2020, *Information Technology - Programming Languages - C++*.
16 This OpenMP API specification refers to ISO/IEC 14882:2020 as C++20.
- 17 ● ISO/IEC 14882:2023, *Information Technology - Programming Languages - C++*.
18 This OpenMP API specification refers to ISO/IEC 14882:2023 as C++23.
- 19 ● ISO/IEC 1539:1980, *Information Technology - Programming Languages - Fortran*.
20 This OpenMP API specification refers to ISO/IEC 1539:1980 as Fortran 77.
- 21 ● ISO/IEC 1539:1991, *Information Technology - Programming Languages - Fortran*.
22 This OpenMP API specification refers to ISO/IEC 1539:1991 as Fortran 90.
- 23 ● ISO/IEC 1539-1:1997, *Information Technology - Programming Languages - Fortran*.
24 This OpenMP API specification refers to ISO/IEC 1539-1:1997 as Fortran 95.
- 25 ● ISO/IEC 1539-1:2004, *Information Technology - Programming Languages - Fortran*.
26 This OpenMP API specification refers to ISO/IEC 1539-1:2004 as Fortran 2003.
- 27 ● ISO/IEC 1539-1:2010, *Information Technology - Programming Languages - Fortran*.
28 This OpenMP API specification refers to ISO/IEC 1539-1:2010 as Fortran 2008.
- 29 ● ISO/IEC 1539-1:2018, *Information Technology - Programming Languages - Fortran*.
30 This OpenMP API specification refers to ISO/IEC 1539-1:2018 as Fortran 2018. While
31 future versions of the OpenMP specification are expected to address the following features,
32 currently their use may result in [unspecified behavior](#).
- 33 – Assumed-type dummy argument
- 34 ● Where this OpenMP API specification refers to C, C++ or Fortran, reference is made to the
35 [base 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 [ICV scope](#). Possible [ICV scopes](#) are: [global](#); [device](#); [implicit task](#); and [data environment](#). If an [ICV scope](#) is [global](#) then one copy of the [ICV](#) exists for the whole [OpenMP program](#). If an [ICV scope](#) is [device](#) then one copy of the [ICV](#) exists for the [current device](#). If an [ICV scope](#) is [implicit task](#) then a distinct copy of the [ICV](#) exists for each [implicit task](#). If an [ICV scope](#) is [data environment](#) then a distinct copy of the [ICV](#) exists for the [data environment](#) of each [task](#), unless otherwise specified. 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 [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 [ICV 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 active 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

ICV	Scope	Description
<i>bind-var</i>	data environment	Controls the binding of 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
<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>	data environment	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>free-agent-thread-limit-var</i>	data environment	Controls the maximum number of free-agent threads that may execute tasks in the contention group in parallel
<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 active 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>nteamsv-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

ICV	Scope	Description
<i>place-assignment-var</i>	implicit task	Controls the places to which threads are bound
<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>structured-thread-limit-var</i>	data environment	Controls the maximum number of structured threads that may execute tasks in the contention group in parallel
<i>target-offload-var</i>	global	Controls the offloading behavior
<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	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 current 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 native threads

Cross References

- Team Generator Types, see [Section 21.3.10](#)

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	See below
<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>
<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>free-agent-thread-limit-var</i>	OMP_THREAD_LIMIT, OMP_THREADS_RESERVE	See below
<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>nteamsv-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-assignment-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>structured-thread-limit-var</i>	OMP_THREAD_LIMIT, OMP_THREADS_RESERVE	See below
<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>

ICV	Environment Variable	Initial Value
<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

If an **ICV** has an associated **environment variable** and that **ICV** neither has **global ICV scope** nor is **default-device-var** then the **ICV** has a set of associated **device-specific environment variables** that extend the associated **environment variable** with the following syntax:

`<ENVIRONMENT VARIABLE>_ALL`

or

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

where `<ENVIRONMENT VARIABLE>` is the associated **environment variable** and `<device>` is the **device number** as specified in the **device clause** (see [Section 14.2](#)); the semantic and precedence is described in [Chapter 3](#).

Semantics

- The initial value of *available-devices-var* is the set of all **accessible devices** that are also **supported devices**.
- The initial value of *dyn-var* is **implementation defined** if the implementation supports dynamic adjustment of the number of **threads**; otherwise, the initial value is *false*.
- The initial value of *free-agent-thread-limit-var* is one less than the initial value of *thread-limit-var*.
- The initial value of *structured-thread-limit-var* is the initial value of *thread-limit-var*.
- If *target-offload-var* is **mandatory** and the number of available **non-host devices** is zero then *default-device-var* is initialized to **omp_invalid_device**. Otherwise, the initial value is an **implementation defined** non-negative integer that is less than or, if *target-offload-var* is not **mandatory**, equal to **omp_get_initial_device()**.
- The value of the *nthreads-var ICV* is a list.
- The value of the *bind-var ICV* is a list.

1 The [host device](#) and [non-host device ICVs](#) are initialized before any [construct](#) or [routine](#) executes.
2 After the initial values are assigned, the values of any [OpenMP environment variables](#) that were set
3 by the user are read and the associated [ICVs](#) are modified accordingly. If no [device number](#) is
4 specified on the [device-specific environment variable](#) then the value is applied to all [non-host](#)
5 [devices](#).

6 **Cross References**

- 7 • [OMP_AFFINITY_FORMAT](#), see [Section 3.2.5](#)
- 8 • [OMP_ALLOCATOR](#), see [Section 3.5.1](#)
- 9 • [OMP_AVAILABLE_DEVICES](#), see [Section 3.2.7](#)
- 10 • [OMP_CANCELLATION](#), see [Section 3.2.6](#)
- 11 • [OMP_DEBUG](#), see [Section 3.4.1](#)
- 12 • [OMP_DEFAULT_DEVICE](#), see [Section 3.2.8](#)
- 13 • [OMP_DISPLAY_AFFINITY](#), see [Section 3.2.4](#)
- 14 • [OMP_DYNAMIC](#), see [Section 3.1.1](#)
- 15 • [OMP_MAX_ACTIVE_LEVELS](#), see [Section 3.1.4](#)
- 16 • [OMP_MAX_TASK_PRIORITY](#), see [Section 3.2.11](#)
- 17 • [OMP_NUM_TEAMS](#), see [Section 3.6.1](#)
- 18 • [OMP_NUM_THREADS](#), see [Section 3.1.2](#)
- 19 • [OMP_PLACES](#), see [Section 3.1.5](#)
- 20 • [OMP_PROC_BIND](#), see [Section 3.1.6](#)
- 21 • [OMP_SCHEDULE](#), see [Section 3.2.1](#)
- 22 • [OMP_STACKSIZE](#), see [Section 3.2.2](#)
- 23 • [OMP_TARGET_OFFLOAD](#), see [Section 3.2.9](#)
- 24 • [OMP_TEAMS_THREAD_LIMIT](#), see [Section 3.6.2](#)
- 25 • [OMP_THREAD_LIMIT](#), see [Section 3.1.3](#)
- 26 • [OMP_TOOL](#), see [Section 3.3.1](#)
- 27 • [OMP_TOOL_LIBRARIES](#), see [Section 3.3.2](#)
- 28 • [OMP_WAIT_POLICY](#), see [Section 3.2.3](#)

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 [routines](#) retrieve or modify ICVs with [data environment ICV scope](#) in the [data environment](#) of their [binding task set](#).

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>free-agent-thread-limit-var</i>	(none)	(none)
<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>
<i>place-assignment-var</i>	(none)	(none)
<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>structured-thread-limit-var</i>	(none)	(none)
<i>target-offload-var</i>	(none)	(none)
<i>team-generator-var</i>	(none)	(none)

ICV	Ways to Modify Value	Ways to Retrieve Value
<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 `omp_get_proc_bind` routine retrieves the value of the first element of this list.
- The value of the *nthreads-var ICV* is a list. The `omp_set_num_threads` routine sets the value of the first element of this list, and the `omp_get_max_threads` routine retrieves the value of the first element of this list.
- Detailed values in the *place-partition-var ICV* are retrieved using the listed routines.
- 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 14.3](#)
- `omp_get_active_level`, see [Section 19.2.18](#)
- `omp_get_affinity_format`, see [Section 19.3.9](#)
- `omp_get_cancellation`, see [Section 19.2.8](#)
- `omp_get_default_allocator`, see [Section 19.13.7](#)
- `omp_get_default_device`, see [Section 19.7.4](#)
- `omp_get_dynamic`, see [Section 19.2.7](#)
- `omp_get_level`, see [Section 19.2.15](#)
- `omp_get_max_active_levels`, see [Section 19.2.14](#)
- `omp_get_max_task_priority`, see [Section 19.5.1](#)
- `omp_get_max_teams`, see [Section 19.4.4](#)
- `omp_get_max_threads`, see [Section 19.2.3](#)

- 1 • `omp_get_num_procs`, see [Section 19.7.1](#)
- 2 • `omp_get_num_threads`, see [Section 19.2.2](#)
- 3 • `omp_get_partition_num_places`, see [Section 19.3.6](#)
- 4 • `omp_get_partition_place_nums`, see [Section 19.3.7](#)
- 5 • `omp_get_place_num_procs`, see [Section 19.3.3](#)
- 6 • `omp_get_place_proc_ids`, see [Section 19.3.4](#)
- 7 • `omp_get_proc_bind`, see [Section 19.3.1](#)
- 8 • `omp_get_schedule`, see [Section 19.2.10](#)
- 9 • `omp_get_supported_active_levels`, see [Section 19.2.12](#)
- 10 • `omp_get_teams_thread_limit`, see [Section 19.4.6](#)
- 11 • `omp_get_thread_limit`, see [Section 19.2.11](#)
- 12 • `omp_get_thread_num`, see [Section 19.2.4](#)
- 13 • `omp_in_final`, see [Section 19.5.3](#)
- 14 • `omp_set_affinity_format`, see [Section 19.3.8](#)
- 15 • `omp_set_default_allocator`, see [Section 19.13.6](#)
- 16 • `omp_set_default_device`, see [Section 19.7.3](#)
- 17 • `omp_set_dynamic`, see [Section 19.2.6](#)
- 18 • `omp_set_max_active_levels`, see [Section 19.2.13](#)
- 19 • `omp_set_num_teams`, see [Section 19.4.3](#)
- 20 • `omp_set_num_threads`, see [Section 19.2.1](#)
- 21 • `omp_set_schedule`, see [Section 19.2.9](#)
- 22 • `omp_set_teams_thread_limit`, see [Section 19.4.5](#)

23 2.4 How the Per-Data Environment ICVs Work

24 When a [task construct](#), a [parallel construct](#) or a [teams construct](#) is encountered, each
25 generated [task](#) inherits the values of the [ICVs with data environment ICV scope](#) from the [ICV](#)
26 values of the [generating task](#), unless otherwise specified.

27 When a [parallel construct](#) is encountered, the value of each [ICV with implicit task ICV scope](#)
28 is inherited from the [binding implicit task](#) of the [generating task](#) unless otherwise specified.

1 When a **task construct** is encountered, the generated **task** inherits the value of *nthreads-var* from
2 the *nthreads-var* value of the **generating task**. If a **parallel construct** is encountered on which a
3 **num_threads clause** is specified with a *nthreads* list of more than one list item, the value of
4 *nthreads-var* for the generated **implicit tasks** is the list obtained by deletion of the first item of the
5 *nthreads* list. Otherwise, when a **parallel construct** is encountered, if the *nthreads-var* list of
6 the **generating task** contains a single element, the generated **implicit tasks** inherit that list as the
7 value of *nthreads-var*; if the *nthreads-var* list of the **generating task** contains multiple elements, the
8 generated **implicit tasks** inherit the value of *nthreads-var* as the list obtained by deletion of the first
9 element from the *nthreads-var* value of the **generating task**. The *bind-var ICV* is handled in the
10 same way as the *nthreads-var ICV*, except that an override list cannot be specified through the
11 **proc_bind clause** of an encountered **parallel construct**.

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

15 When a **target task** executes an **inactive target region**, the generated **initial task** uses the values of the
16 *ICVs* with **data environment ICV scope** from the **data environment** of the **task** that encountered the
17 **target construct**, unless otherwise specified.

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

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

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

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

31 If a **target construct**, **teams construct**, or **parallel construct** is encountered, the
32 *team-generator-var ICV* for the **data environments** of the generated **implicit tasks** is instead set to
33 the value of the appropriate **team** generator type as specified in [Section 21.3.10](#).

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

37 Cross References

- 38 • Team Generator Types, see [Section 21.3.10](#)

2.5 ICV Override Relationships

Table 2.4 shows the override relationships among [construct clauses](#) and [ICVs](#). The table only lists [ICVs](#) that can be overridden by a [clause](#).

TABLE 2.4: ICV Override Relationships

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

If a [schedule clause](#) specifies a [modifier](#) then that [modifier](#) overrides any [modifier](#) that is specified in the *run-sched-var* ICV.

If *bind-var* is not set to *false* then the [proc_bind clause](#) overrides the value of the first element of the *bind-var* ICV; otherwise, the [proc_bind clause](#) has no effect.

Cross References

- `allocate` clause, see [Section 7.6](#)
- `allocator` clause, see [Section 7.4](#)
- `num_teams` clause, see [Section 11.3.1](#)
- `num_threads` clause, see [Section 11.2.2](#)
- `proc_bind` clause, see [Section 11.2.4](#)
- `schedule` clause, see [Section 12.6.3](#)
- `thread_limit` clause, see [Section 14.3](#)

3 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.

3.1 Parallel Region Environment Variables

This section defines environment variables that affect the operation of `parallel` regions.

3.1.1 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 11.2](#)
- *dyn-var* ICV, see [Table 2.1](#)
- `omp_get_dynamic`, see [Section 19.2.7](#)
- `omp_set_dynamic`, see [Section 19.2.6](#)

3.1.2 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 11.2.1](#) 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 3.1.4](#) for details.

11 Example:

```
12 setenv OMP_NUM_THREADS 4,3,2
```

13 Cross References

- 14 • **OMP_MAX_ACTIVE_LEVELS**, see [Section 3.1.4](#)
- 15 • **num_threads** clause, see [Section 11.2.2](#)
- 16 • **parallel** directive, see [Section 11.2](#)
- 17 • *nthreads-var* ICV, see [Table 2.1](#)
- 18 • **omp_set_num_threads**, see [Section 19.2.1](#)

19 3.1.3 OMP_THREAD_LIMIT

20 The **OMP_THREAD_LIMIT** environment variable sets the number of [threads](#) to use for a
21 [contention group](#) by setting the *thread-limit-var* ICV. The value of this environment variable must
22 be a positive integer. The behavior of the program is [implementation defined](#) if the requested value
23 of **OMP_THREAD_LIMIT** is greater than the number of [threads](#) an implementation can support, or
24 if the value is not a positive integer.

25 Cross References

- 26 • *thread-limit-var* ICV, see [Table 2.1](#)

27 3.1.4 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](#)

3.1.5 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 $\langle lower-bound \rangle : \langle length \rangle : \langle stride \rangle$ notation to represent the following list of numbers: “ $\langle lower-bound \rangle, \langle lower-bound \rangle + \langle stride \rangle, \dots, \langle lower-bound \rangle + (\langle length \rangle - 1) * \langle stride \rangle$.” When $\langle stride \rangle$ 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 3.1](#) 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*($\langle num_places \rangle$) or *abstract_name*($\langle num_places \rangle : \langle stride \rangle$), where $\langle num_places \rangle$ denotes the length of the [place list](#) and $\langle stride \rangle$ 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 3.1: Predefined Abstract Names for **OMP_PLACES**

Abstract Name	Meaning
threads	Each place corresponds to a single hardware thread on the device.

table continued on next page

Abstract Name	Meaning
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.

```

⟨list⟩ = ⟨p-list⟩ | ⟨aname⟩
⟨p-list⟩ = ⟨p-interval⟩ | ⟨p-list⟩,⟨p-interval⟩
⟨p-interval⟩ = ⟨place⟩:⟨len⟩:⟨stride⟩ | ⟨place⟩:⟨len⟩ | ⟨place⟩ | !⟨place⟩
⟨place⟩ = {⟨res-list⟩} | ⟨res⟩
⟨res-list⟩ = ⟨res-interval⟩ | ⟨res-list⟩,⟨res-interval⟩
⟨res-interval⟩ = ⟨res⟩:⟨num-places⟩:⟨stride⟩ | ⟨res⟩:⟨num-places⟩ | ⟨res⟩ | !⟨res⟩
⟨aname⟩ = ⟨word⟩(⟨num-places⟩:⟨stride⟩) | ⟨word⟩(⟨num-places⟩) | ⟨word⟩
⟨word⟩ = sockets | cores | ll_caches | numa_domains
        | threads | <implementation-defined abstract name>
⟨res⟩ = non-negative integer
⟨num-places⟩ = positive integer
⟨stride⟩ = integer
⟨len⟩ = positive integer

```


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 3.1.6 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 *threads* between *places*, *thread affinity* is
23 enabled, and the *initial thread* is bound to the first place in the *place-partition-var* ICV prior to the
24 first *active parallel region*. An *initial thread* that is created by a `teams` construct is bound to the first
25 place in its *place-partition-var* ICV before it begins execution of the associated *structured block*.

26 If the environment variable is set to `true`, the thread affinity policy is implementation defined but
27 must conform to the previous paragraph. The behavior of the program is implementation defined if
28 the value in the `OMP_PROC_BIND` environment variable is not `true`, `false`, or a comma
29 separated list of `primary`, `close`, or `spread`. The behavior is also implementation defined if
30 an initial thread cannot be bound to the first place in the *place-partition-var* ICV.

31 The `OMP_PROC_BIND` environment variable sets the *max-active-levels-var* ICV to the number of
32 active levels of parallelism that the implementation supports if the `OMP_PROC_BIND` environment
33 variable is set to a comma-separated list of more than one element. The value of the
34 *max-active-level-var* ICV may be overridden by setting `OMP_MAX_ACTIVE_LEVELS`. See
35 [Section 3.1.4](#) 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 3.1.4](#)
- 6 • `proc_bind` clause, see [Section 11.2.4](#)
- 7 • `parallel` directive, see [Section 11.2](#)
- 8 • `teams` directive, see [Section 11.3](#)
- 9 • Controlling OpenMP Thread Affinity, see [Section 11.2.3](#)
- 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 19.3.1](#)

14 3.2 Program Execution Environment Variables

15 This section defines environment variables that affect program execution.

16 3.2.1 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 12.6.3](#) 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 12.6.3](#)
- 7 • `run-sched-var` ICV, see [Table 2.1](#)

8 3.2.2 OMP_STACKSIZE

9 The `OMP_STACKSIZE` environment variable controls the size of the stack for [threads](#), by setting
10 the value of the [stacksize-var ICV](#). The environment variable does not control the size of the stack
11 for an [initial thread](#). Whether this environment variable also controls the size of the stack of [native](#)
12 [threads](#) is [implementation defined](#). The value of this environment variable takes the form `size[unit]`,
13 where:

- 14 • `size` is a positive integer that specifies the size of the stack for [threads](#).
- 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 3.2.3 OMP_WAIT_POLICY

31 The `OMP_WAIT_POLICY` environment variable provides a hint to an OpenMP implementation
32 about the desired behavior of waiting [native threads](#) by setting the [wait-policy-var ICV](#). A
33 [compliant 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:

35 `active` | `passive`

1 The **active** value specifies that waiting **native threads** should mostly be active, consuming
2 **processor** cycles, while waiting. A **compliant implementation** may, for example, make waiting
3 **native threads** spin. The **passive** value specifies that waiting **native threads** should mostly be
4 passive, not consuming processor cycles, while waiting. For example, a **compliant implementation**
5 may make waiting **native threads** yield the processor to other **native threads** or go to sleep. The
6 details of the **active** and **passive** behaviors are **implementation defined**. The behavior of the
7 program is **implementation defined** if the value of **OMP_WAIT_POLICY** is neither **active** nor
8 **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 3.2.4 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 3.2](#) 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:

```
36 nesting_level= 1,   thread_num= 0,   thread_affinity= 0,1  
37 nesting_level= 1,   thread_num= 1,   thread_affinity= 2,3
```

Cross References

- `OMP_AFFINITY_FORMAT`, see [Section 3.2.5](#)
- Controlling OpenMP Thread Affinity, see [Section 11.2.3](#)
- `affinity-format-var` ICV, see [Table 2.1](#)
- `display-affinity-var` ICV, see [Table 2.1](#)

3.2.5 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 [thread affinity](#) information. The value of this environment variable is case sensitive and leading and trailing [white space](#) 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 3.2: Available Field Types for Formatting OpenMP Thread Affinity Information

Short Name	Long Name	Meaning
<code>t</code>	<code>team_num</code>	The value returned by <code>omp_get_team_num()</code> .
<code>T</code>	<code>num_teams</code>	The value returned by <code>omp_get_num_teams()</code> .
<code>L</code>	<code>nesting_level</code>	The value returned by <code>omp_get_level()</code> .
<code>n</code>	<code>thread_num</code>	The value returned by <code>omp_get_thread_num()</code> .

table continued on next page

table continued from previous page

Short Name	Long Name	Meaning
N	<code>num_threads</code>	The value returned by <code>omp_get_num_threads()</code> .
a	<code>ancestor_tnum</code>	The value returned by <code>omp_get_ancestor_thread_num(level)</code> , where <code>level</code> is <code>omp_get_level()</code> minus 1.
H	<code>host</code>	The name for the host device on which the OpenMP program is running.
P	<code>process_id</code>	The process identifier used by the implementation.
i	<code>native_thread_id</code>	The native thread identifier used by the implementation.
A	<code>thread_affinity</code>	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 11.2.3](#)
- 13 • `affinity-format-var` ICV, see [Table 2.1](#)
- 14 • `omp_get_ancestor_thread_num`, see [Section 19.2.16](#)
- 15 • `omp_get_level`, see [Section 19.2.15](#)
- 16 • `omp_get_num_teams`, see [Section 19.4.1](#)
- 17 • `omp_get_num_threads`, see [Section 19.2.2](#)

- `omp_get_thread_num`, see [Section 19.2.4](#)

- `omp_get_thread_num`, see [Section 19.2.4](#)

3.2.6 OMP_CANCELLATION

The `OMP_CANCELLATION` environment variable sets the initial value of the *cancel-var* ICV. The value of this environment variable must be one of the following:

`true|false`

If the environment variable is set to `true`, the effects of the `cancel` construct and of cancellation points are enabled (i.e., cancellation is enabled). If the environment variable is set to `false`, cancellation is disabled and the `cancel` construct and cancellation points are effectively ignored. The behavior of the program is implementation defined if `OMP_CANCELLATION` is set to neither `true` nor `false`.

Cross References

- `cancel` directive, see [Section 17.2](#)
- *cancel-var* ICV, see [Table 2.1](#)

3.2.7 OMP_AVAILABLE_DEVICES

The `OMP_AVAILABLE_DEVICES` environment variable sets the *available-devices-var* ICV and determines the available non-host `devices` and their `device` numbers by permitting selection of `devices` from the set of supported `accessible devices` and by ordering them. This ICV is initialized before any other ICV that uses a `device` number, depends on the number of available `devices`, or permits device-specific environment variables. After the *available-devices-var* ICV is initialized, only those `devices` that the ICV identifies are available and the `omp_get_num_devices` routine returns the number of `devices` stored in the ICV.

The value of this environment variable must be a comma-separated list. Each item is either a trait specification as specified in the following or `*`. A `*` expands to all accessible and `supported devices` while a trait specification expands to a possibly empty set of accessible and `supported devices` for which the specification is fulfilled. After expansion, further selection via an optional array subscript syntax and removal of `devices` that appear in previous items, each item contains an unordered set of `devices`. A consecutive unique `device` number is then assigned to each `device` in the sets, starting with `device` number zero, where the `device` number of the first `device` in an item is the total number of `devices` in all previous items.

Traits are specified by the case-insensitive trait name followed by the argument in parentheses. The permitted traits are `kind` (*kind-name*), `isa` (*isa-name*), `arch` (*arch-name*), and `vendor` (*vendor-name*), where the names are as specified in [Section 8.1](#) and the [OpenMP Additional Definitions document](#); the *kind-name* `host` is not permitted. Multiple traits can be combined using the binary operators `&&` and `|` to require both or either trait, respectively.

1 Parentheses can be used for grouping, but are optional except that **&&** and **||** may not appear in the
2 same grouping level. The unary **!** operator inverts the meaning of the immediately following trait
3 or parenthesized group.

4 Each trait specification or ***** yields a (possibly zero-sized) array of non-host **devices** with the lowest
5 array element, if it exists, having index zero. The C/C++ syntax `[index]` can be used to select an
6 element and the **array section** syntax for C/C++ as specified in [Section 4.2.5](#) can be used to specify
7 a subset of elements. Any array element specified by the subscript that is outside the bounds of the
8 array resulting from the trait specification or ***** is silently excluded.

9 **Cross References**

- 10 • Device Directives and Clauses, see [Chapter 14](#)
- 11 • *available-devices-var* ICV, see [Table 2.1](#)

12 **3.2.8 OMP_DEFAULT_DEVICE**

13 The **OMP_DEFAULT_DEVICE** environment variable sets the initial value of the *default-device-var*
14 **ICV**. The value of this environment variable must be a comma-separated list, each item being either
15 a non-negative integer value that denotes the **device** number, a trait specification with an optional
16 subscript selector, or one of the following case-insensitive string literals: **initial** to specify the
17 **host device**, **invalid** to specify the **device** number **omp_invalid_device**, or **default** to
18 set the **ICV** as if this environment variable was not specified (see [Section 1.3](#)).

19 The trait specification is as described for **OMP_AVAILABLE_DEVICES** (see [Section 3.2.7](#)), except
20 that in addition the trait *device_num(device number)* may be specified, **host** is permitted as
21 *kind-name*. The **device** numbers yielded by the trait specification are sorted in ascending order by
22 **device** number; the array-element syntax as described in **OMP_AVAILABLE_DEVICES** can be
23 used to select an element from the set. If an item is an empty set, non-existing element, or does not
24 evaluate to an available **device**, the next item is evaluated; otherwise, the *default-device-var* ICV is
25 set to the first value of the set. However, **initial**, **invalid**, and **default** always match. If
26 none of the list items match, the *default-device-var* **ICV** is set to **omp_invalid_device**.

27 **Cross References**

- 28 • Device Directives and Clauses, see [Chapter 14](#)
- 29 • *default-device-var* ICV, see [Table 2.1](#)

30 **3.2.9 OMP_TARGET_OFFLOAD**

31 The **OMP_TARGET_OFFLOAD** environment variable sets the initial value of the *target-offload-var*
32 **ICV**. Its value must be one of the following:

33 **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 14](#)
- Device Memory Routines, see [Section 19.8](#)
- *target-offload-var* ICV, see [Table 2.1](#)

3.2.10 OMP_THREADS_RESERVE

The **OMP_THREADS_RESERVE** environment variable controls the number of [reserved threads](#) in each [contention group](#) by setting the initial value of the *structured-thread-limit-var* and the *free-agent-thread-limit-var* ICVs [structured parallelism](#),

The **OMP_THREADS_RESERVE** environment variable can be defined using a non-negative integer or an unordered list of reservations. Each reservation specifies a [thread-reservation type](#), for which the possible values are listed in [Table 3.3](#). The [reservation type](#) may be appended with one non-negative number in parentheses, that is, *reservation_type(<num-threads>)*, where *<num-threads>* denotes the number of [threads](#) to reserve for that [reservation type](#). If only a non-negative integer is provided, this number denotes the number of [threads](#) to reserve for [structured parallelism](#). If only one [reservation type](#) is provided, and its *<num-threads>* is not specified, the number of [threads](#) to reserve is *thread-limit-var* if the [reservation type](#) is **structured**, or *thread-limit-var* minus 1 if the [reservation type](#) is **free_agent**.

TABLE 3.3: Reservation Types for **OMP_THREADS_RESERVE**

Reservation Type	Meaning	Default Value
Number of structured	Threads reserved for structured threads .	1
Number of free_agent	Threads reserved for free-agent threads .	0

The **OMP_THREADS_RESERVE** environment variable sets the initial value of the *structured-thread-limit-var* and the *free-agent-thread-limit-var* ICVs according to [Algorithm 3.1](#).

The following grammar describes the values accepted for the **OMP_THREADS_RESERVE** environment variable.

Algorithm 3.1 Initial Values of the *structured-thread-limit-var* and *free-agent-thread-limit-var* ICVs

let *structured-reserve* be the number of **threads** to reserve for **structured threads**;
let *free-agent-reserve* be the number of **threads** to reserve for **free-agent threads**;
let *threads-reserve* be the sum of *structured-reserve* and *free-agent-reserve*;
if (*structured-reserve* < 1) **then** *structured-reserve* = 1;
if (*free-agent-reserve* = *thread-limit-var*) **then** *free-agent-reserve* = *free-agent-reserve* - 1;
if (*threads-reserve* ≤ *thread-limit-var*) **then**
 structured-thread-limit-var = *thread-limit-var* - *free-agent-reserve*;
 free-agent-thread-limit-var = *thread-limit-var* - *structured-reserve*;
else behavior is **implementation defined**

$\langle \text{reserve} \rangle \models \langle \text{res-list} \rangle \mid \langle \text{res-type} \rangle \mid \langle \text{res-num} \rangle$
 $\langle \text{res-list} \rangle \models \langle \text{res} \rangle \mid \langle \text{res-list} \rangle, \langle \text{res} \rangle$
 $\langle \text{res} \rangle \models \langle \text{res-type} \rangle (\langle \text{res-num} \rangle)$
 $\langle \text{res-type} \rangle \models \text{structured} \mid \text{free_agent}$
 $\langle \text{res-num} \rangle \models \text{non-negative integer}$

1 Examples:

```
2 setenv OMP_THREADS_RESERVE 4  
3 setenv OMP_THREADS_RESERVE "structured(4)"  
4 setenv OMP_THREADS_RESERVE "structured"  
5 setenv OMP_THREADS_RESERVE "structured(2), free_agent(2)"
```

6 where the first two definitions correspond to the same reservation for **structured parallelism**, the
7 third definition reserves all available **threads** for **structured parallelism**, and the last one reserves
8 **threads** for both **structured parallelism** and **free-agent threads**.

9 Cross References

- 10 • **threadset** clause, see [Section 13.4](#)
- 11 • **parallel** directive, see [Section 11.2](#)
- 12 • *free-agent-thread-limit-var* ICV, see [Table 2.1](#)
- 13 • *structured-thread-limit-var* ICV, see [Table 2.1](#)

3.2.11 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](#)

3.3 OMPT Environment Variables

This section defines environment variables that affect operation of the OMPT tool interface.

3.3.1 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 20](#)
- *tool-var* ICV, see [Table 2.1](#)

3.3.2 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

1 initialized or if `ompt_start_tool` returns `NULL`, an OpenMP implementation will consider
2 libraries in the `tool-libraries-var` list in a left-to-right order. The OpenMP implementation will
3 search the list for a library that meets two criteria: it can be dynamically loaded on the current
4 device and it defines the symbol `ompt_start_tool`. If an OpenMP implementation finds a
5 suitable library, no further libraries in the list will be considered.

6 Example:

```
7 % setenv OMP_TOOL_LIBRARIES libtoolXY64.so:/usr/local/lib/  
8 libtoolXY32.so
```

9 Cross References

- 10 • OMPT Interface, see [Chapter 20](#)
- 11 • `tool-libraries-var` ICV, see [Table 2.1](#)
- 12 • `ompt_start_tool`, see [Section 20.2.1](#)

13 3.3.3 OMP_TOOL_VERBOSE_INIT

14 The `OMP_TOOL_VERBOSE_INIT` environment variable sets the `tool-verbose-init-var` ICV, which
15 controls whether an OpenMP implementation will verbosely log the registration of a tool. The
16 value of this environment variable must be one of the following:

17 `disabled | stdout | stderr | <filename>`

18 If `OMP_TOOL_VERBOSE_INIT` is set to any value other than case insensitive `disabled`,
19 `stdout`, or `stderr`, the value is interpreted as a filename and the OpenMP runtime will try to
20 log to a file with prefix `filename`. If the value is interpreted as a filename, whether it is case
21 sensitive is implementation defined. If opening the logfile fails, the output will be redirected to
22 `stderr`. If `OMP_TOOL_VERBOSE_INIT` is not defined, the default value for `tool-verbose-init-var`
23 is `disabled`. Support for logging to `stdout` or `stderr` is implementation defined. Unless
24 `tool-verbose-init-var` is `disabled`, the OpenMP runtime will log the steps of the tool activation
25 process defined in [Section 20.2.2](#) to a file with a name that is constructed using the provided
26 filename prefix. The format and detail of the log is implementation defined. At a minimum, the log
27 will contain one of the following:

- 28 • That the `tool-var` ICV is disabled;
- 29 • An indication that a tool was available in the address space at program launch; or
- 30 • The path name of each tool in `OMP_TOOL_LIBRARIES` that is considered for dynamic
31 loading, whether dynamic loading was successful, and whether the `ompt_start_tool`
32 function is found in the loaded library.

33 In addition, if an `ompt_start_tool` function is called the log will indicate whether or not the
34 tool will use the OMPT interface.

1 Example:

```
2 % setenv OMP_TOOL_VERBOSE_INIT disabled  
3 % setenv OMP_TOOL_VERBOSE_INIT STDERR  
4 % setenv OMP_TOOL_VERBOSE_INIT omp_load.log
```

5 Cross References

- 6 • OMPT Interface, see [Chapter 20](#)
- 7 • *tool-verbose-init-var* ICV, see [Table 2.1](#)

8 3.4 OMPD Environment Variables

9 This section defines environment variables that affect operation of the OMPD tool interface.

10 3.4.1 OMP_DEBUG

11 The `OMP_DEBUG` environment variable sets the *debug-var* ICV, which controls whether an
12 OpenMP runtime collects information that an OMPD library may need to support a tool. The value
13 of this environment variable must be one of the following:

14 **enabled** | **disabled**

15 If `OMP_DEBUG` is set to any value other than **enabled** or **disabled** then the behavior is
16 implementation defined.

17 Example:

```
18 % setenv OMP_DEBUG enabled
```

19 Cross References

- 20 • Enabling Runtime Support for OMPD, see [Section 21.2.1](#)
- 21 • OMPD Interface, see [Chapter 21](#)
- 22 • *debug-var* ICV, see [Table 2.1](#)

3.5 Memory Allocation Environment Variables

This section defines environment variables that affect memory allocations.

3.5.1 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.

$$\begin{aligned} \langle \text{allocator} \rangle & \models \langle \text{predef-allocator} \rangle \mid \langle \text{predef-mem-space} \rangle \mid \langle \text{predef-mem-space} \rangle : \langle \text{traits} \rangle \\ \langle \text{traits} \rangle & \models \langle \text{trait} \rangle = \langle \text{value} \rangle \mid \langle \text{trait} \rangle = \langle \text{value} \rangle, \langle \text{traits} \rangle \\ \langle \text{predef-allocator} \rangle & \models \textit{one of the predefined allocators from Table 7.3} \\ \langle \text{predef-mem-space} \rangle & \models \textit{one of the predefined memory spaces from Table 7.1} \\ \langle \text{trait} \rangle & \models \textit{one of the allocator trait names from Table 7.2} \\ \langle \text{value} \rangle & \models \textit{one of the allowed values from Table 7.2} \mid \textit{non-negative integer} \\ & \mid \langle \text{predef-allocator} \rangle \end{aligned}$$

The *value* can be an integer only if the *trait* accepts a numerical value, for the **fb_data** *trait* the *value* can only be *predef-allocator*. If the value of this environment variable is not a predefined allocator, then a new allocator with the given predefined memory space and optional traits is created and set as the *def-allocator-var* ICV. If the new allocator cannot be created, the *def-allocator-var* ICV will be set to **omp_default_mem_alloc**.

Example:

```
setenv OMP_ALLOCATOR omp_high_bw_mem_alloc
setenv OMP_ALLOCATOR omp_large_cap_mem_space:alignment=16,\
pinned=true
setenv OMP_ALLOCATOR omp_high_bw_mem_space:pool_size=1048576,\
fallback=allocator_fb,fb_data=omp_low_lat_mem_alloc
```

Cross References

- Memory Allocators, see [Section 7.2](#)
- *def-allocator-var* ICV, see [Table 2.1](#)

3.6 Teams Environment Variables

This section defines environment variables that affect the operation of **teams** regions.

3.6.1 OMP_NUM_TEAMS

The **OMP_NUM_TEAMS** environment variable sets the maximum number of teams created by a **teams** construct by setting the *nteams-var* ICV. The value of this environment variable must be a positive integer. The behavior of the program is implementation defined if the requested value of **OMP_NUM_TEAMS** is greater than the number of teams that an implementation can support, or if the value is not a positive integer.

Cross References

- **teams** directive, see [Section 11.3](#)
- *nteams-var* ICV, see [Table 2.1](#)

3.6.2 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 11.3](#)
- *teams-thread-limit-var* ICV, see [Table 2.1](#)

3.7 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 19.15](#)). 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

1 environment variable other than **true**, **false**, and **verbose**, the displayed information is
2 unspecified.

3 Example:

```
4 | % setenv OMP_DISPLAY_ENV true
```

5 For the output of the above example, see [Section 19.15](#).

6 **Cross References**

- 7 • Environment Display Routine, see [Section 19.15](#)

4 Directive and Construct Syntax

This chapter describes the syntax of **directives** and **clauses** and their association with **base language** code. **Directives** are specified with various **base language** mechanisms that allow compilers to ignore the **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 **directives** and 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 regions** (that arise from **executable directives**) are as follows:

C++

- A **throw** executed inside a **region** that arises from a **thread-limiting construct** must cause execution to resume within the same **region**, and the same **thread** that threw the exception must catch it. If the **directive** also has the **exception-aborting property** 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 has the **pure property**.
- 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 region** will result in **unspecified behavior**.

Fortran

4.1 Directive Format

This section defines several categories of **directives** and **constructs**. **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 **directive**:

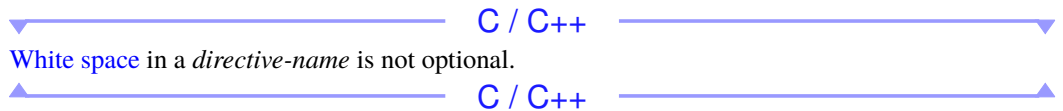
```
directive-specifier [ [ , ] clause [ [ , ] clause ] ... ]
```

The *directive-specifier* is:

```
directive-name
```

or for argument-modified directives:

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

C / C++

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

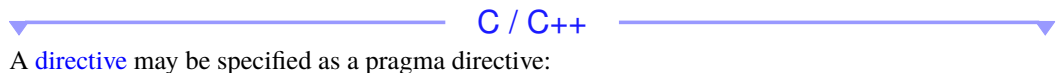
Some **directives** specify a paired **end directive**, where the *directive-name* of the paired **end directive** is:

- 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**.

C / C++

A **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** 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 a *directive-specification* of **#pragma** and **_Pragma** pragmas are subject to macro expansion.

C / C++

C / C++

In C23 and later versions or C++11 and later versions, a *directive* may be specified as a C/C++ attribute specifier:

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

C++

or

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

C++

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(...) {}
```

1 The use of **omp** as the attribute namespace of an attribute specifier, or as the optional namespace
2 qualifier within a **sequence** attribute, is reserved for OpenMP **directives** that are defined in this
3 specification. The use of **omp_x** as the attribute namespace of an attribute specifier, or as the
4 optional namespace qualifier within a **sequence** attribute, is reserved for implementation-defined
5 extensions to the OpenMP **directives**.

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

▲ C / C++ ▲
▼ C / C++ ▼

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

▲ C / C++ ▲
▼ C++ ▼

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

▲ C++ ▲
▼ Fortran ▼

13 A **directive** for Fortran is specified with a stylized comment as follows:

14 *sentinel directive-specification*

15 All **directives** must begin with a **directive sentinel**. The format of a sentinel differs between fixed
16 form and free form source files, as described in [Section 4.1.1](#) and [Section 4.1.2](#). In order to simplify
17 the presentation, free form is used for the syntax of **directives** for Fortran throughout this document,
18 except as noted.

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

▲ Fortran ▲

22 A **directive** may be categorized as one of the following:

- 23 • **metadirective**
- 24 • **declarative directive**
- 25 • **executable directive**
- 26 • **informational directive**
- 27 • **utility directive**
- 28 • **subsidiary directive**

1 Base language code can be associated with directives. The association of a directive can be
2 categorized as:

- 3 • none
- 4 • block-associated directive
- 5 • loop-nest-associated directive
- 6 • loop-sequence-associated directive
- 7 • declaration-associated directive
- 8 • delimited directive
- 9 • separating directive

C / C++

10 A declarative directive that is declaration-associated may alternatively be expressed as an attribute
11 specifier:

```
12 | [[ omp :: decl( directive-specification ) ]]
```

C++

13 or

```
14 | [[ using omp : decl( directive-specification ) ]]
```

C++

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

C / C++

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

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

```
25 | directive
```

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

C / C++

```
27 | directive  
28 |   structured-block
```

C / C++

Fortran

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

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

Fortran

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

Fortran

The associated *structured-block* of a block-associated **directives** can be a **DO CONCURRENT** loop where it is explicitly allowed.

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

Fortran

C / C++

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

```
declaration-associated-specification
```

where *declaration-associated-specification* is either:

```
directive
  function-definition-or-declaration
```

or:

```
directive
declaration-associated-specification
```

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

C / C++

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.

If a **directive** appears in the specification part of a module then the behavior is as if that **directive** appears in the specification part of any **compilation unit** that references the module with a **USE** statement unless otherwise specified.

Fortran

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

```
directive  
  base-language-code  
end-directive
```

Separating **directives** are used to split statements contained in a **structured block** that is associated with a **construct** (the **separated construct**) into multiple **structured block sequences**. If the **separated construct** is a **loop-nest-associated construct** then any **separating directives** divide the loop body of the innermost **associated loop** into **structured block sequences**. Otherwise, the **separating directives** divide the associated **structured block** into **structured block sequences**.

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

```
structured-block-sequence  
directive  
structured-block-sequence  
[directive  
structured-block-sequence ...]
```

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

C / C++

Formations that result from **directives** that are specified as attribute specifiers that use the **directive** attribute are specified as follows. If the **directive** has an association of none, the resulting formation is an *attribute-declaration* if the **directive** is not executable and it consists of the attribute specifier and a null statement (i.e., “;”) if the **directive** is executable. For a block-associated **directive** or **loop-nest-associated directive**, the resulting formation consists of the attribute specifier and a **structured block** to which the specifier applies. If the **directives** are separating or delimited then the resulting formation is as previously specified except the attribute specifier for each **directive**, including the **end** directive, applies to a null statement.

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

C / C++

Restrictions

Restrictions to **directive** format are as follows:

- Orphaned separating **directives** are prohibited. That is, the separating **directives** must appear within the **structured block** associated with the same **construct** with which it is associated and must not be encountered elsewhere in the **region** of that associated **construct**.

- A **stand-alone directive** may be placed only at a point where a **base language** executable statement is allowed.

Fortran

- **Directives** may not appear in the **WHERE**, **FORALL**, or **DO CONCURRENT** constructs.
- A **declarative directive** must be specified in the specification part after all **USE**, **IMPORT** and **IMPLICIT** statements.

Fortran

C / 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 or C++ attribute specifiers, the specified attributes must be applied to the function as permitted by the **base language**.

C / 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

4.1.1 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.

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

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

```
8 c23456789  
9 !$omp parallel do shared(a,b,c)  
10  
11 c$omp parallel do  
12 c$omp+shared(a,b,c)  
13  
14 c$omp paralleldoshared(a,b,c)
```

15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
Fortran
Fortran

4.1.2 Free Source Form Directives

The following sentinels are recognized in free form source files:

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

The **!\$omp** sentinel is reserved for OpenMP **directives** that are defined in this specification. The **!\$omp_x** sentinel is reserved for **implementation defined** extensions to the OpenMP **directives**.

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

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

One or more blanks or horizontal tabs are optional to separate adjacent keywords in *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

4.2 Clause Format

This section defines the format and categories of OpenMP **clauses**. **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. Some **clauses** form natural groupings that have similar semantic effect and so are frequently specified as a clause grouping. A *clause-specification* specifies each **clause** in a *directive-specification* where *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 **omp_x_** 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.

The first *clause-argument-specification* is required unless otherwise explicitly specified while additional ones are only permitted on **clauses** that explicitly allow them. When the first one is omitted, the syntax is simply:

```
clause-name
```

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 [clauses](#), *clause-argument-list* is an OpenMP argument list, which is a comma-separated
6 list of a specific kind of list items (see [Section 4.2.1](#)), 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 a **directive**. These **properties** are summarized in Table 4.1, along with the respective
 7 default **properties** for **clauses**, arguments and **modifiers**.

TABLE 4.1: 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 Some **clause** properties determine the **constituent directives** to which they apply when specified on
 13 **combined directives** and **composite directives**. A **clause** with the all-constituents property applies to
 14 all **constituent directives** of any **combined directive** or **composite directive** on which it is specified.
 15 Unless otherwise specified, a **clause** has the all-constituents property. That is, the all-constituents
 16 property is a default **clause** property. A **clause** with the once-for-all-constituents property applies to
 17 the **directive** once, before any of the **constituent directives** are applied. A **clause** with the
 18 innermost-leaf property applies to the innermost **constituent directive** to which it may be applied. A
 19 **clause** with the outermost-leaf property applies to the outermost **constituent directive** to which it
 20 may be applied. A **clause** with the all-privatizing property applies to all **constituent directives** that
 21 permit the **clause** and to which a data-sharing attribute **clause** that may create a private copy of the
 22 same list item is applied.

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

25

26 **Note** – In this example, *clause-specification* is **depend (inout : d)**, *clause-name* is **depend**
 27 and *clause-argument-specification* is **inout : d**. The **depend clause** has an argument for which

1 *argument-name* is *locator-list*, which syntactically is the OpenMP locator list **d** in the example.
2 Similarly, the **depend clause** accepts a simple **modifier** with the name *task-dependence-type*.
3 Syntactically, *task-dependence-type* is the keyword **inout** in the example.

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

5
6 The **clauses** that a **directive** accepts may form sets. These sets may imply restrictions on their use
7 on that **directive** or may otherwise capture properties for the **clauses** on the **directive**. While specific
8 properties may be defined for a **clause** set on a particular **directive**, the following clause-set
9 properties have general meanings and implications as indicated by the restrictions below: required,
10 unique, and exclusive.

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

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

17 Restrictions

18 Restrictions to **clauses** and **clause** sets are as follows:

- 19 • A required **clause** for a **directive** must appear on the **directive**.
- 20 • A unique **clause** for a **directive** may appear at most once on the **directive**.
- 21 • An exclusive **clause** for a **directive** must not appear if a **clause** with a different *clause-name*
22 also appears on the **directive**.
- 23 • An ultimate **clause** for a **directive** must be the lexically last **clause** to appear on the **directive**.
- 24 • If a **clause** set has the required property, at least one **clause** in the set must be present on the
25 **directive** for which the **clause** set is specified.
- 26 • If a **clause** is a member of a set that has the unique property for a **directive** then the **clause** has
27 the unique property for that **directive** regardless of whether it has the unique property when it
28 is not part of such a set.
- 29 • If one **clause** of a **clause** set with the exclusive property appears on a **directive**, no other
30 **clauses** with a different *clause-name* in that set may appear on the **directive**.
- 31 • A required argument must appear in the *clause-specification*, unless otherwise specified.
- 32 • A unique argument may appear at most once in a *clause-argument-specification*.
- 33 • An exclusive argument must not appear if an argument with a different *argument-name*
34 appears in the *clause-argument-specification*.
- 35 • A required **modifier** must appear in the *clause-argument-specification*.

- 1 • A unique **modifier** may appear at most once in a *clause-argument-specification*.
- 2 • An exclusive **modifier** must not appear if a **modifier** with a different *modifier-name* also
- 3 appears in the *clause-argument-specification*.
- 4 • If a **clause** is pre-modified, an ultimate **modifier** must be the last **modifier** in a
- 5 *clause-argument-specification* in which any **modifier** appears.
- 6 • If a **clause** is post-modified, an ultimate **modifier** must be the first **modifier** in a
- 7 *clause-argument-specification* in which any **modifier** appears.
- 8 • A **modifier** that is an expression must neither lexically match the name of a simple **modifier**
- 9 defined for the **clause** that is an OpenMP keyword nor *modifier-name parenthesized-tokens*,
- 10 where *modifier-name* is the *modifier-name* of a complex **modifier** defined for the **clause** and
- 11 *parenthesized-tokens* is a token sequence that starts with (and ends with) .
- 12 • A constant argument or parameter must be a compile-time constant.
- 13 • A positive argument or parameter must be greater than zero; a non-negative argument or
- 14 parameter must be greater than or equal to zero.
- 15 • A region-invariant argument or parameter must have the same value throughout any given
- 16 execution of the **construct** or, for **declarative directives**, execution of the function or
- 17 subroutine with which the declaration is associated.

18 **Cross References**

- 19 • Directive Format, see [Section 4.1](#)
- 20 • OpenMP Argument Lists, see [Section 4.2.1](#)
- 21 • OpenMP Stylized Expressions, see [Section 5.2](#)
- 22 • OpenMP Types and Identifiers, see [Section 5.1](#)

23 **4.2.1 OpenMP Argument Lists**

24 The OpenMP API defines several kinds of lists, each of which can be used as syntactic instances of

25 **clause** arguments. A list of any OpenMP type consists of a comma-separated collection of one or

26 more expressions of that OpenMP type. A **variable** list consists of a comma-separated collection of

27 one or more *variable list items*. An extended list consists of a comma-separated collection of one or

28 more *extended list items*. A locator list consists of a comma-separated collection of one or more

29 *locator list items*. A parameter list consists of a comma-separated collection of one or more

30 *parameter list items*. A type-name list consists of a comma-separated collection of one or more

31 *type-name list items*. A directive-name list consists of a comma-separated collection of one or more

32 *directive-name list items*, each of which is the *directive-name* of some OpenMP **directive**. A

33 **directive** specification list consists of a comma-separated collection of one or more

34 *directive-specification list items*, each of which is an OpenMP *directive-specification*. A foreign

35 runtime preference list consists of a comma-separated collection of one or more *foreign-runtime list*

1 *items*, each of which is an OpenMP *foreign-runtime* identifier. An OpenMP operation list consists
2 of a comma-separated collection of one or more *OpenMP operation list items*, each of which is an
3 OpenMP operation defined in [Section 4.2.3](#).

▼ C / C++ ▼

4 A *variable list item* is a [variable](#) or an [array section](#). An *extended list item* is a *variable list item* or a
5 function name. A *locator list item* is any lvalue expression including [variables](#), [array sections](#), and
6 reserved locators. A *parameter list item* is the name of a function parameter. A *type-name list item*
7 is a type name.

▲ C / C++ ▲

▼ Fortran ▼

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

- 9 • a [variable](#) that is not coindexed and that is not a substring;
- 10 • an [array section](#) that is not coindexed and that does not contain an element that is a substring;
- 11 • a *named constant*;
- 12 • an associate name that may appear in a [variable](#) definition context; or
- 13 • a common block name (enclosed in slashes).

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

18 A *named constant* as a *list item* can appear only in [clauses](#) where it is explicitly allowed.

19 When a named common block appears in an OpenMP argument list, it has the same meaning and
20 restrictions as if every explicit member of the common block appeared in the list. An explicit
21 member of a common block is a [variable](#) that is named in a **COMMON** statement that specifies the
22 common block name and is declared in the same scoping unit in which the [clause](#) appears. Named
23 common blocks do not include the blank common block.

24 Although [variables](#) in common blocks can be accessed by use association or host association,
25 common block names cannot. As a result, a common block name specified in a [clause](#) must be
26 declared to be a common block in the same scoping unit in which the [clause](#) appears. **construct**.

▲ Fortran ▲

Restrictions

The restrictions to OpenMP lists are as follows:

- Unless otherwise specified, OpenMP list items must be directive-wide unique, i.e., a list item can only appear once in one OpenMP list of all arguments, [clauses](#), and [modifiers](#) of the [directive](#).
- All list items must be visible, according to the scoping rules of the [base language](#).
- The *directive-specifier* and the [clauses](#) in a *directive-specification* item must not be comma-separated.

C

- Unless otherwise specified, a [variable](#) that is part of an [aggregate variable](#) must not be a [variable](#) list item or an extended list item.

C

C++

- Unless otherwise specified, a [variable](#) that is part of an [aggregate variable](#) must not be a [variable](#) list item or an extended 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 an [aggregate variable](#) must not be a [variable](#) list item or an extended list item.

Fortran

4.2.2 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](#).

4.2.3 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 5.1](#)). In general, the format of an OpenMP operation is the following:

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

C / C++

4.2.4 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 [containing array](#) of the list item or a [containing array](#) of one of its [named pointers](#).
- The type of the expression upon which a shape-operator is applied must be a pointer type.

C++

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

C++

C / C++

4.2.5 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 ▲

4.2.6 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 13.6.1](#)
- 19 • **depend** clause, see [Section 16.9.5](#)
- 20 • **from** clause, see [Section 6.9.2](#)
- 21 • **map** clause, see [Section 6.8.3](#)
- 22 • **to** clause, see [Section 6.9.1](#)

23 **4.3 Conditional Compilation**

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

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



5 4.3.1 Fixed Source Form Conditional Compilation Sentinels

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

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

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


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

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

18

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

```
21 c23456789  
22 !$ 10 iam = omp_get_thread_num() +  
23 !$ & index  
24  
25 #ifdef _OPENMP  
26 10 iam = omp_get_thread_num() +  
27 & index  
28 #endif
```



4.3.2 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):

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

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

4.4 *directive-name-modifier* Modifier

Modifiers

Name	Modifies	Type	Properties
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: directive-name	unique

Clauses

`acq_rel`, `acquire`, `adjust_args`, `affinity`, `align`, `aligned`, `allocate`, `allocator`, `append_args`, `apply`, `at`, `atomic_default_mem_order`, `bind`, `capture`, `collapse`, `collector`, `combiner`, `compare`, `copyin`, `copyprivate`, `default`, `defaultmap`, `depend`, `destroy`, `detach`, `device`, `device_type`, `dist_schedule`, `doacross`, `dynamic_allocators`, `enter`, `exclusive`, `fail`, `filter`, `final`, `firstprivate`, `from`, `full`, `grainsize`, `has_device_addr`, `hint`, `if`, `in_reduction`, `inbranch`, `inclusive`, `indirect`, `induction`, `inductor`, `init`, `initializer`, `interop`, `is_device_ptr`, `lastprivate`, `linear`, `link`, `local`, `map`, `match`, `memscope`, `mergeable`, `message`, `no_omp`, `no_omp_routines`, `no_parallelism`, `nocontext`, `nogroup`, `nontemporal`, `notinbranch`, `novariants`, `nowait`, `num_tasks`, `num_teams`, `num_threads`, `order`, `ordered`, `otherwise`, `partial`, `permutation`, `priority`, `proc_bind`, `read`, `reduction`, `relaxed`, `release`, `reverse_offload`, `safelen`, `safesync`, `schedule`, `seq_cst`, `severity`, `simd`, `simdlen`, `sizes`, `task_reduction`, `thread_limit`, `threads`, `threadset`, `to`, `unified_address`, `unified_shared_memory`, `uniform`, `untied`, `update`, `update`, `use`, `use_device_addr`, `use_device_ptr`, `uses_allocators`, `weak`, `when`, `write`

Semantics

The *directive-name-modifier* is a universal modifier that can be used on any OpenMP clause. The *directive-name* identifies the construct or constituent construct to which the clause applies. If *directive-name* is that of a combined or composite construct, then the leaf constructs to which the clause applies are determined as specified in [Section 18.2](#). 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.

Restrictions

Restrictions to the *directive-name-modifier* modifier are as follows:

- The *directive-name-modifier* must specify the *directive-name* of the construct or of a constituent construct of the *directive-specification* on which the clause appears.

Cross References

- `acq_rel` clause, see [Section 16.8.1.1](#)
- `acquire` clause, see [Section 16.8.1.2](#)
- `adjust_args` clause, see [Section 8.5.2](#)
- `affinity` clause, see [Section 13.6.1](#)
- `align` clause, see [Section 7.3](#)
- `aligned` clause, see [Section 6.11](#)
- `allocate` clause, see [Section 7.6](#)

- 1 • **allocator** clause, see [Section 7.4](#)
- 2 • **append_args** clause, see [Section 8.5.3](#)
- 3 • **apply** clause, see [Section 10.6](#)
- 4 • **at** clause, see [Section 9.2](#)
- 5 • **atomic_default_mem_order** clause, see [Section 9.5.1.1](#)
- 6 • **bind** clause, see [Section 12.8.1](#)
- 7 • **capture** clause, see [Section 16.8.3.1](#)
- 8 • **collapse** clause, see [Section 5.4.3](#)
- 9 • **collector** clause, see [Section 6.5.18](#)
- 10 • **combiner** clause, see [Section 6.5.14](#)
- 11 • **compare** clause, see [Section 16.8.3.2](#)
- 12 • **copyin** clause, see [Section 6.7.1](#)
- 13 • **copyprivate** clause, see [Section 6.7.2](#)
- 14 • **default** clause, see [Section 6.4.1](#)
- 15 • **defaultmap** clause, see [Section 6.8.6](#)
- 16 • **depend** clause, see [Section 16.9.5](#)
- 17 • **destroy** clause, see [Section 4.6](#)
- 18 • **detach** clause, see [Section 13.6.2](#)
- 19 • **device** clause, see [Section 14.2](#)
- 20 • **device_type** clause, see [Section 14.1](#)
- 21 • **dist_schedule** clause, see [Section 12.7.1](#)
- 22 • **doacross** clause, see [Section 16.9.6](#)
- 23 • **dynamic_allocators** clause, see [Section 9.5.1.2](#)
- 24 • **enter** clause, see [Section 6.8.4](#)
- 25 • **exclusive** clause, see [Section 6.6.2](#)
- 26 • **fail** clause, see [Section 16.8.3.3](#)
- 27 • **filter** clause, see [Section 11.6.1](#)
- 28 • **final** clause, see [Section 13.3](#)
- 29 • **firstprivate** clause, see [Section 6.4.4](#)

- 1 • **from** clause, see [Section 6.9.2](#)
- 2 • **full** clause, see [Section 10.2.1](#)
- 3 • **grainsize** clause, see [Section 13.7.1](#)
- 4 • **has_device_addr** clause, see [Section 6.4.9](#)
- 5 • **hint** clause, see [Section 16.1.2](#)
- 6 • **if** clause, see [Section 4.5](#)
- 7 • **in_reduction** clause, see [Section 6.5.11](#)
- 8 • **inbranch** clause, see [Section 8.7.1.1](#)
- 9 • **inclusive** clause, see [Section 6.6.1](#)
- 10 • **indirect** clause, see [Section 8.8.3](#)
- 11 • **induction** clause, see [Section 6.5.12](#)
- 12 • **inductor** clause, see [Section 6.5.17](#)
- 13 • **init** clause, see [Section 15.1.2](#)
- 14 • **initializer** clause, see [Section 6.5.15](#)
- 15 • **interop** clause, see [Section 8.6.1](#)
- 16 • **is_device_ptr** clause, see [Section 6.4.7](#)
- 17 • **lastprivate** clause, see [Section 6.4.5](#)
- 18 • **linear** clause, see [Section 6.4.6](#)
- 19 • **link** clause, see [Section 6.8.5](#)
- 20 • **local** clause, see [Section 6.13](#)
- 21 • **map** clause, see [Section 6.8.3](#)
- 22 • **match** clause, see [Section 8.5.1](#)
- 23 • **memscope** clause, see [Section 16.8.4](#)
- 24 • **mergeable** clause, see [Section 13.2](#)
- 25 • **message** clause, see [Section 9.3](#)
- 26 • **no_openmp** clause, see [Section 9.6.1.4](#)
- 27 • **no_openmp_routines** clause, see [Section 9.6.1.6](#)
- 28 • **no_parallelism** clause, see [Section 9.6.1.7](#)
- 29 • **nocontext** clause, see [Section 8.6.3](#)

- 1 • **nogroup** clause, see [Section 16.7](#)
- 2 • **nontemporal** clause, see [Section 11.5.1](#)
- 3 • **notinbranch** clause, see [Section 8.7.1.2](#)
- 4 • **novariants** clause, see [Section 8.6.2](#)
- 5 • **nowait** clause, see [Section 16.6](#)
- 6 • **num_tasks** clause, see [Section 13.7.2](#)
- 7 • **num_teams** clause, see [Section 11.3.1](#)
- 8 • **num_threads** clause, see [Section 11.2.2](#)
- 9 • **order** clause, see [Section 11.4](#)
- 10 • **ordered** clause, see [Section 5.4.4](#)
- 11 • **otherwise** clause, see [Section 8.4.2](#)
- 12 • **partial** clause, see [Section 10.2.2](#)
- 13 • **permutation** clause, see [Section 10.4.1](#)
- 14 • **priority** clause, see [Section 13.5](#)
- 15 • **proc_bind** clause, see [Section 11.2.4](#)
- 16 • **read** clause, see [Section 16.8.2.1](#)
- 17 • **reduction** clause, see [Section 6.5.9](#)
- 18 • **relaxed** clause, see [Section 16.8.1.3](#)
- 19 • **release** clause, see [Section 16.8.1.4](#)
- 20 • **reverse_offload** clause, see [Section 9.5.1.3](#)
- 21 • **safelen** clause, see [Section 11.5.2](#)
- 22 • **safesync** clause, see [Section 11.2.5](#)
- 23 • **schedule** clause, see [Section 12.6.3](#)
- 24 • **seq_cst** clause, see [Section 16.8.1.5](#)
- 25 • **severity** clause, see [Section 9.4](#)
- 26 • **simd** clause, see [Section 16.10.3.2](#)
- 27 • **simdlen** clause, see [Section 11.5.3](#)
- 28 • **sizes** clause, see [Section 10.1.1](#)
- 29 • **task_reduction** clause, see [Section 6.5.10](#)

- 1 • **thread_limit** clause, see [Section 14.3](#)
- 2 • **threads** clause, see [Section 16.10.3.1](#)
- 3 • **threadset** clause, see [Section 13.4](#)
- 4 • **to** clause, see [Section 6.9.1](#)
- 5 • **unified_address** clause, see [Section 9.5.1.4](#)
- 6 • **unified_shared_memory** clause, see [Section 9.5.1.5](#)
- 7 • **uniform** clause, see [Section 6.10](#)
- 8 • **untied** clause, see [Section 13.1](#)
- 9 • **update** clause, see [Section 16.8.2.2](#)
- 10 • **update** clause, see [Section 16.9.3](#)
- 11 • **use** clause, see [Section 15.1.3](#)
- 12 • **use_device_addr** clause, see [Section 6.4.10](#)
- 13 • **use_device_ptr** clause, see [Section 6.4.8](#)
- 14 • **uses_allocators** clause, see [Section 7.8](#)
- 15 • **weak** clause, see [Section 16.8.3.4](#)
- 16 • **when** clause, see [Section 8.4.1](#)
- 17 • **write** clause, see [Section 16.8.2.3](#)

4.5 if Clause

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

Arguments

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

Modifiers

Name	Modifies	Type	Properties
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: directive-name	unique

Directives

[cancel](#), [parallel](#), [simd](#), [target](#), [target data](#), [target enter data](#), [target exit data](#), [target update](#), [task](#), [taskloop](#), [teams](#)

Semantics

The effect of the **if** clause depends on the **construct** to which it is applied. If the **construct** is not a **combined construct** or a **composite construct** then the effect is described in the section that describes that **construct**.

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*.

Cross References

- **cancel** directive, see [Section 17.2](#)
- **parallel** directive, see [Section 11.2](#)
- **simd** directive, see [Section 11.5](#)
- **target** directive, see [Section 14.8](#)
- **target data** directive, see [Section 14.5](#)
- **target enter data** directive, see [Section 14.6](#)
- **target exit data** directive, see [Section 14.7](#)
- **target update** directive, see [Section 14.9](#)
- **task** directive, see [Section 13.6](#)
- **taskloop** directive, see [Section 13.7](#)
- **teams** directive, see [Section 11.3](#)

4.6 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>

Modifiers

Name	Modifies	Type	Properties
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: directive-name	unique

Directives

depobj, **interop**

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

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

9 **Restrictions**

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

15 **Cross References**

- 16 • **depobj** directive, see [Section 16.9.4](#)
- 17 • **interop** directive, see [Section 15.1](#)

5 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 the [base program](#) of an [OpenMP program](#):

- [OpenMP programs](#) must not declare names that begin with the `omp_` or `omp_x_` prefix, as these are reserved for the OpenMP implementation.

C++

- [OpenMP programs](#) must not declare a namespace with the `omp` or `omp_x` names, as these are reserved for the OpenMP implementation.

C++

5.1 OpenMP Types and Identifiers

An OpenMP identifier is a special identifier for use within [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 a [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 expression of OpenMP logical type is a scalar expression. This document uses *true* as a
5 generic term for a non-zero integer value and *false* as a generic term for an integer value of zero.

C / C++

Fortran

6 Any expression of OpenMP logical type is a scalar logical expression. This document uses *true* as a
7 generic term for a logical value of `.TRUE.` and *false* as a generic term for a logical value of
8 `.FALSE.`

Fortran

9 The OpenMP integer type supports integer variables and expressions in any [base language](#).

C / C++

10 Any OpenMP integer expression is an integer expression.

C / C++

Fortran

11 Any OpenMP integer expression is a scalar integer expression.

Fortran

12 The OpenMP string type supports character string variables and expressions in any [base language](#).

C / C++

13 Any OpenMP string expression is an expression of type qualified or unqualified `const char *`
14 or `char *` pointing to a null-terminated character string.

C / C++

Fortran

15 Any OpenMP string expression is a character string of default kind.

Fortran

16 OpenMP function identifiers support [procedure](#) names in any [base language](#). Regardless of the [base](#)
17 [language](#), any OpenMP function identifier is the name of a [procedure](#) as a [base language](#) identifier.

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

C / C++

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 15.1.1](#)
- OpenMP Reduction and Induction Identifiers, see [Section 6.5.1](#)
- `mapper` modifier, see [Section 6.8.2](#)

5.2 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 Collector Expressions, see [Section 6.5.2.4](#)
- OpenMP Combiner Expressions, see [Section 6.5.2.1](#)
- OpenMP Inductor Expressions, see [Section 6.5.2.3](#)
- OpenMP Initializer Expressions, see [Section 6.5.2.2](#)

5.3 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++

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

C / C++

The remainder of this section covers OpenMP context-specific **structured blocks** that conform to specific syntactic forms and restrictions that are required for certain block-associated **directives**.

Restrictions

Restrictions to **structured blocks** are as follows:

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

C / C++

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

C / C++

C++

- **throw**, **co_await**, **co_yield** and **co_return** must not violate the entry/exit criteria of **structured blocks**.

C++

Fortran

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

Fortran

5.3.1 OpenMP Allocator Structured Blocks

Fortran

An OpenMP *allocator structured-block* is a context-specific **structured block** that is associated with an **allocators directive**. It consists of *allocate-stmt*, where *allocate-stmt* is a Fortran **ALLOCATE** statement. For an **allocators directive**, the paired **end directive** is optional.

Fortran

Cross References

- `allocators` directive, see [Section 7.7](#)

5.3.2 OpenMP Function Dispatch Structured Blocks

An OpenMP *function dispatch structured block* is a context-specific [structured block](#) that is associated with a [dispatch directive](#). It 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 8.6](#)

5.3.3 OpenMP Atomic Structured Blocks

An OpenMP *atomic structured block* is a context-specific [structured block](#) that is associated with an [atomic](#) directive. The form of an atomic [structured block](#) depends on the atomic semantics that the [directive](#) enforces.

C / C++

Any instance of any *atomic structured block* in which any statement is enclosed in braces remains an instance of the same kind of *atomic structured block*.

C / C++

Fortran

Enclosing any instance of any *atomic structured block* in the pair of **BLOCK** and **END BLOCK** remains an instance of the same kind of *atomic structured block*, in which case the paired **end** directive is optional.

Fortran

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 either $<$ 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.**
- $equalop$ is **==**, **.EQ.**, or **.EQV.**
- The order operation, $ordop$, is one of **<**, **.LT.**, **>**, or **.GT.**
- **==** or **.EQ.** comparisons are performed by comparing the physical representation of operand values for equality after the usual conversions as described in the [base language](#), while ignoring padding bits, if any.
- **.EQV.** comparisons are performed as described in the [base language](#).
- 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 $equalop$ 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

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

▼ C / C++ ▼

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

```
x = expr;
```

▲ C / C++ ▲

▼ Fortran ▼

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

```
x = expr
```

▲ Fortran ▲

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

▼ C / C++ ▼

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

```
x++;  
x--;  
++x;  
--x;  
x binop= expr;  
x = x binop expr;  
x = expr binop x;
```

▲ C / C++ ▲

▼ Fortran ▼

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

```
x = x operator expr  
x = expr operator x  
x = intrinsic-procedure-name (x, expr-list)  
x = intrinsic-procedure-name (expr-list, x)
```

▲ Fortran ▲

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

C / C++

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

```
x = expr ordop x ? expr : x;  
x = x ordop expr ? expr : x;  
x = x == e ? d : x;
```

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

```
if (expr ordop x) x = expr;  
if (x ordop expr) x = expr;  
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) x = d  
if (x equalop e) then; x = d; end if  
if (x ordop expr) x = expr  
if (x ordop expr) then; x = expr; end if  
if (expr ordop x) x = expr  
if (expr ordop x) then; x = expr; end if
```

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:

```

1 { v = x; cond-update-stmt }
2 { cond-update-stmt v = x; }
3 if(x == e) x = d; else v = x;
4 { r = x == e; if(r) x = d; }
5 { r = x == e; if(r) x = d; else v = x; }

```

C / C++

Fortran

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

```

7 statement
8 capture-statement

```

9 or

```

10 capture-statement
11 statement

```

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

```

13 v = x

```

14 If *statement* is *write-statement* as specified above then it is a *write-capture-atomic structured block*.
15 If *statement* is *update-statement* as specified above then it is an *update-capture-atomic structured block* and may be used in **atomic constructs** that enforce **atomic captured update** semantics. If
16 *statement* is *conditional-update-statement* as specified above then it is a
17 *conditional-update-capture-atomic structured block*. In addition, for a
18 *conditional-update-capture-atomic structured block*, *statement* can have the following form:

```

20 x = expr

```

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

```

23 if (x equalop e) then
24     x = d
25 else
26     v = x
27 end if

```

28 or

```

29 r = x equalop e
30 if (r) x = d

```

31 or

```

1  r = x equalop e
2  if (r) then
3      x = d
4  else
5      v = x
6  endif

```

Fortran

Restrictions

Restrictions to OpenMP atomic **structured block** are as follows:

C / C++

- In forms where *e* is assigned it must be an lvalue.
- *r* must be of integral type.
- During the execution of an **atomic region**, multiple syntactic occurrences of *x* must designate the same **storage location**.
- During the execution of an **atomic region**, multiple syntactic occurrences of *r* must designate the same **storage location**.
- During the execution of an **atomic region**, multiple syntactic occurrences of *expr* must evaluate to the same value.
- None of *v*, *x*, *d* and *expr* (as applicable) may access the **storage location** designated by any other symbol in the list.
- In forms that capture the original value of *x* in *v*, *v* and *e* may not refer to, or access, the same **storage location**.
- *binop*, *binop=*, *ordop*, *==*, *++*, and *--* are not overloaded operators.
- The expression *x binop expr* must be numerically equivalent to *x binop (expr)*. This requirement is satisfied if the operators in *expr* have precedence greater than *binop*, or by using parentheses around *expr* or subexpressions of *expr*.
- The expression *expr binop x* must be numerically equivalent to *(expr) binop x*. This requirement is satisfied if the operators in *expr* have precedence equal to or greater than *binop*, or by using parentheses around *expr* or subexpressions of *expr*.
- The expression *x ordop expr* must be numerically equivalent to *x ordop (expr)*. This requirement is satisfied if the operators in *expr* have precedence greater than *ordop*, or by using parentheses around *expr* or subexpressions of *expr*.
- The expression *expr ordop x* must be numerically equivalent to *(expr) ordop x*. This requirement is satisfied if the operators in *expr* have precedence equal to or greater than *ordop*, or by using parentheses around *expr* or subexpressions of *expr*.

- 1 • The expression $x == e$ must be numerically equivalent to $x == (e)$. This requirement is
2 satisfied if the operators in e have precedence equal to or greater than $==$, or by using
3 parentheses around e or subexpressions of e .

← C / C++ →

← Fortran →

- 4 • x must not have the **ALLOCATABLE** attribute.
- 5 • During the execution of an **atomic region**, multiple syntactic occurrences of x must
6 designate the same **storage location**.
- 7 • During the execution of an **atomic region**, multiple syntactic occurrences of r must
8 designate the same **storage location**.
- 9 • During the execution of an **atomic region**, multiple syntactic occurrences of $expr$ must
10 evaluate to the same value.
- 11 • None of v , x , d , r , $expr$, and $expr$ -list (as applicable) may access the same **storage location** as
12 any other symbol in the list.
- 13 • In forms that capture the original value of x in v , v may not access the same **storage location**
14 as e .
- 15 • If *intrinsic-procedure-name* refers to **IAND**, **IOR**, or **IEOR**, exactly one expression must
16 appear in $expr$ -list.
- 17 • The expression $x operator expr$ must be, depending on its type, either mathematically or
18 logically equivalent to $x operator (expr)$. This requirement is satisfied if the operators in $expr$
19 have precedence greater than $operator$, or by using parentheses around $expr$ or
20 subexpressions of $expr$.
- 21 • The expression $expr operator x$ must be, depending on its type, either mathematically or
22 logically equivalent to $(expr) operator x$. This requirement is satisfied if the operators in $expr$
23 have precedence equal to or greater than $operator$, or by using parentheses around $expr$ or
24 subexpressions of $expr$.
- 25 • The expression $x equalop e$ must be, depending on its type, either mathematically or logically
26 equivalent to $x equalop (e)$. This requirement is satisfied if the operators in e have precedence
27 equal to or greater than $equalop$, or by using parentheses around e or subexpressions of e .
- 28 • *intrinsic-procedure-name* must refer to the intrinsic procedure name and not to other program
29 entities.
- 30 • $operator$ must refer to the intrinsic operator and not to a user-defined operator.
- 31 • All assignments must be intrinsic assignments.

← Fortran →

Cross References

- 32 • **atomic** directive, see [Section 16.8.5](#)
33

5.4 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.

5.4.1 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
```

or

```
{  
  loop-nest  
}
```

C / C++

or

C++

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

A range-based **for** loop is equivalent to a regular **for** loop using iterators, as defined in the [base language](#). A range-based **for** loop has no iteration [variable](#).

C++

or

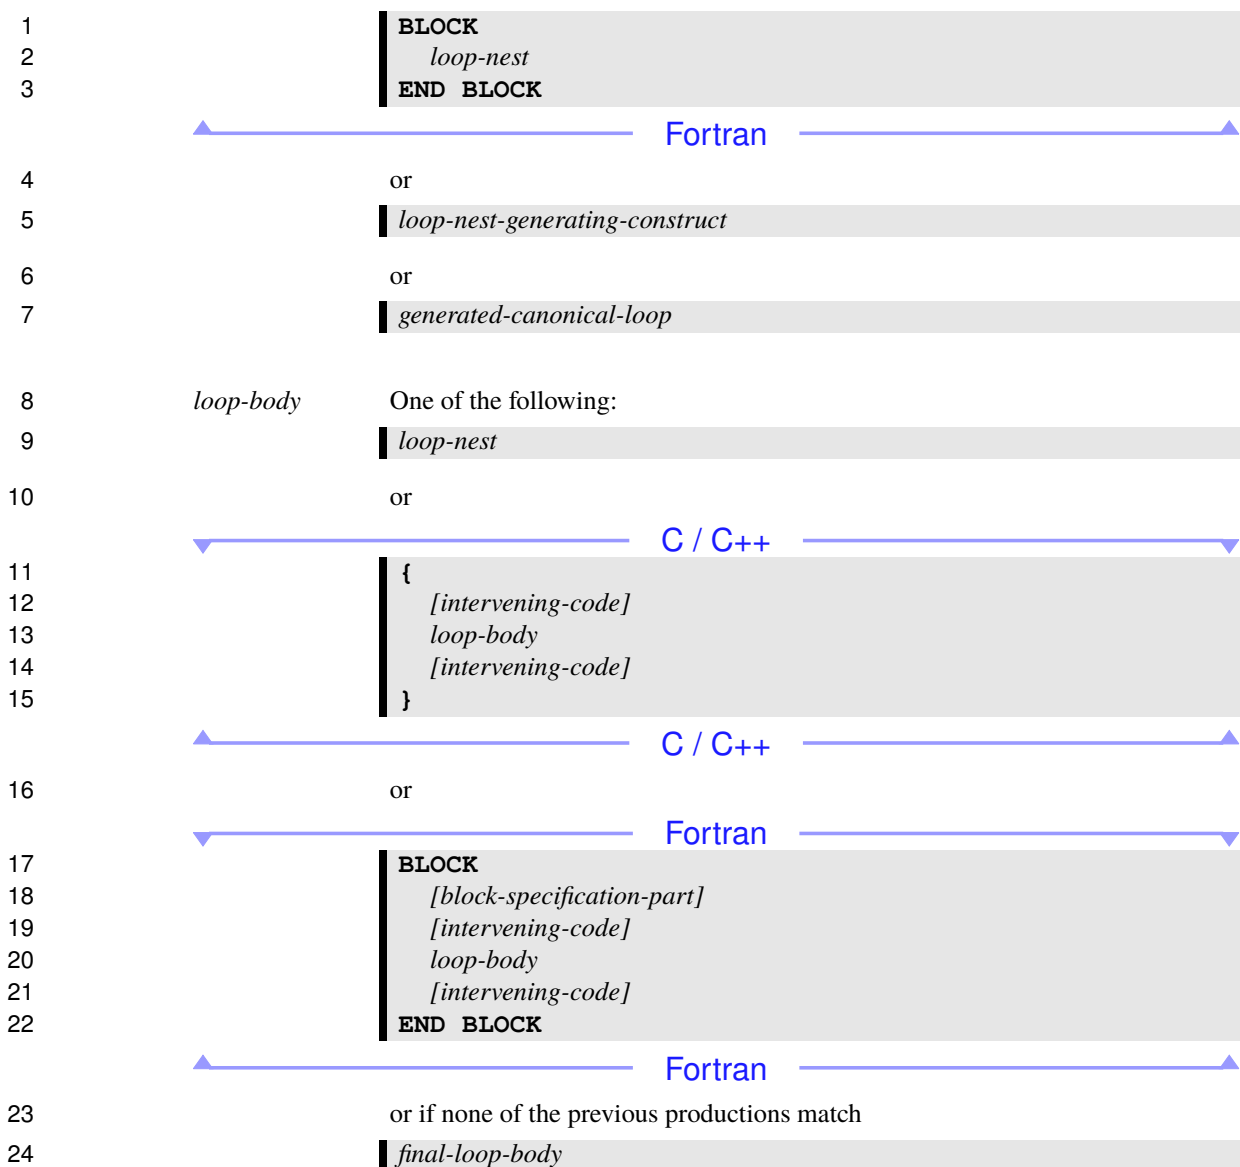
Fortran

```
DO [ label ] var = lb , ub [ , incr ]  
  [intervening-code]  
  loop-body  
  [intervening-code]  
[ label ] END DO
```

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

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

or



25 *loop-nest-generating-construct*
 26 A [loop-transforming construct](#) that generates a [canonical loop nest](#), which may
 27 be a [canonical loop sequence](#) that contains exactly one [canonical loop nest](#).

1
2
3

4
5

6
7
8

9
10

11
12
13
14
15
16
17
18
19
20

21
22
23
24

25
26
27

generated-canonical-loop

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

intervening-code

C / C++

A non-empty sequence of **structured blocks** or declarations, referred to as **intervening code**. It must not contain iteration statements, **continue** statements or **break** statements that apply to the enclosing loop.

C / C++

Fortran

A non-empty **structured block sequence**, referred to as **intervening code**. It must not contain:

- loops;
- **CYCLE** statements;
- **EXIT** statements;
- array expressions;
- array references with a vector subscript;
- assignment statements where the target is an array object;
- references to elemental procedures with an array actual argument; or
- references to procedures where the actual argument is an array that is not simply contiguous and the corresponding dummy argument has the **CONTIGUOUS** attribute or is an explicit-shape or assumed-size array.

Fortran

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

final-loop-body

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

C / C++

1 *init-expr* One of the following:
 2 *var = lb*
 3 *integer-type var = lb*
 4 *pointer-type var = lb*
 5 *random-access-iterator-type var = lb*

C
C
C++
C++

6 *test-expr* One of the following:
 7 *var relational-op ub*
 8 *ub relational-op var*

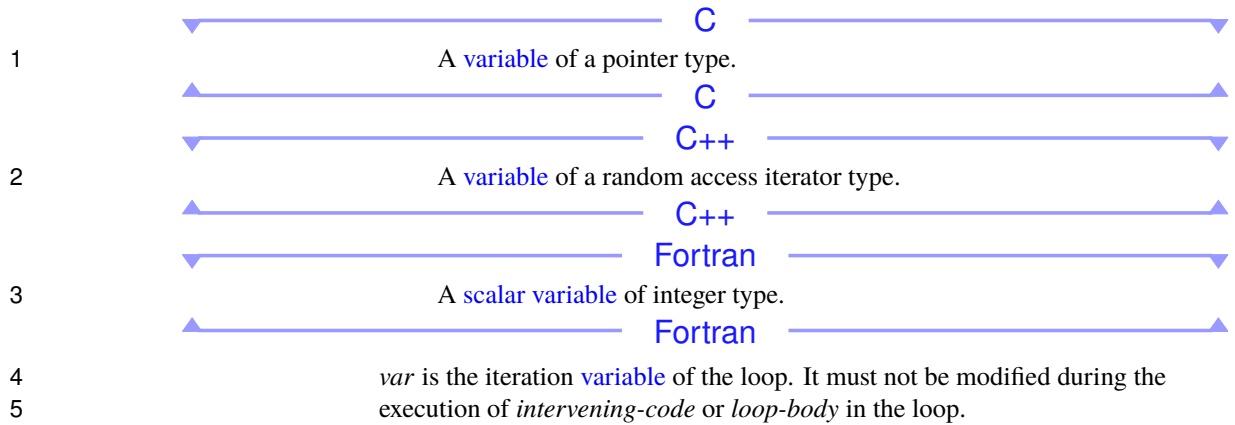
9 *relational-op* One of the following:
 10 <
 11 <=
 12 >
 13 >=
 14 !=

15 *incr-expr* One of the following:
 16 ++*var*
 17 *var*++
 18 -- *var*
 19 *var* --
 20 *var* += *incr*
 21 *var* -= *incr*
 22 *var* = *var* + *incr*
 23 *var* = *incr* + *var*
 24 *var* = *var* - *incr*
 25 The value of *incr*, respectively 1 and -1 for the increment and decrement
 26 operators, is the increment of the loop.

C / C++

27 *var* One of the following:
 28 A [variable](#) of a signed or unsigned integer type.

C / C++



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 [directive](#), the expressions are evaluated before the

13 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
- 24
- 25
- 26
- 27
- 28
- 29
- 30
- 31
- If *test-expr* is of the form *var relational-op b* and *relational-op* is $<$ or $<=$ then *incr-expr* must cause *var* to increase on each iteration of the loop. If *test-expr* is of the form *var relational-op b* and *relational-op* is $>$ or $>=$ then *incr-expr* must cause *var* to decrease on each iteration of the loop. Increase and decrease are using the order induced by *relational-op*.
 - If *test-expr* is of the form *ub relational-op var* and *relational-op* is $<$ or $<=$ then *incr-expr* must cause *var* to decrease on each iteration of the loop. If *test-expr* is of the form *ub relational-op var* and *relational-op* is $>$ or $>=$ then *incr-expr* must cause *var* to increase on each iteration of the loop. Increase and decrease are using the order induced by *relational-op*.
 - If *relational-op* is $!=$ then *incr-expr* must cause *var* to always increase by 1 or always 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 6.2](#)
- 14 • Canonical Loop Sequence Form, see [Section 5.4.6](#)
- 15 • Loop-Transforming Constructs, see [Chapter 10](#)

16 **5.4.2 OpenMP Loop-Iteration Spaces and Vectors**

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

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

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

29 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*, where all of the **associated loops** have a **loop iteration variable**. The OpenMP **loop-iteration vector 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 vector space**.

Alternatively, OpenMP defines a special keyword `omp_cur_iteration` that represents the current [logical iteration](#). It enables identification of relative points in the [logical iteration space](#) with:

`omp_cur_iteration` [\pm *logical_offset*]

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

The iterations of some number of outer [associated loops](#) can be collapsed into one larger [logical iteration space](#) that is the [collapsed 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](#). The number of times that any [intervening code](#) between any two [collapsed loops](#) 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 [collapsed logical iteration](#). If the [iteration count](#) of any loop is zero and that loop does not enclose the [intervening code](#), the behavior is unspecified.

5.4.3 collapse Clause

Name: <code>collapse</code>	Properties: once-for-all-constituents, unique
------------------------------------	--

Arguments

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

Modifiers

Name	Modifies	Type	Properties
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: directive-name	unique

Directives

[distribute](#), [do](#), [for](#), [loop](#), [simd](#), [taskloop](#)

Semantics

The [collapse clause](#) associates one or more loops of a [canonical loop nest](#) 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 [work distribution](#) 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 5.4.4](#)
- **distribute** directive, see [Section 12.7](#)
- **do** directive, see [Section 12.6.2](#)
- **for** directive, see [Section 12.6.1](#)
- **loop** directive, see [Section 12.8](#)
- **simd** directive, see [Section 11.5](#)
- **taskloop** directive, see [Section 13.7](#)

5.4.4 ordered Clause

Name: ordered	Properties: once-for-all-constituents, unique
-----------------------------	--

Arguments

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

Modifiers

Name	Modifies	Type	Properties
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: directive-name	unique

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** nest 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 a **perfectly nested loop**.

- 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**.

C++

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

C++

Cross References

- **collapse** clause, see [Section 5.4.3](#)
- **linear** clause, see [Section 6.4.6](#)
- **do** directive, see [Section 12.6.2](#)
- **for** directive, see [Section 12.6.1](#)
- **simd** directive, see [Section 11.5](#)
- **tile** directive, see [Section 10.1](#)

5.4.5 Consistent Loop Schedules

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

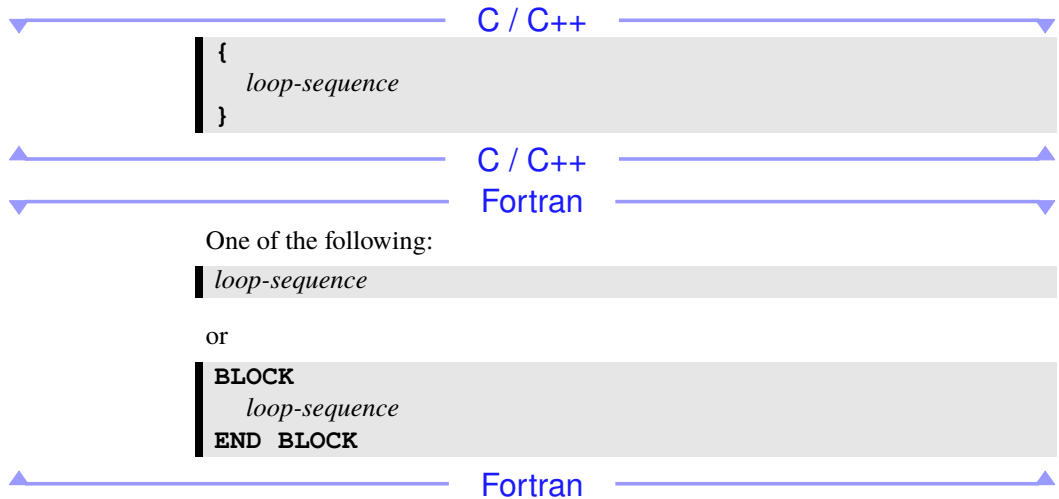
Two **loop-nest-associated constructs** have consistent schedules if all of the following conditions hold:

- The **constructs** have the same *directive-name*;
- The **regions** that correspond to the two **constructs** have the same **binding region**;
- The **constructs** have the same reproducible schedule;
- The **associated loop** nests have identical **logical iteration vector spaces**; and
- The **associated loop** nests are either both rectangular loops or both **non-rectangular loops**.

5.4.6 Canonical Loop Sequence Form

A structured-block has [canonical loop sequence form](#) if it conforms to *canonical-loop-sequence* in the following grammar:

canonical-loop-sequence



loop-sequence A [structured block sequence](#) with executable statements that match *canonical-loop-sequence*, *loop-sequence-generating-construct*, or *loop-nest* (a [canonical loop nest](#) as defined in [Section 5.4.1](#)). The loops must be [bounds-independent loops](#) with respect to *canonical-loop-sequence*.

loop-transforming-construct

A [loop-transforming construct](#) that generates a [canonical loop sequence](#) or [canonical loop nest](#).

The [loop sequence length](#) and consecutive order of [canonical loop nests](#) matched by *loop-nest* ignore how they are nested in *canonical-loop-sequence* or *loop-sequence*.

Cross References

- `looprange` clause, see [Section 5.4.7](#)
- Canonical Loop Nest Form, see [Section 5.4.1](#)
- Loop-Transforming Constructs, see [Chapter 10](#)

5.4.7 looprange Clause

Name: looprange	Properties: unique
------------------------	---------------------------

Arguments

Name	Type	Properties
<i>first</i>	expression of OpenMP integer type	constant, positive
<i>count</i>	expression of OpenMP integer type	constant, positive, ultimate

Directives

fuse

Semantics

For a [loop-sequence-associated construct](#), the **looprange clause** determines the [canonical loop nests](#) of the [associated loop sequence](#) that are affected by the [directive](#). The [affected loop nests](#) are the *count* consecutive [canonical loop nests](#) that begin with the [canonical loop nest](#) specified by the *first* argument.

For all [directives](#) on which the **looprange clause** may appear, if the [clause](#) is not specified then the effect is as if the [clause](#) was specified with a value equal to the [loop sequence lengths](#) of the [canonical loop sequence](#).

Restrictions

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

- $first + count - 1$ must not evaluate to a value greater than the [loop sequence length](#) of the associated [canonical loop sequence](#).

Cross References

- **fuse** directive, see [Section 10.5](#)
- Canonical Loop Sequence Form, see [Section 5.4.6](#)

1

Part II

2

Directives and Clauses

6 Data Environment

This chapter presents [directives](#) and [clauses](#) for controlling [data environments](#). These [directives](#) and [clauses](#) include the [data-environment attribute clauses](#), which explicitly determine the [data-environment attributes](#) of [list items](#) specified in a [list](#) argument. 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 clause](#). Additional restrictions apply to the use of these [clause sets](#) on [directives](#) that accept any members of them.

[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](#). Any [clause](#) that indicates a [variable](#) is private in that scope is a [privatization clause](#).

[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](#).

6.1 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 6.1.1](#) describes the [data-sharing attribute](#) rules for [variables](#) referenced in a [construct](#).
- [Section 6.1.2](#) describes the [data-sharing attribute](#) rules for [variables](#) referenced in a [region](#), but outside any [construct](#).

6.1.1 Variables Referenced in a Construct

A [variable](#) that is referenced in a [construct](#) can have a [predetermined data-sharing attribute](#), an [explicitly determined data-sharing attribute](#), or an [implicitly determined data-sharing attribute](#), 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

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

Fortran

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

Fortran

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

Fortran

- 8 • **variables** and common blocks (in Fortran) that appear as arguments in **threadprivate**
9 **directives** or **variables** with the **_Thread_local** (in C) or **thread_local** (in C/C++)
10 storage-class specifier are threadprivate.
- 11 • **Variables** and common blocks (in Fortran) that appear as arguments in **groupprivate**
12 **directives** are **groupprivate variables**.
- 13 • Variables and common blocks (in Fortran) that appear as **list items** in **local clauses** on
14 **declare target** directives are **device local variables**.

C

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

C

C++

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

C++

C / C++

- 19 • Objects with dynamic storage duration are shared.

C / C++

- 20 • The **loop iteration variable** in the **associated loop** of a **simd** construct with just one
21 **associated loop** is linear with a *linear-step* that is the increment of the **associated loop**.
- 22 • The **loop iteration variable** in the **associated loops** of a **simd** construct with multiple
23 **associated loops** are lastprivate.
- 24 • The **loop iteration variable** in any **associated loop** of a **loop** construct is lastprivate.

- 1 • The **loop iteration variable** in any **associated loop** of a **loop-nest-associated directive** is
2 otherwise private.

▼ C++ ▼

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

▲ C++ ▲

▼ Fortran ▼

- 4 • **Loop iteration variables** inside **parallel**, **teams**, or **task-generating constructs** are private
5 in the innermost such **construct** that encloses the loop.

- 6 • Implied-do, **FORALL** and **DO CONCURRENT** indices are private.

▲ Fortran ▲

▼ C / C++ ▼

- 7 • **Variables** with **static storage duration** that are declared in a scope inside the **construct** are
8 shared.

- 9 • If a **list item** in a **has_device_addr** clause or in a **map** clause on the **target** construct
10 has a **base pointer**, and the **base pointer** is a **scalar variable** that does not appear in a **map**
11 clause on the **construct**, the **base pointer** is firstprivate.

- 12 • If a **list item** in a **reduction** or **in_reduction** clause on the **construct** has a **base**
13 **pointer** then the **base pointer** is private.

- 14 • Static data members are shared.

- 15 • The **__func__** variable and similar function-local predefined **variables** are shared.

▲ C / C++ ▲

▼ Fortran ▼

- 16 • **Assumed-size arrays** and *named constants* are shared in **constructs** that are not **data-mapping**
17 **constructs**.

- 18 • *Named constants* are firstprivate in **target** constructs.

- 19 • An associate name that may appear in a **variable** definition context is shared if its association
20 occurs outside of the **construct** and otherwise it has the same **data-sharing attribute** as the
21 selector with which it is associated.

▲ Fortran ▲

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

- 26 • The **loop iteration variable** in any **associated loop** of a **loop-nest-associated directive** may be
27 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 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 6.4.

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 data-sharing attributes 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 6.4.1).
- 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 6.8) 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.

An **OpenMP 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 6.4.4](#).

6.1.2 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

6.2 threadprivate Directive

Name: <code>threadprivate</code>	Association: none
Category: <code>declarative</code>	Properties: <code>pure</code>

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.

C++

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

C++

1 The content of a **threadprivate variable** can change across a **task scheduling point** if the executing
2 **thread** switches to another **task** that modifies the **variable**. For more details on **task** scheduling, see
3 **Section 1.3** and **Chapter 13**.

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

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

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

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

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

C / C++

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

C / C++

Fortran

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

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

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

Fortran

C++

21 The order in which any constructors for different **threadprivate variables** of **class type** are called is
22 unspecified. The order in which any destructors for different **threadprivate variables** of **class type**
23 are called is unspecified. A **variable** that is part of an **aggregate variable** may appear in a
24 **threadprivate directive** only if it is a static data member of a C++ class.

C++

Restrictions

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

- 27 • A **thread** must not reference the copy of another **thread** of a **threadprivate variable**.
- 28 • A **threadprivate variable** must not appear as the **base variable** of a **list item** in any **clause**
29 except for the **copyin** and **copyprivate** clauses.
- 30 • An **OpenMP program** in which an **untied task** accesses threadprivate storage is
31 non-conforming.

C / C++

- 1 • Each **list item** must be a file-scope, namespace-scope, or static block-scope **variable**.
- 2 • No **list item** may have an incomplete type.
- 3 • The address of a **threadprivate variable** must not be an address constant.
- 4 • If the value of a **variable** referenced in an explicit initializer of a **threadprivate variable** is
5 modified prior to the first reference to any instance of the **threadprivate variable**, the behavior
6 is unspecified.
- 7 • A **threadprivate directive** for file-scope **variables** must appear outside any definition or
8 declaration, and must lexically precede all references to any of the **variables** in its *list*.
- 9 • A **threadprivate directive** for namespace-scope **variables** must appear outside any
10 definition or declaration other than the namespace definition itself and must lexically precede
11 all references to any of the **variables** in its *list*.
- 12 • Each **variable** in the list of a **threadprivate directive** at file, namespace, or class scope
13 must refer to a **variable** declaration at file, namespace, or class scope that lexically precedes
14 the **directive**.
- 15 • A **threadprivate directive** for a static block-scope **variable** must appear in the scope of
16 the **variable** and not in a nested scope. The **directive** must lexically precede all references to
17 any of the **variables** in its *list*.
- 18 • Each **variable** in the *list* of a **threadprivate directive** in block scope must refer to a
19 **variable** declaration in the same scope that lexically precedes the **directive**. The **variable** must
20 have **static storage duration**.
- 21 • If a **variable** is specified in a **threadprivate directive** in one **compilation unit**, it must be
22 specified in a **threadprivate directive** in every **compilation unit** in which it is declared.

C / C++

C++

- 23 • A **threadprivate directive** for static class member **variables** must appear in the class
24 definition, in the same scope in which the member **variables** are declared, and must lexically
25 precede all references to any of the **variables** in its *list*.
- 26 • A **threadprivate variable** must not have an incomplete type or a reference type.
- 27 • A **threadprivate variable** with class type must have:
 - 28 – An accessible, unambiguous default constructor in the case of default initialization
29 without a given initializer;
 - 30 – An accessible, unambiguous constructor that accepts the given argument in the case of
31 direct initialization; and

- 1 – An accessible, unambiguous copy constructor in the case of copy initialization with an
2 explicit initializer.

C++
Fortran

- 3 • Each **list item** must be a named variable or a named common block; a named common block
4 must appear between slashes.
- 5 • The **list** argument must not include any coarrays or associate names.
- 6 • The **threadprivate** directive must appear in the declaration section of a scoping unit in
7 which the common block or **variable** is declared.
- 8 • If a **threadprivate** directive that specifies a common block name appears in one
9 **compilation unit**, then such a directive must also appear in every other **compilation unit** that
10 contains a **COMMON** statement that specifies the same name. It must appear after the last such
11 **COMMON** statement in the **compilation unit**.
- 12 • If a **threadprivate** variable or a threadprivate common block is declared with the **BIND**
13 attribute, the corresponding C entities must also be specified in a **threadprivate**
14 **directive** in the C program.
- 15 • A **variable** may only appear as an argument in a **threadprivate** directive in the scope in
16 which it is declared. It must not be an element of a common block or appear in an
17 **EQUIVALENCE** statement.
- 18 • A **variable** that appears as an argument in a **threadprivate** directive must be declared in
19 the scope of a module or have the **SAVE** attribute, either explicitly or implicitly.
- 20 • The effect of an access to a **threadprivate** variable in a **DO CONCURRENT** construct is
21 unspecified.

Fortran

22 **Cross References**

- 23 • **copyin** clause, see [Section 6.7.1](#)
- 24 • **order** clause, see [Section 11.4](#)
- 25 • *dyn-var* ICV, see [Table 2.1](#)
- 26 • Determining the Number of Threads for a **parallel** Region, see [Section 11.2.1](#)

6.3 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](#) is a [loop-collapsing construct](#), 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 items](#). Inside the [construct](#), all references to the [original list items](#) are replaced by references to the [new list items](#) received by the [task](#) or [SIMD lane](#).

If the [construct](#) is a [loop-collapsing construct](#) then, within the same [collapsed logical iteration](#) of the [collapsed loops](#), the same [new list item](#) replaces all references to the [original list item](#). For any two [collapsed iterations](#), if the references to the [original list item](#) are replaced by the same [new list item](#) then the [collapsed 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 list 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 [new 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**, **SELECT RANK** or **SELECT TYPE**
28 **construct**. If the construct association is established prior to a **parallel region**, the association
29 between the 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 for
2 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:

▼ C++ ▼

- 6 • A **variable** of **class type** (or array thereof) that is privatized requires an accessible,
7 unambiguous default constructor for the **class type**.
- 8 • A **variable** that is privatized must not have the **constexpr** specifier unless it is of **class type**
9 with a **mutable** member. This restriction does not apply to the **firstprivate** clause.

▲ C++ ▲

▼ 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.
- 12 • A **variable** that is privatized must not have an incomplete type or be a reference to an
13 incomplete type.

▲ C / C++ ▲

▼ Fortran ▼

- 14 • **Variable** 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
19 where the corresponding dummy argument is a coarray.
- 20 • **Assumed-size arrays** must not be privatized.
- 21 • An optional dummy argument that is not present must not appear as a **list item** in a
22 **privatization clause** or be privatized as a result of an **implicitly determined data-sharing**
23 **attribute** or **predetermined data-sharing attribute**.

▲ Fortran ▲

6.4 Data-Sharing Attribute Clauses

Several [constructs](#) accept [clauses](#) that allow a user to control the [data-sharing attributes](#) of [variables](#) referenced in the [construct](#). Not all of the [clauses](#) listed in this section are valid on all [directives](#). The set of [clauses](#) that is valid on a particular [directive](#) is described with the [directive](#). The [reduction clauses](#) are explained in [Section 6.5](#).

A [list item](#) may be specified in both [firstprivate](#) and [lastprivate](#) clauses.

C++

If a [variable](#) referenced in a [data-sharing attribute clause](#) has a type derived from a template and the [OpenMP 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

6.4.1 default Clause

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

Arguments

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

Modifiers

Name	Modifies	Type	Properties
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: directive-name	<code>unique</code>

Directives

[parallel](#), [task](#), [taskloop](#), [teams](#)

Semantics

The [default clause](#) determines the [implicitly determined data-sharing attributes](#) of certain [variables](#) that are referenced in the [construct](#), in accordance with the rules given in [Section 6.1.1](#).

If *data-sharing-attribute* is not **none**, the [data-sharing attributes](#) 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 an **explicitly determined data-sharing attribute**.

C / C++

- If *data-sharing-attribute* is **firstprivate** or **private**, each **variable** with **static storage duration** that is declared in a namespace or global scope, is referenced in the **construct**, and does not have a **predetermined data-sharing attribute** must have an **explicitly determined data-sharing attribute**.

C / C++

Cross References

- **parallel** directive, see [Section 11.2](#)
- **task** directive, see [Section 13.6](#)
- **taskloop** directive, see [Section 13.7](#)
- **teams** directive, see [Section 11.3](#)

6.4.2 shared Clause

Name: shared	Properties: data-environment attribute, data-sharing attribute
----------------------------	---

Arguments

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

Modifiers

Name	Modifies	Type	Properties
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: directive-name	unique

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 list item** 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 11.2](#)
- `task` directive, see [Section 13.6](#)
- `taskloop` directive, see [Section 13.7](#)
- `teams` directive, see [Section 11.3](#)

6.4.3 private Clause

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

Arguments

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

Modifiers

Name	Modifies	Type	Properties
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: directive-name	unique

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 6.3](#). 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 6.3.

Cross References

- **distribute** directive, see Section 12.7
- **do** directive, see Section 12.6.2
- **for** directive, see Section 12.6.1
- **loop** directive, see Section 12.8
- **parallel** directive, see Section 11.2
- **scope** directive, see Section 12.2
- **sections** directive, see Section 12.3
- **simd** directive, see Section 11.5
- **single** directive, see Section 12.1
- **target** directive, see Section 14.8
- **task** directive, see Section 13.6
- **taskloop** directive, see Section 13.7
- **teams** directive, see Section 11.3
- List Item Privatization, see Section 6.3

6.4.4 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>

Modifiers

Name	Modifies	Type	Properties
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: directive-name	unique

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 6.4.3, except as noted. In addition, the **new list item** is initialized from the **original list item**. 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 ▼

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

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

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

Fortran

Restrictions

4 Restrictions to the **firstprivate** clause are as follows:

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

C++

- 26 • A **variable** of **class type** (or array thereof) that appears in a **firstprivate** clause requires
27 an accessible, unambiguous copy constructor for the **class type**.
- 28 • If the **original list item** in a **firstprivate** clause on a **work-distribution** construct has a
29 reference type then it must bind to the same object for all **threads** in the **binding thread set** of
30 the **work-distribution region**.

C++

Fortran

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

Fortran

Cross References

- **private** clause, see [Section 6.4.3](#)
- **distribute** directive, see [Section 12.7](#)
- **do** directive, see [Section 12.6.2](#)
- **for** directive, see [Section 12.6.1](#)
- **parallel** directive, see [Section 11.2](#)
- **scope** directive, see [Section 12.2](#)
- **sections** directive, see [Section 12.3](#)
- **single** directive, see [Section 12.1](#)
- **target** directive, see [Section 14.8](#)
- **task** directive, see [Section 13.6](#)
- **taskloop** directive, see [Section 13.7](#)
- **teams** directive, see [Section 11.3](#)

6.4.5 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
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: directive-name	unique

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 6.4.3. In addition, when a **lastprivate** clause without the **conditional** modifier appears on a **directive** and the **list item** is not a **loop 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 a **loop iteration variable** of one of the **associated loops**, if sequential execution of the associated **structured block** 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 **structured block**.

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 a **loop 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 **loop iteration variable** of an **associated loop** and that would not be assigned a value during sequential execution of the **canonical 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.

Otherwise, if the **lastprivate** clause is used on a **construct** to which the **nowait** or the **nogroup** clauses are applied, accesses to the **original list item** may create a data race. To avoid

1 this data race, if an assignment to the **original list item** occurs then synchronization must be inserted
2 to ensure that the assignment completes and the **original list item** is flushed to **memory**.

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

6 Restrictions

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

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

▼ C++ ▼

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

▲ C++ ▲

▼ Fortran ▼

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

▲ Fortran ▲

Cross References

- **private** clause, see [Section 6.4.3](#)
- **distribute** directive, see [Section 12.7](#)
- **do** directive, see [Section 12.6.2](#)
- **for** directive, see [Section 12.6.1](#)
- **loop** directive, see [Section 12.8](#)
- **sections** directive, see [Section 12.3](#)
- **simd** directive, see [Section 11.5](#)
- **taskloop** directive, see [Section 13.7](#)

6.4.6 linear Clause

Name: <code>linear</code>	Properties: data-environment attribute, data-sharing attribute, privatization, innermost-leaf, 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
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: directive-name	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 6.4.3, 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.

When a **linear** clause is specified on a **loop-collapsing construct**, the value of the **new list item** on each **collapsed iteration** corresponds to the value of the **original list item** before entering the **construct** plus the logical number of the iteration times *linear-step*. The value that corresponds to the sequentially last **collapsed iteration** of the **collapsed loops** is assigned to the **original list item**.

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

Restrictions

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

- Only a **loop iteration variable** of an **associated loop** may appear as a **list item** in a **linear** clause if a **reduction** clause with the **inscan** modifier also appears on the **construct**.
- A *linear-modifier* may be specified as **ref** or **uval** only on a **declare simd** directive.
- For a **linear** clause that appears on a **loop-nest-associated directive**, the difference between the value of a **list item** at the end of a **collapsed iteration** and its value at the beginning of the **collapsed iteration** must be equal to *linear-step*.
- If *linear-modifier* is **uval** for a **list item** in a **linear** clause that is specified on a **declare simd** directive and the **list item** is modified during a call to the **SIMD** version of the **procedure**, the **OpenMP program** must not depend on the value of the **list item** upon return from the **procedure**.
- If *linear-modifier* is **uval** for a **list item** in a **linear** clause that is specified on a **declare simd** directive, the **OpenMP program** must not depend on the storage of the argument in the **procedure** being the same as the storage of the corresponding argument at the callsite.

C

- All **list items** must be of integral or pointer type.

1

- If specified, *linear-modifier* must be **val**.

C

C++

2

- If *linear-modifier* is not **ref**, all **list items** must be of integral or pointer type, or must be a reference to an integral or pointer type.

3

4

- If *linear-modifier* is **ref** or **uval**, all **list items** must be of a reference type.

5

- 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**.

6

7

- 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**.

8

9

10

C++

Fortran

11

- If *linear-modifier* is not **ref**, all **list items** must be of type **integer**.

12

- If *linear-modifier* is **ref** or **uval**, all **list items** must be dummy arguments without the **VALUE** attribute.

13

14

- **List items** must not be **variables** that have the **POINTER** attribute.

15

- If *linear-modifier* is not **ref** and a **list item** has the **ALLOCATABLE** attribute, the allocation status of the **list item** in the last **collapsed iteration** must be allocated upon exit from that **collapsed iteration**.

16

17

18

- If *linear-modifier* is **ref**, **list items** must be polymorphic **variables**, assumed-shape arrays, or **variables** with the **ALLOCATABLE** attribute.

19

20

- 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**.

21

22

23

24

- A common block name must not appear in a **linear** clause.

Fortran

Cross References

- **private** clause, see [Section 6.4.3](#)
- **declare simd** directive, see [Section 8.7](#)
- **do** directive, see [Section 12.6.2](#)
- **for** directive, see [Section 12.6.1](#)
- **simd** directive, see [Section 11.5](#)
- **taskloop** directive, see [Section 13.7](#)

6.4.7 is_device_ptr Clause

Name: is_device_ptr	Properties: data-environment attribute, data-sharing attribute, innermost-leaf
----------------------------	---

Arguments

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

Modifiers

Name	Modifies	Type	Properties
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: directive-name	unique

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](#) is 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 6.4.9](#)
- `dispatch` directive, see [Section 8.6](#)
- `target` directive, see [Section 14.8](#)

6.4.8 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>

Modifiers

Name	Modifies	Type	Properties
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: directive-name	unique

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 6.8.3](#)), 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](#) or is accessible on the [target device](#).

Cross References

- `target data` directive, see [Section 14.5](#)

6.4.9 `has_device_addr` Clause

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

Arguments

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

Modifiers

Name	Modifies	Type	Properties
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: directive-name	unique

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 14.8](#)

6.4.10 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>

Modifiers

Name	Modifies	Type	Properties
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: directive-name	unique

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 no `corresponding list item` is 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 14.5](#)

6.5 Reduction and Induction Clauses and Directives

The [reduction clauses](#) and [induction clause](#) 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. The [induction clause](#) can be used to express [induction operations](#) in a loop.

6.5.1 OpenMP Reduction and Induction Identifiers

The syntax of OpenMP reduction and induction identifiers is defined as follows:

C

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

An induction identifier is either an *identifier* or one of the following operators: `+` and `*`.

C

C++

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

An induction identifier is either an *id-expression* or one of the following operators: `+` and `*`.

C++

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`.

An induction identifier is either a [base language](#) identifier, or a user-defined operator, or one of the following operators: `+` and `*`.

Fortran

6.5.2 OpenMP Reduction and Induction Expressions

A [reduction expression](#) is an [OpenMP stylized expression](#) that is relevant to [reduction clauses](#). An [induction expression](#) is an [OpenMP stylized expression](#) that is relevant to the [induction clause](#).

Restrictions

Restrictions to [reduction expressions](#) and [induction expressions](#) are as follows:

- If execution of a [reduction expression](#) or [induction expression](#) results in the execution of a [construct](#) or an OpenMP API call, the behavior is unspecified.

C / C++

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

C / C++

Fortran

- Any generic identifier, defined operation, defined assignment, or specific procedure used in a **reduction expression** or **induction expression** must be resolvable to a **procedure** with an explicit interface that has only scalar dummy arguments.
- Any **procedure** used in a **reduction expression** or **induction expression** must not have any alternate returns appear in the argument list.
- Any **procedure** called in the **region** of a **reduction expression** or **induction expression** must be pure and may not reference any host-associated or use-associated **variables** nor any **variables** in a common block.
- A **declare target directive** must be specified for any **procedure** that can be accessed through any **reduction expression** or **induction expression** that corresponds to an identifier that is used in a **target region**.

Fortran

6.5.2.1 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

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

Fortran

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

1 If a generic name is used in a [combiner expression](#) and the [list item](#) in the corresponding [reduction](#)
2 [clause](#) is an array or [array section](#), it is resolved to the specific procedure that is elemental or only
3 has scalar dummy arguments.

Fortran

4 **Restrictions**

5 Restrictions to [combiner expressions](#) are as follows:

- 6 • The only [variables](#) allowed in a [combiner expression](#) are `omp_in` and `omp_out`.

Fortran

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

Fortran

8 **6.5.2.2 OpenMP Initializer Expressions**

9 If the initialization of the private copies of reduction [list items](#) is not determined *a priori*, the syntax
10 of an [initializer expression](#) is as follows:

C

```
11 | omp_priv = initializer
```

C

12 or

C++

```
13 | omp_priv initializer
```

C++

14 or

C / C++

```
15 | function-name (argument-list)
```

C / C++

16 or

Fortran

```
17 | omp_priv = expression
```

18 or

```
19 | subroutine-name (argument-list)
```

Fortran

20 In the definition of an [initializer expression](#), the `omp_priv` special [variable](#) identifier refers to the
21 storage to be initialized. The special [variable](#) identifier `omp_orig` can be used in an [initializer](#)
22 [expression](#) to refer to the storage of the [original list item](#) to be reduced. The number of times that an
23 [initializer expression](#) is evaluated and the order of these evaluations are unspecified.

C / C++

1 If an **initializer expression** is a function name with an argument list, it is evaluated by calling the
2 function with the specified argument list. Otherwise, an **initializer expression** specifies how
3 **omp_priv** is declared and initialized.

C / C++

Fortran

4 If an **initializer expression** is a subroutine name with an argument list, it is evaluated by calling the
5 subroutine with the specified argument list. If an **initializer expression** is an assignment statement,
6 the **initializer expression** is evaluated by executing the assignment statement.

Fortran

C

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

C

C++

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

C++

Fortran

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

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

Fortran

Restrictions

17 Restrictions to **initializer expressions** are as follows:

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

C

- 20 • If an **initializer expression** is a function name with an argument list, one of the arguments
21 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

6.5.2.3 OpenMP Inductor Expressions

An **inductor expression** specifies how an **induction operation** determines a new value of the **induction variable** from its previous value and a **step expression**.

Fortran

An **inductor expression** is an assignment statement or a subroutine name followed by an argument list.

Fortran

In the definition of an **inductor** expression, **omp_var** is a special **variable** identifier that refers to storage of the type of the **induction variable** to which the **induction operation** applies, and **omp_step** is a special **variable** identifier that refers to the **step expression** of the **induction operation**. If the **list item** is an array or **array section**, the identifier to which **omp_var** corresponds refers to an array element.

Fortran

If the **inductor expression** is a subroutine name with an argument list, the **inductor expression** is evaluated by calling the subroutine with the specified argument list. If the **inductor expression** is an assignment statement, the **inductor expression** is evaluated by executing the assignment statement.

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

Fortran

Restrictions

Restrictions to **inductor expressions** are as follows:

- The only **variables** allowed in an **inductor expression** are **omp_var** and **omp_step**.

Fortran

- Any selectors in the designator of **omp_var** and **omp_step** must be *component selectors*.

Fortran

6.5.2.4 OpenMP Collector Expressions

A [collector expression](#) evaluates to the value of the [collective step expression](#) of a [collapsed iteration](#). In the definition of a [collector expression](#), `omp_step` is a special [variable](#) identifier that refers to the [step expression](#), and `omp_idx` is a special [variable](#) identifier that refers to the [collapsed iteration](#).

Restrictions

Restrictions to [collector expressions](#) are as follows:

- The only [variables](#) allowed in a [collector expression](#) are `omp_step` and `omp_idx`.

6.5.3 Implicitly Declared OpenMP Reduction Identifiers

C / C++

Table 6.1 lists each reduction identifier that is implicitly declared at every scope and its 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 6.5.6](#).

TABLE 6.1: Implicitly Declared C/C++ Reduction Identifiers

Identifier	Initializer	Combiner
<code>+</code>	<code>omp_priv = 0</code>	<code>omp_out += omp_in</code>
<code>*</code>	<code>omp_priv = 1</code>	<code>omp_out *= omp_in</code>
<code>&</code>	<code>omp_priv = ~ 0</code>	<code>omp_out &= omp_in</code>
<code> </code>	<code>omp_priv = 0</code>	<code>omp_out = omp_in</code>
<code>^</code>	<code>omp_priv = 0</code>	<code>omp_out ^= omp_in</code>
<code>&&</code>	<code>omp_priv = 1</code>	<code>omp_out = omp_in && omp_out</code>
<code> </code>	<code>omp_priv = 0</code>	<code>omp_out = omp_in omp_out</code>
<code>max</code>	<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>
<code>min</code>	<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>

C / C++

Fortran

Table 6.2 lists each reduction identifier that is implicitly declared for numeric and logical types and its semantic initializer value. The actual initializer value is that value as expressed in the data type of the reduction [list item](#).

TABLE 6.2: Implicitly Declared Fortran Reduction Identifiers

Identifier	Initializer	Combiner
<code>+</code>	<code>omp_priv = 0</code>	<code>omp_out = omp_in + omp_out</code>
<code>*</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>
<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

6.5.4 Implicitly Declared OpenMP Induction Identifiers

C / C++

Table 6.3 lists each induction identifier that is implicitly declared at every scope for arithmetic types and its corresponding [inductor expression](#) and [collector expression](#).

TABLE 6.3: Implicitly Declared C/C++ Induction Identifiers

Identifier	Inductor Expression	Collector Expression
+	<code>omp_var = omp_var + omp_step</code>	<code>omp_step * omp_idx</code>
*	<code>omp_var = omp_var * omp_step</code>	<code>pow(omp_step, omp_idx)</code>



1 Table 6.4 lists each induction identifier that is implicitly declared for numeric types and its
2 corresponding [inductor expression](#) and [collector expression](#).


TABLE 6.4: Implicitly Declared Fortran Induction Identifiers

Identifier	Inductor Expression	Collector Expression
+	<code>omp_var = omp_var + omp_step</code>	<code>omp_step * omp_idx</code>
*	<code>omp_var = omp_var * omp_step</code>	<code>omp_step ** omp_idx</code>



3 6.5.5 Properties Common to Reduction and induction 4 Clauses

5 The [list items](#) that appear in a [reduction clause](#) or [induction clause](#) may include [array sections](#)
6 and [array elements](#).



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

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

13 If the reduction or induction identifier is qualified then a qualified name lookup is used to find the
14 declaration.

1 If the reduction or induction identifier is unqualified then an *argument-dependent name lookup*
2 must be performed using the type of each **list item**.

▲ C++ ▼

3 If a **list item** is an array or **array section**, it will be treated as if a **reduction clause** or **induction**
4 **clause** would be applied to each separate element of the array or **array section**.

5 If a **list item** is an **array section**, the elements of any copy of the **array section** will be stored
6 contiguously.

▼ Fortran ▲

7 If the **original list item** has the **POINTER** attribute, any copies of the **list item** are associated with
8 private targets.

▲ Fortran ▼

9 Restrictions

10 Restrictions common to **reduction clauses** and **induction clauses** are as follows:

- 11 • Any array element must be specified at most once in all **list items** on a **directive**.
- 12 • For a reduction or induction identifier declared in a **declare reduction** or a **declare**
13 **induction directive**, the **directive** must appear before its use in a **reduction clause** or
14 **induction clause**.
- 15 • If a **list item** is an **array section**, it must specify contiguous storage, it cannot be a zero-length
16 **array section** and its **base expression** must be a **base language** identifier.
- 17 • If a **list item** is an **array section** or an array element, accesses to the elements of the array
18 outside the specified **array section** or array element result in **unspecified behavior**.

▼ C / C++ ▲

- 19 • The type of a **list item** that appears in a **reduction clause** must be valid for the reduction
20 identifier. The type of a **list item** and of the **step expression** that appear in an **induction**
21 **clause** must be valid for the induction identifier.
- 22 • A **list item** that appears in a **reduction clause** or **induction clause** must not be
23 **const**-qualified.
- 24 • The reduction or induction identifier for any **list item** must be unambiguous and accessible.

▲ C / C++ ▼

▼ Fortran ▲

- 25 • The type, type parameters and rank of a **list item** that appears in a **reduction clause** must be
26 valid for the **combiner expression** and the **initializer expression**. The type, type parameters
27 and rank of a **list item** and of the **step expression** that appear in an **induction clause** must
28 be valid for the **inductor expression**.

- A [list item](#) that appears in a reduction or [induction clause](#) must be definable.
- A procedure pointer must not appear in a [reduction clause](#) or [induction clause](#).
- A pointer with the **INTENT (IN)** attribute must not appear in a [reduction clause](#) or [induction 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](#) or [inductor expression](#) must be associated at entry to the [construct](#) that contains the [reduction clause](#) or [induction 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 a [combiner expression](#), [initializer expression](#), or [inductor expression](#) must be in the allocated state at entry to the [construct](#) that contains the [reduction clause](#) or [induction 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 or induction identifier is defined in a [declare reduction](#) or [declare induction directive](#), that [directive](#) must be in the same subprogram, or accessible by host or use association.
- If the reduction or induction identifier is a user-defined operator, the same explicit interface for that operator must be accessible at the location of the [declare reduction](#) or [declare induction directive](#) that defines the reduction or induction identifier.
- If the reduction or induction identifier is defined in a [declare reduction](#) or [declare induction directive](#), any procedure referenced in the [initializer](#), [combiner](#), [inductor](#), or [collector clause](#) must be an intrinsic function, or must have an explicit interface where the same explicit interface is accessible as at the [declare reduction](#) or [declare induction directive](#).

Fortran

6.5.6 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](#).

C++

If the type is of [class type](#) and the reduction identifier is implicitly declared, then it must provide the operator as described in [Section 6.5.5](#) as well as one of:

- A default constructor and an assignment operator that accepts a type that can be implicitly constructed from an integer expression.

```
template<typename T>
requires(T&& t) {
    T();
    t = 0;
};
```

- A single-argument constructor that accepts a type that can be implicitly constructed from an integer expression.

```
template<typename T>
requires() {
    T(0);
};
```

The first of these that matches will be used, with the initializer value being passed to the assignment operator or constructor.

C++

Any copies of a [list item](#) associated with the reduction are initialized with the initializer value of the reduction identifier. Any copies are combined using the combiner associated with the reduction identifier.

Execution Model Events

The *reduction-begin event* occurs before a [task](#) begins to perform loads and stores that belong to the implementation of a reduction and the *reduction-end event* occurs after the [task](#) has completed loads and stores associated with the reduction. If a [task](#) participates in multiple reductions, each reduction may be bracketed by its own pair of *reduction-begin/reduction-end events* or multiple reductions may be bracketed by a single pair of *events*. The interval defined by a pair of *reduction-begin/reduction-end events* may not contain a [task scheduling point](#).

Tool Callbacks

A [thread](#) dispatches a registered `ompt_callback_reduction` with `ompt_sync_region_reduction` in its *kind* argument and `ompt_scope_begin` as its *endpoint* argument for each occurrence of a *reduction-begin event* in that [thread](#). Similarly, a [thread](#) dispatches a registered `ompt_callback_reduction` with `ompt_sync_region_reduction` in its *kind* argument and `ompt_scope_end` as its *endpoint* argument for each occurrence of a *reduction-end event* in that [thread](#). These [callbacks](#) occur in the context of the [task](#) that performs the reduction and has the type signature `ompt_callback_sync_region_t`.

Restrictions

Restrictions common to [reduction clauses](#) are as follows:

C

- For a **max** or **min** reduction, the type of the [list item](#) must be an allowed arithmetic data type: **char**, **int**, **float**, **double**, or **_Bool**, possibly modified with **long**, **short**, **signed**, or **unsigned**.

C

C++

- For a **max** or **min** reduction, the type of the [list item](#) must be an allowed arithmetic data type: **char**, **wchar_t**, **int**, **float**, **double**, or **bool**, possibly modified with **long**, **short**, **signed**, or **unsigned**.

C++

Cross References

- [ompt_callback_sync_region_t](#), see [Section 20.5.2.13](#)
- [ompt_scope_endpoint_t](#), see [Section 20.4.4.11](#)
- [ompt_sync_region_t](#), see [Section 20.4.4.14](#)

6.5.7 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 6.5.5](#) and [Section 6.5.6](#), 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](#).

6.5.8 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 6.5.5](#) and [Section 6.5.6](#), 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.

6.5.9 reduction Clause

Name: <code>reduction</code>	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
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: directive-name	unique

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 6.5.7](#) and [Section 6.5.8](#). 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

1 **private** clause had been used. For the **taskloop** construct, private copies are created
2 according to the rules of the **reduction scoping** clause. For the **teams** construct, one or more
3 private copies of each **list item** are created for the **initial task** of each **team** in the **league**, as if the
4 **private** clause had been used. For the **loop** construct, private copies are created and used in the
5 **construct** according to the description and restrictions in Section 6.3. At the end of a **region** that
6 corresponds to a **construct** for which the **reduction** clause was specified, the **original list item** is
7 updated by combining its original value with the final value of each of the private copies, using the
8 combiner of the specified *reduction-identifier*.

9 If the **inscan** *reduction-modifier* is present, a **scan computation** is performed over updates to the
10 **list item** performed in each **logical iteration** of the **associated loops** (see Section 6.6). The **list items**
11 are privatized in the **construct** according to the description and restrictions in Section 6.3. At the
12 end of the **region**, each **original list item** is assigned the value described in Section 6.6.

13 If the **task** *reduction-modifier* is present for a **parallel** or **worksharing construct**, then each **list**
14 **item** is privatized according to the description and restrictions in Section 6.3, and an unspecified
15 number of additional private copies may be created to support **task** reductions. Any copies
16 associated with the reduction are initialized before they are accessed by the **tasks** that participate in
17 the reduction, which include all **implicit tasks** in the corresponding **region** and all participating
18 **explicit tasks** that specify an **in_reduction** clause (see Section 6.5.11). After the end of the
19 **region**, the **original list item** contains the result of the reduction.

20 Restrictions

21 Restrictions to the **reduction** clause are as follows:

- 22 • All restrictions common to all **reduction** clauses, as listed in Section 6.5.5 and Section 6.5.6,
23 apply to this clause.
- 24 • A **list item** that appears in a **reduction** clause on a **worksharing** construct must be shared
25 in the **parallel** region to which the **worksharing** region binds.
- 26 • If an **array section** or array element appears as a **list item** in a **reduction** clause on a
27 **worksharing** construct, all **threads** of the **team** must specify the same **storage location**.
- 28 • Each **list item** specified with the **inscan** *reduction-modifier* must appear as a **list item** in an
29 **inclusive** or **exclusive** clause on a **scan** directive enclosed by the **construct**.
- 30 • If the **inscan** *reduction-modifier* is specified, a **reduction** clause without the **inscan**
31 *reduction-modifier* must not appear on the same **construct**.
- 32 • A **reduction** clause with the **task** *reduction-modifier* may only appear on a **parallel**
33 **construct** or a **worksharing** construct, or a **combined** construct or a **composite** construct for
34 which any of the aforementioned **constructs** is a **constituent construct** and neither **simd** nor
35 **loop** are constituent constructs.
- 36 • A **reduction** clause with the **inscan** *reduction-modifier* may only appear on a
37 **worksharing-loop** construct or a **simd** construct, or a **combined** construct or a **composite**

1 `construct` for which any of the aforementioned `constructs` is a `constituent construct` and
2 `distribute` is not a `constituent construct`.

- 3 • The `inscan reduction-modifier` must not be specified on a `construct` for which the
4 `ordered` or `schedule` clause is specified.
- 5 • A `list item` that appears in a `reduction clause` of the innermost enclosing `worksharing`
6 `construct` or `parallel construct` must not be accessed in an `explicit task` generated by a
7 `construct` for which an `in_reduction clause` over the same `list item` does not appear.
- 8 • The `task reduction-modifier` must not appear in a `reduction clause` if the `nowait`
9 `clause` is specified on the same `construct`.

C / C++

- 10 • If a `list item` in a `reduction clause` on a `worksharing construct` has a reference type then it
11 must bind to the same object for all `threads` of the `team`.
- 12 • If a `list item` in a `reduction clause` on a `worksharing construct` is an `array section` or an
13 array element then the `base pointer` must point to the same `variable` for all `thread` of the `team`.
- 14 • A `variable` of `class type` (or array thereof) that appears in a `reduction clause` with the
15 `inscan reduction-modifier` requires an accessible, unambiguous default constructor for the
16 `class type`; the number of calls to it while performing the `scan computation` is unspecified.
- 17 • A `variable` of `class type` (or array thereof) that appears in a `reduction clause` with the
18 `inscan reduction-modifier` requires an accessible, unambiguous copy assignment operator
19 for the `class type`; the number of calls to it while performing the `scan computation` is
20 unspecified.

C / C++

Cross References

- 21 • `ordered` clause, see [Section 5.4.4](#)
- 22
- 23 • `private` clause, see [Section 6.4.3](#)
- 24 • `schedule` clause, see [Section 12.6.3](#)
- 25 • `do` directive, see [Section 12.6.2](#)
- 26 • `for` directive, see [Section 12.6.1](#)
- 27 • `loop` directive, see [Section 12.8](#)
- 28 • `parallel` directive, see [Section 11.2](#)
- 29 • `scan` directive, see [Section 6.6](#)
- 30 • `scope` directive, see [Section 12.2](#)
- 31 • `sections` directive, see [Section 12.3](#)
- 32 • `simd` directive, see [Section 11.5](#)

- **taskloop** directive, see [Section 13.7](#)
- **teams** directive, see [Section 11.3](#)
- List Item Privatization, see [Section 6.3](#)

6.5.10 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
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: directive-name	unique

Directives

taskgroup

Semantics

The **task_reduction** clause is a [reduction scoping clause](#), as described in [Section 6.5.7](#), 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 6.5.5](#) and [Section 6.5.6](#), apply to this [clause](#).

Cross References

- **taskgroup** directive, see [Section 16.4](#)

6.5.11 `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
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: directive-name	unique

Directives

target, **task**, **taskloop**

Semantics

The `in_reduction` clause is a `reduction participating clause`, as described in Section 6.5.8, 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.

For the `task` construct, the generated `task` becomes the participating `task`. For each `list item`, a private copy may be created as if the `private` clause had been used.

For the `target` construct, the `target` `task` becomes the participating `task`. For each `list item`, a private copy may be created in the `data environment` of the `target` `task` as if the `private` clause had been used. This private copy will be implicitly mapped into the `device data environment` of the `target device`, if the `target device` is not the `parent device`.

At the end of the `task region`, if a private copy was created its value is combined with a copy created by a `reduction scoping clause` or with the `original list item`.

Restrictions

Restrictions to the `in_reduction` clause are as follows:

- All restrictions common to all `reduction` clauses, as listed in Section 6.5.5 and Section 6.5.6, apply to this clause.
- A `list item` that appears in a `task_reduction` clause or a `reduction` clause with `task` as the `reduction-modifier` that is specified on a `construct` that corresponds to a `region` in which the `region` of the participating `task` is a `closely nested region` must match each `list item`. The `construct` that corresponds to the innermost enclosing `region` that meets this condition must specify the same `reduction-identifier` for the matching `list item` as the `in_reduction` clause.

Cross References

- `target` directive, see Section 14.8
- `task` directive, see Section 13.6
- `taskloop` directive, see Section 13.7

6.5.12 induction Clause

Name: <code>induction</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>induction-identifier</i>	<i>list</i>	OpenMP induction identifier	required, ultimate
<i>step-modifier</i>	<i>list</i>	Complex, name: step Arguments: <i>induction-step</i> expression of induction-step type (region-invariant)	required
<i>induction-modifier</i>	<i>list</i>	Keyword: relaxed , strict	default
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: directive-name	unique

Directives

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

Semantics

The **induction** clause provides a superset of the functionality provided by the **private** clause. A **list item** that appears in an **induction** clause is subject to the **private** clause semantics described in Section 6.4.3, except as otherwise specified.

When an **induction** clause is specified on a **loop-nest-associated directive** and the **strict induction-modifier** is present, the value of the **new list item** at the beginning of each **collapsed iteration** is determined by the closed form of the **induction operation**. The value of the **original list item** at the end of the last **collapsed iteration** is the result of applying the **inductor expression** to the value of the **new list item** at the beginning of that **collapsed iteration**. When the **relaxed induction-modifier** is present, the implementation may assume that the value of the **new list item** at the end of the previous **collapsed iteration**, if executed by the same **task** or **SIMD lane**, is the value determined by the closed form of the **induction operation**. When an **induction-modifier** is not specified, the behavior is as if the **relaxed induction-modifier** is present.

The value of the **new list item** at the end of the last **collapsed iteration** is assigned to the **original list item**.

If the **construct** is a **worksharing-loop construct** with the **nowait** clause present and the **original list item** is shared in the **enclosing context**, access to the **original list item** after the **construct** may create a data race. To avoid this data race, user code must insert synchronization.

The **induction-identifier** must match a previously declared induction identifier of the same name and type for each of the **list items** and for the **induction-step-expr**. This match is done by means of a name lookup in the **base language**.

Restrictions

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

- All restrictions listed in Section 6.5.5 apply to this clause.
- The **induction-step** must not be an array or array section.
- If an array section or array element appears as a **list item** in an **induction** clause on a **worksharing construct**, all **threads** of the **team** must specify the same **storage location**.

▼ C / C++ ▼

- If a **list item** in an **induction** clause on a **worksharing construct** has a reference type and the **original list item** is shared in the **enclosing context** then it must bind to the same object for all **threads** of the **team**.
- If a **list item** in an **induction** clause on a **worksharing construct** is an array section or an array element and the **original list item** is shared in the **enclosing context** then the **base pointer** must point to the same **variable** for all **threads** of the **team**.

▲ C / C++ ▲

Cross References

- **private** clause, see [Section 6.4.3](#)
- **distribute** directive, see [Section 12.7](#)
- **do** directive, see [Section 12.6.2](#)
- **for** directive, see [Section 12.6.1](#)
- **simd** directive, see [Section 11.5](#)
- **taskloop** directive, see [Section 13.7](#)
- List Item Privatization, see [Section 6.3](#)

6.5.13 declare reduction Directive

Name: <code>declare reduction</code>	Association: none
Category: declarative	Properties: pure

Arguments

`declare reduction`(*reduction-specifier*)

Name	Type	Properties
<i>reduction-specifier</i>	OpenMP reduction specifier	<i>default</i>

Clauses

[combiner](#), [initializer](#)

Additional information

The syntax *reduction-identifier* : *typename-list* : *combiner-expr*, where *combiner* is an OpenMP combiner expression, may alternatively be used for *reduction-specifier*. The [combiner clause](#) must not be specified if this syntax is used. This syntax has been [deprecated](#).

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
```

where *reduction-identifier* is a reduction identifier and *typename-list* is a type-name [list](#).

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++

1 The **declare reduction directive** can also appear at the locations in a program where a static
2 data member could be declared. In this case, the visibility and accessibility of the declaration are
3 the same as those of a static data member declared at the same location in the program.

C++

4 The **enclosing context** of the *combiner-expr* specified by the **combiner clause** and of the
5 *initializer-expr* that is specified by the **initializer clause** is that of the **declare**
6 **reduction directive**. The *combiner-expr* and the *initializer-expr* must be correct in the **base**
7 **language** as if they were the body of a function defined at the same location in the program.

Fortran

8 If a type with deferred or assumed length type parameter is specified in a **declare reduction**
9 **directive**, the *reduction-identifier* of that **directive** can be used in a **reduction clause** with any
10 **variable** of the same type and the same kind parameter, regardless of the length type parameters
11 with which the **variable** is declared.

12 If the *reduction-identifier* is the same as the name of a user-defined operator or an extended
13 operator, or the same as a generic name that is one of the allowed intrinsic procedures, and if the
14 operator or procedure name appears in an accessibility statement in the same module, the
15 accessibility of the corresponding **declare reduction directive** is determined by the
16 accessibility attribute of the statement.

17 If the *reduction-identifier* is the same as a generic name that is one of the allowed intrinsic
18 procedures and is accessible, and if it has the same name as a derived type in the same module, the
19 accessibility of the corresponding **declare reduction directive** is determined by the
20 accessibility of the generic name according to the **base language**.

Fortran

Restrictions

21 Restrictions to the **declare reduction directive** are as follows:

- 22 • A *reduction-identifier* may not be re-declared in the current scope for the same type or for a
23 type that is compatible according to the **base language** rules.
- 24 • The *typename-list* must not declare new types.

C / C++

- 25 • A type name in a **declare reduction directive** cannot be a function type, an array type,
26 a reference type, or a type qualified with **const**, **volatile** or **restrict**.

C / C++

Fortran

- 27 • If the length type parameter is specified for a type, it must be a constant, a colon (:) or an
28 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

- **combiner** clause, see [Section 6.5.14](#)
- **initializer** clause, see [Section 6.5.15](#)
- OpenMP Combiner Expressions, see [Section 6.5.2.1](#)
- OpenMP Initializer Expressions, see [Section 6.5.2.2](#)
- OpenMP Reduction and Induction Identifiers, see [Section 6.5.1](#)

6.5.14 combiner Clause

Name: combiner	Properties: unique, required
------------------------------	-------------------------------------

Arguments

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

Modifiers

Name	Modifies	Type	Properties
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: directive-name	unique

Directives

declare reduction

Semantics

This **clause** specifies *combiner-expr* as the **combiner expression** for a **user-defined reduction**.

Cross References

- **declare reduction** directive, see [Section 6.5.13](#)
- OpenMP Combiner Expressions, see [Section 6.5.2.1](#)

6.5.15 initializer Clause

Name: initializer	Properties: unique
---------------------------------	---------------------------

Arguments

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

Modifiers

Name	Modifies	Type	Properties
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: directive-name	unique

Directives

declare reduction

Semantics

This [clause](#) specifies *initializer-expr* as the [initializer expression](#) for a user-defined-reduction.

Cross References

- **declare reduction** directive, see [Section 6.5.13](#)
- OpenMP Initializer Expressions, see [Section 6.5.2.2](#)

6.5.16 declare induction Directive

Name: declare induction	Association: none
Category: declarative	Properties: pure

Arguments

declare induction(*induction-specifier*)

Name	Type	Properties
<i>induction-specifier</i>	OpenMP induction specifier	<i>default</i>

Clauses

[collector](#), [inductor](#)

Semantics

The [declare induction directive](#) declares an *induction-identifier* that can be used in an [induction clause](#) as a user-defined-induction. The [directive](#) argument *induction-specifier* uses the following syntax:

```
induction-identifier : type-specifier-list
type-specifier-list := type-specifier | type-specifier , type-specifier-list
type-specifier     := typename-list | typename-pair
typename-pair     := ( type , type )
```

where *induction-identifier* is an induction identifier and *typename-list* is a type-name [list](#).

The *induction-identifier* identifies the [declare induction directive](#). The *induction-identifier* can be used in an [induction clause](#) that lists [induction variables](#) of the types specified in the *typename-list*, with corresponding [step expressions](#) of the same type if the *type-specifier-list* item uses the form that specifies only one *type*. If the *type-specifier-list* item uses the *typename-pair*

1 form then the *induction-identifier* can be used in an **induction clause** that lists that pair, in
2 which case the **induction variable** must be of the first type specified in the *typename-pair* while the
3 corresponding **step expression** must be of the second type in the *typename-pair*.

4 The visibility and accessibility of a user-defined-induction are the same as those of a **variable**
5 declared at the same location in the program.

C++

6 The **declare induction directive** can also appear at the locations in a program where a static
7 data member could be declared. In this case, the visibility and accessibility of the declaration are
8 the same as those of a static data member declared at the same location in the program.

C++

9 The **enclosing context** of the **inductor expression** specified by the **inductor clause** and of the
10 **collector expression** specified by the **collector clause** is that of the **declare induction**
11 **directive**. The **inductor expression** and the **collector expression** must be correct in the **base language**
12 as if they were the body of a function defined at the same location in the program.

Fortran

13 If the *induction-identifier* is the same as the name of a user-defined operator or an extended
14 operator, or the same as a generic name that is one of the allowed intrinsic procedures, and if the
15 operator or procedure name appears in an accessibility statement in the same module, the
16 accessibility of the corresponding **declare induction directive** is determined by the
17 accessibility attribute of the statement.

18 If the *induction-identifier* is the same as a generic name that is one of the allowed intrinsic
19 procedures and is accessible, and if it has the same name as a derived type in the same module, the
20 accessibility of the corresponding **declare induction directive** is determined by the
21 accessibility of the generic name according to the **base language**.

Fortran

22 Restrictions

23 Restrictions to the **declare induction directive** are as follows:

- 24 • A *induction-identifier* may not be re-declared in the current scope for the same type or for a
25 type that is compatible according to the **base language** rules.
- 26 • The *typename-list* must not declare new types.

C / C++

- 27 • A type name in a **declare induction directive** cannot be a function type, an array type,
28 a reference type, or a type qualified with **const**, **volatile** or **restrict**.

C / C++

Cross References

- `collector` clause, see [Section 6.5.18](#)
- `inductor` clause, see [Section 6.5.17](#)
- OpenMP Collector Expressions, see [Section 6.5.2.4](#)
- OpenMP Inductor Expressions, see [Section 6.5.2.3](#)
- OpenMP Reduction and Induction Identifiers, see [Section 6.5.1](#)

6.5.17 inductor Clause

Name: <code>inductor</code>	Properties: unique, required
------------------------------------	-------------------------------------

Arguments

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

Modifiers

Name	Modifies	Type	Properties
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: directive-name	unique

Directives

[declare induction](#)

Semantics

This [clause](#) specifies *inductor-expr* as the [inductor expression](#) for a [user-defined induction](#).

Cross References

- `declare induction` directive, see [Section 6.5.16](#)
- OpenMP Inductor Expressions, see [Section 6.5.2.3](#)

6.5.18 collector Clause

Name: <code>collector</code>	Properties: unique, required
-------------------------------------	-------------------------------------

Arguments

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

Modifiers

Name	Modifies	Type	Properties
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: directive-name	unique

Directives

[declare induction](#)

Semantics

This [clause](#) specifies *collector-expr* as the [collector expression](#) for a [user-defined induction](#), which ensures that a [collector](#) is available for use in the closed form of the [induction operation](#).

Cross References

- [declare induction](#) directive, see [Section 6.5.16](#)
- OpenMP Collector Expressions, see [Section 6.5.2.4](#)

6.6 scan Directive

Name: scan Category: subsidiary	Association: separating Properties: pure
--	---

Separated directives

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

Clauses

[exclusive](#), [inclusive](#)

Clause set

Properties: unique, required, exclusive	Members: exclusive , inclusive
--	---

Semantics

The [scan directive](#) is a [subsidiary directive](#) that 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 [collapsed iteration](#), and the [scan phase](#) ensures that any statement that reads the [list item](#) uses the result of the [scan computation](#) for that [collapsed iteration](#). Thus, the [scan directive](#) specifies that a [scan computation](#) updates each [list item](#) on each [collapsed iteration](#) of the enclosing [canonical loop nest](#) that is associated with the [separated construct](#).

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

1 preceding **structured block sequence** and, thus, the **directive** specifies that an **exclusive scan**
2 **computation** is performed for each **list item** of *list*.

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

- 8 • if $N = 1$, a_1
- 9 • if $N > 1$, $op(\text{PRESUM}_{\text{last}}(op, a_1, \dots, a_j), \text{PRESUM}_{\text{last}}(op, a_k, \dots, a_N))$,
10 $1 \leq j + 1 = k \leq N$.

11 At the beginning of the **input phase** of each **collapsed iteration**, the **new list item** is initialized with
12 the value of the **initializer expression** of the *reduction-identifier* specified by the **reduction**
13 **clause** on the **separated construct**. The **update value** of a **new list item** is, for a given **collapsed**
14 **iteration**, the value of the **new list item** on completion of its **input phase**.

15 Let *orig-val* be the value of the **original list item** on entry to the **separated construct**. Let *combiner*
16 be the **combiner expression** for the *reduction-identifier* specified by the **reduction clause** on the
17 **construct**. Let u_i be the **update value** of a **list item** for **collapsed iteration** *i*. For **list items** that appear
18 in an **inclusive clause** on the **scan directive**, at the beginning of the **scan phase** for **collapsed**
19 **iteration** *i* the **new list item** is assigned the result of the operation $\text{PRESUM}_{\text{last}}(\text{combiner}, \text{orig-val},$
20 $u_0, \dots, u_i)$. For **list items** that appear in an **exclusive clause** on the **scan directive**, at the
21 beginning of the **scan phase** for **collapsed iteration** $i = 0$ the **list item** is assigned the value *orig-val*,
22 and at the beginning of the **scan phase** for **collapsed iteration** $i > 0$ the **list item** is assigned the
23 result of the operation $\text{PRESUM}_{\text{last}}(\text{combiner}, \text{orig-val}, u_0, \dots, u_{i-1})$.

24 For **list items** that appear in an **inclusive clause**, at the end of the **separated construct**, the
25 **original list item** is assigned the private copy from the last **collapsed iteration** of the **associated**
26 **loops** of the **separated construct**. For **list items** that appear in an **exclusive clause**, let *k* be the
27 last **collapsed iteration** of the **associated loops** of the **separated construct**. At the end of the
28 **separated construct**, the **original list item** is assigned the result of the operation $\text{PRESUM}_{\text{last}}(\text{combiner},$
29 $\text{orig-val}, u_0, \dots, u_k)$.

30 Restrictions

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

- 32 • A **separated construct** must have at most one **scan directive** as a **separating directive**.
- 33 • The **associated loops** of the **directive** to which the **scan directive** is associated must all be
34 **perfectly nested loops**.
- 35 • Each **list item** that appears in the **inclusive** or **exclusive clause** must appear in a
36 **reduction clause** with the **inscan** modifier on the **separated construct**.
- 37 • Each **list item** that appears in a **reduction clause** with the **inscan** modifier on the
38 **separated construct** must appear in a **clause** on the **scan separating directive**.

- Cross-iteration dependences across different [collapsed 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 item](#) that appear in the [inclusive](#) or [exclusive](#) clause must not be modified in the [scan phase](#).

Cross References

- [exclusive](#) clause, see [Section 6.6.2](#)
- [inclusive](#) clause, see [Section 6.6.1](#)
- [reduction](#) clause, see [Section 6.5.9](#)
- [do](#) directive, see [Section 12.6.2](#)
- [for](#) directive, see [Section 12.6.1](#)
- [simd](#) directive, see [Section 11.5](#)

6.6.1 inclusive Clause

Name: inclusive	Properties: innermost-leaf, unique
--	---

Arguments

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

Modifiers

Name	Modifies	Type	Properties
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: directive-name	unique

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](#) and [array elements](#).

Cross References

- [scan](#) directive, see [Section 6.6](#)

6.6.2 exclusive Clause

Name: <code>exclusive</code>	Properties: innermost-leaf, unique
-------------------------------------	---

Arguments

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

Modifiers

Name	Modifies	Type	Properties
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: directive-name	unique

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](#) and [array elements](#).

Cross References

- **scan** directive, see [Section 6.6](#)

6.7 Data Copying Clauses

This section describes the [copyin](#) clause and the [copyprivate](#) clause. These two [clauses](#) support copying data values from [private variables](#) or [threadprivate variables](#) of an [implicit task](#) or [thread](#) to the corresponding [variables](#) of other [implicit tasks](#) or [threads](#) in the [team](#).

6.7.1 copyin Clause

Name: <code>copyin</code>	Properties: outermost-leaf, data copying
----------------------------------	---

Arguments

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

Modifiers

Name	Modifies	Type	Properties
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: directive-name	unique

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, the copy of each **thread** 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 11.2](#)
- **threadprivate** directive, see [Section 6.2](#)

6.7.2 copyprivate Clause

Name: copyprivate	Properties: innermost-leaf, end-clause, data copying
--------------------------	---

Arguments

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

Modifiers

Name	Modifies	Type	Properties
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: directive-name	unique

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

- 26 • A common block that appears in a **copyprivate** clause must be threadprivate.
- 27 • 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 6.4.4](#)
- **private** clause, see [Section 6.4.3](#)
- **single** directive, see [Section 12.1](#)

6.8 Data-Mapping Control

This section describes the available mechanisms for controlling how data are mapped to [device data environments](#). It covers [implicitly determined data-mapping attribute](#) rules for [variables](#) referenced in [target constructs](#), [clauses](#) that support [explicitly determined data-mapping attributes](#), and [clauses](#) for mapping [variables](#) with static lifetimes and making procedures available 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.

6.8.1 Implicit Data-Mapping Attribute Rules

When specified, [data-mapping attribute clauses](#) on [target directives](#) determine the [data-mapping attributes](#) for [variables](#) referenced in a [target construct](#). Otherwise, the first matching rule from the following list determines the [implicitly determined data-mapping attribute](#) (or [implicitly determined data-sharing attribute](#)) for [variables](#) referenced in a [target construct](#) that do not have a [predetermined data-sharing attribute](#) according to [Section 6.1.1](#). References to structure elements or array elements are treated as references to the structure or array, respectively, for the purposes of [implicitly determined data-mapping attributes](#) or [implicitly determined data-sharing attributes](#) of [variables](#) referenced 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](#).
- If a [variable](#) is the base variable of a [list item](#) in a [reduction](#), [lastprivate](#) or [linear](#) clause on a [combined target construct](#) then the [list item](#) is treated as if it had appeared in a [map](#) clause with a [map-type](#) of [tofrom](#) if [Section 18.2](#) specifies this behavior.

- 1 • If a **variable** is the **base variable** of a **list item** in an **in_reduction** clause on a **target**
- 2 **construct** then it is treated as if the **list item** had appeared in a **map** clause with a **map-type** of
- 3 **tofrom** and an **always-modifier**.
- 4 • If a **defaultmap** clause is present for the category of the **variable** and specifies an implicit
- 5 behavior other than **default**, the **data-mapping attribute** or **data-sharing attribute** is
- 6 determined by that **clause**.

C++

- 7 • If the **target construct** is within a class non-static member function, and a **variable** is an
- 8 accessible data member of the object for which the non-static data member function is
- 9 invoked, the **variable** is treated as if the **this[:1]** expression had appeared in a **map** clause
- 10 with a **map-type** of **tofrom**. Additionally, if the **variable** is of type pointer or reference to
- 11 pointer, it is also treated as if it is the **base expression** of a **zero-offset assumed-size array** that
- 12 appears in a **map** clause with the **alloc map-type**.
- 13 • If the **this** keyword is referenced inside a **target construct** within a class non-static
- 14 member function, it is treated as if the **this[:1]** expression had appeared in a **map** clause
- 15 with a **map-type** of **tofrom**.

C++

C / C++

- 16 • A **variable** that is of type pointer, but is neither a pointer to function nor (for C++) a pointer
- 17 to a member function, is treated as if it is the **base expression** of a **zero-offset assumed-size**
- 18 **array** that appears in a **map** clause with the **alloc map-type**.

C / C++

C++

- 19 • A **variable** that is of type reference to pointer, but is neither a reference to pointer to function
- 20 nor a reference to a pointer to a member function, is treated as if it is the **base expression** of a
- 21 **zero-offset assumed-size array** that appears in a **map** clause with the **alloc map-type**.

C++

Fortran

- 22 • If a **combined target construct** is associated with a **DO CONCURRENT** loop, a **variable** that has
- 23 **SHARED** locality in the loop is treated as if it had appeared in a **map** clause with a **map-type**
- 24 of **tofrom**.

Fortran

- 25 • If a **variable** is not a **scalar variable** then it is treated as if it had appeared in a **map** clause with
- 26 a **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

- If the above rules do not apply then a [scalar variable](#) is not mapped but instead has an [implicitly determined data-sharing attribute](#) of firstprivate (see [Section 6.1.1](#)).

6.8.2 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](#):

```

C / C++
#pragma omp declare mapper( $T$   $v$ ) map(tofrom:  $v$ )

C / C++
Fortran
!$omp declare mapper( $T$  ::  $v$ ) map(tofrom:  $v$ )

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 6.9.2](#)
- [map](#) clause, see [Section 6.8.3](#)
- [to](#) clause, see [Section 6.9.1](#)

6.8.3 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>always-modifier</i>	<i>locator-list</i>	Keyword: always	map-type-modifying
<i>close-modifier</i>	<i>locator-list</i>	Keyword: close	map-type-modifying
<i>present-modifier</i>	<i>locator-list</i>	Keyword: present	map-type-modifying
<i>self-modifier</i>	<i>locator-list</i>	Keyword: self	map-type-modifying
<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
<i>map-type</i>	<i>locator-list</i>	Keyword: alloc, delete, from, release, to, tofrom	default
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: directive-name	unique

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 data environment of the current task 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 a

1 `map-entering clause`, which can only appear on `constructs` that have the `map-entering property`, if
2 the `map-type` is `to`, `tofrom` or `alloc`. The `map clause` is a `map-exiting clause`, which can only
3 appear on `constructs` that have the `map-exiting property`, if the `map-type` is `from`, `tofrom`,
4 `release` or `delete`.

5 The `list items` that appear in a `map clause` may include `array sections`, `assumed-size arrays`, and
6 `structure elements`. A `list item` in a `map clause` may reference any `iterator-identifier` defined in its
7 `iterator modifier`. A `list item` may appear more than once in the `map clauses` that are specified on
8 the same `directive`.

C / C++

9 If a `list item` is a zero-length `array section` that has a single array subscript, the behavior is as if the
10 `list item` is an `assumed-size array` that is instead mapped with the `alloc map-type`.

C / C++

11 When a `list item` in a `map clause` that is not an `assumed-size array` is mapped on a `map-entering`
12 `construct` and `corresponding storage` is created in the `device data environment` on entry to the `region`,
13 the `list item` becomes a `matchable candidate` with an associated `starting address`, `ending address`,
14 and `base address` that define its `mapped address range` and `extended address range`. The current set
15 of `matchable candidates` consists of any `map clause list item` on the `construct` that is a `matchable`
16 `candidate` and all `matchable candidates` that were previously mapped and are still mapped.

17 A `list item` in a `map clause` that is an `assumed-size array` is treated as if an `array section`, with a `base`
18 `expression`, lower bound and length determined as follows, is substituted in its place if a `matched`
19 `candidate` is found. If the `assumed-size array` is an `array section`, the `base expression` of the
20 substitute `array section` is the same as for the `assumed-size array`; otherwise, the `base expression` is
21 the `assumed-size array`. If the `mapped address range` of a `matchable candidate` includes the first
22 `storage location` of the `assumed-size array`, it is a `matched candidate`. If a `matchable candidate` does
23 not exist for which the `mapped address range` includes the first `storage location` of the `assumed-size`
24 `array`, then a `matchable candidate` is a `matched candidate` if its `extended address range` includes the
25 first `storage location` of the `assumed-size array`. If multiple `matched candidates` exist, an arbitrary
26 one of them is the found `matched candidate`. The lower bound and length of the substitute `array`
27 `section` are set such that its storage is identical to the storage of the found `matched candidate`. If a
28 `matched candidate` is not found then a substitute `array section` is not formed and no further actions
29 that are described in this section are performed for the `list item`.

30 A `list item` that is an `array` or `array section` and for which the map type is `tofrom`, `to`, or `from` is
31 mapped as if the map type decays to `alloc` or, if the `construct` on which the `map clause` appears is
32 `target exit data`, to `release`. If a `list item` is an `array` or `array section`, the array elements
33 become implicitly mapped `list items` with the same `modifiers` (including the original map type) as in
34 the `clause`. If the `array` or `array section` is implicitly mapped and `corresponding storage` exists in the
35 `device data environment` prior to a `task` encountering the `construct` on which the `clauserefmap`
36 `clause` appears, only those array elements that have `corresponding storage` are implicitly mapped.

37 If a `mapper modifier` is not present, the behavior is as if a `mapper modifier` was specified with the
38 `default` parameter. The map behavior of a `list item` in a `map clause` is modified by a visible

1 user-defined mapper (see Section 6.8.7) if the *mapper-identifier* of the *mapper modifier* is defined
2 for a *base language* type that matches the type of the *list item*. Otherwise, the predefined default
3 *mapper* for the type of the *list item* applies. The effect of the *mapper* is to remove the *list item* from
4 the *map clause* and to apply the *clauses* specified in the declared *mapper* to the *construct* on which
5 the *map clause* appears. In the *clauses* applied by the *mapper*, references to *var* are replaced with
6 references to the *list item* and the *map-type* is replaced with a final map type that is determined
7 according to the rules of *map-type decay* (see Section 6.8.7). If any *modifier* with the
8 *map-type-modifying property* appears in the *map clause* then the effect is as if that *map-type*
9 *modifier* appears in each *map clause* specified in the declared *mapper*.

Fortran

10 If a component of a derived type *list item* is a *map clause list item* that results from the predefined
11 default *mapper* for that derived type, and if the derived type component is not an explicit *list item* or
12 the *base expression* of an explicit *list item* in a *map clause* on the *construct*, then:

- 13 • If it has the **POINTER** attribute, it is *attach-ineligible*; and
- 14 • If it has the **ALLOCATABLE** attribute and an allocated allocation status, and it is present in
15 the *device data environment* when the *construct* is encountered, the *map clause* may treat its
16 allocation status as if it is unallocated if the corresponding component does not have
17 allocated storage.

18 If a *list item* in a *map clause* is an associated pointer that is not *attach-ineligible* and the pointer is
19 not the *base pointer* of another *list item* in a *map clause* on the same *construct*, then it is treated as if
20 its pointer target is implicitly mapped in the same *clause*. For the purposes of the *map clause*, the
21 mapped pointer target is treated as if its *base pointer* is the associated pointer.

Fortran

C++

22 If a *list item* has a closure type that is associated with a lambda expression, it is mapped as if it has
23 a *structure* type. For each *variable* that is captured by reference by the lambda expression,
24 references to the *variable* in the function call operator for the *new list item* refer to its *corresponding*
25 *storage* in the *device data environment*, if it exists prior to a *task* encountering the *construct*
26 associated with the *map clause*, and otherwise refer to its original storage. For each pointer that is
27 not a function pointer that is captured by the lambda expression, the behavior is as if the pointer or,
28 for capture by copy, the corresponding pointer member of the closure object is the *base expression*
29 of an *zero-offset assumed-size array* that appears in a *map clause* with the **alloc** *map-type*.

30 If the **this** pointer is captured by a lambda expression in class scope, and a *variable* of the
31 associated closure type is later mapped explicitly or implicitly with its full static type, the behavior
32 is as if the object to which **this** points is also mapped as an *array section*, of length one, for which
33 the *base pointer* is the non-static data member that corresponds to the **this** pointer in the closure
34 object.

C++

1 If a **map** clause with a *present-modifier* appears on a **construct** and on entry to the **region** the
2 **corresponding list item** is not present in the **device data environment**, **runtime error termination** is
3 performed.

4 If a **map-entering clause** has the *self-modifier*, the resulting **mapping operations** are **self maps**.

5 The **map** clauses on a **construct** collectively determine the set of **mappable storage blocks** for that
6 **construct**. All **map** clause **list items** that share storage or have the same **containing structure** or
7 **containing array** result in a single **mappable storage block** that contains the storage of the **list items**.
8 The storage for each other **map** clause **list item** becomes a distinct **mappable storage block**.

9 For each **mappable storage block** that is determined by the **map** clauses on a **map-entering**
10 **construct**, on entry to the **region** the following sequence of steps occurs as if performed as a single
11 **atomic operation**:

- 12 1. If a **corresponding storage block** is not present in the **device data environment** then:
 - 13 a) A **corresponding storage block**, which may share storage with the **original storage**
14 **block**, is created in the **device data environment** of the **target device**;
 - 15 b) The **corresponding storage block** receives a reference count that is initialized to zero.
16 This reference count also applies to any part of the **corresponding storage block**.
- 17 2. The reference count of the **corresponding storage block** is incremented by one.
- 18 3. For each **map** clause **list item** on the **construct** that is contained by the **mappable storage**
19 **block**:
 - 20 a) If the reference count of the **corresponding storage block** is one, a **new list item** with
21 language-specific attributes derived from the **original list item** is created in the
22 **corresponding storage block**. The reference count of the **new list item** is always equal to
23 the reference count of its storage.
 - 24 b) If the reference count of the **corresponding list item** is one or if the *always-modifier* is
25 specified, and if the *map-type* is **to** or **tofrom**, the **corresponding list item** is updated
26 as if the **list item** appeared in a **to** clause on a **target update** directive.

27 If the effect of the **map** clauses on a **construct** would assign the value of an **original list item** to a
28 **corresponding list item** more than once, then an implementation is allowed to ignore additional
29 assignments of the same value to the **corresponding list item**.

30 In all cases on entry to the **region**, concurrent reads or updates of any part of the **corresponding list**
31 **item** must be synchronized with any update of the **corresponding list item** that occurs as a result of
32 the **map** clause to avoid data races.

33 For **map** clauses on **map-entering constructs**, if any **list item** has a **base pointer** for which a
34 **corresponding pointer** exists in the **device data environment** after all **mappable storage blocks** are
35 mapped, and either a **new list item** or the **corresponding pointer** is created in the **device data**
36 **environment** on entry to the **region**, then **pointer attachment** is performed and the **corresponding**

1 pointer becomes an attached pointer to the corresponding list item via corresponding base pointer
2 initialization.

3 The original list item and corresponding list item may share storage such that writes to either item
4 by one task followed by a read or write of the other list item by another task without intervening
5 synchronization can result in data races. They are guaranteed to share storage if the mapping
6 operation is a self map, if the map clause appears on a target construct that corresponds to an
7 inactive target region, if it appears on a mapping-only construct that applies to the device data
8 environment of the host device, or if the corresponding list item has an attached pointer that shares
9 storage with its original pointer.

10 For each mappable storage block that is determined by the map clauses on a map-exiting construct,
11 and for which corresponding storage is present in the device data environment, on exit from the
12 region the following sequence of steps occurs as if performed as a single atomic operation:

- 13 1. For each map clause list item that is contained by the mappable storage block:
 - 14 a) If the reference count of the corresponding list item is one or if the *always-modifier* is
15 specified, and if the *map-type* is **from** or **tofrom**, the original list item is updated as if
16 the list item appeared in a **from** clause on a **target update** directive.
 - 17 2. If the *map-type* is not **delete** and the reference count of the corresponding storage block is
18 finite then the reference count is decremented by one.
 - 19 3. If the *map-type* is **delete** and the reference count of the corresponding storage block is
20 finite then the reference count is set to zero.
 - 21 4. If the reference count of the corresponding storage block is zero, all storage to which that
22 reference count applies is removed from the device data environment.

23 If the effect of the map clauses on a construct would assign the value of a corresponding list item to
24 an original list item more than once, then an implementation is allowed to ignore additional
25 assignments of the same value to the original list item.

26 In all cases on exit from the region, concurrent reads or updates of any part of the original list item
27 must be synchronized with any update of the original list item that occurs as a result of the map
28 clause to avoid data races.

29 If a single contiguous part of the original storage of a list item that results from an implicitly
30 determined data-mapping attribute has corresponding storage in the device data environment prior
31 to a task encountering the construct on which the map clause appears, only that part of the original
32 storage will have corresponding storage in the device data environment as a result of the map clause.

33 If a list item with an implicitly determined data-mapping attribute does not have any corresponding
34 storage in the device data environment prior to a task encountering the construct associated with the
35 map clause, and one or more contiguous parts of the original storage are either list items or base
36 pointers to list items that are explicitly mapped on the construct, only those parts of the original
37 storage will have corresponding storage in the device data environment as a result of the map
38 clauses on the construct.

C / C++

1 If a **new list item** is created then the **new list item** will have the same static type as the **original list**
2 **item**, and language-specific attributes of the **new list item**, including size and alignment, are
3 determined by that type.

C / C++

C++

4 If **corresponding storage** that differs from the **original storage** is created in a **device data**
5 **environment**, all **new list items** that are created in that **corresponding storage** are default initialized.
6 Default initialization for **new list items** of **class type**, including their data members, is performed as
7 if with an implicitly-declared default constructor and as if non-static data member initializers are
8 ignored.

9 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
10 the **device data environment** that corresponds to the object referenced by the **original list item**. The
11 effect is as if the object were mapped through a pointer with an **array section** of length one and
12 elements of type *T*.

C++

Fortran

13 If a **new list item** is created then the **new list item** will have the same type, type parameter, and rank
14 as the **original list item**. The **new list item** inherits all default values for the type parameters from
15 the **original list item**.

16 If the allocation status of an **original list item** that has the **ALLOCATABLE** attribute is changed
17 while a **corresponding list item** is present in the **device data environment**, the allocation status of the
18 **corresponding list item** is unspecified until entry to a **region** that corresponds to a **map-entering**
19 **construct** that maps the **list item** with a **map clause** for which the **always-modifier** is specified.

Fortran

20 The **close-modifier** is a hint that the **corresponding storage** should be close to the **target device**.

21 If a **map-entering clause** specifies a **self map** for a **list item** then **runtime error termination** is
22 performed if any of the following is true:

- 23 • The **original list item** is not accessible and cannot be made accessible from the **device**;
- 24 • The **corresponding list item** is present prior to a **task** encountering the **construct** on which the
25 **clause** appears, and the **corresponding storage** differs from the **original storage**; or
- 26 • The **list item** is a pointer that would be assigned a different value as a result of **pointer**
27 **attachment**.

Execution Model Events

28 The **target-map event** occurs in a **thread** that executes the outermost **region** that corresponds to an
29 encountered **device construct** with a **map clause**, after the **target-task-begin event** for the **device**
30 **construct** and before any **mapping operations** are performed.

1 The *target-data-op-begin* event occurs before a thread initiates a data operation on the target device
2 that is associated with a **map** clause, in the outermost region that corresponds to the encountered
3 construct.

4 The *target-data-op-end* event occurs after a thread initiates a data operation on the target device
5 that is associated with a **map** clause, in the outermost region that corresponds to the encountered
6 construct.

7 Tool Callbacks

8 A thread dispatches one or more registered **ompt_callback_target_map** or
9 **ompt_callback_target_map_emi** callbacks for each occurrence of a *target-map* event in
10 that thread. The callback occurs in the context of the target task and has type signature
11 **ompt_callback_target_map_t** or **ompt_callback_target_map_emi_t**,
12 respectively.

13 A thread dispatches a registered **ompt_callback_target_data_op_emi** callback with
14 **ompt_scope_begin** as its endpoint argument for each occurrence of a *target-data-op-begin*
15 event in that thread. Similarly, a thread dispatches a registered
16 **ompt_callback_target_data_op_emi** callback with **ompt_scope_end** as its endpoint
17 argument for each occurrence of a *target-data-op-end* event in that thread. These callbacks have
18 type signature **ompt_callback_target_data_op_emi_t**.

19 A thread dispatches a registered **ompt_callback_target_data_op** callback for each
20 occurrence of a *target-data-op-end* event in that thread. The callback occurs in the context of the
21 target task and has type signature **ompt_callback_target_data_op_t**.

22 Restrictions

23 Restrictions to the **map** clause are as follows:

- 24 • Two list items of the **map** clauses on the same construct must not share original storage
25 unless one of the following is true: they are the same list item, one is the containing structure
26 of the other, at least one is an assumed-size array, or at least one is implicitly mapped due to
27 the list item also appearing in a **use_device_addr** clause.
- 28 • If the same list item appears more than once in **map** clauses on the same construct, the **map**
29 clauses must specify the same *mapper* modifier.
- 30 • A variable that is a **groupprivate** variable or a device local variable must not appear as a list
31 item in a **map** clause.
- 32 • If a list item is an array section, it must specify contiguous storage.
- 33 • If an expression that is used to form a list item in a **map** clause contains an iterator identifier,
34 the list item instances that would result from different values of the iterator must not have the
35 same containing array and must not have base pointers that share original storage.
- 36 • If multiple list items are explicitly mapped on the same construct and have the same
37 containing array or have base pointers that share original storage, and if any of the list items

1 do not have **corresponding list items** that are present in the **device data environment** prior to a
2 **task** encountering the **construct**, then the **list items** must refer to the same array elements of
3 either the **containing array** or the **implicit array** of the **base pointers**.

- 4 • If any part of the **original storage** of a **list item** that is explicitly mapped by a **map clause** has
5 **corresponding storage** in the **device data environment** prior to a **task** encountering the
6 **construct** associated with the **map clause**, all of the **original storage** must have **corresponding**
7 **storage** in the **device data environment** prior to the **task** encountering the **construct**.
- 8 • If an array appears as a **list item** in a **map clause** and it has **corresponding storage** in the
9 **device data environment**, the **corresponding storage** must correspond to a single **mappable**
10 **storage block** that was previously mapped.
- 11 • If a **list item** is an element of a **structure**, and a different element of the **structure** has a
12 **corresponding list item** in the **device data environment** prior to a **task** encountering the
13 **construct** associated with the **map clause**, then the **list item** must also have a **corresponding**
14 **list item** in the **device data environment** prior to the **task** encountering the **construct**.
- 15 • Each **list item** must have a **mappable type**.
- 16 • If a **mapper modifier** appears in a **map clause**, the type on which the specified **mapper**
17 operates must match the type of the **list items** in the **clause**.
- 18 • **Handles** for **memory spaces** and **memory allocators** must not appear as **list items** in a **map**
19 **clause**.
- 20 • If a **list item** is an **assumed-size array**, multiple **matched candidates** must not exist unless they
21 are subobjects of the same **containing structure**.
- 22 • If a **list item** is an **assumed-size array**, the **map-type** must be **alloc**.
- 23 • If a **list item** appears in a **map clause** with the **self-modifier**, any other **list item** in a **map**
24 **clause** on the same **construct** that has the same **base variable** or **base pointer** must also be
25 specified with the **self-modifier**.

C++

- 26 • If a **list item** has a polymorphic **class type** and its static type does not match its dynamic type,
27 the behavior is unspecified if the **map clause** is specified on a **map-entering construct** and a
28 **corresponding list item** is not present in the **device data environment** prior to a **task**
29 encountering the **construct**.
- 30 • No type mapped through a reference may contain a reference to its own type, or any
31 references to types that could produce a cycle of references.

C++

C / C++

- 32 • A **list item** cannot be a **variable** that is a member of a **structure** of a union type.
- 33 • A bit-field cannot appear in a **map clause**.

- A pointer that has a [corresponding pointer](#) that is an [attached pointer](#) must not be modified for the duration of the lifetime of the [list item](#) to which the [corresponding pointer](#) is attached in the [device data environment](#).

C / C++

Fortran

- 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](#).
- If a [list item](#) of a [map clause](#) is an allocatable [variable](#) or is the subobject of an allocatable [variable](#), the [original list item](#) may not be allocated, deallocated or reshaped while the [corresponding list item](#) has allocated storage.
- A pointer that has a [corresponding pointer](#) that is an [attached pointer](#) and is associated with a given pointer target must not become associated with a different pointer target for the duration of the lifetime of the [list item](#) to which the [corresponding pointer](#) is attached in the [device data environment](#).
- If an [array section](#) is mapped and the size of the [array section](#) is smaller than that of the whole array, the behavior of referencing the whole array in a [target region](#) is unspecified.
- A [list item](#) must not be a complex part designator.

Fortran

Cross References

- `declare mapper` directive, see [Section 6.8.7](#)
- `target` directive, see [Section 14.8](#)
- `target data` directive, see [Section 14.5](#)
- `target enter data` directive, see [Section 14.6](#)
- `target exit data` directive, see [Section 14.7](#)
- `target update` directive, see [Section 14.9](#)
- Array Sections, see [Section 4.2.5](#)
- `iterator` modifier, see [Section 4.2.6](#)
- `mapper` modifier, see [Section 6.8.2](#)
- `ompt_callback_target_data_op_emi_t` and `ompt_callback_target_data_op_t`, see [Section 20.5.2.25](#)
- `ompt_callback_target_map_emi_t` and `ompt_callback_target_map_t`, see [Section 20.5.2.27](#)

6.8.4 enter Clause

Name: enter	Properties: data-environment attribute, data-mapping attribute
--------------------	---

Arguments

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

Modifiers

Name	Modifies	Type	Properties
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: directive-name	unique

Directives

declare target

Semantics

The **enter** clause is a data-mapping attribute clause.

If a procedure name appears in an **enter** clause in the same compilation unit in which the definition of the procedure occurs then a device-specific version of the procedure is created for all device to which the directive of the clause applies.

▼ C / C++ ▼

If a variable appears in an **enter** clause in the same compilation unit in which the definition of the variable occurs then a corresponding list item to the original list item is created in the device data environment of all devices to which the directive of the clause applies.

▲ C / C++ ▲

▼ Fortran ▼

If a variable that is host associated appears in an **enter** clause then a corresponding list item to the original list item is created in the device data environment of all devices to which the directive of the clause applies.

▲ Fortran ▲

If a variable appears in an **enter** clause then the corresponding list item in the device data environment of each device to which the directive of the clause applies is initialized once, in the manner specified by the OpenMP program, but at an unspecified point in the OpenMP program prior to the first reference to that list item. The list item is never removed from those device data environments, as if its reference count was initialized to positive infinity.

Restrictions

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

- Each list item must have a mappable type.
- Each list item must have static storage duration.

Cross References

- `declare target` directive, see [Section 8.8.1](#)

6.8.5 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>

Modifiers

Name	Modifies	Type	Properties
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: directive-name	unique

Directives

`declare target`

Semantics

The `link` clause supports compilation of `device procedures` that refer to `variables` with `static storage duration` that appear as `list items` in the clause. The `declare target` directive on which the clause appears does not map the `list items`. Instead, they are mapped according to the data-mapping rules described in [Section 6.8](#).

Restrictions

Restrictions to the `link` clause are as follows:

- Each `list item` must have a `mappable type`.
- Each `list item` must have `static storage duration`.

Cross References

- `declare target` directive, see [Section 8.8.1](#)
- Data-Mapping Control, see [Section 6.8](#)

6.8.6 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, self, 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
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: directive-name	unique

Directives

target

Semantics

The **defaultmap** clause controls the **implicitly determined data-mapping attributes** or **implicitly determined data-sharing attributes** of certain **variables** that are referenced in a **target** construct, in accordance with the rules given in **Section 6.8.1**. 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 **aggregate variables**. 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 a *map-type* of **alloc** and the *present-modifier*. If *implicit-behavior* is **self**, the attribute is a **data-mapping attribute** determined by an implicit **map** clause with a *map-type* of **alloc** and the *self-modifier*. If *implicit-behavior* is **none** then no **implicitly determined data-mapping attributes** or **implicitly determined 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 attribute** 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 14.8](#)
- Implicit Data-Mapping Attribute Rules, see [Section 6.8.1](#)

6.8.7 declare mapper Directive

Name: <code>declare mapper</code> Category: <code>declarative</code>	Association: none Properties: <code>pure</code>
---	--

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* argument declares the `mapper` using the following syntax:

C / C++

```
[ mapper-identifier : ] type var
```

C / C++

Fortran

```
[ mapper-identifier : ] type :: var
```

Fortran

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 [data-motion clause](#) later in the [OpenMP program](#). The visibility and accessibility of this declaration are the same as those of a [variable](#) declared at the same location in the [OpenMP 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 [data-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 6.8.3](#) and [Section 6.9](#).

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 6.5](#). [Table 6.5](#) 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 6.5: 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++

The **declare mapper** directive can also appear at locations in the **OpenMP program** at which a static data member could be declared. In this case, the visibility and accessibility of the declaration are the same as those of a static data member declared at the same location in the **OpenMP program**.

C++

Restrictions

Restrictions to the **declare mapper** directive are as follows:

- No instance of *type* can be mapped as part of the **mapper**, either directly or indirectly through another **base language** type, except the instance *var* that is passed as the **list item**. If a set of **declare mapper** directives results in a cyclic definition then the behavior is unspecified.
- The *type* must not declare a new **base language** type.
- At least one **map clause** that maps *var* or at least one element of *var* is required.
- **List items** in **map clauses** on the **declare mapper** directive may only refer to the declared **variable** *var* and entities that could be referenced by a **procedure defined** at the same location.
- Neither the **release** or **delete map-type** may be specified on any **map clause**.
- If a **mapper-modifier** is specified for a **map clause**, its parameter must be **default**.
- Multiple **declare mapper** directives that specify the same **mapper-identifier** for the same **base language** type or for compatible **base language** types, according to the **base language** rules, may not appear in the same scope.

C

- *type* must be a **struct** or **union** type.

C

C++

- *type* must be a **struct**, **union**, or **class** type.
- If *type* is **struct** or **class**, it must not be derived from any virtual base class.

C++

Fortran

- *type* must not be an intrinsic type, an abstract type, or a parameterized derived type.

Fortran

Cross References

- **map** clause, see [Section 6.8.3](#)

6.9 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*. How the other `devices`, which are the *target device*, are determined is defined by the `construct` specification. Each `data-motion clause` specifies a `data-motion attribute` relative to the *target 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 list item` and the `original list item`. 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 list item` or the `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` `mapper modifier`. The effect of a `data-motion clause` on a `list item` is modified by a visible `user-defined mapper` if a `mapper modifier` is specified with a `mapper-identifier` 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 `compatible map type` with respect to the `data-motion attribute` of 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** of *locator-list* must have a **mappable type**.
- 10 • If an array appears as a **list item** in a **data-motion clause** and it has **corresponding storage** in
11 the **device data environment**, the **corresponding storage** must correspond to a single
12 **mappable storage block** that was previously mapped.
- 13 • If a **mapper modifier** appears in a **data-motion clause**, the specified **mapper** must operate on a
14 type that matches either the type or array element type of each **list item** in the **clause**.
- 15

Fortran

- 16 • The association status of a **list item** that is a pointer must not be undefined unless it is a
17 structure component and it results from a predefined default **mapper**.

Fortran

Cross References

- 18 • **device** clause, see [Section 14.2](#)
- 19 • **from** clause, see [Section 6.9.2](#)
- 20 • **to** clause, see [Section 6.9.1](#)
- 21 • **declare mapper** directive, see [Section 6.8.7](#)
- 22 • **target update** directive, see [Section 14.9](#)
- 23 • Array Sections, see [Section 4.2.5](#)
- 24 • Array Shaping, see [Section 4.2.4](#)
- 25 • **iterator** modifier, see [Section 4.2.6](#)
- 26

6.9.1 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>locator-list</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
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: directive-name	unique

Directives

target update

Semantics

The **to** clause is a [data-motion clause](#) that specifies movement to the [target devices](#) from the [encountering device](#) so the [corresponding list items](#) are the [assigned list items](#) and the [compatible map types](#) are **to** and **tofrom**.

▼ C++ ▲

A list item for which a [mapper](#) does not exist is ignored if it has [static storage duration](#) and either it has the **constexpr** specifier or it is a non-mutable member of a [structure](#) that has the **constexpr** specifier.

▲ C++ ▼

Cross References

- **target update** directive, see [Section 14.9](#)
- **iterator** modifier, see [Section 4.2.6](#)

6.9.2 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>locator-list</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
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: directive-name	unique

Directives

target update

Semantics

The **from** clause is a **data-motion clause** that specifies movement from the **target devices** to the **encountering device** so the **original list items** are the **assigned list items** and the **compatible map types** are **from** and **tofrom**.

▼ C ▼

A list item for which a **mapper** does not exist is ignored if it has the **const** specifier or if it is a member of a **structure** that has the **const** specifier.

▲ C ▲

▼ C++ ▼

A list item for which a **mapper** does not exist is ignored if it has the **const** or **constexpr** specifier or if it is a non-mutable member of a **structure** that has the **const** or **constexpr** specifier.

▲ C++ ▲

Cross References

- `target update` directive, see [Section 14.9](#)
- `iterator` modifier, see [Section 4.2.6](#)

6.10 uniform Clause

Name: <code>uniform</code>	Properties: data-environment attribute
-----------------------------------	---

Arguments

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

Modifiers

Name	Modifies	Type	Properties
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: directive-name	unique

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 8.7](#)

6.11 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
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: directive-name	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 each **list item** must be an array or pointer type.

C

C++

- The type of each **list item** must be an array, pointer, reference to array, or reference to pointer type.

C++

Fortran

- Each **list item** must be an array.

Fortran

Cross References

- `declare simd` directive, see [Section 8.7](#)
- `simd` directive, see [Section 11.5](#)

6.12 `groupprivate` Directive

Name: <code>groupprivate</code>	Association: none
Category: <code>declarative</code>	Properties: <code>pure</code>

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*.
- If any **list item** is a namespace-scope **variable**, the **directive** 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 the *list*.
- Each **variable** in the *list* of a **groupprivate** directive at file, namespace, or class scope must refer to a **variable** declaration at file, namespace, or class scope that lexically precedes the **directive**.
- If any **list item** is a static block-scope **variable**, the **directive** must appear in the scope of the **variable** and not in a nested scope and must lexically precede all references to any of the **variables** in the *list*.

- 1 • Each **variable** in the *list* of a **groupprivate** directive in block scope must have **static storage duration** and must refer to a **variable** declaration in the same scope that lexically precedes the **directive**.
- 2
- 3
- 4 • If a **variable** is specified in a **groupprivate** directive in one **compilation unit**, it must be specified in a **groupprivate** directive in every **compilation unit** in which it is declared.
- 5



- 6 • If any **list item** is a static class member variable, the **directive** must appear in the class definition, in the same scope in which the member **variable** is declared, and must lexically precede all references the **variable**.
- 7
- 8
- 9 • A **groupprivate variable** must not have an incomplete type or a reference type.



- 10 • Each **list item** must be a named **variable** or a named common block; a named common block must appear between slashes.
- 11
- 12 • The *list* argument must not include any coarrays or associate names.
- 13 • The **groupprivate directive** must appear in the declaration section of a scoping unit in which the common block or **variable** is declared.
- 14
- 15 • If a **groupprivate directive** that specifies a common block name appears in one **compilation unit**, then such a **directive** must also appear in every other **compilation unit** that contains a **COMMON** statement that specifies the same name. Each such **directive** must appear after the last such **COMMON** statement in that **compilation unit**.
- 16
- 17
- 18
- 19 • If a **groupprivate variable** or a **groupprivate common block** is declared with the **BIND** attribute, the corresponding C entities must also be specified in a **groupprivate directive** in the C program.
- 20
- 21
- 22 • A **variable** may only appear as an argument in a **groupprivate 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.
- 23
- 24
- 25 • 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.
- 26
- 27 • The effect of an access to a **groupprivate variable** in a **DO CONCURRENT** construct is unspecified.
- 28



29 **Cross References**

- 30 • **device_type** clause, see [Section 14.1](#)

6.13 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>

Modifiers

Name	Modifies	Type	Properties
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: directive-name	unique

Directives

declare target

Semantics

The **local clause** specifies that a reference to a **list item** on a given **device** will refer to a copy of the **list item** that is a **device local variable** and is in **memory** associated with the **device**.

Cross References

- **declare target** directive, see [Section 8.8.1](#)

7 Memory Management

This chapter defines [directives](#), [clauses](#) and related concepts for managing [memory](#) used by [OpenMP programs](#).

7.1 Memory Spaces

OpenMP [memory spaces](#) represent storage resources where [variables](#) can be stored and retrieved. Table [7.1](#) 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 7.1: 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.

7.2 Memory Allocators

OpenMP [memory allocators](#) can be used by an [OpenMP 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 a [memory allocator](#).

The behavior of the allocation process can be affected by the [allocator traits](#) that the user specifies. Table 7.2 shows the allowed [allocator traits](#), their possible values and the default value of each [trait](#).

TABLE 7.2: Allocator Traits

Allocator trait	Allowed values	Default value
<code>sync_hint</code>	<code>contended, uncontended, serialized, private</code>	<code>contended</code>
<code>alignment</code>	Positive integer powers of 2	1 byte
<code>access</code>	<code>all, memspace, device, cgroup, pteam, thread</code>	<code>memspace</code>
<code>pool_size</code>	Any positive integer	Implementation defined
<code>fallback</code>	<code>default_mem_fb, null_fb, abort_fb, allocator_fb</code>	See below
<code>fb_data</code>	An allocator handle	(none)
<code>pinned</code>	<code>true, false</code>	<code>false</code>
<code>partition</code>	<code>environment, nearest, blocked, 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, multiple</code>	<code>single</code>
<code>atomic_scope</code>	<code>all, device</code>	<code>device</code>
<code>part_size</code>	Positive integer value	Implementation defined

1 The **sync_hint trait** describes the expected manner in which multiple **threads** may use the
2 **allocator**. The values and their descriptions are:

- 3 • **contended**: high contention is expected on the **allocator**; that is, many **tasks** are expected
4 to request allocations simultaneously;
- 5 • **uncontended**: low contention is expected on the **allocator**; that is, few **task** are expected to
6 request allocations simultaneously;
- 7 • **serialized**: one **task** at a time will request allocations with the **allocator**. Requesting two
8 allocations simultaneously when specifying **serialized** results in **unspecified behavior**;
9 and
- 10 • **private**: the same **thread** will execute **all tasks** that request allocations with the **allocator**.
11 Requesting an allocation from **tasks** that different **threads** execute, simultaneously or not,
12 when specifying **private** results in **unspecified behavior**.

13 Allocated **memory** will be byte aligned to at least the value specified for the **alignment trait** of
14 the **allocator**. Some **directives** and API routines can specify additional requirements on alignment
15 beyond those described in this section.

16 The **access trait** defines the *access group* of **tasks** that may access **memory** that is allocated by a
17 **memory allocator**. If the value is **all**, the access group consists of **all tasks** that execute on all
18 available **devices**. If the value is **memspace**, the access group consists of **all tasks** that execute on
19 all **devices** that are associated with the **allocator**. if the value is **device**, the access group consists
20 of **all tasks** that execute on the **device** where the allocation was requested. If the value is **cgroup**,
21 the access group consists of **all tasks** in the same **contention group** as the **task** that requested the
22 allocation. If the value is **pteam**, the access group consists of all **current team tasks** of the
23 innermost enclosing parallel **region** in which the allocation was requested. If the value is **thread**,
24 the access group consists of **all tasks** that are executed by the same **thread** that executed the
25 allocation request. **Memory** returned by the **allocator** will be **memory** accessible by **all tasks** in the
26 same access group as the **task** that requested the allocation. Attempts to access this **memory** from a
27 **task** that is not in same access group results in **unspecified behavior**.

28 The total amount of storage in bytes that an **allocator** can use for allocation requests from **tasks** in
29 the same access group is limited by the **pool_size trait**. Requests that would result in using more
30 storage than **pool_size** will not be fulfilled by the **allocator**.

31 The **fallback trait** specifies how the **memory allocator** behaves when it cannot fulfill an
32 allocation request. If the **fallback trait** is set to **null_fb**, the **allocator** returns the value zero if
33 it fails to allocate the **memory**. If the **fallback trait** is set to **abort_fb**, the behavior is as if an
34 **error directive** for which *sev-level* is **fatal** and *action-time* is **execution** is encountered if
35 the allocation fails. If the **fallback trait** is set to **allocator_fb** then when an allocation fails
36 the request will be delegated to the **allocator** specified in the **fb_data trait**. If the **fallback trait**
37 is set to **default_mem_fb** then when an allocation fails another allocation will be tried in
38 **omp_default_mem_space**, which assumes all **allocator traits** to be set to their default values
39 except for **fallback trait**, which will be set to **null_fb**. The default value for the **fallback**

1 **trait** is **null_fb** for any **allocator** that is associated with a **target memory space**. Otherwise, the
2 default value is **default_mem_fb**.

3 All **memory** that is allocated with an **allocator** for which the **pinned trait** is specified as **true**
4 must remain in the same storage resource at the same location for its entire lifetime. If
5 **pin_device** is also specified then the allocation must be allocated in that **device**.

6 The **partition trait** describes the partitioning of allocated **memory** over the storage resources
7 represented by the **memory space** associated with the **allocator**. The partitioning will be done in
8 parts with a minimum size that is **implementation defined**. The values are:

- 9 • **environment**: the placement of allocated **memory** is determined by the execution
10 environment;
- 11 • **nearest**: allocated **memory** is placed in the storage resource that is nearest to the **thread**
12 that requests the allocation;
- 13 • **blocked**: allocated **memory** is partitioned into parts of approximately the same size with at
14 most one part per storage resource; and
- 15 • **interleaved**: allocated **memory** parts are distributed in a round-robin fashion across the
16 storage resources such that the size of each part is the value of the **part_size trait** except
17 possibly the last part, which can be smaller.

18 The **part_size trait** specifies the size of the parts allocated over the storage resources for some
19 of the partition **trait** policies. The actual value of the **trait** might be rounded up to an
20 **implementation defined** value to comply with hardware restrictions of the storage resources.

21 If the **preferred_device trait** is specified then storage resources of the specified **device** are
22 preferred to fulfill the allocation.

23 If the value of the **target_access trait** is **single** then data from this **allocator** cannot be
24 accessed on two different **devices** unless, for any given **host device** access, the entry and exit of the
25 target **region** in which any accesses occur either both precede or both follow the **host device** access
26 in **happens-before order**. Additionally, for any two **target regions** that may access data from this
27 **allocator** and execute on distinct **devices**, the entry and exit of one of the **regions** must precede those
28 of the other in **happens-before order**. If the value of the **target_access trait** is **multiple** then
29 accesses of data from this **allocator** from different **devices** may be arbitrarily interleaved, provided
30 that synchronization ensures data races do not occur.

31 If the value of the **atomic_scope trait** is **all** then all storage locations of data from this
32 **allocator** have an **atomic scope** that consists of all **threads** on the devices associated with the
33 **allocator**. If the value is **device** then all storage locations have an **atomic scope** that consists of all
34 **threads** on the **device** on which the **atomic operation** is performed.

35 Table 7.3 shows the list of predefined **memory allocators** and their **associated memory spaces**. The
36 predefined **memory allocators** have default values for their **allocator traits** unless otherwise
37 specified.

TABLE 7.3: 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

If any operation of the [base language](#) causes a reallocation of a [variable](#) that is allocated with a [memory allocator](#) then that [memory allocator](#) will be used to deallocate the current [memory](#) and to allocate the new [memory](#). For any allocatable subcomponents, the [allocator](#) that is used for the 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](#).

7.3 align Clause

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

Arguments

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

Modifiers

Name	Modifies	Type	Properties
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: directive-name	unique

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 7.5](#)
- Memory Allocators, see [Section 7.2](#)

7.4 allocator Clause

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

Arguments

Name	Type	Properties
<i>allocator</i>	expression of <code>allocator_handle</code> type	<i>default</i>

Modifiers

Name	Modifies	Type	Properties
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: <code>directive-name</code>	<code>unique</code>

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 7.5](#)
- Memory Allocators, see [Section 7.2](#)
- `def-allocator-var` ICV, see [Table 2.1](#)

7.5 `allocate` Directive

Name: <code>allocate</code> Category: declarative	Association: none Properties: pure
--	---

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

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 7.3](#)
- **allocator** clause, see [Section 7.4](#)
- Memory Allocators, see [Section 7.2](#)

7.6 allocate Clause

Name: allocate	Properties: all-privatizing
------------------------------	------------------------------------

Arguments

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

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: <i>allocator</i> expression of <code>allocator_handle</code> type (<i>default</i>)	unique
<i>align-modifier</i>	<i>list</i>	Complex, name: align Arguments: <i>alignment</i> expression of integer type (constant, positive)	unique
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: directive-name	unique

Directives

allocators, **distribute**, **do**, **for**, **parallel**, **scope**, **sections**, **single**, **target**, **task**, **taskgroup**, **taskloop**, **teams**

Semantics

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*

1 argument. The *allocator-complex-modifier* and *align-modifier* have the same syntax and semantics
2 for the **allocate** clause as the **allocator** and **align** clauses have for the **allocate**
3 **directive**.

4 For allocations that arise from this **clause**, the **null_fb** value of the fallback **allocator trait**
5 behaves as if the **abort_fb** had been specified.

6 **Restrictions**

7 Restrictions to the **allocate clause** are as follows:

- 8 • For any list item that is specified in the **allocate clause** on a **directive** other than the
9 **allocators directive**, a data-sharing attribute **clause** that may create a private copy of that
10 list item must be specified on the same **directive**.
- 11 • For **task**, **taskloop** or **target** directives, allocation requests to **memory allocators** with
12 the **access trait** set to **thread** result in **unspecified behavior**.
- 13 • **allocate clauses** that appear on a **target** construct or on **constructs** in a **target region**
14 must specify an *allocator-simple-modifier* or *allocator-complex-modifier* unless a
15 **requires directive** with the **dynamic_allocators clause** is present in the same
16 **compilation unit**.

17 **Cross References**

- 18 • **align** clause, see [Section 7.3](#)
- 19 • **allocator** clause, see [Section 7.4](#)
- 20 • **allocators** directive, see [Section 7.7](#)
- 21 • **distribute** directive, see [Section 12.7](#)
- 22 • **do** directive, see [Section 12.6.2](#)
- 23 • **for** directive, see [Section 12.6.1](#)
- 24 • **parallel** directive, see [Section 11.2](#)
- 25 • **scope** directive, see [Section 12.2](#)
- 26 • **sections** directive, see [Section 12.3](#)
- 27 • **single** directive, see [Section 12.1](#)
- 28 • **target** directive, see [Section 14.8](#)
- 29 • **task** directive, see [Section 13.6](#)
- 30 • **taskgroup** directive, see [Section 16.4](#)
- 31 • **taskloop** directive, see [Section 13.7](#)
- 32 • **teams** directive, see [Section 11.3](#)
- 33 • **Memory Allocators**, see [Section 7.2](#)

7.7 allocators Construct

Name: <code>allocators</code>	Association: block (allocator structured block)
Category: <code>executable</code>	Properties: <code>default</code>

Clauses

`allocate`

Semantics

The `allocators` construct specifies that `memory allocators` are used for certain `variables` that are allocated by the associated `allocate-stmt`. The list items that appear in an `allocate` clause may include structure elements. If a `variable` that is to be allocated appears as a list item in an `allocate` clause on the directive, an `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 7.6](#)
- Memory Allocators, see [Section 7.2](#)
- OpenMP Allocator Structured Blocks, see [Section 5.3.1](#)

7.8 uses_allocators Clause

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

Arguments

Name	Type	Properties
<code>allocator</code>	expression of <code>allocator_handle</code> type	<code>default</code>

Modifiers

Name	Modifies	Type	Properties
<i>mem-space</i>	<i>allocator</i>	Complex, name: memspace Arguments: memspace-handle expression of memspace_handle type (<i>default</i>)	default
<i>traits-array</i>	<i>allocator</i>	Complex, name: traits Arguments: traits variable of alloctrait array type (<i>default</i>)	default
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: directive-name	unique

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 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 14.8](#)
- Memory Allocators, see [Section 7.2](#)
- Memory Spaces, see [Section 7.1](#)
- `omp_destroy_allocator`, see [Section 19.13.5](#)
- `omp_init_allocator`, see [Section 19.13.3](#)

8 Variant Directives

This chapter defines [directives](#) and related concepts to support the seamless adaption of [OpenMP programs](#) to [OpenMP contexts](#).

8.1 OpenMP Contexts

At any point in an [OpenMP program](#), an [OpenMP context](#) exists that defines [traits](#) that describe the active [constructs](#), the execution [devices](#), functionality supported by the implementation and available dynamic values. The [traits](#) are grouped into [trait sets](#). The defined [trait sets](#) are: the [construct trait set](#); the [device trait set](#); the [target device trait set](#); the [implementation trait set](#); and the [dynamic trait set](#). [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 8.2](#).

The [construct trait set](#) is composed of the [directive](#) names, each being a [trait](#), of all enclosing [constructs](#) at that point in the [OpenMP program](#) up to a [target construct](#). [Combined constructs](#) and [composite constructs](#) are added to the set as distinct [constructs](#) in the same nesting order specified by the original [constructs](#). The [dispatch construct](#) is added to the [construct trait set](#) only for the [target-call](#) of the associated function dispatch [structured block](#). The [construct trait 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 [OpenMP 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 [construct trait set](#) as c_1 for any generated [SIMD](#) versions so the total size of the [trait set](#) is increased by one.
2. For [procedures](#) that are determined to be [function variants](#) by a [declare variant directive](#), the [trait selectors](#) c_1, \dots, c_M of the [construct selector set](#) are added in the same order to the beginning of the [construct trait set](#) as c_1, \dots, c_M so the total size of the [trait set](#) is increased by M .
3. For [procedures](#) that are determined to be [target variants](#) by a [declare target directive](#), the [target trait](#) is added to the beginning of the [construct trait set](#) as c_1 so the total size of the [trait set](#) is increased by one.

1 The *simd* trait is a [clause-list trait](#) that is defined with [properties](#) that match the [clauses](#) that can be
2 specified on the [declare simd directive](#) with the same names and semantics. The *simd* trait
3 defines at least the *simdlen* [property](#) and one of the *inbranch* or *notinbranch* [properties](#). Traits in the
4 [construct trait set](#) other than *simd* are [non-property traits](#).

5 The [device trait set](#) includes [traits](#) that define the characteristics of the [device](#) being targeted by the
6 compiler at that point in the [OpenMP program](#). For each [target device](#) that the implementation
7 supports, a [target device trait set](#) exists that defines the characteristics of that [device](#). At least the
8 following [traits](#) must be defined for the [device trait set](#) and all [target device trait sets](#):

- 9 • The *kind(kind-list)* [name-list trait](#) specifies the general kind of the [device](#). Each member of
10 *kind-list* is a *kind-name*, for which the following values are defined:
 - 11 – *host*, which specifies that the [device](#) is the [host device](#);
 - 12 – *nohost*, which specifies that the [device](#) is not the [host device](#); and
 - 13 – the values defined in the [OpenMP Additional Definitions document](#).
- 14 • The *isa(isa-list)* [name-list trait](#) specifies the Instruction Set Architectures supported by the
15 [device](#). Each member of *isa-list* is an *isa-name*, for which the accepted values are
16 [implementation defined](#).
- 17 • The *arch(arch-list)* [name-list trait](#) specifies the architectures supported by the [device](#). Each
18 member of *arch-list* is an *arch-name*, for which the accepted values are [implementation](#)
19 [defined](#).

20 The [target device trait set](#) also defines the following [trait](#):

- 21 • The *device_num* [trait](#) specifies the *device number* of the [device](#).

22 The [implementation trait set](#) includes [traits](#) that describe the functionality supported by the OpenMP
23 implementation at that point in the [OpenMP program](#). At least the following [traits](#) can be defined:

- 24 • The *vendor(vendor-list)* [name-list trait](#), which specifies the vendor identifiers of the
25 implementation. Each member of *vendor-list* is a *vendor-name*, for which the defined values
26 are in the [OpenMP Additional Definitions document](#).
- 27 • The *extension(extension-list)* [name-list trait](#), which specifies vendor-specific extensions to the
28 OpenMP specification. Each member of *extension-list* is an *extension-name*, for which the
29 accepted values are [implementation defined](#).
- 30 • A *requires(requires-list)* [clause-list trait](#), for which the [properties](#) are the [clauses](#) that have
31 been supplied to the [requires directive](#) prior to the program point as well as
32 [implementation defined](#) implicit requirements.

33 Implementations can define additional [traits](#) in the [device trait set](#), [target device trait set](#) and
34 [implementation trait set](#); these [traits](#) are [extension traits](#).

35 The [dynamic trait set](#) includes [traits](#) that define the dynamic [properties](#) of an [OpenMP program](#) at a
36 point in its execution. The *data state* [trait](#) in the [dynamic trait set](#) refers to the complete data state of
37 the [OpenMP program](#) that may be accessed at runtime.

8.2 Context Selectors

Context selectors are used to define the **properties** that can match an **OpenMP context**. OpenMP defines different **trait selector sets**, each of which contains different **trait 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 **clause**.

5 The **construct selector set** defines the **traits** in the **construct trait set** that should be active in the
6 **OpenMP context**. Each **trait selector** that can be defined in the **construct selector set** is the
7 *directive-name* of a **context-matching construct**. Each *trait-property* of the **simd trait selector** is a
8 *trait-property-clause*. The syntax is the same as for a valid **clause** of the **declare simd** directive
9 and the restrictions on the **clauses** from that **directive** apply. The **construct selector set** is an
10 ordered list c_1, \dots, c_N .

11 The **device selector set** and **implementation selector set** define the **traits** that should be
12 active in the corresponding **trait set** of the **OpenMP context**. The **target_device selector set**
13 defines the **traits** that should be active in the **target device trait set** for the **device** that the specified
14 **device_num trait selector** identifies. The same **traits** that are defined in the corresponding **trait**
15 **sets** can be used as **trait selectors** with the same **properties**. The **kind trait selector** of the **device**
16 **selector set** and **target_device selector set** can also specify the value **any**, which is as if no
17 **kind trait selector** was specified. If a **device_num trait selector** does not appear in the
18 **target_device selector set** then a **device_num trait selector** that specifies the value of the
19 *default-device-var* **ICV** is implied. For the **device_num trait selector** of the **target_device**
20 **selector set**, a single *trait-property-expression* must be specified. For the
21 **atomic_default_mem_order trait selector** of the **implementation selector set**, a single
22 *trait-property* must be specified as an identifier equal to one of the valid arguments to the
23 **atomic_default_mem_order** clause on the **requires** directive. For the **requires trait**
24 **selector** of the **implementation selector set**, each *trait-property* is a *trait-property-clause*. The
25 syntax is the same as for a valid **clause** of the **requires** directive and the restrictions on the
26 **clauses** from that **directive** apply.

27 The **user selector set** defines the **condition trait selector** that provides additional user-defined
28 conditions. The **condition trait selector** contains a single *trait-property-expression* that must
29 evaluate to *true* for the **trait selector** to be true. Any non-constant *trait-property-expression* that is
30 evaluated to determine the suitability of a variant is evaluated according to the *data state trait* in the
31 **dynamic trait set** of the **OpenMP context**. The **user selector set** is dynamic if the **condition**
32 **trait selector** is present and the expression in the **condition trait selector** is not a constant
33 expression; otherwise, it is static.

34 All parts of a **context selector** define the static part of the **context selector** except the following
35 parts, which define the dynamic part of the **context selector**:

- 36 • Its **user selector set** if it is dynamic; and
- 37 • Its **target_device selector set**.

38 For the **match** clause of a **declare variant** directive, any argument of the **base function** that
39 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 [context](#)
4 [selector](#) 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 [trait selectors](#) to be specified. Each specified *trait-property* for
9 these [implementation defined trait selectors](#) should be a *trait-property-extension*. Implementations
10 can ignore specified [trait selectors](#) that are not those described in this section.

11 Restrictions

12 Restrictions to [context selectors](#) are as follows:

- 13 ● Each *trait-property* may only be specified once in a [trait selector](#) other than those in the
14 [construct selector set](#).
- 15 ● Each *trait-set-selector-name* may only be specified once.
- 16 ● Each *trait-selector-name* may only be specified once.
- 17 ● A *trait-score* cannot be specified in [traits](#) from the [construct selector set](#), the [device](#)
18 [selector set](#) or the [target_device 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 returned by [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 [context selector](#) appears unless
24 the [directive](#) is a [declare variant directive](#) and the [variable](#) is an argument of the
25 associated [base function](#).
- 26 ● If *trait-property any* is specified in the **kind** *trait-selector* of the [device selector set](#) or
27 the [target_device selector sets](#), no other *trait-property* may be specified in the same
28 [selector set](#).
- 29 ● For a *trait-selector* that corresponds to a [name-list trait](#), at least one *trait-property* must be
30 specified.
- 31 ● For a *trait-selector* that corresponds to a [non-property trait](#), no *trait-property* may be
32 specified.
- 33 ● For the [requires trait selector](#) of the [implementation selector set](#), at least one
34 *trait-property* must be specified.

8.3 Matching and Scoring Context Selectors

A **context selector** is compatible with an **OpenMP context** if the following conditions are satisfied:

- All **trait selectors** in its **user selector set** are true;
- All **traits** and **trait properties** that are defined by **trait selectors** in the **target_device selector set** are active in the **target device trait set** for the **device** that is identified by the **device_num trait selector**;
- All **traits** and **trait properties** that are defined by **trait selectors** in its **construct selector set**, its **device selector set** and its **implementation selector set** are active in the corresponding **trait sets** of the **OpenMP context**;
- For each **trait selector** in the **context selector**, its **properties** are a subset of the **properties** of the corresponding **trait** of the **OpenMP context**;
- **Trait selectors** in its **construct selector set** appear in the same relative order as their corresponding **traits** in the **construct trait set** of the **OpenMP context**; and
- No specified **implementation defined trait selector** is ignored by the implementation.

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

- The **simdlen (N)** **property** of the **trait 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 **trait 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 **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 **trait selectors** in the same order are used;
2. The **kind**, **arch**, and **isa trait 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 trait 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 **trait selectors** allowed by the implementation are **implementation defined**;
5. Other **trait selectors** are given a value of zero; and

- 1 6. A **context selector** that is a strict subset of another **context selector** has a score of zero. For
2 other **context selectors**, the final score is the sum of the values of all specified **trait selectors**
3 plus 1.

4 **8.4 Metadirectives**

5 A **metadirective** is a **directive** that can specify multiple **directive variants** of which one may be
6 conditionally selected to replace the **metadirective** based on the **enclosing context**. A **metadirective**
7 is replaced by a **nothing directive** or one of the **directive variants** specified by the **when clauses**
8 or the **otherwise clause**. If no **otherwise clause** is specified the effect is as if one was
9 specified without an associated **directive variant**.

10 The **OpenMP context** for a given **metadirective** is defined according to **Section 8.1**. The order of
11 **clauses** that appear on a **metadirective** is significant and, if specified, **otherwise** must be the last
12 **clause** specified on a **metadirective**.

13 **Replacement candidates** for a **metadirective** are ordered according to the following rules in
14 decreasing precedence:

- 15 • A **candidate** is before another one if the score associated with the **context selector** of the
16 corresponding **when clause** is higher.
- 17 • A **candidate** that was explicitly specified is before one that was implicitly specified.
- 18 • **Candidates** are ordered according to the order in which they lexically appear on the
19 **metadirective**.

20 The list of **dynamic replacement candidates** is the prefix of the sorted list of **replacement candidates**
21 up to and including the first **candidate** for which the corresponding **when** or **otherwise clause**
22 has a **static context selector**. The first **dynamic replacement candidate** for which the corresponding
23 **when** or **otherwise clause** has a **compatible context selector**, according to the matching rules
24 defined in **Section 8.3**, replaces the **metadirective**.

25 **Restrictions**

26 Restrictions to **metadirectives** are as follows:

- 27 • Replacement of the **metadirective** with the **directive variant** associated with any of the
28 **dynamic replacement candidates** must result in a **conforming program**.
- 29 • Insertion of user code at the location of a **metadirective** must be allowed if the first **dynamic**
30 **replacement candidate** does not have a **static context selector**.
- 31 • If the list of **dynamic replacement candidates** has multiple items then all items must be
32 **executable directives**.

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.
- A **metadirective** is pure if and only if all **directive variants** specified for it are pure.

Fortran

8.4.1 when Clause

Name: when	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
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: directive-name	unique

Directives

begin metadirective, metadirective

Semantics

The specified *directive-variant* is a **replacement candidate** for the **metadirective** on which the **clause** is specified if the static part of the **context selector** specified by *context-selector* is compatible with the **OpenMP context** according to the matching rules defined in [Section 8.3](#). 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 **OpenMP 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** trait selector.

C / C++

- *directive-variant* must not specify a **begin declare variant** directive.

C / C++

Cross References

- **begin metadirective** directive, see [Section 8.4.4](#)
- **metadirective** directive, see [Section 8.4.3](#)
- **nothing** directive, see [Section 9.7](#)
- Context Selectors, see [Section 8.2](#)

8.4.2 otherwise Clause

Name: otherwise	Properties: unique, ultimate
-------------------------------	-------------------------------------

Arguments

Name	Type	Properties
<i>directive-variant</i>	directive-specification	optional, unique

Modifiers

Name	Modifies	Type	Properties
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: directive-name	unique

Directives

begin metadirective, **metadirective**

Semantics

The **otherwise clause** is treated as a **when clause** with the specified **directive variant**, if any, and a **static context selector** that is always compatible and has a score lower than the scores associated with any other **directive variant**.

Restrictions

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

- *directive-variant* must not specify a **metadirective**.

C / C++

- *directive-variant* must not specify a **begin declare variant** directive.

C / C++

Cross References

- **when** clause, see [Section 8.4.1](#)
- **begin metadirective** directive, see [Section 8.4.4](#)
- **metadirective** directive, see [Section 8.4.3](#)

8.4.3 metadirective

Name: metadirective Category: meta	Association: none Properties: pure
---	---

Clauses

[otherwise](#), [when](#)

Semantics

The [metadirective](#) specifies [metadirective](#) semantics.

Cross References

- **otherwise** clause, see [Section 8.4.2](#)
- **when** clause, see [Section 8.4.1](#)
- Metadirectives, see [Section 8.4](#)

8.4.4 begin metadirective

Name: begin metadirective 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 a [directive](#) that has a paired [end directive](#) or must be the [nothing directive](#).

Cross References

- **otherwise** clause, see [Section 8.4.2](#)
- **when** clause, see [Section 8.4.1](#)
- **nothing** directive, see [Section 9.7](#)
- Metadirectives, see [Section 8.4](#)

8.5 Declare Variant Directives

[Declare variant directives](#) declare [base functions](#) to have the specified [function variant](#). The [context selector](#) specified by *context-selector* in the **match** clause is associated with the [function variant](#).

The [OpenMP context](#) for a direct call to a given [base function](#) is defined according to [Section 8.1](#). 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 8.3](#) then the [function 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 **match** clause has a [static context selector](#).

The first [dynamic replacement candidate](#) for which the corresponding **match** clause has a [compatible context selector](#) 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 [variant substitution](#) occurs is [implementation defined](#).

▲ C++ ▼

For indirect function calls that can be determined to call a particular [base function](#), whether [variant substitution](#) occurs is unspecified.

Any differences that the specific [OpenMP context](#) requires in the prototype of the [function 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 [procedures](#) 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 [procedure](#) 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 [procedure](#) 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 may only be specified if the [dispatch trait 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](#) may not be specified for virtual, defaulted or deleted functions.
- [Declare variant directives](#) may not be specified for constructors or destructors.
- [Declare variant directives](#) may not be specified for immediate functions.
- The [procedure](#) 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 8.5.5](#)
- [declare variant](#) directive, see [Section 8.5.4](#)
- [Context Selectors](#), see [Section 8.2](#)
- [OpenMP Contexts](#), see [Section 8.1](#)

8.5.1 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>

Modifiers

Name	Modifies	Type	Properties
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: directive-name	unique

Directives

begin declare variant, **declare variant**

Semantics

The *context-selector* argument of the **match** clause specifies the **context selector** to use to determine if a specified **function variant** is a **replacement candidate** for the specified **base function** in a given **OpenMP 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 each call site to the **base function** according to the **base language** rules.

Cross References

- **begin declare variant** directive, see [Section 8.5.5](#)
- **declare variant** directive, see [Section 8.5.4](#)
- Context Selectors, see [Section 8.2](#)

8.5.2 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
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: directive-name	unique

Directives

declare variant

Semantics

The `adjust_args` clause specifies how to adjust the arguments of the `base function` when a specified `function variant` is selected for replacement. For each `adjust_args` clause that is present on the selected `function variant`, the adjustment operation specified by the `adjust-op modifier` is applied to each argument specified in the `clause` before being passed to the selected `function variant`. If the `adjust-op modifier` is `nothing`, the argument is passed to the selected `function 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 they are not already `device pointers`. If the `current task` has the `is_device_ptr property` for a given argument in its `interoperability requirement set`, 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 the `clause` with a `need_device_ptr adjust-op` must be of type `C_PTR` in the dummy argument declaration of the `function variant`.

Fortran

Cross References

- `declare variant` directive, see [Section 8.5.4](#)

8.5.3 append_args Clause

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

Arguments

Name	Type	Properties
<code>append-op-list</code>	list of OpenMP operation list item type	<code>default</code>

Modifiers

Name	Modifies	Type	Properties
<code>directive-name-modifier</code>	<code>all arguments</code>	Keyword: <code>directive-name</code>	unique

Directives

`declare variant`

Semantics

The `append_args` clause specifies additional arguments to pass in the call when a specified `function variant` 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 `function variant` 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 *interop-type*, the behavior is as if the corresponding `clauses` would also be part of the `interop` construct and those `properties` are removed from the `interoperability requirement set`.

This argument is destroyed after the call to the selected `function variant` returns, as if an `interop` construct with a `destroy` clause was used with the same `clauses` that were used to initialize the argument.

Cross References

- `init` clause, see [Section 15.1.2](#)
- `declare variant` directive, see [Section 8.5.4](#)
- `interop` directive, see [Section 15.1](#)
- Interoperability Requirement Set, see [Section 15.2](#)
- OpenMP Operations, see [Section 4.2.3](#)

8.5.4 declare variant Directive

Name: <code>declare variant</code> Category: declarative	Association: declaration Properties: pure
---	--

Arguments

`declare variant` (*[base-name:]variant-name*)

Name	Type	Properties
<i>base-name</i>	identifier of function type	optional
<i>variant-name</i>	identifier of function type	<i>default</i>

Clauses

[adjust_args](#), [append_args](#), [match](#)

Semantics

The [declare variant directive](#) specifies declare variant semantics for a single [replacement candidate](#). *variant-name* identifies the [function variant](#) while *base-name* identifies the [base function](#).

▼ **C** ▼
Any expressions in the [match clause](#) are interpreted as if they appeared in the scope of arguments of the [base function](#).

▲ **C** ▲
▼ **C++** ▼
variant-name and any expressions in the [match clause](#) are interpreted as if they appeared at the scope of the trailing return type of the [base function](#).

The [function variant](#) is determined by [base language](#) standard name lookup rules ([basic.lookup]) of *variant-name* using the argument types at the call site after [implementation defined](#) changes have 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 8.5.2](#)
- **append_args** clause, see [Section 8.5.3](#)
- **match** clause, see [Section 8.5.1](#)
- Declare Variant Directives, see [Section 8.5](#)

C / C++

8.5.5 begin declare variant Directive

Name: <code>begin declare variant</code>	Association: delimited (declaration-definition-seq)
Category: declarative	Properties: <i>default</i>

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.

1 If a **declare variant directive** appears between a **begin declare variant directive** and its
2 paired **end directive**, the effective **context selectors** of the outer **directive** are appended to the
3 **context selector** of the inner **directive** to form the effective **context selector** of the inner **directive**. If
4 a *trait-set-selector* is present on both **directives**, the *trait-selector* list of the outer **directive** is
5 appended to the *trait-selector* list of the inner **directive** after equivalent *trait-selectors* have been
6 removed from the outer list. Restrictions that apply to explicitly specified **context selectors** also
7 apply to effective **context selectors** constructed through this process.

8 The symbol name of a function definition that appears between a **begin declare variant**
9 **directive** and its paired **end directive** is determined through the **base language** rules after the name
10 of the function has been augmented with a string that is determined according to the effective
11 **context selector** of the **begin declare variant directive**. The symbol names of two
12 definitions of a function are considered to be equal if and only if their effective **context selectors** are
13 equivalent.

14 If the **context selector** of a **begin declare variant directive** contains **traits** in the *device* or
15 *implementation* set that are known never to be compatible with an **OpenMP context** during the
16 current compilation, the **preprocessed code** that follows the **begin declare variant**
17 **directive** up to its paired **end directive** is elided.

18 Any expressions in the **match clause** are interpreted at the location of the **directive**.

19 Restrictions

20 The restrictions to **begin declare variant directive** are as follows:

- 21 • **match clause** must not contain a **simd trait selector**.
- 22 • Two **begin declare variant directives** and their paired **end directives** must either
23 encompass disjoint source ranges or be perfectly nested.

▼ C++ ▼

- 24 • A **match clause** must not contain a **dynamic context selector** that references the **this**
25 pointer.
- 26 • If an expression in the **context selector** that appears in **match clause** references the **this**
27 pointer, the **base function** must be a non-static member function.

▲ C++ ▲

28 Cross References

- 29 • **match clause**, see [Section 8.5.1](#)
- 30 • **Declare Variant Directives**, see [Section 8.5](#)

▲ C / C++ ▲

8.6 dispatch Construct

Name: <code>dispatch</code>	Association: block (function dispatch structured block)
Category: <code>executable</code>	Properties: <code>context-matching</code>

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 `function variant` if `variant substitution` occurs.

Properties added to the `interoperability requirement set` can be removed by the effect of other `directives` (see [Section 15.2](#)) 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 `function variant` 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 16.9.5](#)
- **device** clause, see [Section 14.2](#)
- **interop** clause, see [Section 8.6.1](#)
- **is_device_ptr** clause, see [Section 6.4.7](#)
- **nocontext** clause, see [Section 8.6.3](#)
- **novariants** clause, see [Section 8.6.2](#)
- **nowait** clause, see [Section 16.6](#)
- Interoperability Requirement Set, see [Section 15.2](#)
- OpenMP Function Dispatch Structured Blocks, see [Section 5.3.2](#)

8.6.1 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>

Modifiers

Name	Modifies	Type	Properties
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: directive-name	unique

Directives

[dispatch](#)

Semantics

The [interop clause](#) specifies additional arguments to pass to the [function variant](#) 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 8.6](#)

8.6.2 novariants Clause

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

Arguments

Name	Type	Properties
<i>do-not-use-variant</i>	expression of OpenMP logical type	<i>default</i>

Modifiers

Name	Modifies	Type	Properties
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: directive-name	unique

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 8.6](#)

8.6.3 nocontext Clause

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

Arguments

Name	Type	Properties
<i>do-not-update-context</i>	expression of OpenMP logical type	<i>default</i>

Modifiers

Name	Modifies	Type	Properties
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: directive-name	unique

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 trait 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 8.6](#)

8.7 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 **procedure** declaration or definition enables the creation of corresponding **SIMD** versions of the associated **procedure** 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 a *step-complex-modifier* as *linear-step*.

▼ C / C++ —————

The expressions that appear in the **clauses** of each **directive** are evaluated in the scope of the arguments of the **procedure** declaration or definition.

▲ C / C++ —————

▼ C++ —————

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

▲ C++ —————

Restrictions

Restrictions to the **declare simd** directive are as follows:

- The **procedure** body must be a **structured block**.
- The execution of the **procedure**, when called from a **SIMD loop**, may not result in the execution of any **constructs** except for **atomic constructs** and **ordered constructs** on which the **simd clause** is specified.
- The execution of the **procedure** may not have any side effects that would alter its execution for concurrent iterations of a **SIMD chunk**.

C / C++

- If the **procedure** 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 **procedure** may not contain calls to the **longjmp** or **setjmp** functions.

C / C++

C++

- The **procedure** may not 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 an explicit interface and a **declare simd directive** is also specified for the definition of the **procedure** then the two **declare simd directives** must specify equivalent **clauses**.
- **Procedures** pointers may not be used to access versions created by the **declare simd directive**.

Fortran

Cross References

- **aligned** clause, see [Section 6.11](#)
- **linear** clause, see [Section 6.4.6](#)
- **reduction** clause, see [Section 6.5.9](#)
- **simdlen** clause, see [Section 11.5.3](#)
- **uniform** clause, see [Section 6.10](#)

8.7.1 *branch* Clauses

Clause groups

Properties: unique, exclusive	Members: Clauses inbranch , notinbranch
--------------------------------------	---

Directives

[declare simd](#)

Semantics

The *branch clause group* defines a set of [clauses](#) that indicate if a [procedure](#) can be assumed to be or not to be encountered in a branch. If neither [clause](#) is specified, then the [procedure](#) may or may not be called from inside a conditional statement of the calling context.

Cross References

- **declare simd** directive, see [Section 8.7](#)

8.7.1.1 *inbranch* Clause

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

Arguments

Name	Type	Properties
<code>inbranch</code>	expression of OpenMP logical type	constant, optional

Modifiers

Name	Modifies	Type	Properties
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: directive-name	unique

Directives

[declare simd](#)

Semantics

If *inbranch* evaluates to true, the **inbranch** clause specifies that the **procedure** will always be called from inside a conditional statement of the calling context. If *inbranch* evaluates to false, the **procedure** may be called other than from inside a conditional statement. If *inbranch* is not specified, the effect is as if *inbranch* evaluates to true.

Cross References

- **declare simd** directive, see [Section 8.7](#)

8.7.1.2 notinbranch Clause

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

Arguments

Name	Type	Properties
<i>notinbranch</i>	expression of OpenMP logical type	constant, optional

Modifiers

Name	Modifies	Type	Properties
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: directive-name	unique

Directives

declare simd

Semantics

If *notinbranch* evaluates to true, the **notinbranch** clause specifies that the **procedure** will never be called from inside a conditional statement of the calling context. If *notinbranch* evaluates to false, the **procedure** may be called from inside a conditional statement. If *notinbranch* is not specified, the effect is as if *notinbranch* evaluates to true.

Cross References

- **declare simd** directive, see [Section 8.7](#)

8.8 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 it generates different versions, and whether it calls a different version in a **target** region from the version that it calls outside a **target** region, are **implementation defined**.

1 To facilitate `device` usage, OpenMP defines rules that implicitly specify `declare target directives` for
2 `procedures` and `variables`. The remainder of this section defines those rules as well as restrictions
3 that apply to all `declare target directives`.

C++

4 If a `variable` with `static storage duration` has the `constexpr` specifier and is not a `groupprivate`
5 `variable` then the `variable` is treated as if it had appeared as a `list item` in an `enter` clause on a
6 `declare target directive`.

C++

7 If a `variable` with `static storage duration` that is not a `device local variable` (including not a
8 `groupprivate variable`) is declared in a `device procedure` then the `variable` is treated as if it had
9 appeared as a `list item` in an `enter` clause on a `declare target directive`.

10 If a `procedure` is referenced outside of any `reverse-offload region` in a `procedure` that appears as a
11 `list item` in an `enter` clause on a `non-host declare target directive` then the name of the referenced
12 `procedure` is treated as if it had appeared in an `enter` clause on a `declare target directive`.

C / C++

13 If a `variable` with `static storage duration` or a function (except `lambda` for C++) is referenced in the
14 initializer expression list of a `variable` with `static storage duration` that appears as a `list item` in an
15 `enter` or `local` clause on a `declare target directive` then the name of the referenced `variable` or
16 `procedure` is treated as if it had appeared in an `enter` clause on a `declare target directive`.

C / C++

Fortran

17 If a `declare target directive` has a `device_type` clause then any enclosed internal
18 `procedure` cannot contain any `declare target directives`. The enclosing `device_type`
19 `clause` implicitly applies to internal `procedures`.

Fortran

20 A reference to a `device local variable` that has `static storage duration` inside a `device procedure` is
21 replaced with a reference to the copy of the `variable` for the `device`. Otherwise, a reference to a
22 `variable` that has `static storage duration` in a `device procedure` is replaced with a reference to a
23 corresponding `variable` in the `device data environment`. If the corresponding `variable` does not exist
24 or the `variable` does not appear in an `enter` or `link` clause on a `declare target directive`, the
25 behavior is unspecified.

Execution Model Events

26 The *target-global-data-op* event occurs when an `original list item` is associated with a
27 `corresponding list item` on a `device` as a result of a `declare target directive`; the event occurs before
28 the first access to the `corresponding list item`.
29

1 Tool Callbacks

2 A [thread](#) dispatches a registered `ompt_callback_target_data_op_callback`, or a registered
3 `ompt_callback_target_data_op_emi_callback` with `ompt_scope_beginend` as its
4 endpoint argument for each occurrence of a *target-global-data-op event* in that [thread](#). These
5 [callbacks](#) have type signature `ompt_callback_target_data_op_t` or
6 `ompt_callback_target_data_op_emi_t`, respectively.

7 Restrictions

8 Restrictions to any [declare target directive](#) are as follows:

- 9 • The same [list item](#) must not explicitly appear in both an [enter clause](#) on one [declare target](#)
10 [directive](#) and a [link](#) or [local clause](#) on another [declare target directive](#).
- 11 • The same [list item](#) must not explicitly appear in both a [link clause](#) on one [declare target](#)
12 [directive](#) and a [local clause](#) on another [declare target directive](#).
- 13 • If a [variable](#) appears in a [enter clause](#) on the [declare target directive](#), its initializer must not
14 refer to a [variable](#) that appears in a [link clause](#) on a [declare target directive](#).

15 Cross References

- 16 • [enter clause](#), see [Section 6.8.4](#)
- 17 • [link clause](#), see [Section 6.8.5](#)
- 18 • `begin declare target` directive, see [Section 8.8.2](#)
- 19 • `declare target` directive, see [Section 8.8.1](#)
- 20 • `target` directive, see [Section 14.8](#)
- 21 • `ompt_callback_target_data_op_emi_t` and
22 `ompt_callback_target_data_op_t`, see [Section 20.5.2.25](#)

23 8.8.1 declare target Directive

24	Name: <code>declare target</code> Category: declarative	Association: none Properties: device , declare target , pure
----	--	---

25 Arguments

26 `declare target` (*extended-list*)

27	Name	Type	Properties
	<i>extended-list</i>	list of extended list item type	optional

28 Clauses

29 [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 / 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** 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 / 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 **list 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 **groupprivate variable** must not appear in any **enter clauses** or **link clauses**.

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.

- 1 • If the **directive** is specified in an interface block for a **procedure**, it must match a **declare**
2 **target directive** in the definition of the **procedure**, including the **device_type** clause if
3 present.
- 4 • If an external **procedure** is a type-bound **procedure** of a derived type and the **directive** is
5 specified in the definition of the external **procedure**, it must appear in the interface block that
6 is accessible to the derived-type definition.
- 7 • If any **procedure** is declared via a **procedure** declaration statement that is not in the
8 type-bound **procedure** part of a derived-type definition, any **declare target directive**
9 with the **procedure** name must appear in the same specification part.
- 10 • The **directive** must appear in the declaration section of a scoping unit in which the common
11 block or **variable** is declared.
- 12 • If a **declare target directive** that specifies a common block name appears in one
13 program unit, then such a **directive** must also appear in every other program unit that contains
14 a **COMMON** statement that specifies the same name, after the last such **COMMON** statement in
15 the program unit.
- 16 • If a **list item** is declared with the **BIND** attribute, the corresponding C entities must also be
17 specified in a **declare target directive** in the C program.
- 18 • A **variable** can only appear in a **declare target directive** in the scope in which it is
19 declared. It must not be an element of a common block or appear in an **EQUIVALENCE**
20 statement.
- 21 • A **variable** that appears in a **declare target directive** must be declared in the Fortran
22 scope of a module or have the **SAVE** attribute, either explicitly or implicitly.

Fortran

Cross References

- 23 • **device_type** clause, see [Section 14.1](#)
- 24
- 25 • **enter** clause, see [Section 6.8.4](#)
- 26 • **indirect** clause, see [Section 8.8.3](#)
- 27 • **link** clause, see [Section 6.8.5](#)
- 28 • **local** clause, see [Section 6.13](#)
- 29 • Declare Target Directives, see [Section 8.8](#)

8.8.2 begin declare target Directive

Name: <code>begin declare target</code>	Association: delimited (declaration-definition-seq)
Category: <code>declarative</code>	Properties: <code>device</code> , <code>declare target</code>

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 `procedure` names of any `variable` or `procedure` 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++

Additionally, the implicit *extended-list* and *local-list* consist of the `variable` and `procedure` names of any `variable` or `procedure` declarations at namespace or class scope that appear in the delimited code region, including the `operator ()` member function of the resulting closure type of any lambda expression that is defined in the delimited code region.

C++

The delimited code region may contain `declare target directives`. If a `device_type clause` is present on the contained `declare target directive`, then its argument determines which versions are made available. If a `list item` appears both in an implicit and explicit `list`, the explicit `list` determines which versions are made available.

Restrictions

Restrictions to the `begin declare target` directive are as follows:

- The function names of overloaded functions or template functions may only be specified within an implicit *extended-list*.
- If a *lambda declaration and definition* appears between a `begin declare target directive` and the paired `end directive`, all `variables` that are captured by the lambda expression must also appear in an `enter clause`.
- A module `export` or `import` statement may not appear between a `begin declare target directive` and the paired `end directive`.

C++

Cross References

- `device_type` clause, see [Section 14.1](#)
- `enter` clause, see [Section 6.8.4](#)
- `indirect` clause, see [Section 8.8.3](#)
- Declare Target Directives, see [Section 8.8](#)

▲ C / C++ ▼

8.8.3 indirect Clause

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

Arguments

Name	Type	Properties
<i>invoked-by-fptr</i>	expression of OpenMP logical type	constant, optional

Modifiers

Name	Modifies	Type	Properties
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: <code>directive-name</code>	unique

Directives

`begin declare target`, `declare target`

Semantics

If *invoked-by-fptr* evaluates to true, any `procedure` that appear in an `enter` clause on the `directive` on which the `indirect` clause is specified may be called with an `indirect device invocation`. If the *invoked-by-fptr* does not evaluate to true, any `procedures` that appear in an `enter` clause on the `directive` may not be called with an `indirect device invocation`. Unless otherwise specified by an `indirect` clause, `procedures` may not be called with an `indirect device invocation`. If the `indirect` clause is specified and *invoked-by-fptr* is not specified, the effect of the `clause` is as if *invoked-by-fptr* evaluates to true.

▲ C / C++ ▼

If a `procedure` appears in the implicit `enter` clause of a `begin declare target` directive and in the `enter` clause of a `declare target` directive that is contained in the delimited code region of the `begin declare target` directive, and if an `indirect` clause appears on both `directives`, then the `indirect` clause on the `begin declare target` directive has no effect or that `procedure`.

▲ C / C++ ▼

1
2
3
4
5
6
7

Restrictions

Restrictions to the [indirect clause](#) are as follows:

- If *invoked-by-fptr* evaluates to true, a [device_type clause](#) must not appear on the same [directive](#) unless it specifies **any** for its *device-type-description*.

Cross References

- **begin declare target** directive, see [Section 8.8.2](#)
- **declare target** directive, see [Section 8.8.1](#)

9 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 an [executable directive](#).

9.1 error Directive

Name: <code>error</code> Category: utility	Association: none 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 9.2](#)
- **message** clause, see [Section 9.3](#)
- **severity** clause, see [Section 9.4](#)
- **ompt_callback_error_t**, see [Section 20.5.2.30](#)

9.2 at Clause

Name: at	Properties: unique
------------------------	---------------------------

Arguments

Name	Type	Properties
<i>action-time</i>	Keyword: compilation , execution	<i>default</i>

Modifiers

Name	Modifies	Type	Properties
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: directive-name	unique

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 9.1](#)

9.3 message Clause

Name: message	Properties: unique
----------------------	------------------------------------

Arguments

Name	Type	Properties
<i>msg-string</i>	expression of string type	<i>default</i>

Modifiers

Name	Modifies	Type	Properties
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: directive-name	unique

Directives

[error](#), [parallel](#)

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

- If the *action-time* is **compilation**, *msg-string* must be a constant expression.

Cross References

- **error** directive, see [Section 9.1](#)
- **parallel** directive, see [Section 11.2](#)

9.4 severity Clause

Name: severity	Properties: unique
-----------------------	------------------------------------

Arguments

Name	Type	Properties
<i>sev-level</i>	Keyword: fatal , warning	<i>default</i>

Modifiers

Name	Modifies	Type	Properties
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: directive-name	unique

Directives

[error](#), [parallel](#)

Semantics

The **severity** clause determines the action that the implementation performs if an error is encountered with respect to the **directive** on which the **clause** appears. 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 9.1](#)
- **parallel** directive, see [Section 11.2](#)

9.5 requires Directive

Name: <code>requires</code> Category: <code>informational</code>	Association: none 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

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

- A **requires** directive may not appear lexically after a **context selector** in which any **clause** of the **requires** directive is used.

▼ C ▼

- The **requires** directive may only appear at file scope.

▲ C ▲

▼ C++ ▼

- The **requires** directive may only appear at file or namespace scope.

▲ 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.

9.5.1 *requirement* Clauses

Clause groups

<p>Properties: required, unique</p>	<p>Members:</p> <p>Clauses <code>atomic_default_mem_order</code>, <code>dynamic_allocators</code>, <code>reverse_offload</code>, <code>self_maps</code>, <code>unified_address</code>, <code>unified_shared_memory</code></p>
--	--

Directives

requires

Semantics

The *requirement clause group* defines a *clause set* that indicates the requirements that a program requires the implementation to support. 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 group*. If the implementation does not support the requirement then it must perform *compile-time error termination*.

Restrictions

- All *compilation units* of a program that contain *declare target directives*, *device constructs* or *device procedures* must specify the same set of requirements that are defined by *clauses* with the *device global requirement property* in the *requirement clause group*.

Cross References

- **requires** directive, see [Section 9.5](#)

9.5.1.1 `atomic_default_mem_order` Clause

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

Arguments

Name	Type	Properties
<i>memory-order</i>	Keyword: <code>acq_rel</code> , <code>acquire</code> , <code>relaxed</code> , <code>release</code> , <code>seq_cst</code>	<i>default</i>

Modifiers

Name	Modifies	Type	Properties
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: <code>directive-name</code>	unique

Directives

`requires`

Semantics

The `atomic_default_mem_order` clause specifies the default memory ordering behavior for `atomic` constructs that an implementation must provide. The effect is as if its argument appears as a clause on any `atomic` construct that does not specify a *memory-order* clause.

Restrictions

Restrictions to the `atomic_default_mem_order` clause are as follows:

- All `requires` directives in the same compilation unit that specify the `atomic_default_mem_order` requirement must specify the same argument.
- Any directive that specifies the `atomic_default_mem_order` clause must not appear lexically after any `atomic` construct on which a *memory-order* clause is not specified.

Cross References

- *memory-order* Clauses, see Section 16.8.1
- `atomic` directive, see Section 16.8.5
- `requires` directive, see Section 9.5

9.5.1.2 `dynamic_allocators` Clause

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

Arguments

Name	Type	Properties
<i>required</i>	expression of OpenMP logical type	constant, optional

Modifiers

Name	Modifies	Type	Properties
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: directive-name	unique

Directives

requires

Semantics

If *required* evaluates to true, the **dynamic_allocators** clause removes certain restrictions on the use of **memory allocators** in **target** regions. Specifically, **allocators** (including the default **allocator** that is specified by the *def-allocator-var* ICV) may be used in a **target** region or in an **allocate** clause on a **target** construct without specifying the **uses_allocators** clause on the **target** construct. Additionally, the implementation must support calls to the **omp_init_allocator** and **omp_destroy_allocator** API routines in **target** regions. If *required* is not specified, the effect is as if *required* evaluates to true.

Cross References

- **allocate** clause, see [Section 7.6](#)
- **uses_allocators** clause, see [Section 7.8](#)
- **requires** directive, see [Section 9.5](#)
- **target** directive, see [Section 14.8](#)
- *def-allocator-var* ICV, see [Table 2.1](#)
- **omp_destroy_allocator**, see [Section 19.13.5](#)
- **omp_init_allocator**, see [Section 19.13.3](#)

9.5.1.3 reverse_offload Clause

Name: reverse_offload	Properties: unique, device global requirement
-------------------------------------	--

Arguments

Name	Type	Properties
<i>required</i>	expression of OpenMP logical type	constant, optional

Modifiers

Name	Modifies	Type	Properties
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: directive-name	unique

Directives

requires

Semantics

If *required* evaluates to true, the `reverse_offload` clause requires an implementation to guarantee that if a `target` construct specifies a `device` clause in which the `ancestor` `devie-modifier` appears, the `target` region can execute on the `parent device` of an enclosing `target` region. If *required* is not specified, the effect is as if *required* evaluates to true.

Restrictions

Restrictions to the `reverse_offload` clause are as follows:

C / C++
<ul style="list-style-type: none"> Any <code>directive</code> that specifies a <code>reverse_offload</code> clause must appear lexically before any <code>device</code> constructs or <code>device</code> procedures.
C / C++

Cross References

- `device` clause, see [Section 14.2](#)
- `requires` directive, see [Section 9.5](#)
- `target` directive, see [Section 14.8](#)
- Declare Target Directives, see [Section 8.8](#)

9.5.1.4 unified_address Clause

Name: <code>unified_address</code>	Properties: unique, device global requirement
---	--

Arguments

Name	Type	Properties
<i>required</i>	expression of OpenMP logical type	constant, optional

Modifiers

Name	Modifies	Type	Properties
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: <code>directive-name</code>	unique

Directives

requires

Semantics

If *required* evaluates to true, 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 exempt from the **unified_address** requirement, following the restrictions of locality to a given execution context, **thread** or **contention group**. If *required* is not specified, the effect is as if *required* evaluates to true.

Restrictions

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

C / C++

- Any **directive** that specifies a **unified_address** clause must appear lexically before any **device constructs** or **device procedures**.

C / C++

Cross References

- **is_device_ptr** clause, see [Section 6.4.7](#)
- **requires** directive, see [Section 9.5](#)
- **target** directive, see [Section 14.8](#)
- Declare Target Directives, see [Section 8.8](#)

9.5.1.5 unified_shared_memory Clause

Name: <code>unified_shared_memory</code>	Properties: unique, device global requirement
---	--

Arguments

Name	Type	Properties
<i>required</i>	expression of OpenMP logical type	constant, optional

Modifiers

Name	Modifies	Type	Properties
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: directive-name	unique

Directives

requires

Semantics

If *required* evaluates to true, the **unified_shared_memory** clause requires the implementation to guarantee that all **devices** share **memory** that is generally accessible to all **threads**.

The **unified_shared_memory** clause implies the **unified_address** requirement, inheriting all of its behaviors.

The implementation must guarantee that storage locations in **memory** are accessible to **threads** on all **accessible devices**, except for **memory** that is local to a specific execution context and exempt from the **unified_address** requirement (see Section 9.5.1.4). Every **device address** that refers to storage allocated through **OpenMP API routines** is a valid **host pointer** that may be dereferenced and may be used as a **host address**. Values stored into **memory** by one **device** may not be visible to another **device** until synchronization establishes a **happens-before order** between the **memory** accesses.

The use of **declare target** directives in an OpenMP program is optional for referencing **variables** with **static storage duration** in **device procedures**.

Any data object that results from the declaration of a **variable** that has **static storage duration** is treated as if it is mapped with a **persistent self map** at the beginning of the program to the **device data environments** of all **target devices** if:

- The **variable** is not a **device local variable**;
- The **variable** is not listed in an **enter** clause on a **declare target** directive; and
- The **variable** is referenced in a **device procedure**.

If *required* is not specified, the effect is as if *required* evaluates to true.

Restrictions

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

C / C++

- Any **directive** that specifies a **unified_shared_memory** clause must appear lexically before any **device constructs** or **device procedures**.

C / C++

Cross References

- **requires** directive, see Section 9.5
- **target** directive, see Section 14.8
- Declare Target Directives, see Section 8.8

9.5.1.6 self_maps Clause

Name: <code>self_maps</code>	Properties: unique, device global requirement
-------------------------------------	--

Arguments

Name	Type	Properties
<i>required</i>	expression of OpenMP logical type	constant, optional

Modifiers

Name	Modifies	Type	Properties
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: directive-name	unique

Directives

requires

Semantics

If *required* evaluates to true, the **self_maps** clause implies the **unified_shared_memory** clause, inheriting all of its behaviors. Additionally, **map-entering** clauses in the **compilation unit** behave as if all resulting **mapping operations** are **self maps**, and all **corresponding list items** created by the **enter** clauses specified by **declare target** directives in the **compilation unit** share storage with the **original list items**.

Restrictions

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

C / C++
<ul style="list-style-type: none">Any directive that specifies a self_maps clause must appear lexically before any device constructs or device procedures.

Cross References

- requires** directive, see [Section 9.5](#)
- target** directive, see [Section 14.8](#)
- Declare Target Directives, see [Section 8.8](#)

9.6 Assumption Directives

Different **assumption directives** facilitate definition of assumptions for a scope that is appropriate to each **base language**. The **assumption scope** of a particular format is defined in the section that defines that **directive**. If the invariants do not hold at runtime, the behavior is unspecified.

9.6.1 *assumption* Clauses

Clause groups

Properties: required, unique	Members: Clauses absent , contains , holds , no_openmp , no_openmp_constructs , no_openmp_routines , no_parallelism
-------------------------------------	---

Directives

[assume](#), [assumes](#), [begin assumes](#)

Semantics

The *assumption clause group* defines a [clause set](#) that indicate the invariants that a program ensures the implementation can exploit.

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 construct](#) or [composite construct](#)) one of its [leaf constructs](#) has the same *directive-name* as one of the list items.

Restrictions

The restrictions to *assumption clauses* are as follows:

- A *directive-name* list item must not specify a [combined directive](#) or a [composite directive](#).
- A *directive-name* list item must not specify a [directive](#) that is a [declarative directive](#), an [informational directive](#), or a [metadirective](#).

Cross References

- [assume](#) directive, see [Section 9.6.3](#)
- [assumes](#) directive, see [Section 9.6.2](#)
- [begin assumes](#) directive, see [Section 9.6.4](#)

9.6.1.1 absent Clause

Name: absent	Properties: unique
-------------------------------------	---------------------------

Arguments

Name	Type	Properties
<i>directive-name-list</i>	list of directive-name list item type	<i>default</i>

Modifiers

Name	Modifies	Type	Properties
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: directive-name	unique

Directives

[assume](#), [assumes](#), [begin assumes](#)

Semantics

The [absent clause](#) specifies that the program guarantees that no [construct](#) that match a *directive-name* list item are encountered in the [assumption scope](#).

Cross References

- **assume** directive, see [Section 9.6.3](#)
- **assumes** directive, see [Section 9.6.2](#)
- **begin assumes** directive, see [Section 9.6.4](#)

9.6.1.2 contains Clause

Name: contains	Properties: unique
------------------------------	---------------------------

Arguments

Name	Type	Properties
<i>directive-name-list</i>	list of directive-name list item type	<i>default</i>

Modifiers

Name	Modifies	Type	Properties
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: directive-name	unique

Directives

[assume](#), [assumes](#), [begin assumes](#)

Semantics

The [contains clause](#) specifies that [constructs](#) that match the *directive-name* list items are likely to be encountered in the [assumption scope](#).

Cross References

- **assume** directive, see [Section 9.6.3](#)
- **assumes** directive, see [Section 9.6.2](#)
- **begin assumes** directive, see [Section 9.6.4](#)

9.6.1.3 holds Clause

Name: holds	Properties: unique
--------------------	---------------------------

Arguments

Name	Type	Properties
<i>hold-expr</i>	expression of OpenMP logical type	<i>default</i>

Modifiers

Name	Modifies	Type	Properties
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: directive-name	unique

Directives

assume, **assumes**, **begin assumes**

Semantics

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.

Cross References

- **assume** directive, see [Section 9.6.3](#)
- **assumes** directive, see [Section 9.6.2](#)
- **begin assumes** directive, see [Section 9.6.4](#)

9.6.1.4 no_openmp Clause

Name: no_openmp	Properties: unique
------------------------	---------------------------

Arguments

Name	Type	Properties
<i>can_assume</i>	expression of OpenMP logical type	constant, optional

Modifiers

Name	Modifies	Type	Properties
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: directive-name	unique

Directives

assume, **assumes**, **begin assumes**

Semantics

If *can_assume* evaluates to true, the **no_openmp** clause guarantees that no OpenMP related code is 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++

Cross References

- **assume** directive, see [Section 9.6.3](#)
- **assumes** directive, see [Section 9.6.2](#)
- **begin assumes** directive, see [Section 9.6.4](#)

9.6.1.5 no_openmp_constructs Clause

Name: no_openmp_constructs	Properties: unique
-----------------------------------	--------------------

Arguments

Name	Type	Properties
<i>can_assume</i>	expression of OpenMP logical type	constant, optional

Modifiers

Name	Modifies	Type	Properties
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: directive-name	unique

Directives

assume, **assumes**, **begin assumes**

Semantics

If *can_assume* evaluates to true, the **no_openmp_constructs** clause guarantees that no **constructs** are encountered in the **assumption scope**.

Cross References

- **assume** directive, see [Section 9.6.3](#)
- **assumes** directive, see [Section 9.6.2](#)
- **begin assumes** directive, see [Section 9.6.4](#)

9.6.1.6 no_openmp_routines Clause

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

Arguments

Name	Type	Properties
<i>can_assume</i>	expression of OpenMP logical type	constant, optional

Modifiers

Name	Modifies	Type	Properties
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: directive-name	unique

Directives

[assume](#), [assumes](#), [begin assumes](#)

Semantics

If *can_assume* evaluates to true, the [no_openmp_routines](#) clause guarantees that no OpenMP API routines are executed in the assumption scope.

Cross References

- [assume](#) directive, see [Section 9.6.3](#)
- [assumes](#) directive, see [Section 9.6.2](#)
- [begin assumes](#) directive, see [Section 9.6.4](#)

9.6.1.7 no_parallelism Clause

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

Arguments

Name	Type	Properties
<i>can_assume</i>	expression of OpenMP logical type	constant, optional

Modifiers

Name	Modifies	Type	Properties
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: directive-name	unique

Directives

[assume](#), [assumes](#), [begin assumes](#)

Semantics

If *can_assume* evaluates to true, the [no_parallelism](#) clause guarantees that no [tasks](#) (explicit or implicit) will be generated and that no [simd constructs](#) will be executed in the assumption scope.

Cross References

- **assume** directive, see [Section 9.6.3](#)
- **assumes** directive, see [Section 9.6.2](#)
- **begin assumes** directive, see [Section 9.6.4](#)

9.6.2 **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](#).

Fortran

Referencing a module that has an **assumes** directive in its specification part does not have the effect as if the **assumes** directive appeared in the specification part of the referencing scope.

Fortran

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

9.6.3 `assume` Directive

Name: <code>assume</code> Category: <code>informational</code>	Association: <code>block</code> Properties: <code>pure</code>
---	--

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`.



9.6.4 `begin assumes` Directive

Name: <code>begin assumes</code> Category: <code>informational</code>	Association: <code>delimited (declaration-definition-seq)</code> Properties: <code>default</code>
--	--

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.



9.7 `nothing` Directive

Name: <code>nothing</code> Category: <code>utility</code>	Association: <code>none</code> Properties: <code>pure, loop-transforming</code>
--	--

Clauses

`apply`

Loop Modifiers for the `apply` Clause

<i>loop-modifier</i>	Number of Generated Loops	Description
<code>identity</code> (<i>default</i>)	1	the copy of the associated loop

1 **Semantics**
2 The **nothing directive** has no effect on the execution of the **OpenMP program** unless otherwise
3 specified by the **apply clause**.

4 If the **nothing directive** immediately precedes a **canonical loop nest** then it forms a
5 **loop-transforming construct**. It associates with the outermost loop and generates one loop that has
6 the same **logical iterations** in the same order as the **associated loop**.

7 **Restrictions**

8

- The **apply clause** can be specified if and only if the **nothing directive** forms a
9 **loop-transforming construct**.

10 **Cross References**

11

- **apply clause**, see [Section 10.6](#)
- **Loop-Transforming Constructs**, see [Chapter 10](#)
- **Metadirectives**, see [Section 8.4](#)

13

10 Loop-Transforming Constructs

A [loop-transforming construct](#) replaces itself, including its [associated loop](#) (see [Section 5.4.1](#)) or [associated loop sequence](#) (see [Section 5.4.6](#)), with a [structured block](#) that may be another loop nest or loop sequence. If the replacement of a [loop-transforming construct](#) is another loop nest or sequence, that loop nest or sequence, possibly as part of an enclosing loop nest or sequence, may be associated with another [loop-nest-associated directive](#) or [loop-sequence-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](#).

A [loop-sequence-transforming construct](#) generates a [canonical loop sequence](#). The [canonical loop nests](#) that are before the [affected loop nests](#) as specified by the [looprange](#) clause are prepended to the generated [canonical loop nest](#), and the loop nests trailing the [affected loop nests](#) are appended to the generated [canonical loop nest](#).

All [generated loops](#) have [canonical loop nest](#) form, unless otherwise specified. Loop iteration variables of [generated loops](#) are always private in the innermost enclosing [parallelism-generating construct](#).

At the beginning of each [logical iteration](#), the [loop iteration variable](#) or the [variable](#) declared by *range-decl* has the value that it would have if the [associated loop](#) was not associated with any [directive](#). After the execution of the [loop-transforming construct](#), the [loop iteration variables](#) of any of its [associated loops](#) have the values that they would have without the [loop-transforming directive](#).

Restrictions

The following restrictions apply to [loop-transforming constructs](#):

- The replacement of a [loop-transforming construct](#) with its [generated loop nests](#) or [generated loop sequences](#) must result in a [conforming program](#).

Cross References

- [nothing](#) directive, see [Section 9.7](#)
- Canonical Loop Nest Form, see [Section 5.4.1](#)

10.1 tile Construct

Name: <code>tile</code> Category: <code>executable</code>	Association: loop nest Properties: <code>pure</code> , <code>loop-transforming</code> , <code>simdizable</code>
--	--

Clauses

`apply`, `sizes`

Loop Modifiers for the `apply` Clause

<i>loop-modifier</i>	Number of Generated Loops	Description
<code>grid</code>	n	the grid loops g_1, \dots, g_n
<code>intratile</code>	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 loops`.
- 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-nest-associated directive`.

Cross References

- `apply` clause, see [Section 10.6](#)
- `sizes` clause, see [Section 10.1.1](#)

10.1.1 sizes Clause

Name: <code>sizes</code>	Properties: unique, required
---------------------------------	-------------------------------------

Arguments

Name	Type	Properties
<i>size-list</i>	list of OpenMP integer expression type	constant, positive

Modifiers

Name	Modifies	Type	Properties
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: directive-name	unique

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 10.1](#)

10.2 unroll Construct

Name: <code>unroll</code> Category: <code>executable</code>	Association: loop nest Properties: <code>pure</code> , <code>loop-transforming</code> , <code>simdizable</code>
--	--

Clauses

`apply`, `full`, `partial`

Clause set

Properties: <code>exclusive</code>	Members: <code>full</code> , <code>partial</code>
---	--

Loop Modifiers for the `apply` Clause

<i>loop-modifier</i>	Number of Generated Loops	Description
unrolled (<i>default</i>)	1	the grid loop g_1 of the tiling step

Semantics

The **unroll** construct is associated with one loop, which is unrolled according to its specified clauses. 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.

If the **apply** clause is specified on **construct** without a *loop-modifier*, the effect is as if **unrolled** 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 10.6](#)
- **full** clause, see [Section 10.2.1](#)
- **partial** clause, see [Section 10.2.2](#)

10.2.1 full Clause

Name: full	Properties: unique
--------------------------	---------------------------

Arguments

Name	Type	Properties
<i>fully_unroll</i>	expression of OpenMP logical type	constant, optional

Modifiers

Name	Modifies	Type	Properties
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: directive-name	unique

Directives

unroll

Semantics

If *fully_unroll* evaluates to true, 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* **associated iterations** and in their **logical iteration** order. If *fully_unroll* evaluates to false, the **full** clause has no effect. If *fully_unroll* is not specified, the effect is as if *fully_unroll* evaluates to true.

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 10.2](#)

10.2.2 partial Clause

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

Arguments

Name	Type	Properties
<i>unroll-factor</i>	expression of integer type	optional, constant, positive

Modifiers

Name	Modifies	Type	Properties
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: directive-name	unique

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 10.2](#)

10.3 reverse Construct

Name: <code>reverse</code> Category: <code>executable</code>	Association: loop nest Properties: <code>pure</code> , <code>loop-transforming</code> , <code>simdizable</code>
---	--

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 10.6](#)

10.4 interchange Construct

Name: <code>interchange</code> Category: <code>executable</code>	Association: loop nest Properties: <code>pure</code> , <code>loop-transforming</code> , <code>simdizable</code>
---	--

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 loops**.

Cross References

- **apply** clause, see [Section 10.6](#)
- **permutation** clause, see [Section 10.4.1](#)

10.4.1 permutation Clause

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

Arguments

Name	Type	Properties
<i>permutation-list</i>	list of OpenMP integer expression type	constant, positive

Modifiers

Name	Modifies	Type	Properties
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: directive-name	unique

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 10.4](#)

10.5 fuse Construct

Name: <code>fuse</code> Category: executable	Association: loop sequence Properties: pure , loop-transforming , simdizable
---	---

Clauses

[looprange](#)

Loop Modifiers for the `apply` Clause

<i>loop-modifier</i>	Number of Generated Loops	Description
fused (<i>default</i>)	1	the fused loop

Semantics

The **fuse** construct merges the **affected loop nests** specified by the **looprange** clause into a single **canonical loop nest** where execution of each **logical iteration** of the **generated loop** executes a **logical iteration** of each **affected loop nest**.

Let ℓ^1, \dots, ℓ^n be the **affected loop nests** with m^1, \dots, m^n **logical iterations** each, and i_j^k the j^{th} **logical iteration** of loop ℓ^k . Let i_j^k be an empty iteration if $j \geq m^k$. Let m_{\max} be the number of **logical iterations** of the **affected loop nest** with the most **logical iterations**. The loop generated by the **fuse** construct has m_{\max} **logical iterations**, where execution of the j^{th} **logical iteration** executes the **logical iterations** i_j^1, \dots, i_j^n , in that order.

Cross References

- **looprange** clause, see [Section 5.4.7](#)

10.6 `apply` Clause

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

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: fused , grid , identity , interchanged , intratile , reversed , unrolled	default
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: directive-name	unique

Directives

interchange, **nothing**, **reverse**, **tile**, **unroll**

Semantics

The **apply** clause applies [loop-transforming constructs](#), specified by the *applied-directives* list, to the [generated loops](#) of a [loop-transforming construct](#). The *loop-modifier* specifies to which [generated loops](#) the [directives](#) are applied. An applied [loop-transforming construct](#) 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. 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 construct](#).
- 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* modifier must not be omitted.

Cross References

- **interchange** directive, see [Section 10.4](#)
- **metadirective** directive, see [Section 8.4.3](#)
- **nothing** directive, see [Section 9.7](#)
- **reverse** directive, see [Section 10.3](#)
- **tile** directive, see [Section 10.1](#)
- **unroll** directive, see [Section 10.2](#)

11 Parallelism Generation and Control

This chapter defines `constructs` for generating and controlling parallelism.

11.1 `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`.

11.2 `parallel` Construct

Name: <code>parallel</code> Category: <code>executable</code>	Association: block Properties: <code>parallelism-generating</code> , <code>team-generating</code> , <code>cancellable</code> , <code>thread-limiting</code> , <code>context-matching</code>
--	--

Clauses

`allocate`, `copyin`, `default`, `firstprivate`, `if`, `message`, `num_threads`, `private`, `proc_bind`, `reduction`, `safesync`, `severity`, `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` is formed to execute the `parallel` region (see Section 11.2.1 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 formed, 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

1 number of **threads** in the **team**. A **thread** may obtain its own **thread number** by a call to the
2 **omp_get_thread_num** library routine.

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

9 The implementation may cause any **thread** to suspend execution of its **implicit task** at a **task**
10 **scheduling point**, and to switch to execution of any **explicit task** generated by any of the **threads** in
11 the **team**, before eventually resuming execution of the **implicit task** (for more details see
12 Chapter 13).

13 An implicit **barrier** occurs at the end of a **parallel region**. After the end of a **parallel region**,
14 only the **primary thread** of the **team** resumes execution of the enclosing **task region**.

15 If a **thread** in a **team** that is executing a **parallel region** encounters another **parallel**
16 **directive**, it forms a new **team**, according to the rules in Section 11.2.1, and it becomes the **primary**
17 **thread** of that new **team**.

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

22 Execution Model Events

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

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

28 If a new **native thread** is created for the **team** that executes the **parallel region** upon
29 encountering the **construct**, a *native-thread-begin event* occurs as the first **event** in the context of the
30 new **thread** prior to the *implicit-task-begin event*.

31 **Events** associated with implicit **barriers** occur at the end of a **parallel region**. Section 16.3.2
32 describes **events** associated with implicit **barriers**.

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

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

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

4 **Tool Callbacks**

5 A **thread** dispatches a registered **ompt_callback_parallel_begin** **callback** for each
6 occurrence of a *parallel-begin event* in that **thread**. The **callback** occurs in the **task** that encounters
7 the **parallel** **construct**. This **callback** has the type signature
8 **ompt_callback_parallel_begin_t**. In the dispatched **callback**,
9 (*flags & ompt_parallel_team*) evaluates to *true*.

10 A **thread** dispatches a registered **ompt_callback_implicit_task** **callback** with
11 **ompt_scope_begin** as its *endpoint* argument for each occurrence of an *implicit-task-begin*
12 **event** in that **thread**. Similarly, a **thread** dispatches a registered
13 **ompt_callback_implicit_task** **callback** with **ompt_scope_end** as its *endpoint*
14 argument for each occurrence of an *implicit-task-end event* in that **thread**. The **callbacks** occur in
15 the context of the **implicit task** and have type signature **ompt_callback_implicit_task_t**.
16 In the dispatched **callback**, (*flags & ompt_task_implicit*) evaluates to *true*.

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

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

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

27 **Cross References**

- 28 • **allocate** clause, see [Section 7.6](#)
- 29 • **copyin** clause, see [Section 6.7.1](#)
- 30 • **default** clause, see [Section 6.4.1](#)
- 31 • **firstprivate** clause, see [Section 6.4.4](#)
- 32 • **if** clause, see [Section 4.5](#)
- 33 • **message** clause, see [Section 9.3](#)
- 34 • **num_threads** clause, see [Section 11.2.2](#)
- 35 • **private** clause, see [Section 6.4.3](#)
- 36 • **proc_bind** clause, see [Section 11.2.4](#)

- 1 • **reduction** clause, see [Section 6.5.9](#)
- 2 • **safesync** clause, see [Section 11.2.5](#)
- 3 • **severity** clause, see [Section 9.4](#)
- 4 • **shared** clause, see [Section 6.4.2](#)
- 5 • **omp_get_thread_num**, see [Section 19.2.4](#)
- 6 • **ompt_callback_implicit_task_t**, see [Section 20.5.2.11](#)
- 7 • **ompt_callback_parallel_begin_t**, see [Section 20.5.2.3](#)
- 8 • **ompt_callback_parallel_end_t**, see [Section 20.5.2.4](#)
- 9 • **ompt_callback_thread_begin_t**, see [Section 20.5.2.1](#)
- 10 • **ompt_callback_thread_end_t**, see [Section 20.5.2.2](#)
- 11 • **ompt_scope_endpoint_t**, see [Section 20.4.4.11](#)
- 12 • Determining the Number of Threads for a **parallel** Region, see [Section 11.2.1](#)

11.2.1 Determining the Number of Threads for a **parallel** Region

When execution encounters a **parallel** directive, the value of the **if** clause or the first item of the *nthreads* list of the **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-expression* of an **if** clause or in an element of the *nthreads* list of a **num_threads** clause of a **parallel** construct causes an implicit reference to the **variable** in all enclosing constructs. The *if-expression* and the *nthreads* list items 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 *nthreads* list items or an *if-expression* occur is also unspecified.

When a **thread** encounters a **parallel** construct, the number of **threads** is determined according to Algorithm 11.1.

Cross References

- **if** clause, see [Section 4.5](#)
- **num_threads** clause, see [Section 11.2.2](#)
- **parallel** directive, see [Section 11.2](#)
- *dyn-var* ICV, see [Table 2.1](#)
- *max-active-levels-var* ICV, see [Table 2.1](#)

Algorithm 11.1 Determine Number of Threads

let *ThreadsBusy* be the number of **threads** currently executing **tasks** in this **contention group**;
let *StructuredThreadsBusy* be the number of **structured 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 the first item of the *nthreads* list;
else let *ThreadsRequested* = value of the first element of *nthreads-var*;
let *ThreadsAvailable* = $\min(\text{thread-limit-var} - \text{ThreadsBusy}, \text{structured-thread-limit-var} - \text{StructuredThreadsBusy}) + 1$;
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**

- *nthreads-var* ICV, see [Table 2.1](#)
- *thread-limit-var* ICV, see [Table 2.1](#)

11.2.2 num_threads Clause

Name: num_threads	Properties: unique
--------------------------	---------------------------

Arguments

Name	Type	Properties
<i>nthreads</i>	list of OpenMP integer expression type	positive

Modifiers

Name	Modifies	Type	Properties
<i>prescriptiveness</i>	<i>nthreads</i>	Keyword: strict	default
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: directive-name	unique

Directives

parallel

Semantics

The **num_threads** clause specifies the desired number of **threads** to execute a **parallel** region. Algorithm 11.1 determines the number of **threads** that execute the **parallel** region. If *prescriptiveness* is specified as **strict** and an implementation determines that Algorithm 11.1 would always result in a number of **threads** other than the value of the first item of the *nthreads* list then **compile-time error termination** may be performed in which case the effect of any **message** clause associated with the **directive** is **implementation defined**. Otherwise, if *prescriptiveness* is specified as **strict** and Algorithm 11.1 would result in a number of **threads** other than the value of the first item of the *nthreads* list then **runtime error termination** is performed. In both **error termination** scenarios, the effect is as if an **error** **directive** has been encountered on which any specified **message** and **severity** clauses and an **at** clause with **execution** as *action-time* are specified.

Cross References

- **at** clause, see [Section 9.2](#)
- **message** clause, see [Section 9.3](#)
- **parallel** directive, see [Section 11.2](#)

11.2.3 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 **threads** to **places** within the input **place partition**, as defined in the following paragraph. If the **parallel** directive has a **proc_bind** clause then the **thread affinity** 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**.

If the **encountering thread** is a **free-agent thread** that is executing an **explicit task** that was created in an **implicit parallel region**, the input **place partition** for all **thread affinity** policies is the value of the *place-partition-var* ICV of the **initial task**. If the **encountering thread** is a **free-agent thread** that is executing an **explicit task** that was created in an **explicit parallel region**, the input **place partition** for all **thread affinity** policies is the input **place partition** of that **parallel region**. If the **encountering thread** is not a **free-agent thread**, the input **place partition** for all **thread affinity** policies is the value of the *place-partition-var* ICV of its **binding implicit task**.

Under the **primary** and **close** **thread affinity** policies, the *place-partition-var* ICV of each **implicit task** is assigned the input **place partition**. As discussed below, under the **spread thread affinity** policy, the *place-partition-var* ICV of each **implicit task** is derived from the value of the input **place partition**.

The *place-assignment-var* ICV is a list of positions. Each position is assigned to a group that is derived based on the **thread affinity** policy that applies to the **parallel** directive as described below. A set of **places** is assigned to each group and its positions; if more than one **place** is assigned to a group, the positions assigned to the group are associated with the **places** in round robin fashion with wrap-around, starting with the first **place** that is assigned to the group. Each **thread** assigned to the **team** is bound to the **place** that is associated with the group that includes the position that equals its **thread number**. That is, each **thread** of the **team** is assigned to the position of the *place-assignment-var* that corresponds to its **thread number**.

Free-agent threads that execute **tasks** bound to the **team** are assigned to the first position of the *place-assignment-var* that has not been assigned to any other **thread** and are bound to a **place** that is associated with that position. If another **OpenMP thread** is bound to that **place**, the **place** to which the **free-agent thread** is bound is **implementation defined**.

The assignment of positions to groups that determines the *place-assignment-var* ICV uses the following symbols:

- T : the number of **threads** in the **team**;
- P : the number of **places** in the input **place partition**;
- L : the value of the *thread-limit-var* ICV;
- NG : the total number of groups;
- BT : the below **thread** count, which is equal to $\lfloor T/NG \rfloor$;

- 1 • AT : the above **thread** count, which is equal to $\lceil T/NG \rceil$;
- 2 • ET : the excess **thread** count, which is equal to $T \bmod NG$;
- 3 • g_i : a member of the set of groups, g_0, \dots, g_{NG-1} ; and
- 4 • l_j : the group assigned to position j ;

5 The *place-assignment-var ICV* consists of L positions. Thus, each **thread affinity** policy determines
 6 the composition of each group g_i by assigning position j to one of them for each j ,
 7 $j = 0, \dots, L - 1$.

8 Under the **primary thread affinity** policy, $NG = 1$ and all positions are assigned to a single
 9 group, g_0 . The **place** assigned to g_0 is the **place** to which the **encountering thread** is assigned. Thus,
 10 all positions of the *place-assignment-var ICV* are associated with the same **place** as the **primary**
 11 **thread**.

12 The **close thread affinity** policy sets NG to P . Each **place** in the input **place partition** is assigned
 13 to one group, starting with the **place** to which the **encountering thread** is assigned, which is
 14 assigned to g_0 . The **place** assigned to group g_i is then the next **place** in the **place partition** of the one
 15 assigned to group g_{i-1} with wrap around with respect to the input **place partition**.

16 The purpose of the **spread thread affinity** policy is to create a sparse distribution for a **team** of T
 17 **threads** among the P **places** of the parent's **place partition**. A sparse distribution is achieved by first
 18 subdividing the parent partition into T subpartitions if $T \leq P$ (in which case $NG = T$), or P
 19 subpartitions if $T > P$ (in which case $NG = P$). The subpartitions are determined as follows:

- 20 • $T \leq P$: The input **place partition** is split into T subpartitions, where each subpartition
 21 contains $\lfloor P/T \rfloor$ or $\lceil P/T \rceil$ consecutive **places**; if $P \bmod T$ is not zero, which subpartitions
 22 contain $\lceil P/T \rceil$ **places** is **implementation defined**;
- 23 • $T > P$: The input **place partition** is split into P subpartitions, each with a single **place**.

24 In either case, a subpartition is assigned to each group. The subpartition that is assigned to group
 25 g_0 is the one that includes the **place** on which the **encountering thread** is executing. The
 26 subpartition that is assigned to group g_i is the one that includes the next **place** to those in the
 27 subpartition assigned to group g_{i-1} , with wrap around with respect to the input **place partition**. The
 28 *place-partition-var ICV* of each **implicit task** is set to the subpartition associated with the group to
 29 which its corresponding position is assigned. Thus, the subpartitioning is not only a mechanism for
 30 achieving a sparse distribution, it also defines a subset of **places** for a **thread** to use when creating a
 31 nested **parallel region**.

32 Both the **close** and the **spread thread affinity** policies assign the values of the
 33 *place-assignment-var ICV* as follows:

- 34 • For positions from 0 up to $T - 1$: The positions are partitioned into NG sets of consecutive
 35 positions, ET of which have AT positions and $NG - ET$ of which have BT positions
 36 (when ET is not zero, which sets have which count is **implementation defined** unless the
 37 **thread affinity** policy is **close** and $T < P$, in which case the first T groups are assigned the

- 1 sets with *AT* positions) and the sets are assigned to each group, with the first set, which starts
 2 with position 0, assigned to the first group, g_0 , and with each successive set i , which starts
 3 with the position immediately after the last position in the set assigned to group g_{i-1} ,
 4 assigned to the next group g_i ;
- 5 • If $ET \neq 0$, for the positions from T up to $(AT * NG) - 1$: Each of these positions is
 6 assigned to a group g_i that received fewer than AT positions in the above step such that each
 7 such g_i is assigned AT positions (which positions are assigned to which group is
 8 [implementation defined](#));
 - 9 • For the remaining positions from $AT * NG$ up to L : Each position is assigned to a group in
 10 round robin fashion, starting with g_0 .

11 The determination of whether the affinity request can be fulfilled is [implementation defined](#). If it
 12 cannot be fulfilled, then the affinity of [threads](#) in the [team](#) is [implementation defined](#).

13

14 **Note** – Wrap around is needed if the end of a [place partition](#) is reached before all [thread](#)
 15 assignments are done. For example, wrap around may be needed in the case of `close` and $T \leq P$,
 16 if the [primary thread](#) is assigned to a [place](#) other than the first [place](#) in the [place partition](#). In this
 17 case, [thread 1](#) is assigned to the [place](#) after the [place](#) of the [primary thread](#), [thread 2](#) is assigned to
 18 the [place](#) after that, and so on. The end of the [place partition](#) may be reached before all [threads](#) are
 19 assigned. In this case, assignment of [threads](#) is resumed with the first [place](#) in the [place partition](#).
 20

21 Cross References

- 22 • `proc_bind` clause, see [Section 11.2.4](#)
- 23 • `parallel` directive, see [Section 11.2](#)
- 24 • `bind-var` ICV, see [Table 2.1](#)
- 25 • `place-partition-var` ICV, see [Table 2.1](#)

26 11.2.4 `proc_bind` Clause

27 Name: <code>proc_bind</code>	Properties: unique
--	---------------------------

28 Arguments

29 Name	Type	Properties
<i>affinity-policy</i>	Keyword: <code>close</code> , <code>primary</code> , <code>spread</code>	<i>default</i>

30 Modifiers

31 Name	Modifies	Type	Properties
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: <code>directive-name</code>	unique

Directives

`parallel`

Semantics

The `proc_bind` clause specifies the mapping of threads to places within the current place partition, that is, within the place 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 11.2.3

Cross References

- `parallel` directive, see Section 11.2
- Controlling OpenMP Thread Affinity, see Section 11.2.3
- *place-partition-var* ICV, see Table 2.1

11.2.5 safesync Clause

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

Arguments

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

Modifiers

Name	Modifies	Type	Properties
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: directive-name	unique

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. 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. If the *width* argument is not specified, the behavior is as if the *width* argument is one.

Cross References

- `parallel` directive, see Section 11.2

11.3 teams Construct

Name: <code>teams</code> Category: <code>executable</code>	Association: <code>block</code> Properties: <code>parallelism-generating, team-generating, thread-limiting, context-matching</code>
---	--

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 teams` 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 formed `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` construct 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 generated for the corresponding **teams region**.

Upon generation 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 a new **native thread** is created for the **league of teams** that executes the **teams region** upon encountering the **construct**, 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** completes 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 **initial thread** that uses the **native thread** as the last **event** prior to destruction of the **native 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 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**, **(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 each *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 each *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 **regions** that may be strictly nested inside the **teams** region.

Cross References

- **allocate** clause, see [Section 7.6](#)
- **default** clause, see [Section 6.4.1](#)
- **firstprivate** clause, see [Section 6.4.4](#)
- **if** clause, see [Section 4.5](#)
- **num_teams** clause, see [Section 11.3.1](#)
- **private** clause, see [Section 6.4.3](#)
- **reduction** clause, see [Section 6.5.9](#)
- **shared** clause, see [Section 6.4.2](#)
- **thread_limit** clause, see [Section 14.3](#)
- **distribute** directive, see [Section 12.7](#)
- **parallel** directive, see [Section 11.2](#)
- **target** directive, see [Section 14.8](#)
- **omp_get_num_teams**, see [Section 19.4.1](#)
- **omp_get_team_num**, see [Section 19.4.2](#)
- **ompt_callback_implicit_task_t**, see [Section 20.5.2.11](#)
- **ompt_callback_parallel_begin_t**, see [Section 20.5.2.3](#)
- **ompt_callback_parallel_end_t**, see [Section 20.5.2.4](#)
- **ompt_callback_thread_begin_t**, see [Section 20.5.2.1](#)
- **ompt_callback_thread_end_t**, see [Section 20.5.2.2](#)

11.3.1 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>upper-bound</i>	OpenMP integer expression	positive, ultimate, unique
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: directive-name	unique

Directives

teams

Semantics

The **num_teams** clause specifies the bounds on the number of **teams** formed 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** formed 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 *nteams-var* ICV is greater than zero, the effect is as if *upper-bound* was specified as an **implementation defined** value greater than zero but less than or equal to the value of the *nteams-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 11.3](#)

11.4 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
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: directive-name	unique

Directives

distribute, **do**, **for**, **loop**, **simd**

Semantics

The **order** clause specifies an *ordering* of execution for the **collapsed iterations** of a **loop-collapsing construct**. If *ordering* is **concurrent**, different **collapsed iterations** 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 **native threads** to participate in the parallel execution of any **collapsed iterations**.

The *order-modifier* on the **order** clause affects the schedule specification for the purpose of determining its consistency with other schedules (see [Section 5.4.5](#)). 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, the **atomic** 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 **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 12.7](#)
- **do** directive, see [Section 12.6.2](#)
- **for** directive, see [Section 12.6.1](#)
- **loop** directive, see [Section 12.8](#)
- **simd** directive, see [Section 11.5](#)

11.5 simd Construct

Name: <code>simd</code> Category: <code>executable</code>	Association: loop nest Properties: <code>parallelism-generating</code> , <code>context-matching</code> , <code>simdizable</code> , <code>pure</code>
--	---

Separating directives

`scan`

Clauses

`aligned`, `collapse`, `if`, `induction`, `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 `collapsed iterations` concurrently by using `SIMD instructions`. At the beginning of each `collapsed iteration`, the loop iteration `variable` or the `variable` declared by `range-decl` of each `collapsed loop` has the value that it would have if the set of the `collapsed loops` was executed sequentially. The number of `collapsed 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 with an `if-expression` that 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` may be encountered during execution of a `simd` region.
- If an `order` clause that specifies `concurrent` appears on a `simd` directive, the `safelen` clause must not also appear.

▼ C / C++ ▼

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

▲ C / C++ ▲

- No exceptions can be raised in the **simd** region.
- The only random access iterator types that are allowed for the **collapsed loops** are pointer types.

Cross References

- **aligned** clause, see [Section 6.11](#)
- **collapse** clause, see [Section 5.4.3](#)
- **if** clause, see [Section 4.5](#)
- **induction** clause, see [Section 6.5.12](#)
- **lastprivate** clause, see [Section 6.4.5](#)
- **linear** clause, see [Section 6.4.6](#)
- **nontemporal** clause, see [Section 11.5.1](#)
- **order** clause, see [Section 11.4](#)
- **private** clause, see [Section 6.4.3](#)
- **reduction** clause, see [Section 6.5.9](#)
- **safelen** clause, see [Section 11.5.2](#)
- **simdlen** clause, see [Section 11.5.3](#)
- **scan** directive, see [Section 6.6](#)

11.5.1 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>

Modifiers

Name	Modifies	Type	Properties
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: directive-name	unique

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 11.5](#)

11.5.2 safelen Clause

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

Arguments

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

Modifiers

Name	Modifies	Type	Properties
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: directive-name	unique

Directives

simd

Semantics

The **safelen** clause specifies that no two concurrent iterations within a [SIMD chunk](#) can have a distance in the [collapsed iteration space](#) that is greater than or equal to the value given in the [clause](#).

Cross References

- **simd** directive, see [Section 11.5](#)

11.5.3 simdlen Clause

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

Arguments

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

Modifiers

Name	Modifies	Type	Properties
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: directive-name	unique

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 `collapsed iterations` to be executed concurrently. When the `simdlen` clause appears on a `declare simd construct`, if a `SIMD` version of the associated `procedure` is created, *length* corresponds to the number of concurrent arguments of the `procedure`.

Cross References

- `declare simd` directive, see [Section 8.7](#)
- `simd` directive, see [Section 11.5](#)

11.6 masked Construct

Name: <code>masked</code> Category: <code>executable</code>	Association: <code>block</code> Properties: <code>thread-limiting</code>
--	---

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` argument 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`.

1 Tool Callbacks

2 A [thread](#) dispatches a registered `ompt_callback_masked` [callback](#) with
3 `ompt_scope_begin` as its *endpoint* argument for each occurrence of a *masked-begin event* in
4 that [thread](#). Similarly, a [thread](#) dispatches a registered `ompt_callback_masked` [callback](#) with
5 `ompt_scope_end` as its *endpoint* argument for each occurrence of a *masked-end event* in that
6 [thread](#). These [callbacks](#) occur in the context of the [task](#) executed by the current [thread](#) and have the
7 type signature `ompt_callback_masked_t`.

8 Cross References

- 9 • `filter` clause, see [Section 11.6.1](#)
- 10 • `ompt_callback_masked_t`, see [Section 20.5.2.12](#)
- 11 • `ompt_scope_endpoint_t`, see [Section 20.4.4.11](#)

12 11.6.1 filter Clause

13 Name: <code>filter</code>	Properties: unique
------------------------------	--------------------

14 Arguments

15 Name	Type	Properties
<code>thread_num</code>	expression of integer type	<i>default</i>

16 Modifiers

17 Name	Modifies	Type	Properties
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: <code>directive-name</code>	unique

18 Directives

19 [masked](#)

20 Semantics

21 If `thread_num` specifies the [thread number](#) of the current [thread](#) in the [current team](#) then the
22 [filter clause](#) selects the current [thread](#). If the [filter clause](#) is not specified, the effect is as if
23 the [clause](#) is specified with `thread_num` equal to zero, so that the [filter clause](#) selects the
24 [primary thread](#). The use of a [variable](#) in a `thread_num` argument expression causes an implicit
25 reference to the [variable](#) in all enclosing constructs.

26 Cross References

- 27 • `masked` directive, see [Section 11.6](#)

12 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 a **worksharing construct** if the **binding thread set** is a **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 with *do_not_synchronize* specified as true, in which case 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.

If a **work-distribution construct** is a **partitioned construct** then all user code encountered in the **region**, but not in a **nested region** that is not a **closely nested region**, is executed by one **thread** from the **binding thread set**.

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 **team** must be the same for every **thread** in the **team**.

Fortran

- A **variable** must not be private within a **teams** or **parallel region** if it has either **LOCAL_INIT** or **SHARED** locality in a **DO CONCURRENT** loop that is associated with a **work-distribution construct**, where the **teams** or **parallel region** is a binding region of the corresponding **work-distribution region**.
- If a **variable** is accessed in more than one iteration of a **DO CONCURRENT** loop that is associated with the loop directive and at least one of the accesses modifies the **variable**, the **variable** must have locality specified in the **DO CONCURRENT** loop.

Fortran

12.1 single Construct

Name: <code>single</code> Category: <code>executable</code>	Association: <code>block</code> Properties: <code>work-distribution</code> , <code>team-executed</code> , <code>partitioned</code> , <code>worksharing</code> , <code>thread-limiting</code>
--	--

Clauses

`allocate`, `copyprivate`, `firstprivate`, `nowait`, `private`

Clause set

Properties: <code>exclusive</code>	Members: <code>copyprivate</code> , <code>nowait</code>
---	--

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` does not specify otherwise.

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`.

Cross References

- `allocate` clause, see [Section 7.6](#)
- `copyprivate` clause, see [Section 6.7.2](#)
- `firstprivate` clause, see [Section 6.4.4](#)
- `nowait` clause, see [Section 16.6](#)
- `private` clause, see [Section 6.4.3](#)
- `ompt_callback_work_t`, see [Section 20.5.2.5](#)
- `ompt_scope_endpoint_t`, see [Section 20.4.4.11](#)
- `ompt_work_t`, see [Section 20.4.4.16](#)

12.2 scope Construct

Name: <code>scope</code> Category: executable	Association: block Properties: work-distribution , team-executed , 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 does not specify otherwise.

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 7.6](#)
- [firstprivate](#) clause, see [Section 6.4.4](#)
- [nowait](#) clause, see [Section 16.6](#)
- [private](#) clause, see [Section 6.4.3](#)
- [reduction](#) clause, see [Section 6.5.9](#)
- [ompt_callback_work_t](#), see [Section 20.5.2.5](#)
- [ompt_scope_endpoint_t](#), see [Section 20.4.4.11](#)
- [ompt_work_t](#), see [Section 20.4.4.16](#)

12.3 sections Construct

Name: sections Category: executable	Association: block Properties: work-distribution , team-executed , partitioned , 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 does not specify otherwise.

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

Tool Callbacks

A **thread** dispatches a registered **ompt_callback_work** callback with **ompt_scope_begin** as its *endpoint* argument and **ompt_work_sections** as its *work_type* argument for each occurrence of a *sections-begin* event in that **thread**. Similarly, a **thread** dispatches a registered **ompt_callback_work** callback with **ompt_scope_end** as its *endpoint* argument and **ompt_work_sections** as its *work_type* argument for each occurrence of a *sections-end* event in that **thread**. The **callbacks** occur in the context of the **implicit task**. The **callbacks** have type signature **ompt_callback_work_t**.

Cross References

- **allocate** clause, see [Section 7.6](#)
- **firstprivate** clause, see [Section 6.4.4](#)
- **lastprivate** clause, see [Section 6.4.5](#)
- **nowait** clause, see [Section 16.6](#)
- **private** clause, see [Section 6.4.3](#)
- **reduction** clause, see [Section 6.5.9](#)
- **section** directive, see [Section 12.3.1](#)
- **ompt_callback_dispatch_t**, see [Section 20.5.2.6](#)
- **ompt_callback_work_t**, see [Section 20.5.2.5](#)
- **ompt_scope_endpoint_t**, see [Section 20.4.4.11](#)
- **ompt_work_t**, see [Section 20.4.4.16](#)

12.3.1 section Directive

Name: <code>section</code> Category: <code>subsidiary</code>	Association: separating Properties: <code>default</code>
---	---

Separated directives

`sections`

Semantics

The `section` directive splits a `structured block sequence` that is associated with a `sections` construct into two `structured block sequences`.

Execution Model Events

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

Tool Callbacks

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

Cross References

- `sections` directive, see [Section 12.3](#)

Fortran

12.4 workshare Construct

Name: <code>workshare</code> Category: <code>executable</code>	Association: block Properties: <code>work-distribution</code> , <code>team-executed</code> , <code>partitioned</code> , <code>worksharing</code>
---	--

Clauses

`nowait`

Binding

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

1 **Semantics**

2 The **workshare** construct divides the execution of the associated **structured block** into separate
 3 **units of work** and causes the **threads** of the **team** to share the work such that each **unit of work** is
 4 executed only once by one **thread**, in the context of its **implicit task**. An implicit **barrier** occurs at
 5 the end of a **workshare region** if a **nowait** clause does not specify otherwise.

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

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

- 12 • For array expressions within each statement, including transformational array intrinsic
 13 functions that compute scalar values from arrays:
 - 14 – Evaluation of each element of the array expression, including any references to
 15 elemental functions, is a **unit of work**.
 - 16 – Evaluation of transformational array intrinsic functions may be subdivided into any
 17 number of **units of work**.
- 18 • For array assignment statements, assignment of each element is a **unit of work**.
- 19 • For scalar assignment statements, each assignment operation is a **unit of work**.
- 20 • For **WHERE** statements or constructs, evaluation of the mask expression and the masked
 21 assignments are each a **unit of work**.
- 22 • For **FORALL** statements or constructs, evaluation of the mask expression, expressions
 23 occurring in the specification of the iteration space, and the masked assignments are each a
 24 **unit of work**.
- 25 • For **atomic** constructs, **critical** constructs, and **parallel** constructs, the **construct** is
 26 a **unit of work**. A new **team** executes the statements contained in a **parallel** construct.
- 27 • If none of the rules above apply to a portion of a statement in the **structured block**, then that
 28 portion is a **unit of work**.

29 The transformational array intrinsic functions are **MATMUL**, **DOT_PRODUCT**, **SUM**, **PRODUCT**,
 30 **MAXVAL**, **MINVAL**, **COUNT**, **ANY**, **ALL**, **SPREAD**, **PACK**, **UNPACK**, **RESHAPE**, **TRANSPOSE**,
 31 **EOSHIFT**, **CSHIFT**, **MINLOC**, and **MAXLOC**.

32 The **units of work** are assigned to the **threads** that execute a **workshare region** such that each **unit**
 33 **of work** is executed once.

34 If an array expression in the **structured block** references the value, association status, or allocation
 35 status of private **variables**, the value of the expression is undefined, unless the same value would be
 36 computed by every **thread**.

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

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

5 Execution Model Events

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

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

10 Tool Callbacks

11 A **thread** dispatches a registered **ompt_callback_work callback** with **ompt_scope_begin**
 12 as its *endpoint* argument and **ompt_work_workshare** as its *work_type* argument for each
 13 occurrence of a *workshare-begin event* in that **thread**. Similarly, a **thread** dispatches a registered
 14 **ompt_callback_work callback** with **ompt_scope_end** as its *endpoint* argument and
 15 **ompt_work_workshare** as its *work_type* argument for each occurrence of a *workshare-end*
 16 **event** in that **thread**. The **callbacks** occur in the context of the **implicit task**. The **callbacks** have type
 17 signature **ompt_callback_work_t**.

18 Restrictions

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

- 20 • The only OpenMP **constructs** that may be **closely nested constructs** of a **workshare**
 21 **construct** are the **atomic**, **critical**, and **parallel** **constructs**.
- 22 • **Base language** statements that are encountered inside a **workshare construct** but that are
 23 not enclosed within a **parallel** or **atomic construct** that is nested inside the
 24 **workshare construct** must consist of only the following:
 - 25 – array assignments;
 - 26 – scalar assignments;
 - 27 – **FORALL** statements;
 - 28 – **FORALL** constructs;
 - 29 – **WHERE** statements;
 - 30 – **WHERE** constructs; and
 - 31 – **BLOCK** constructs that are **strictly structured blocks** associated with **directives**.
- 32 • All array assignments, scalar assignments, and masked array assignments that are
 33 encountered inside a **workshare construct** but are not nested inside a **parallel construct**
 34 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 16.6](#)
- `atomic` directive, see [Section 16.8.5](#)
- `critical` directive, see [Section 16.2](#)
- `parallel` directive, see [Section 11.2](#)
- `ompt_callback_work_t`, see [Section 20.5.2.5](#)
- `ompt_scope_endpoint_t`, see [Section 20.4.4.11](#)
- `ompt_work_t`, see [Section 20.4.4.16](#)



12.5 coexecute Construct

Name: <code>coexecute</code>	Association: block
Category: <code>executable</code>	Properties: <code>work-distribution</code> , <code>partitioned</code>

Binding

The `binding region` is the innermost enclosing `teams region`. The `binding thread set` is the set of `initial threads` executing the enclosing `teams region`.

Semantics

The `coexecute construct` divides the execution of the associated `structured block` into separate `units of work` and causes the `threads` of the `binding thread set` to share the work such that each `unit of work` is executed only once by one `thread`, in the context of its `implicit task`. No `implicit barrier` occurs at the end of a `coexecute region`.

An implementation must enforce ordering of statements that is required to maintain Fortran semantics. For example, the effects of each statement within the `structured block` must appear to occur before the execution of the subsequent statements, and the evaluation of the right hand side of an assignment must appear to complete prior to the effects of assigning to the left hand side.

The statements in the `coexecute construct` are divided into `units of work` as follows:

- For array expressions within each statement, including transformational array intrinsic functions that compute scalar values from arrays:
 - Evaluation of each element of the array expression, including any references to pure elemental `procedures`, is a `unit of work`.

– Evaluation of transformational array intrinsic functions may be subdivided into any number of **units of work**.

- For array assignment statements, assignment of each element is a **unit of work**.
- For scalar assignment statements, each assignment operation is a **unit of work**.

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

The **units of work** are assigned to the **binding thread set** that execute a **coexecute region** such that each **unit of work** is executed once.

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

Execution Model Events

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

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

Tool Callbacks

A **thread** dispatches a registered **ompt_callback_work callback** with **ompt_scope_begin** as its *endpoint* argument and **ompt_work_coexecute** as its *work_type* argument for each occurrence of a *coexecute-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_coexecute** as its *work_type* argument for each occurrence of a *coexecute-end* event in that **thread**. The **callbacks** occur in the context of the **implicit task**. The **callbacks** have type signature **ompt_callback_work_t**.

Restrictions

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

- The **coexecute construct** must be a **closely nested construct** inside a **teams construct**.
- No **explicit region** may be nested inside a **coexecute region**.
- Base language statements that are encountered inside a **coexecute** must consist of only the following:
 - array assignments;
 - scalar assignments; and
 - calls to pure and elemental **procedures**.

- All array assignments and scalar assignments that are encountered inside a **coexecute construct** must be intrinsic assignments.
- The **construct** must not contain any calls to **procedures** that are not pure and elemental.
- If a **threadprivate variable** or **groupprivate variable** is referenced inside a **coexecute region**, the behavior is unspecified.

Cross References

- **target** directive, see [Section 14.8](#)
- **teams** directive, see [Section 11.3](#)
- **ompt_callback_work_t**, see [Section 20.5.2.5](#)

Fortran

12.6 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 **associated 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 **collapsed iterations** will be executed in parallel by **threads** in the **team** in the context of their **implicit tasks**. The **collapsed iterations** are distributed across **threads** that already are assigned to 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 execution of the **collapsed iterations** of a given **chunk** is consistent with their sequential order.

At the beginning of each **collapsed iteration**, the loop iteration **variable** or the **variable** declared by *range-decl* of each **collapsed loop** has the value that it would have if the **collapsed loops** were executed sequentially.

The **schedule kind** is reproducible if one of the following conditions is true:

- The **order clause** is specified with the **reproducible order-modifier modifier**; or
- The **schedule clause** is specified with **static** as the *kind* argument but not with the **simd ordering-modifier** and the **order clause** is not specified with the **unconstrained order-modifier**.

OpenMP programs can only depend on which **thread** executes a particular **collapsed iteration** if the **schedule kind** is reproducible. Schedule reproducibility also determines the consistency with the execution of **constructs** with the same **schedule kind**.

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 [collapsed 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 [collapsed 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 kind](#) as shown in Table 12.1.

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 12.1: `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>	static
<code>ompt_work_loop_dynamic</code>	dynamic
<code>ompt_work_loop_guided</code>	guided
<code>ompt_work_loop_other</code>	implementation defined

Restrictions

Restrictions to the [worksharing-loop construct](#) are as follows:

- The [associated iteration space](#) 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 3.2.1](#)
- `nowait` clause, see [Section 16.6](#)
- `order` clause, see [Section 11.4](#)

- 1 • **schedule** clause, see [Section 12.6.3](#)
- 2 • **do** directive, see [Section 12.6.2](#)
- 3 • **for** directive, see [Section 12.6.1](#)
- 4 • Consistent Loop Schedules, see [Section 5.4.5](#)
- 5 • **ompt_callback_work_t**, see [Section 20.5.2.5](#)
- 6 • **ompt_scope_endpoint_t**, see [Section 20.4.4.11](#)
- 7 • **ompt_work_t**, see [Section 20.4.4.16](#)

C / C++

12.6.1 for Construct

<p>Name: <code>for</code></p> <p>Category: executable</p>	<p>Association: loop nest</p> <p>Properties: work-distribution, team-executed, partitioned, worksharing, worksharing-loop, cancellable, context-matching</p>
---	--

Separating directives

[scan](#)

Clauses

[allocate](#), [collapse](#), [firstprivate](#), [induction](#), [lastprivate](#), [linear](#), [nowait](#), [order](#), [ordered](#), [private](#), [reduction](#), [schedule](#)

Semantics

The [for](#) construct is a [worksharing-loop](#) construct.

Cross References

- **allocate** clause, see [Section 7.6](#)
- **collapse** clause, see [Section 5.4.3](#)
- **firstprivate** clause, see [Section 6.4.4](#)
- **induction** clause, see [Section 6.5.12](#)
- **lastprivate** clause, see [Section 6.4.5](#)
- **linear** clause, see [Section 6.4.6](#)
- **nowait** clause, see [Section 16.6](#)
- **order** clause, see [Section 11.4](#)
- **ordered** clause, see [Section 5.4.4](#)

- **private** clause, see [Section 6.4.3](#)
- **reduction** clause, see [Section 6.5.9](#)
- **schedule** clause, see [Section 12.6.3](#)
- **scan** directive, see [Section 6.6](#)
- Worksharing-Loop Constructs, see [Section 12.6](#)

C / C++

Fortran

12.6.2 do Construct

<p>Name: do Category: executable</p>	<p>Association: loop Properties: work-distribution, team-executed, partitioned, worksharing, worksharing-loop, cancellable, context-matching</p>
---	---

Separating directives

[scan](#)

Clauses

[allocate](#), [collapse](#), [firstprivate](#), [induction](#), [lastprivate](#), [linear](#), [nowait](#), [order](#), [ordered](#), [private](#), [reduction](#), [schedule](#)

Semantics

The [do construct](#) is a [worksharing-loop](#) construct.

Cross References

- **allocate** clause, see [Section 7.6](#)
- **collapse** clause, see [Section 5.4.3](#)
- **firstprivate** clause, see [Section 6.4.4](#)
- **induction** clause, see [Section 6.5.12](#)
- **lastprivate** clause, see [Section 6.4.5](#)
- **linear** clause, see [Section 6.4.6](#)
- **nowait** clause, see [Section 16.6](#)
- **order** clause, see [Section 11.4](#)
- **ordered** clause, see [Section 5.4.4](#)
- **private** clause, see [Section 6.4.3](#)
- **reduction** clause, see [Section 6.5.9](#)

- **schedule** clause, see [Section 12.6.3](#)
- **scan** directive, see [Section 6.6](#)
- Worksharing-Loop Constructs, see [Section 12.6](#)

Fortran

12.6.3 **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
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: directive-name	unique

Directives

do, **for**

Semantics

The **schedule** clause specifies how [collapsed iterations](#) of a [worksharing-loop construct](#) are divided into [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 variables](#) 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**, [chunks](#) of increasing [collapsed iteration](#) numbers are assigned to the [threads](#) of the [team](#) in a round-robin fashion in the order of the [thread number](#). Each [chunk](#) includes *chunk_size* [collapsed iterations](#), except possibly for the [chunk](#) that contains the sequentially last iteration, which may have fewer iterations. If *chunk_size* is not specified, the [collapsed iteration space](#) is divided into [chunks](#) that are approximately equal in size, and at most one [chunk](#) is distributed to each [thread](#).

1 If the *kind* argument is **dynamic**, each **thread** executes a **chunk**, then requests another **chunk**, until
2 no **chunks** remain to be assigned. Each **chunk** contains *chunk_size* **collapsed iterations**, except for
3 the **chunk** that contains the sequentially last iteration, which may have fewer iterations. If
4 *chunk_size* is not specified, it defaults to 1.

5 If the *kind* argument is **guided**, each **thread** executes a **chunk**, then requests another **chunk**, until
6 no **chunks** remain to be assigned. For a *chunk_size* of 1, the size of each **chunk** is proportional to
7 the number of unassigned **collapsed iterations** divided by the number of **threads** in the **team**,
8 decreasing to 1. For a *chunk_size* with value $k > 1$, the size of each **chunk** is determined in the
9 same way, with the restriction that the **chunks** do not contain fewer than k **collapsed iterations**
10 (except for the **chunk** that contains the sequentially last iteration, which may have fewer than k
11 iterations). If *chunk_size* is not specified, it defaults to 1.

12 If the *kind* argument is **auto**, the decision regarding scheduling is **implementation defined**. If the
13 **schedule** clause is not specified on a **worksharing-loop construct** then the effect is as if the
14 **schedule** clause was specified with **auto** as its *kind* argument.

15 If the *kind* argument is **runtime**, the decision regarding scheduling is deferred until runtime, and
16 the behavior is as if the **clause** specifies *kind*, *chunk-size* and *ordering-modifier* as set in the
17 *run-sched-var ICV*. If the **schedule** clause explicitly specifies any **modifiers** then they override
18 any corresponding **modifiers** that are specified in the *run-sched-var ICV*.

19 If the **simd chunk-modifier** is specified and the **canonical loop nest** is associated with a **SIMD**
20 **construct**, $new_chunk_size = \lceil chunk_size / simd_width \rceil * simd_width$ is the *chunk_size* for
21 all **chunks** except the first and last **chunks**, where *simd_width* is an **implementation defined** value.
22 The first **chunk** will have at least *new_chunk_size* **collapsed iterations** except if it is also the last
23 **chunk**. The last **chunk** may have fewer **collapsed iterations** than *new_chunk_size*. If the **simd**
24 **chunk-modifier** is specified and the **canonical loop nest** is not associated with a **SIMD construct**, the
25 **modifier** is ignored.

26
27 **Note** – For a **team** of p **threads** and **collapsed loops** of n **collapsed iterations**, let $\lceil n/p \rceil$ be the
28 integer q that satisfies $n = p * q - r$, with $0 \leq r < p$. One **compliant implementation** of the
29 **static schedule kind** (with no specified *chunk_size*) would behave as though *chunk_size* had
30 been specified with value q . Another **compliant implementation** would assign q **collapsed iterations**
31 to the first $p - r$ **threads**, and $q - 1$ **collapsed iterations** to the remaining r **threads**. This illustrates
32 why a **conforming program** must not rely on the details of a particular implementation.

33 A **compliant implementation** of the **guided schedule kind** with a *chunk_size* value of k would
34 assign $q = \lceil n/p \rceil$ **collapsed iterations** to the first available **thread** and set n to the larger of $n - q$
35 and $p * k$. It would then repeat this process until q is greater than or equal to the number of
36 remaining **collapsed iterations**, at which time the remaining iterations form the final **chunk**.
37 Another **compliant implementation** could use the same method, except with $q = \lceil n/(2p) \rceil$, and set
38 n to the larger of $n - q$ and $2 * p * k$.
39

1 If the **monotonic** *ordering-modifier* is specified then each **thread** executes the **chunks** that it is
 2 assigned in increasing **collapsed iteration** order. When the **nonmonotonic** *ordering-modifier* is
 3 specified then **chunks** may be assigned to **threads** in any order and the behavior of an application
 4 that depends on any execution order of the **chunks** is unspecified. If an *ordering-modifier* is not
 5 specified, the effect is as if the **monotonic** *ordering-modifier* is specified if the *kind* argument is
 6 **static** or an **ordered** clause is specified on the **construct**; otherwise, the effect is as if the
 7 **nonmonotonic** *ordering-modifier* is specified.

8 Restrictions

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

- 10 • The **schedule** clause cannot be specified if any of the **collapsed loops** is a **non-rectangular**
 11 **loop**.
- 12 • The value of the *chunk_size* expression must be the same for all **threads** in the **team**.
- 13 • If **runtime** or **auto** is specified for *kind*, *chunk_size* must not be specified.
- 14 • The **nonmonotonic** *ordering-modifier* cannot be specified if an **ordered** clause is
 15 specified on the same **construct**.

16 Cross References

- 17 • **ordered** clause, see [Section 5.4.4](#)
- 18 • **do** directive, see [Section 12.6.2](#)
- 19 • **for** directive, see [Section 12.6.1](#)
- 20 • *run-sched-var* ICV, see [Table 2.1](#)

21 12.7 distribute Construct

22 Name: distribute	Association: loop nest
Category: executable	Properties: work-distribution , partitioned

23 Clauses

24 **allocate**, **collapse**, **dist_schedule**, **firstprivate**, **induction**, **lastprivate**,
 25 **order**, **private**

26 Binding

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

Semantics

The **distribute** construct specifies that the **collapsed iterations** will be executed by the **initial teams** in the context of their **implicit tasks**. The **collapsed 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 **collapsed iteration**, the loop iteration **variable** or the **variable** declared by *range-decl* of each **collapsed loop** has the value that it would have if the set of **collapsed 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** modifier; or
- The **dist_schedule** clause is specified with **static** as the *kind* argument and the **order** clause is not specified with the **unconstrained order-modifier**.

OpenMP programs can only depend on which **team** executes a particular **collapsed 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 *distribute-begin* event occurs after an **initial task** encounters a **distribute** construct but before the **task** starts to execute the **structured block** of the **distribute region**.

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

The *distribute-chunk-begin* event occurs for each scheduled **chunk** of a **distribute region** before execution of any **collapsed iteration**.

Tool Callbacks

A **thread** dispatches a registered **ompt_callback_work** callback with **ompt_scope_begin** as its *endpoint* argument and **ompt_work_distribute** as its *work_type* argument for each occurrence of a *distribute-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_distribute** as its *work_type* argument for each occurrence of a *distribute-end* event in that **thread**. The **callbacks** occur in the context of the **implicit task**. The **callbacks** have type signature **ompt_callback_work_t**.

A **thread** dispatches a registered **ompt_callback_dispatch** callback for each occurrence of a *distribute-chunk-begin* event in that **thread**. The **callback** occurs in the context of the **initial task**. The **callback** has type signature **ompt_callback_dispatch_t**.

Restrictions

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

- The **associated iteration space** must be the same for all **teams** in the **league**.
- The **region** that corresponds to the **distribute** construct must be a **strictly nested region** of a **teams region**.
- A list item may appear in a **firstprivate** or **lastprivate** clause, but not in both.
- The **conditional** *lastprivate-modifier* must not be specified.
- All list items that appear in an **induction** clause must be **private variables** in the **enclosing context**.

Cross References

- **allocate** clause, see [Section 7.6](#)
- **collapse** clause, see [Section 5.4.3](#)
- **dist_schedule** clause, see [Section 12.7.1](#)
- **firstprivate** clause, see [Section 6.4.4](#)
- **induction** clause, see [Section 6.5.12](#)
- **lastprivate** clause, see [Section 6.4.5](#)
- **order** clause, see [Section 11.4](#)
- **private** clause, see [Section 6.4.3](#)
- **teams** directive, see [Section 11.3](#)
- Consistent Loop Schedules, see [Section 5.4.5](#)
- **ompt_callback_work_t**, see [Section 20.5.2.5](#)
- **ompt_work_t**, see [Section 20.4.4.16](#)

12.7.1 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

Modifiers

Name	Modifies	Type	Properties
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: directive-name	unique

Directives

distribute

Semantics

The **dist_schedule** clause specifies how **collapsed iterations** of a **distribute** construct are divided into **chunks**, and how these **chunks** are distributed among the **teams** of the **league**. If **chunk_size** is not specified, the **collapsed 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, **collapsed iterations** are divided into chunks of **chunk_size** iterations. The **chunk_size** expression is evaluated using the **original list items** of any **variables** that become **private variables** 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 their **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 **collapsed loops** is a **non-rectangular loop**.

Cross References

- **distribute** directive, see [Section 12.7](#)

12.8 loop Construct

Name: loop Category: executable	Association: loop nest Properties: work-distribution, team-executed, partitioned, 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 **collapsed iterations** 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 **collapsed 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 **collapsed iteration**, the loop iteration **variable** or the **variable** declared by *range-decl* of each **collapsed loop** has the value that it would have if the **collapsed loops** were 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 **collapsed iterations** to complete. The **collapsed iterations** are guaranteed to complete before the end of the **teams region**. If the **loop region** does not bind to a **teams region**, all **collapsed iterations** 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**.

Fortran

The **associated loop** may be a **DO CONCURRENT** loop.

Fortran

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 an **associated loop**.
- 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 **construct** then the **bind clause** must be present.
- If a **loop region** binds to a **teams region** or **parallel region**, it must be encountered by all **threads** in the **binding thread set** or by none of them.

Fortran

- If the **associated loop** is a **DO CONCURRENT** loop, neither the **data-sharing attribute clauses** nor the **collapse clause** may be specified.

Fortran

Cross References

- **bind** clause, see [Section 12.8.1](#)
- **collapse** clause, see [Section 5.4.3](#)
- **lastprivate** clause, see [Section 6.4.5](#)
- **order** clause, see [Section 11.4](#)
- **private** clause, see [Section 6.4.3](#)
- **reduction** clause, see [Section 6.5.9](#)
- **teams** directive, see [Section 11.3](#)
- Consistent Loop Schedules, see [Section 5.4.5](#)

12.8.1 bind Clause

Name: bind	Properties: unique
--------------------------	---------------------------

Arguments

Name	Type	Properties
<i>binding</i>	Keyword: parallel , teams , thread	<i>default</i>

Modifiers

Name	Modifies	Type	Properties
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: directive-name	unique

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 a **closely nested construct** of 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 a **strictly nested region** of 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 a **closely nested region** of a **simd region**.

Cross References

- **loop** directive, see [Section 12.8](#)
- **parallel** construct, see [Section 11.2](#)
- **teams** construct, see [Section 11.3](#).

13 Tasking Constructs

This chapter defines [directives](#) and concepts related to [explicit tasks](#).

13.1 `untied` Clause

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

Arguments

Name	Type	Properties
<i>can_change_threads</i>	expression of OpenMP logical type	constant, optional

Modifiers

Name	Modifies	Type	Properties
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: directive-name	unique

Directives

[task](#), [taskloop](#)

Semantics

If *can-change-threads* evaluates to true, the [untied clause](#) specifies that [tasks](#) generated by the [construct](#) on which it appears are [untied tasks](#), which means that any [thread](#) in the [binding thread set](#) can resume the [task region](#) after a suspension. If *can-change-threads* evaluates to false or 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 task](#) or an [included task](#), the [untied clause](#) is ignored and the [task](#) is tied. If *can-change-threads* is not specified, the effect is as if *can-change-threads* evaluates to true.

Cross References

- [task](#) directive, see [Section 13.6](#)
- [taskloop](#) directive, see [Section 13.7](#)

13.2 mergeable Clause

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

Arguments

Name	Type	Properties
<i>can_merge</i>	expression of OpenMP logical type	constant, optional

Modifiers

Name	Modifies	Type	Properties
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: directive-name	unique

Directives

`task`, `taskloop`

Semantics

If *can_merge* evaluates to true, the `mergeable` clause specifies that `tasks` generated by the `construct` on which it appears are `mergeable tasks`. If *can_merge* evaluates to false, the `mergeable` clause specifies that `tasks` generated by the `construct` on which it appears are not `mergeable tasks`. If *can_merge* is not specified, the effect is as if *can_merge* evaluates to true.

Cross References

- `task` directive, see [Section 13.6](#)
- `taskloop` directive, see [Section 13.7](#)

13.3 final Clause

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

Arguments

Name	Type	Properties
<i>finalize</i>	expression of OpenMP logical type	<i>default</i>

Modifiers

Name	Modifies	Type	Properties
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: directive-name	unique

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 included **final 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 13.6](#)
- **taskloop** directive, see [Section 13.7](#)

13.4 threadset Clause

Name: threadset	Properties: unique
------------------------	--------------------

Arguments

Name	Type	Properties
<i>set</i>	Keyword: omp_pool , omp_team	<i>default</i>

Modifiers

Name	Modifies	Type	Properties
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: directive-name	unique

Directives

task, **taskloop**

Semantics

The **threadset** clause specifies the set of **threads** that may execute **tasks** that are generated by the **construct** on which it appears. If the *set* argument is **omp_team**, the generated **tasks** may only be scheduled onto **threads** of the **current team**. If the *set* argument is **omp_pool**, the generated **tasks** may be scheduled onto **unassigned threads** of the current **OpenMP thread pool** in addition to **threads** of the **current team**. If the **threadset** clause is not specified on a **construct** on which it may appear, then the effect is as if the **threadset** clause was specified with **omp_team** as its *set* argument.

If the **encountering task** is a **final task**, the **threadset** clause is ignored.

Cross References

- **task** directive, see [Section 13.6](#)
- **taskloop** directive, see [Section 13.7](#)

13.5 priority Clause

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

Arguments

Name	Type	Properties
<i>priority-value</i>	expression of integer type	constant, non-negative

Modifiers

Name	Modifies	Type	Properties
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: directive-name	unique

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**. An **OpenMP program** that relies on the **task** execution order being determined by the *priority-value* may have **unspecified behavior**.

Cross References

- **task** directive, see [Section 13.6](#)
- **taskloop** directive, see [Section 13.7](#)
- *max-task-priority-var* **ICV**, see [Table 2.1](#)

13.6 task Construct

Name: <code>task</code> Category: <code>executable</code>	Association: block Properties: <code>parallelism-generating</code> , <code>thread-limiting</code> , <code>task-generating</code>
--	---

Clauses

affinity, **allocate**, **default**, **depend**, **detach**, **final**, **firstprivate**, **if**, **in_reduction**, **mergeable**, **priority**, **private**, **shared**, **threadset**, **untied**

Clause set

Properties: exclusive	Members: detach , mergeable
------------------------------	--

1 Binding

2 The **binding thread set** of the **task region** is the set of threads specified in the **threadset clause**.
3 A **task region** binds to the innermost enclosing **parallel region**.

4 Semantics

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

10 The **encountering thread** may immediately execute the **task**, or defer its execution. In the latter case,
11 any **thread** of the current **binding thread set** may be assigned the **task**. **Task completion** of the **task**
12 can be guaranteed using **task synchronization constructs** and **clauses**. If a **task construct** is
13 encountered during execution of an outer task, the generated **task region** that corresponds to this
14 **construct** is not a part of the outer **task region** unless the generated **task** is an **included task**.

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

18 A **thread** that encounters a **task scheduling point** within the **task region** may temporarily suspend
19 the **task region**.

20 The **task construct** includes a **task scheduling point** in the **task region** of its **generating task**,
21 immediately following the generation of the **explicit task**. Each **explicit task region** includes a **task**
22 **scheduling point** at the end of its associated **structured block**.

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

26 When an **if** clause is present on a **task construct** and the **if** clause expression evaluates to *false*,
27 an undeferred task is generated, and the encountering thread must suspend the current task region,
28 for which execution cannot be resumed until execution of the structured block that is associated
29 with the generated task is completed. The use of a variable in an **if** clause expression of a **task**
30 construct causes an implicit reference to the variable in all enclosing constructs. The **if** clause
31 expression is evaluated in the context outside of the **task construct**.

32 Execution Model Events

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

35 Tool Callbacks

36 A thread dispatches a registered **ompt_callback_task_create** callback for each occurrence
37 of a *task-create* event in the context of the encountering task. This callback has the type signature

1 `ompt_callback_task_create_t` and the *flags* argument indicates the task types shown in
 2 Table 13.1.

TABLE 13.1: `ompt_callback_task_create` Callback Flags Evaluation

Operation	Evaluates to true
<code>(flags & ompt_task_explicit)</code>	Always in the dispatched callback
<code>(flags & ompt_task_undeferred)</code>	If the task is an undeferred task
<code>(flags & ompt_task_final)</code>	If the task is a final task
<code>(flags & ompt_task_untied)</code>	If the task is an untied task
<code>(flags & ompt_task_mergeable)</code>	If the task is a mergeable task
<code>(flags & ompt_task_merged)</code>	If the task is a merged task

Cross References

3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21

- `affinity` clause, see [Section 13.6.1](#)
- `allocate` clause, see [Section 7.6](#)
- `default` clause, see [Section 6.4.1](#)
- `depend` clause, see [Section 16.9.5](#)
- `detach` clause, see [Section 13.6.2](#)
- `final` clause, see [Section 13.3](#)
- `firstprivate` clause, see [Section 6.4.4](#)
- `if` clause, see [Section 4.5](#)
- `in_reduction` clause, see [Section 6.5.11](#)
- `mergeable` clause, see [Section 13.2](#)
- `priority` clause, see [Section 13.5](#)
- `private` clause, see [Section 6.4.3](#)
- `shared` clause, see [Section 6.4.2](#)
- `threadset` clause, see [Section 13.4](#)
- `untied` clause, see [Section 13.1](#)
- Task Scheduling, see [Section 13.10](#)
- `omp_fulfill_event`, see [Section 19.11.1](#)
- `ompt_callback_task_create_t`, see [Section 20.5.2.7](#)

13.6.1 affinity Clause

Name: affinity	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: iterator-specifier OpenMP expression (repeatable)	unique
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: directive-name	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.

▼ C / C++ ▼

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

▲ C / C++ ▲

Cross References

- **task** directive, see [Section 13.6](#)
- **iterator** modifier, see [Section 4.2.6](#)

13.6.2 detach Clause

Name: detach	Properties: innermost-leaf, unique
---------------------	------------------------------------

Arguments

Name	Type	Properties
<i>event-handle</i>	variable of event_handle type	<i>default</i>

Modifiers

Name	Modifies	Type	Properties
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: directive-name	unique

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 an **aggregate variable** cannot appear in a **detach** clause.

Fortran

- *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**.

Fortran

Cross References

- **firstprivate** clause, see [Section 6.4.4](#).
- **task** directive, see [Section 13.6](#)

13.7 taskloop Construct

Name: <code>taskloop</code> Category: <code>executable</code>	Association: loop nest Properties: <code>parallelism-generating</code> , <code>task-generating</code>
--	--

Clauses

`allocate`, `collapse`, `default`, `final`, `firstprivate`, `grainsize`, `if`, `in_reduction`, `induction`, `lastprivate`, `mergeable`, `nogroup`, `num_tasks`, `priority`, `private`, `reduction`, `shared`, `threadset`, `untied`

Clause set synchronization-clause

Properties: <code>exclusive</code>	Members: <code>nogroup</code> , <code>reduction</code>
---	---

Clause set granularity-clause

Properties: <code>exclusive</code>	Members: <code>grainsize</code> , <code>num_tasks</code>
---	---

Binding

The `binding thread set` of the `taskloop region` is the set of `threads` specified in the `threadset clause`. A `taskloop region` binds to the innermost enclosing `parallel region`.

Semantics

When a `thread` encounters a `taskloop construct`, the `construct` partitions the `collapsed iterations` 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 identifier 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 identifier and `list items` is present. Thus, the generated `tasks` are participants of the reduction defined by the `task_reduction clause` that was applied to the implicit `taskgroup construct`.

If an `in_reduction clause` is present, the behavior is as if each generated `task` was defined by a `task construct` on which an `in_reduction clause` with the same reduction identifier and `list`



1 `items` is present. Thus, the generated `tasks` are participants of a reduction previously defined by a
2 `reduction scoping clause`.

3 If a `threadset clause` is present, the behavior is as if each generated `task` was defined by a `task`
4 `construct` on which a `threadset clause` with the same set of `threads` is present. Thus, the `binding`
5 `thread set` of the generated `tasks` is the same as that of the `taskloop region`.

6 If no `clause` from the `granularity-clause clause set` is present, the number of loop `tasks` generated
7 and the number of `logical iterations` assigned to these `tasks` is `implementation defined`.

8 At the beginning of each `logical iteration`, the `loop iteration variable` or the `variable` declared by
9 `range-decl` of each `collapsed loop` has the value that it would have if the `collapsed loops` were
10 executed sequentially.

11 When an `if clause` is present and the `if clause` expression evaluates to *false*, `underrred tasks` are
12 generated. The use of a `variable` in an `if clause` expression causes an implicit reference to the
13 `variable` in all enclosing `constructs`.

14  For `firstprivate variables` of class type, the number of invocations of copy constructors that
15 perform the initialization is `implementation defined`.
16 

17 When storage is shared by a `taskloop region`, the programmer must ensure, by adding proper
18 synchronization, that the storage does not reach the end of its lifetime before the `taskloop region`
19 and its `descendent tasks` complete their execution.

19 Execution Model Events

20 The `taskloop-begin event` occurs upon entering the `taskloop region`. A `taskloop-begin` will
21 precede any `task-create events` for the generated `tasks`. The `taskloop-end event` occurs upon
22 completion of the `taskloop region`.

23 `Events` for an implicit `taskgroup region` that surrounds the `taskloop region` are the same as
24 for the `taskgroup construct`.

25 The `taskloop-iteration-begin event` occurs at the beginning of each `logical iteration` of a `taskloop`
26 `region` before an `explicit task` executes the `logical iteration`. The `taskloop-chunk-begin event` occurs
27 before an `explicit task` executes any of its associated `logical iterations` in a `taskloop region`.

28 Tool Callbacks

29 A `thread` dispatches a registered `ompt_callback_work callback` for each occurrence of a
30 `taskloop-begin` and `taskloop-end event` in that `thread`. The `callback` occurs in the context of the
31 `encountering task`. The `callback` has type signature `ompt_callback_work_t`. The `callback`
32 receives `ompt_scope_begin` or `ompt_scope_end` as its `endpoint` argument, as appropriate,
33 and `ompt_work_taskloop` as its `work_type` argument.

34 A `thread` dispatches a registered `ompt_callback_dispatch callback` for each occurrence of a
35 `taskloop-iteration-begin` or `taskloop-chunk-begin event` in that `thread`.

1 The [callback](#) binds to the [explicit task](#) executing the [logical iterations](#). The [callback](#) has type
2 signature `ompt_callback_dispatch_t`.

3 **Restrictions**

4 Restrictions to the [taskloop](#) construct are as follows:

- 5 • The [reduction-modifier](#) must be **default**.
- 6 • The **conditional** [lastprivate-modifier](#) must not be specified.

7 **Cross References**

- 8 • **allocate** clause, see [Section 7.6](#)
- 9 • **collapse** clause, see [Section 5.4.3](#)
- 10 • **default** clause, see [Section 6.4.1](#)
- 11 • **final** clause, see [Section 13.3](#)
- 12 • **firstprivate** clause, see [Section 6.4.4](#)
- 13 • **grainsize** clause, see [Section 13.7.1](#)
- 14 • **if** clause, see [Section 4.5](#)
- 15 • **in_reduction** clause, see [Section 6.5.11](#)
- 16 • **induction** clause, see [Section 6.5.12](#)
- 17 • **lastprivate** clause, see [Section 6.4.5](#)
- 18 • **mergeable** clause, see [Section 13.2](#)
- 19 • **nogroup** clause, see [Section 16.7](#)
- 20 • **num_tasks** clause, see [Section 13.7.2](#)
- 21 • **priority** clause, see [Section 13.5](#)
- 22 • **private** clause, see [Section 6.4.3](#)
- 23 • **reduction** clause, see [Section 6.5.9](#)
- 24 • **shared** clause, see [Section 6.4.2](#)
- 25 • **threadset** clause, see [Section 13.4](#)
- 26 • **untied** clause, see [Section 13.1](#)
- 27 • **task** directive, see [Section 13.6](#)
- 28 • **taskgroup** directive, see [Section 16.4](#)
- 29 • Canonical Loop Nest Form, see [Section 5.4.1](#)
- 30 • `ompt_callback_dispatch_t`, see [Section 20.5.2.6](#)

- `ompt_callback_work_t`, see [Section 20.5.2.5](#)
- `ompt_scope_endpoint_t`, see [Section 20.4.4.11](#)
- `ompt_work_t`, see [Section 20.4.4.16](#)

13.7.1 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
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: directive-name	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 **logical 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 13.7](#)

13.7.2 num_tasks Clause

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

Arguments

Name	Type	Properties
<code>num-tasks</code>	expression of integer type	positive

Modifiers

Name	Modifies	Type	Properties
<i>prescriptiveness</i>	<code>num-tasks</code>	Keyword: strict	unique
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: directive-name	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 `taskloop` region 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 logical 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 13.7](#)

13.8 taskyield Construct

Name: <code>taskyield</code>	Association: none
Category: <code>executable</code>	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 13.10](#)

13.9 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 20.5.2.11](#)
- `ompt_callback_parallel_begin_t`, see [Section 20.5.2.3](#)
- `ompt_callback_parallel_end_t`, see [Section 20.5.2.4](#)
- `ompt_callback_thread_begin_t`, see [Section 20.5.2.1](#)

- 1 • `ompt_callback_thread_end_t`, see [Section 20.5.2.2](#)
- 2 • `ompt_task_flag_t`, see [Section 20.4.4.19](#)
- 3 • `ompt_thread_t`, see [Section 20.4.4.10](#)

4 13.10 Task Scheduling

5 Whenever a [thread](#) reaches a [task scheduling point](#), it may begin or resume execution of a [task](#) from
6 its [schedulable task set](#). An [idle thread](#) is treated as if it is always at a [task scheduling point](#). For
7 other [threads](#), [task scheduling points](#) are implied at the following locations:

- 8 • during the generation of an [explicit task](#);
- 9 • the point immediately following the generation of an [explicit task](#);
- 10 • after the point of completion of the [structured block](#) associated with a [task](#);
- 11 • in a [taskyield](#) region;
- 12 • in a [taskwait](#) region;
- 13 • at the end of a [taskgroup](#) region;
- 14 • in an implicit [barrier](#) region;
- 15 • in an explicit [barrier](#) region;
- 16 • during the generation of a [target](#) region;
- 17 • the point immediately following the generation of a [target](#) region;
- 18 • at the beginning and end of a [target data](#) region;
- 19 • in a [target update](#) region;
- 20 • in a [target enter data](#) region;
- 21 • in a [target exit data](#) region;
- 22 • in the `omp_target_memcpy` routine;
- 23 • in the `omp_target_memcpy_async` routine;
- 24 • in the `omp_target_memcpy_rect` routine;
- 25 • in the `omp_target_memcpy_rect_async` routine;
- 26 • in the `omp_target_memset` routine; and
- 27 • in the `omp_target_memset_async` routine.

28 When a [thread](#) encounters a [task scheduling point](#) it may do one of the following, subject to the [task](#)
29 scheduling constraints specified below:

- 1 • begin execution of a **tied task** in its **schedulable task set** ;
- 2 • resume the suspended **task region** of any **task** to which it is tied;
- 3 • begin execution of an **untied task** in its **schedulable task set** ; or
- 4 • resume the suspended **task region** of any **untied task** in its **schedulable task set** .

5 If more than one of the above choices is available, which one is chosen is unspecified.

6 *Task Scheduling Constraints* are as follows:

- 7 1. If any suspended **tasks** are tied to the **thread** and are not suspended in a **barrier region** , a new
8 explicit **tied task** may be scheduled only if it is a **descendent task** of all of those suspended
9 **tasks** . Otherwise, any new explicit **tied task** may be scheduled.
- 10 2. A **dependent task** shall not start its execution until its **task dependences** are fulfilled.
- 11 3. A **task** shall not be scheduled while another **task** has been scheduled but has not yet
12 completed, if they are **mutually exclusive tasks** .
- 13 4. A **task** shall not start or resume execution on an **unassigned thread** if it would result in the
14 total number of **free-agent threads** in the **OpenMP thread pool** exceeding
15 **free-agent-thread-limit-var** .

16 A program that relies on any other assumption about **task** scheduling is non-conforming.

17

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

22 A program must behave correctly and consistently with all conceivable scheduling sequences that
23 are compatible with the rules above.

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

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

32 The use of threadprivate variables and the use of locks or critical sections in an explicit task with an
33 **if** clause must take into account that when the **if** clause evaluates to *false*, the task is executed
34 immediately, without regard to *Task Scheduling Constraint 2*.
35

1 **Execution Model Events**

2 The *task-schedule* event occurs in a thread when the thread switches tasks at a task scheduling
3 point; no event occurs when switching to or from a merged task.

4 **Tool Callbacks**

5 A thread dispatches a registered **ompt_callback_task_schedule** callback for each
6 occurrence of a *task-schedule* event in the context of the task that begins or resumes. This callback
7 has the type signature **ompt_callback_task_schedule_t**. The argument *prior_task_status*
8 is used to indicate the cause for suspending the prior task. This cause may be the completion of the
9 prior task region, the encountering of a **taskyield** construct, or the encountering of an active
10 cancellation point.

11 **Cross References**

- 12
 - **ompt_callback_task_schedule_t**, see [Section 20.5.2.10](#)

14 Device Directives and Clauses

This chapter defines `constructs` and concepts related to `device` execution.

14.1 `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>

Modifiers

Name	Modifies	Type	Properties
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: directive-name	unique

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 8.8.2](#)
- `declare target` directive, see [Section 8.8.1](#)
- `groupprivate` directive, see [Section 6.12](#)

14.2 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
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: directive-name	unique

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 levels 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 8.6](#)
- `interop` directive, see [Section 15.1](#)
- `target` directive, see [Section 14.8](#)
- `target data` directive, see [Section 14.5](#)
- `target enter data` directive, see [Section 14.6](#)
- `target exit data` directive, see [Section 14.7](#)
- `target update` directive, see [Section 14.9](#)
- `target-offload-var` ICV, see [Table 2.1](#)

14.3 `thread_limit` Clause

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

Arguments

Name	Type	Properties
<code>threadlim</code>	expression of integer type	positive

Modifiers

Name	Modifies	Type	Properties
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: <code>directive-name</code>	unique

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 14.8](#)
- `teams` directive, see [Section 11.3](#)

14.4 Device Initialization

Execution Model Events

The *device-initialize* event occurs in a thread that begins initialization of OpenMP on the device, after OpenMP initialization of the device, 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 a construct to a device until a dispatched `ompt_callback_device_initialize` callback completes.
- No thread may offload execution of a construct to a device after a dispatched `ompt_callback_device_finalize` callback occurs.

Cross References

- `ompt_callback_device_finalize_t`, see [Section 20.5.2.20](#)
- `ompt_callback_device_initialize_t`, see [Section 20.5.2.19](#)
- `ompt_callback_device_load_t`, see [Section 20.5.2.21](#)
- `ompt_callback_device_unload_t`, see [Section 20.5.2.22](#)

14.5 target data Construct

Name: <code>target data</code> Category: <code>executable</code>	Association: <code>block</code> Properties: <code>device</code> , <code>device-affecting</code> , <code>data-mapping</code> , <code>map-entering</code> , <code>map-exiting</code> , <code>mapping-only</code>
---	---

Clauses

`device`, `if`, `map`, `use_device_addr`, `use_device_ptr`

Clause set `data-environment-clause`

Properties: <code>required</code>	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 `if-expression` evaluates to `false`, the `target device` is the `host device`. `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` and `use_device_addr` clauses 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 are associated with a `target enter data` construct, as described in Section 14.6.

The `events` associated with exiting a `target data` region are the same `events` as are associated with a `target exit data` construct, as described in Section 14.7.

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 14.6.

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 14.7.

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 14.2](#)
- **if** clause, see [Section 4.5](#)
- **map** clause, see [Section 6.8.3](#)
- **use_device_addr** clause, see [Section 6.4.10](#)
- **use_device_ptr** clause, see [Section 6.4.8](#)

14.6 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, **ICVs** with **data environment ICV scope**, 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 6.8.3](#)) occur when the
2 [target task](#) executes.

3 When an [if clause](#) is present and *if-expression* evaluates to *false*, the [target device](#) is the [host](#)
4 [device](#).

5 Execution Model Events

6 [Events](#) associated with a [target task](#) are the same as for the [task construct](#) defined in [Section 13.6](#).

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 [target enter](#)
11 [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 13.6](#); (*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 device](#). Similarly, a [thread](#) dispatches
20 a 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 device](#). 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 clauses](#).

30 Cross References

- 31 • [depend clause](#), see [Section 16.9.5](#)
- 32 • [device clause](#), see [Section 14.2](#)
- 33 • [if clause](#), see [Section 4.5](#)
- 34 • [map clause](#), see [Section 6.8.3](#)
- 35 • [nowait clause](#), see [Section 16.6](#)

- **task** directive, see [Section 13.6](#)
- **ompt_callback_target_emi_t** and **ompt_callback_target_t**, see [Section 20.5.2.26](#)

14.7 target exit data Construct

Name: target exit data 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, **ICVs** with **data environment ICV scope**, 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 6.8.3](#)) occur when the **target task** executes.

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

Execution Model Events

Events associated with a **target task** are the same as for the **task construct** defined in [Section 13.6](#).

The *target-exit-data-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-exit-data-begin event* is a *target-task-begin event*.

The *target-exit-data-end event* occurs after all other **events** associated with the **target exit data construct**.

Tool Callbacks

Callbacks associated with **events** for **target tasks** are the same as for the **task construct** defined in [Section 13.6](#); (*flags & ompt_task_target*) always evaluates to *true* in the dispatched **callback**.

A **thread** dispatches a registered **ompt_callback_target** or **ompt_callback_target_emi** **callback** with **ompt_scope_begin** as its *endpoint* argument and **ompt_target_exit_data** or **ompt_target_exit_data_nowait** if the **nowait clause** is present as its *kind* argument for each occurrence of a *target-exit-data-begin event* in that **thread** in the context of the **target task** on the **host device**. Similarly, a **thread** dispatches a registered **ompt_callback_target** or **ompt_callback_target_emi** **callback** with **ompt_scope_end** as its *endpoint* argument and **ompt_target_exit_data** or **ompt_target_exit_data_nowait** if the **nowait clause** is present as its *kind* argument for each occurrence of a *target-exit-data-end event* in that **thread** in the context of the **target task** on the **host device**. These **callbacks** have type signature **ompt_callback_target_t** or **ompt_callback_target_emi_t**, respectively.

Restrictions

Restrictions to the **target exit data construct** are as follows:

- At least one **map clause** must appear on the **directive**.
- All **map clauses** must be **map-exiting clauses**.

Cross References

- **depend** clause, see [Section 16.9.5](#)
- **device** clause, see [Section 14.2](#)
- **if** clause, see [Section 4.5](#)
- **map** clause, see [Section 6.8.3](#)
- **nowait** clause, see [Section 16.6](#)
- **task** directive, see [Section 13.6](#)
- **ompt_callback_target_emi_t** and **ompt_callback_target_t**, see [Section 20.5.2.26](#)

14.8 target Construct

Name: target Category: executable	Association: block Properties: parallelism-generating, team-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 attribute clauses and data-mapping attribute clauses on the target construct, ICVs with data environment ICV scope, 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 6.8.3) 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 if-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 Tool Callbacks

2 Callbacks associated with `events` for `target tasks` are the same as for the `task construct` defined in
3 Section 13.6; `(flags & omp_target)` always evaluates to `true` in the dispatched `callback`.

4 A `thread` dispatches a registered `omp_callback_target` or
5 `omp_callback_target_emi` `callback` with `omp_scope_begin` as its `endpoint`
6 argument and `omp_target` or `omp_target_nowait` if the `nowait` clause is present as its
7 `kind` argument for each occurrence of a `target-begin` event in that `thread` in the context of the `target`
8 `task` on the `host device`. Similarly, a `thread` dispatches a registered `omp_callback_target` or
9 `omp_callback_target_emi` `callback` with `omp_scope_end` as its `endpoint` argument
10 and `omp_target` or `omp_target_nowait` if the `nowait` clause is present as its `kind`
11 argument for each occurrence of a `target-end` event in that `thread` in the context of the `target task` on
12 the `host device`. These `callbacks` have type signature `omp_callback_target_t` or
13 `omp_callback_target_emi_t`, respectively.

14 A `thread` dispatches a registered `omp_callback_target_submit_emi` `callback` with
15 `omp_scope_begin` as its `endpoint` argument for each occurrence of a `target-submit-begin`
16 event in that `thread`. Similarly, a `thread` dispatches a registered
17 `omp_callback_target_submit_emi` `callback` with `omp_scope_end` as its `endpoint`
18 argument for each occurrence of a `target-submit-end` event in that `thread`. These `callbacks` have type
19 signature `omp_callback_target_submit_emi_t`.

20 A `thread` dispatches a registered `omp_callback_target_submit` `callback` for each
21 occurrence of a `target-submit-begin` event in that `thread`. The `callback` occurs in the context of the
22 `target task` and has type signature `omp_callback_target_submit_t`.

23 Restrictions

24 Restrictions to the `target construct` are as follows:

- 25 • `Device-affecting constructs`, other than `target constructs` for which the `ancestor`
26 `device-modifier` is specified, must not be encountered during execution of a `target region`.
- 27 • The result of an `omp_set_default_device`, `omp_get_default_device`, or
28 `omp_get_num_devices` routine called within a `target region` is unspecified.
- 29 • The effect of an access to a `threadprivate variable` in a `target region` is unspecified.
- 30 • If a `list item` in a `map clause` is a `structure` element, any other element of that `structure` that is
31 referenced in the `target construct` must also appear as a `list item` in a `map clause`.
- 32 • A `list item` in a `data-sharing attribute clause` that is specified on a `target construct` must not
33 have the same `base variable` as a `list item` in a `map clause` on the `construct`.
- 34 • A `variable` referenced in a `target region` but not the `target construct` that is not declared
35 in the `target region` must appear in a `declare target directive`.
- 36 • A `map-type` in a `map clause` must be `to`, `from`, `tofrom` or `alloc`.

- If a **device** clause is specified with the **ancestor** *device-modifier*, only the **device**, **firstprivate**, **private**, **defaultmap**, **nowait**, and **map** clauses may appear on the **construct** and no **constructs** or calls to **routines** are allowed inside the corresponding **target region**.
- **Memory allocators** that do not appear in a **uses_allocators** clause cannot appear as an **allocator** in an **allocate** clause or be used in the **target region** unless a **requires** directive with the **dynamic_allocators** clause is present in the same **compilation unit**.
- Any IEEE floating-point exception status flag, halting mode, or rounding mode set prior to a **target region** is unspecified in the **region**.
- Any IEEE floating-point exception status flag, halting mode, or rounding mode set in a **target region** is unspecified upon exiting the **region**.
- An **OpenMP program** must not rely on the value of a function address in a **target region** except for assignments, comparisons to zero and indirect calls.

C / C++

- Upon exit from a **target region**, the value of an **attached pointer** must not be different from the value when entering the **region**.

C / C++

C++

- The run-time type information (RTTI) of an object can only be accessed from the **device** on which it was constructed.
- Invoking a virtual member function of an object on a **device** other than the **device** on which the object was constructed results in **unspecified behavior**, unless the object is accessible and was constructed on the **host device**.
- If an object of polymorphic **class type** is destructed, virtual member functions of any previously existing corresponding objects in other **device data environments** must not be invoked.

C++

Fortran

- An **attached pointer** that is associated with a given pointer target must not be associated with a different pointer target upon exit from a **target region**.
- A reference to a coarray that is encountered on a **non-host device** must not be coindexed or appear as an actual argument to a **procedure** where the corresponding dummy argument is a coarray.
- If the allocation status of a **mapped variable** or a **list item** that appears in a **has_device_addr** clause that has the **ALLOCATABLE** attribute is unallocated on entry to a **target region**, the allocation status of the corresponding **variable** in the **device data environment** must be unallocated upon exiting the **region**.

- If the allocation status of a [mapped variable](#) or a [list item](#) that appears in a [has_device_addr](#) clause that has the **ALLOCATABLE** attribute is allocated on entry to a [target region](#), the allocation status and shape of the corresponding [variable](#) in the [device data environment](#) may not be changed, either explicitly or implicitly, in the [region](#) after entry to it.
- 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 associated upon entry to the [target 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](#).
- An [OpenMP 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 7.6](#)
- **defaultmap** clause, see [Section 6.8.6](#)
- **depend** clause, see [Section 16.9.5](#)
- **device** clause, see [Section 14.2](#)
- **firstprivate** clause, see [Section 6.4.4](#)
- **has_device_addr** clause, see [Section 6.4.9](#)
- **if** clause, see [Section 4.5](#)
- **in_reduction** clause, see [Section 6.5.11](#)
- **is_device_ptr** clause, see [Section 6.4.7](#)
- **map** clause, see [Section 6.8.3](#)
- **nowait** clause, see [Section 16.6](#)
- **private** clause, see [Section 6.4.3](#)
- **thread_limit** clause, see [Section 14.3](#)
- **uses_allocators** clause, see [Section 7.8](#)
- **target data** directive, see [Section 14.5](#)
- **task** directive, see [Section 13.6](#)

- `ompt_callback_target_emi_t` and `ompt_callback_target_t`, see [Section 20.5.2.26](#)
- `ompt_callback_target_submit_emi_t` and `ompt_callback_target_submit_t`, see [Section 20.5.2.28](#)

14.9 target update Construct

Name: <code>target update</code>	Association: none
Category: <code>executable</code>	Properties: <code>parallelism-generating</code> , <code>task-generating</code> , <code>device</code> , <code>device-affecting</code>

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, `ICVs` with `data environment ICV scope`, 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 `data-motion clauses` occur when the `target task` executes. When an `if` clause is present and `if-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 13.6](#).

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**.

Tool Callbacks

Callbacks associated with events for **target tasks** are the same as for the **task construct** defined in [Section 13.6](#); (*flags & ompt_task_target*) always evaluates to *true* in the dispatched **callback**.

A **thread** dispatches a registered **ompt_callback_target** or **ompt_callback_target_emi** **callback** with **ompt_scope_begin** as its *endpoint* argument and **ompt_target_update** or **ompt_target_update_nowait** if the **nowait clause** is present as its *kind* argument for each occurrence of a *target-update-begin* event in that **thread** in the context of the **target task** on the **host device**. Similarly, a **thread** dispatches a registered **ompt_callback_target** or **ompt_callback_target_emi** **callback** with **ompt_scope_end** as its *endpoint* argument and **ompt_target_update** or **ompt_target_update_nowait** if the **nowait clause** is present as its *kind* argument for each occurrence of a *target-update-end* event in that **thread** in the context of the **target task** on the **host device**. These **callbacks** have type signature **ompt_callback_target_t** or **ompt_callback_target_emi_t**, respectively.

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**.

Cross References

- **depend** clause, see [Section 16.9.5](#)
- **device** clause, see [Section 14.2](#)
- **from** clause, see [Section 6.9.2](#)

- 1 • **if** clause, see [Section 4.5](#)
- 2 • **nowait** clause, see [Section 16.6](#)
- 3 • **to** clause, see [Section 6.9.1](#)
- 4 • **task** directive, see [Section 13.6](#)
- 5 • **ompt_callback_target_emi_t** and **ompt_callback_target_t**, see
- 6 [Section 20.5.2.26](#)
- 7 • **ompt_callback_task_create_t**, see [Section 20.5.2.7](#)

15 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 19.12](#)

15.1 `interop` Construct

Name: <code>interop</code>	Association: none
Category: executable	Properties: device

Clauses

[depend](#), [destroy](#), [device](#), [init](#), [nowait](#), [use](#)

Clause set action-clause

Properties: required	Members: destroy , init , use
-----------------------------	--

Binding

The **binding task set** for an **interop region** is the **generating task**. The **interop region** binds to the **region** of the **generating task**.

Semantics

The **interop construct** retrieves interoperability properties from the OpenMP implementation to enable interoperability with **foreign execution contexts**. When an **interop construct** is encountered, the **encountering task** executes the **region**.

For each *action-clause*, the *interop-type* set is the set of *interop-type modifiers* specified for the *clause* if the *clause* is **init** or for the **init clause** that initialized the *interop-var* that is specified for the *clause* if the *clause* is not **init**.

If the *interop-type* set includes **targetsync**, an empty **mergeable task** is generated. If the **nowait clause** is not present on the **construct** then the **task** is also an **included task**. Any **depend clauses** that are present on the **construct** apply to the generated **task**.

The **interop construct** ensures an ordered execution of the generated **task** relative to **foreign tasks** executed in the **foreign execution context** through the foreign synchronization object that is accessible through the **targetsync** property. When the creation of the **foreign task** precedes the encountering of an **interop construct** in happens before order (see Section 1.4.5), the **foreign task** must complete execution before the generated **task** begins execution. Similarly, when the creation of a **foreign task** follows the encountering of an **interop construct** in happens before order, the **foreign task** must not begin execution until the generated **task** completes execution. No ordering is imposed between the **encountering thread** and either **foreign tasks** or OpenMP **tasks** by the **interop construct**.

If the *interop-type* set does not include **targetsync**, the **nowait clause** has no effect.

Restrictions

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

- A **depend clause** can only appear on the **directive** if the *interop-type* includes **targetsync**.
- Each *interop-var* may be specified for at most one *action-clause* of each **interop construct**.

Cross References

- **depend clause**, see Section 16.9.5
- **destroy clause**, see Section 4.6
- **device clause**, see Section 14.2
- **init clause**, see Section 15.1.2
- **nowait clause**, see Section 16.6
- **use clause**, see Section 15.1.3
- Interoperability Routines, see Section 19.12

15.1.1 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.

15.1.2 `init` Clause

Name: <code>init</code>	Properties: innermost-leaf
--------------------------------	-----------------------------------

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>interop-var</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>interop-var</i>	Keyword: target , targetsync	repeatable, required
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: directive-name	unique

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-preference modifier](#) 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 15.1](#)
- OpenMP Foreign Runtime Identifiers, see [Section 15.1.1](#)

15.1.3 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>

Modifiers

Name	Modifies	Type	Properties
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: directive-name	unique

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 15.1](#)

15.2 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](#);
- *nowait*, which specifies that the [construct](#) is asynchronous; and
- *is_device_ptr(list-item)*, which specifies that the *list-item* is a [device pointer](#) in the [construct](#).

The following [directives](#) may add properties to the set:

- **dispatch**.

The following [directives](#) may remove properties from the set:

- **declare variant**.

Cross References

- **dispatch** directive, see [Section 8.6](#)
- Declare Variant Directives, see [Section 8.5](#)

16 Synchronization Constructs and Clauses

A [synchronization construct](#) imposes an order on the completion of code executed by different [threads](#) through synchronizing [flushes](#) that are executed as part of the [region](#) that corresponds to the [construct](#). [Section 1.4.4](#) and [Section 1.4.6](#) describe synchronization through the use of synchronizing [flushes](#) and [atomic operations](#). [Section 16.8.7](#) defines the behavior of synchronizing [flushes](#) that are implied at various other locations in an [OpenMP program](#).

16.1 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 16.1.2](#)
- [atomic](#) directive, see [Section 16.8.5](#)
- [critical](#) directive, see [Section 16.2](#)
- `omp_init_lock_with_hint` and `omp_init_nest_lock_with_hint`, see [Section 19.9.2](#)

16.1.1 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 [synchronization hint](#) constants. The valid constants must include the following, which can be extended with [implementation defined](#) values:

C / C++

```
1 typedef enum omp_sync_hint_t {  
2     omp_sync_hint_none = 0x0,  
3     omp_sync_hint_uncontended = 0x1,  
4     omp_sync_hint_contended = 0x2,  
5     omp_sync_hint_nonspeculative = 0x4,  
6     omp_sync_hint_speculative = 0x8,  
7 } omp_sync_hint_t;
```

C / C++

Fortran

```
8 integer (kind=omp_sync_hint_kind), &  
9     parameter :: omp_sync_hint_none = &  
10        int(Z'0', kind=omp_sync_hint_kind)  
11 integer (kind=omp_sync_hint_kind), &  
12     parameter :: omp_sync_hint_uncontended = &  
13        int(Z'1', kind=omp_sync_hint_kind)  
14 integer (kind=omp_sync_hint_kind), &  
15     parameter :: omp_sync_hint_contended = &  
16        int(Z'2', kind=omp_sync_hint_kind)  
17 integer (kind=omp_sync_hint_kind), &  
18     parameter :: omp_sync_hint_nonspeculative = &  
19        int(Z'4', kind=omp_sync_hint_kind)  
20 integer (kind=omp_sync_hint_kind), &  
21     parameter :: omp_sync_hint_speculative = &  
22        int(Z'8', kind=omp_sync_hint_kind)
```

Fortran

23 [Synchronization hints](#) can be combined by using the + or | operators in C/C++ or the + operator in
24 Fortran. Combining `omp_sync_hint_none` with any other [synchronization hint](#) is equivalent to
25 specifying the other [synchronization hint](#).

26 The intended meaning of each [synchronization hint](#) is:

- 27 • `omp_sync_hint_uncontended`: low contention is expected in this operation, that is,
28 few [threads](#) are expected to perform the operation simultaneously in a manner that requires
29 synchronization;
- 30 • `omp_sync_hint_contended`: high contention is expected in this operation, that is,
31 many [threads](#) are expected to perform the operation simultaneously in a manner that requires
32 synchronization;
- 33 • `omp_sync_hint_speculative`: the programmer suggests that the operation should be
34 implemented using speculative techniques such as transactional memory; and

- `omp_sync_hint_nonspeculative`: the programmer suggests that the operation should not be implemented using speculative techniques such as transactional memory.

Note – Future OpenMP specifications may add additional [synchronization hints](#) to the `sync_hint` type. Implementers are advised to add [implementation defined synchronization hints](#) starting from the most significant bit of the type and to include the name of the implementation in the name of the added [synchronization hint](#) to avoid name conflicts with other OpenMP implementations.

Restrictions

Restrictions to the [synchronization hints](#) are as follows:

- The [synchronization hints](#) `omp_sync_hint_uncontended` and `omp_sync_hint_contended` may not be combined.
- The [synchronization hints](#) `omp_sync_hint_nonspeculative` and `omp_sync_hint_speculative` may not be combined.

16.1.2 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>

Modifiers

Name	Modifies	Type	Properties
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: directive-name	unique

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 16.8.5](#)
- **critical** directive, see [Section 16.2](#)
- Synchronization Hint Type, see [Section 16.1.1](#)

16.2 critical Construct

Name: <code>critical</code> Category: <code>executable</code>	Association: block Properties: thread-limiting , thread-exclusive
--	--

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 [regions](#) that correspond to any **critical construct** of a given name are executed as if only by a single [thread](#) at a time among all [threads](#) associated with the [contention group](#) that execute the [regions](#), 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 [OpenMP 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 in a **hint clause**, the **critical construct** must specify a name.
- The *hint-expr* that is specified in the **hint clause** on each **critical construct** 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 16.1.2](#)
- `ompt_callback_mutex_acquire_t`, see [Section 20.5.2.14](#)
- `ompt_callback_mutex_t`, see [Section 20.5.2.15](#)
- `ompt_mutex_t`, see [Section 20.4.4.17](#)

16.3 Barriers

16.3.1 barrier Construct

Name: <code>barrier</code>	Association: none
Category: <code>executable</code>	Properties: <code>team-executed</code>

Binding

The [binding thread set](#) for a `barrier` region is the [current team](#). A `barrier` region binds to the innermost enclosing [parallel region](#).

Semantics

The `barrier` construct specifies an [explicit barrier](#) at the point at which the `construct` appears. Unless the [binding region](#) is canceled, all [threads](#) of the [team](#) that executes that [binding region](#) must enter the `barrier` region and complete execution of all [explicit tasks](#) bound to that [binding region](#) before any of the [threads](#) continue execution beyond the `barrier`.

The `barrier` region includes an implicit [task scheduling point](#) in the [current task region](#).

Execution Model Events

The *explicit-barrier-begin* event occurs in each [thread](#) that encounters the `barrier` construct on entry to the `barrier` region.

The *explicit-barrier-wait-begin* event occurs when a [task](#) begins an interval of active or passive waiting in a `barrier` region.

The *explicit-barrier-wait-end* event occurs when a [task](#) ends an interval of active or passive waiting and resumes execution in a `barrier` region.

The *explicit-barrier-end* event occurs in each [thread](#) that encounters the `barrier` construct after the `barrier` synchronization on exit from the `barrier` region.

A *cancellation* event occurs if [cancellation](#) is activated at an implicit [cancellation point](#) in a `barrier` region.

1 Tool Callbacks

2 A [thread](#) dispatches a registered [ompt_callback_sync_region](#) [callback](#) with
3 [ompt_sync_region_barrier_explicit](#) as its *kind* argument and [ompt_scope_begin](#)
4 as its *endpoint* argument for each occurrence of an *explicit-barrier-begin* [event](#). Similarly, a [thread](#)
5 dispatches a registered [ompt_callback_sync_region](#) [callback](#) with
6 [ompt_sync_region_barrier_explicit](#) as its *kind* argument and [ompt_scope_end](#) as
7 its *endpoint* argument for each occurrence of an *explicit-barrier-end* [event](#). These [callbacks](#) occur
8 in the context of the [task](#) that encountered the [barrier](#) [construct](#) and have type signature
9 [ompt_callback_sync_region_t](#).

10 A [thread](#) dispatches a registered [ompt_callback_sync_region_wait](#) [callback](#) with
11 [ompt_sync_region_barrier_explicit](#) as its *kind* argument and [ompt_scope_begin](#)
12 as its *endpoint* argument for each occurrence of an *explicit-barrier-wait-begin* [event](#). Similarly, a
13 [thread](#) dispatches a registered [ompt_callback_sync_region_wait](#) [callback](#) with
14 [ompt_sync_region_barrier_explicit](#) as its *kind* argument and [ompt_scope_end](#) as
15 its *endpoint* argument for each occurrence of an *explicit-barrier-wait-end* [event](#). These [callbacks](#)
16 occur in the context of the [task](#) that encountered the [barrier](#) [construct](#) and have type signature
17 [ompt_callback_sync_region_t](#).

18 A [thread](#) dispatches a registered [ompt_callback_cancel](#) [callback](#) with
19 [ompt_cancel_detected](#) as its *flags* argument for each occurrence of a *cancellation* [event](#) in
20 that [thread](#). The [callback](#) occurs in the context of the [encountering](#) [task](#). The [callback](#) has type
21 signature [ompt_callback_cancel_t](#).

22 Restrictions

23 Restrictions to the [barrier](#) [construct](#) are as follows:

- 24 • Each [barrier](#) [region](#) must be encountered by all [threads](#) in a [team](#) or by none at all, unless
25 [cancellation](#) has been requested for the innermost enclosing [parallel](#) [region](#).
- 26 • The sequence of [worksharing](#) [regions](#) and [barrier](#) [regions](#) encountered must be the same
27 for every [thread](#) in a [team](#).

28 Cross References

- 29 • [ompt_callback_cancel_t](#), see [Section 20.5.2.18](#)
- 30 • [ompt_callback_sync_region_t](#), see [Section 20.5.2.13](#)
- 31 • [ompt_scope_endpoint_t](#), see [Section 20.4.4.11](#)
- 32 • [ompt_sync_region_t](#), see [Section 20.4.4.14](#)

33 16.3.2 Implicit Barriers

34 This section describes the [OMPT](#) [events](#) and [tool](#) [callbacks](#) associated with [implicit](#) [barriers](#), which
35 occur at the end of various [regions](#) as defined in the description of the [constructs](#) to which they
36 correspond. [Implicit](#) [barriers](#) are [task](#) [scheduling](#) [points](#). For a description of [task](#) [scheduling](#)
37 [points](#), associated [events](#), and [tool](#) [callbacks](#), see [Section 13.10](#).

Execution Model Events

The *implicit-barrier-begin* event occurs in each **task** that encounters an **implicit barrier** at the beginning of the **implicit barrier region**.

The *implicit-barrier-wait-begin* event occurs when a **task** begins an interval of active or passive waiting in an **implicit barrier region**.

The *implicit-barrier-wait-end* event occurs when a **task** ends an interval of active or waiting and resumes execution of an **implicit barrier region**.

The *implicit-barrier-end* event occurs in a **task** that encounters an **implicit barrier** after the **barrier** synchronization on exit from an **implicit barrier region**.

A *cancellation* event occurs if **cancellation** is activated at an implicit **cancellation point** in an **implicit barrier region**.

Tool Callbacks

A **thread** dispatches a registered **ompt_callback_sync_region** callback for each *implicit-barrier-begin* and *implicit-barrier-end* event. Similarly, a **thread** dispatches a registered **ompt_callback_sync_region_wait** callback for each *implicit-barrier-wait-begin* and *implicit-barrier-wait-end* event. All **callbacks** for **implicit barrier events** execute in the context of the **encountering task** and have type signature **ompt_callback_sync_region_t**.

For the **implicit barrier** at the end of a **worksharing construct**, the *kind* argument is **ompt_sync_region_barrier_implicit_workshare**. For the **implicit barrier** at the end of a **parallel region**, the *kind* argument is **ompt_sync_region_barrier_implicit_parallel**. For a **barrier** at the end of a **teams region**, the *kind* argument is **ompt_sync_region_barrier_teams**. For an extra **barrier** added by an OpenMP implementation, the *kind* argument is **ompt_sync_region_barrier_implementation**.

A **thread** dispatches a registered **ompt_callback_cancel** callback with **ompt_cancel_detected** as its *flags* argument for each occurrence of a *cancellation* event in that **thread**. The **callback** occurs in the context of the **encountering task**. The **callback** has type signature **ompt_callback_cancel_t**.

Restrictions

Restrictions to **implicit barriers** are as follows:

- If a **thread** is in the state **ompt_state_wait_barrier_implicit_parallel**, a call to **ompt_get_parallel_info** may return a pointer to a copy of the data object associated with the **parallel region** rather than a pointer to the associated data object itself. Writing to the data object returned by **ompt_get_parallel_info** when a **thread** is in the state **ompt_state_wait_barrier_implicit_parallel** results in **unspecified behavior**.

Cross References

- `ompt_callback_cancel_t`, see [Section 20.5.2.18](#)
- `ompt_callback_sync_region_t`, see [Section 20.5.2.13](#)
- `ompt_cancel_flag_t`, see [Section 20.4.4.26](#)
- `ompt_scope_endpoint_t`, see [Section 20.4.4.11](#)
- `ompt_sync_region_t`, see [Section 20.4.4.14](#)

16.3.3 Implementation-Specific Barriers

An OpenMP implementation can execute implementation-specific [barriers](#) that the OpenMP specification does not imply; therefore, no execution model [events](#) are bound to them. The implementation can handle these [barriers](#) like [implicit barriers](#) and dispatch all [events](#) as for [implicit barriers](#). Any [callbacks](#) for these [events](#) use `ompt_sync_region_barrier_implementation` — or `ompt_sync_region_barrier`, if the implementation cannot make a distinction — as the *kind* argument when they are dispatched.

16.4 taskgroup Construct

Name: <code>taskgroup</code> Category: executable	Association: <code>block</code> 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](#) dispatches a registered [ompt_callback_sync_region_wait callback](#) with [ompt_sync_region_taskgroup](#) as its *kind* argument and [ompt_scope_end](#) as its *endpoint* argument for each occurrence of a *taskgroup-wait-end event*. These [callbacks](#) occur in the context of the [task](#) that encounters the [taskgroup construct](#) and have type signature [ompt_callback_sync_region_t](#).

Cross References

- [allocate](#) clause, see [Section 7.6](#)
- [task_reduction](#) clause, see [Section 6.5.10](#)
- Task Scheduling, see [Section 13.10](#)
- [ompt_callback_sync_region_t](#), see [Section 20.5.2.13](#)
- [ompt_scope_endpoint_t](#), see [Section 20.4.4.11](#)
- [ompt_sync_region_t](#), see [Section 20.4.4.14](#)

16.5 `taskwait` Construct

Name: <code>taskwait</code> Category: <code>executable</code>	Association: <code>none</code> Properties: <code>default</code>
--	--

Clauses

`depend`, `nowait`

Binding

The `binding thread set` of the `taskwait` region is the `current team`. The `taskwait` region binds to the `current task region`.

Semantics

The `taskwait` construct specifies a wait on the completion of `child tasks` of the `current task`.

If no `depend` clause is present on the `taskwait` construct, the `current task region` is suspended at an implicit `task scheduling point` associated with the `construct`. The `current task region` remains suspended until all `child tasks` that it generated before the `taskwait` region complete execution.

If one or more `depend` clauses are present on the `taskwait` construct and the `nowait` clause is not also present, the behavior is as if these clauses were applied to a `task` construct with an empty associated `structured block` that generates a `mergeable task` and `included task`. Thus, the `current task region` is suspended until the `predecessor tasks` of this `task` complete execution.

If one or more `depend` clauses are present on the `taskwait` construct and the `nowait` clause is also present, the behavior is as if these clauses were applied to a `task` construct with an empty associated `structured block` that generates a `task` for which execution may be deferred. Thus, all `predecessor tasks` of this `task` must complete execution before any subsequently generated `task` that 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` that corresponds 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` that corresponds 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

1 the `dependent task` and before any subsequently generated `task` that depends on the `dependent task`
2 starts its execution.

3 **Tool Callbacks**

4 A `thread` dispatches a registered `ompt_callback_sync_region_callback` with
5 `ompt_sync_region_taskwait` as its *kind* argument and `ompt_scope_begin` as its
6 *endpoint* argument for each occurrence of a *taskwait-begin event* in the `task` that encounters the
7 `taskwait construct`. Similarly, a `thread` dispatches a registered
8 `ompt_callback_sync_region_callback` with `ompt_sync_region_taskwait` as its
9 *kind* argument and `ompt_scope_end` as its *endpoint* argument for each occurrence of a
10 *taskwait-end event* in the `task` that encounters the `taskwait construct`. These `callbacks` occur in
11 the `task` that encounters the `taskwait construct` and have the type signature
12 `ompt_callback_sync_region_t`.

13 A `thread` dispatches a registered `ompt_callback_sync_region_wait_callback` with
14 `ompt_sync_region_taskwait` as its *kind* argument and `ompt_scope_begin` as its
15 *endpoint* argument for each occurrence of a *taskwait-wait-begin event*. Similarly, a `thread`
16 dispatches a registered `ompt_callback_sync_region_wait_callback` with
17 `ompt_sync_region_taskwait` as its *kind* argument and `ompt_scope_end` as its *endpoint*
18 argument for each occurrence of a *taskwait-wait-end event*. These `callbacks` occur in the context of
19 the `task` that encounters the `taskwait construct` and have type signature
20 `ompt_callback_sync_region_t`.

21 A `thread` dispatches a registered `ompt_callback_task_create_callback` for each occurrence
22 of a *taskwait-init event* in the context of the `encountering task`. This `callback` has the type signature
23 `ompt_callback_task_create_t`. In the dispatched `callback`, `(flags &`
24 `ompt_task_taskwait)` always evaluates to true. If the `nowait clause` is not present, `(flags &`
25 `ompt_task_undeferred)` also evaluates to true.

26 A `thread` dispatches a registered `ompt_callback_task_schedule_callback` for each
27 occurrence of a *taskwait-complete event*. This `callback` has the type signature
28 `ompt_callback_task_schedule_t` with `ompt_taskwait_complete` as its
29 *prior_task_status* argument.

30 **Restrictions**

31 Restrictions to the `taskwait construct` are as follows:

- 32 • The `mutexinoutset task-dependence-type` may not appear in a `depend clause` on a
33 `taskwait construct`.
- 34 • If the `task-dependence-type` of a `depend clause` is `depobj` then the depend objects may not
35 represent dependences of the `mutexinoutset` dependence type.
- 36 • The `nowait clause` may only appear on a `taskwait directive` if the `depend clause` is
37 present.

Cross References

- **depend** clause, see [Section 16.9.5](#)
- **nowait** clause, see [Section 16.6](#)
- **task** directive, see [Section 13.6](#)
- **ompt_callback_sync_region_t**, see [Section 20.5.2.13](#)
- **ompt_scope_endpoint_t**, see [Section 20.4.4.11](#)
- **ompt_sync_region_t**, see [Section 20.4.4.14](#)

16.6 **nowait** Clause

Name: <code>nowait</code>	Properties: outermost-leaf, unique, end-clause
----------------------------------	---

Arguments

Name	Type	Properties
<code>do_not_synchronize</code>	expression of OpenMP logical type	optional

Modifiers

Name	Modifies	Type	Properties
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: directive-name	unique

Directives

[dispatch](#), [do](#), [for](#), [interop](#), [scope](#), [sections](#), [single](#), [target](#), [target enter data](#), [target exit data](#), [target update](#), [taskwait](#), [workshare](#)

Semantics

If `do_not_synchronize` evaluates to true, the **nowait clause** overrides any synchronization that would otherwise occur at the end of a [construct](#). It can also specify that an [interoperability requirement set](#) includes the `nowait` [property](#). If `do_not_synchronize` is not specified, the effect is as if `do_not_synchronize` evaluates to true. If `do_not_synchronize` evaluates to false, the effect is as if the **nowait clause** is not specified on the [directive](#).

If the [construct](#) includes an [implicit barrier](#) and `do_not_synchronize` evaluates to true, the **nowait clause** specifies that the [barrier](#) will not occur. If the [construct](#) includes an [implicit barrier](#) and the **nowait** is not specified, the [barrier](#) will occur.

For [constructs](#) that generate a [task](#), if `do_not_synchronize` evaluates to true, the **nowait clause** specifies that the generated [task](#) may be deferred. If the **nowait clause** is not specified on the [directive](#) then the generated [task](#) is an [included task](#) (so it executes synchronously in the context of the [encountering task](#)).

For [constructs](#) that generate an [interoperability requirement set](#), the `nowait` clause adds the *nowait* property to the set if *do-not-synchronize* evaluates to true.

Restrictions

Restrictions to the `nowait` clause are as follows:

- The *do_not_synchronize* argument must evaluate to the same value for all [threads](#) in the [binding thread set](#), if defined for the [construct](#) on which the `nowait` clause appears.
- The *do_not_synchronize* argument must evaluate to the same value for all [tasks](#) in the [binding task set](#), if defined for the [construct](#) on which the `nowait` clause appears.

Cross References

- `dispatch` directive, see [Section 8.6](#)
- `do` directive, see [Section 12.6.2](#)
- `for` directive, see [Section 12.6.1](#)
- `interop` directive, see [Section 15.1](#)
- `scope` directive, see [Section 12.2](#)
- `sections` directive, see [Section 12.3](#)
- `single` directive, see [Section 12.1](#)
- `target` directive, see [Section 14.8](#)
- `target enter data` directive, see [Section 14.6](#)
- `target exit data` directive, see [Section 14.7](#)
- `target update` directive, see [Section 14.9](#)
- `taskwait` directive, see [Section 16.5](#)
- `workshare` directive, see [Section 12.4](#)

16.7 nogroup Clause

Name: <code>nogroup</code>	Properties: outermost-leaf, unique
-----------------------------------	---

Arguments

Name	Type	Properties
<i>do_not_synchronize</i>	expression of OpenMP logical type	optional

Modifiers

Name	Modifies	Type	Properties
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: directive-name	unique

1 **Directives**

2 **taskloop**

3 **Semantics**

4 If *do_not_synchronize* evaluates to true, the **nogroup clause** overrides any implicit **taskgroup**
5 that would otherwise enclose the **construct**. If *do_not_synchronize* evaluates to false, the effect is as
6 if the **nogroup clause** is not specified on the **directive**. If *do_not_synchronize* is not specified, the
7 effect is as if *do_not_synchronize* evaluates to true.

8 **Cross References**

- 9
 - **taskloop** directive, see [Section 13.7](#)

10 **16.8 OpenMP Memory Ordering**

11 This sections describes **constructs** and **clauses** that support ordering of **memory** operations.

12 **16.8.1 *memory-order* Clauses**

13 **Clause groups**

<p>Properties: unique, exclusive, inarguable</p>	<p>Members:</p> <p>Clauses acq_rel, acquire, relaxed, release, seq_cst</p>
---	--

15 **Directives**

16 **atomic**, **flush**

17 **Semantics**

18 The *memory-order clause group* defines a set of **clauses** that indicate the **memory** ordering
19 requirements for the visibility of the effects of the **constructs** on which they may be specified.

20 **Cross References**

- 21
 - **atomic** directive, see [Section 16.8.5](#)

22
 - **flush** directive, see [Section 16.8.6](#)

23
 - OpenMP Memory Consistency, see [Section 1.4.6](#)

16.8.1.1 `acq_rel` Clause

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

Arguments

Name	Type	Properties
<i>use-<code>semantics</code></i>	expression of OpenMP logical type	constant, optional

Modifiers

Name	Modifies	Type	Properties
<i><code>directive-name-modifier</code></i>	<i>all arguments</i>	Keyword: <code>directive-name</code>	unique

Directives

`atomic`, `flush`

Semantics

If *use_*`semantics` evaluates to true, the `acq_rel` clause specifies for the `construct` to use acquire/release memory ordering semantics. If *use_*`semantics` evaluates to false, the effect is as if the `acq_rel` clause is not specified. If *use_*`semantics` is not specified, the effect is as if *use_*`semantics` evaluates to true.

Cross References

- `atomic` directive, see [Section 16.8.5](#)
- `flush` directive, see [Section 16.8.6](#)
- OpenMP Memory Consistency, see [Section 1.4.6](#)

16.8.1.2 `acquire` Clause

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

Arguments

Name	Type	Properties
<i>use_</i> <code>semantics</code>	expression of OpenMP logical type	constant, optional

Modifiers

Name	Modifies	Type	Properties
<i><code>directive-name-modifier</code></i>	<i>all arguments</i>	Keyword: <code>directive-name</code>	unique

Directives

`atomic`, `flush`

Semantics

If *use_semantics* evaluates to true, the **acquire clause** specifies for the **construct** to use acquire **memory** ordering semantics. If *use_semantics* evaluates to false, the effect is as if the **acquire clause** is not specified. If *use_semantics* is not specified, the effect is as if *use_semantics* evaluates to true.

Cross References

- **atomic** directive, see [Section 16.8.5](#)
- **flush** directive, see [Section 16.8.6](#)
- OpenMP Memory Consistency, see [Section 1.4.6](#)

16.8.1.3 relaxed Clause

Name: relaxed	Properties: unique
-----------------------------	---------------------------

Arguments

Name	Type	Properties
<i>use_semantics</i>	expression of OpenMP logical type	constant, optional

Modifiers

Name	Modifies	Type	Properties
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: directive-name	unique

Directives

atomic, **flush**

Semantics

If *use_semantics* evaluates to true, the **relaxed clause** specifies for the **construct** to use relaxed **memory** ordering semantics. If *use_semantics* evaluates to false, the effect is as if the **relaxed clause** is not specified. If *use_semantics* is not specified, the effect is as if *use_semantics* evaluates to true.

Cross References

- **atomic** directive, see [Section 16.8.5](#)
- **flush** directive, see [Section 16.8.6](#)
- OpenMP Memory Consistency, see [Section 1.4.6](#)

16.8.1.4 `release` Clause

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

Arguments

Name	Type	Properties
<i>use_semantics</i>	expression of OpenMP logical type	constant, optional

Modifiers

Name	Modifies	Type	Properties
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: directive-name	unique

Directives

`atomic`, `flush`

Semantics

If *use_semantics* evaluates to true, the `release` clause specifies for the `construct` to use release memory ordering semantics. If *use_semantics* evaluates to false, the effect is as if the `release` clause is not specified. If *use_semantics* is not specified, the effect is as if *use_semantics* evaluates to true.

Cross References

- `atomic` directive, see [Section 16.8.5](#)
- `flush` directive, see [Section 16.8.6](#)
- OpenMP Memory Consistency, see [Section 1.4.6](#)

16.8.1.5 `seq_cst` Clause

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

Arguments

Name	Type	Properties
<i>use_semantics</i>	expression of OpenMP logical type	constant, optional

Modifiers

Name	Modifies	Type	Properties
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: directive-name	unique

Directives

`atomic`, `flush`

Semantics

If *use_semantics* evaluates to true, the **seq_cst** clause specifies for the **construct** to use sequentially consistent **memory** ordering semantics. If *use_semantics* evaluates to false, the effect is as if the **seq_cst** clause is not specified. If *use_semantics* is not specified, the effect is as if *use_semantics* evaluates to true.

Cross References

- **atomic** directive, see [Section 16.8.5](#)
- **flush** directive, see [Section 16.8.6](#)
- OpenMP Memory Consistency, see [Section 1.4.6](#)

16.8.2 atomic Clauses

Clause groups

Properties: unique, exclusive	Members: Clauses read, update, write
--------------------------------------	---

Directives

atomic

Semantics

The *atomic* clause group defines a set of **clauses** that defines the semantics for which a **directive** enforces atomicity. If a **construct** accepts the *atomic* clause group and no member of the **clause group** is specified, the effect is as if the **update** clause is specified.

Cross References

- **atomic** directive, see [Section 16.8.5](#)

16.8.2.1 read Clause

Name: read	Properties: innermost-leaf, unique
--------------------------	---

Arguments

Name	Type	Properties
<i>use_semantics</i>	expression of OpenMP logical type	constant, optional

Modifiers

Name	Modifies	Type	Properties
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: directive-name	unique

Directives

atomic

Semantics

If *use_semantics* evaluates to true, the **read clause** specifies that the **atomic construct** has **atomic read** semantics, which read the value of the **shared variable** atomically. If *use_semantics* evaluates to false, the effect is as if the **read** clause is not specified. If *use_semantics* is not specified, the effect is as if *use_semantics* evaluates to true.

Cross References

- **atomic** directive, see [Section 16.8.5](#)

16.8.2.2 update Clause

Name: update	Properties: innermost-leaf, unique
----------------------------	---

Arguments

Name	Type	Properties
<i>use_semantics</i>	expression of OpenMP logical type	constant, optional

Modifiers

Name	Modifies	Type	Properties
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: directive-name	unique

Directives

atomic

Semantics

If *use_semantics* evaluates to true, the **update clause** specifies that the **atomic construct** has **atomic update** semantics, which read and write the value of the **shared variable** atomically. If *use_semantics* evaluates to false, the effect is as if the **update** is not specified. If *use_semantics* is not specified, the effect is as if *use_semantics* evaluates to true.

Cross References

- **atomic** directive, see [Section 16.8.5](#)

16.8.2.3 write Clause

Name: write	Properties: innermost-leaf, unique
---------------------------	---

Arguments

Name	Type	Properties
<i>use_semantics</i>	expression of OpenMP logical type	constant, optional

Modifiers

Name	Modifies	Type	Properties
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: directive-name	unique

Directives

atomic

Semantics

If *use_semantics* evaluates to true, the **write clause** specifies that the **atomic construct** has **atomic write** semantics, which write the value of the **shared variable** atomically. If *use_semantics* evaluates to false, the effect is as if the **write clause** is not specified. If *use_semantics* is not specified, the effect is as if *use_semantics* evaluates to true.

Cross References

- **atomic** directive, see [Section 16.8.5](#)

16.8.3 *extended-atomic* Clauses

Clause groups

Properties: unique	Members: Clauses capture, compare, fail, weak
---------------------------	--

Directives

atomic

Semantics

The *extended-atomic clause group* defines a set of **clauses** that extend the atomicity semantics specified by members of the *atomic clause group*.

Restrictions

Restrictions to the *extended-atomic clause group* are as follows:

- The **compare clause** may not be specified such that *use_semantics* evaluates to false if the **weak clause** is specified such that *use_semantics* evaluates to true.

Cross References

- *atomic* Clauses, see [Section 16.8.2](#)
- **atomic** directive, see [Section 16.8.5](#)

16.8.3.1 capture Clause

Name: <code>capture</code>	Properties: innermost-leaf, unique
-----------------------------------	---

Arguments

Name	Type	Properties
<i>use_semantics</i>	expression of OpenMP logical type	constant, optional

Modifiers

Name	Modifies	Type	Properties
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: directive-name	unique

Directives

atomic

Semantics

If *use_semantics* evaluates to true, the **capture clause** extends the semantics of the **atomic construct** to have **atomic captured update** semantics, which capture the value of the **shared variable** being updated atomically. If *use_semantics* evaluates to false, the value is not captured. If *use_semantics* is not specified, the effect is as if *use_semantics* evaluates to true.

Cross References

- **atomic** directive, see [Section 16.8.5](#)

16.8.3.2 compare Clause

Name: <code>compare</code>	Properties: innermost-leaf, unique
-----------------------------------	---

Arguments

Name	Type	Properties
<i>use_semantics</i>	expression of OpenMP logical type	constant, optional

Modifiers

Name	Modifies	Type	Properties
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: directive-name	unique

Directives

atomic

Semantics

If *use_semantics* evaluates to true, the **compare clause** extends the semantics of the **atomic construct** with **atomic conditional update** semantics so the **atomic update** is performed conditionally. If *use_semantics* evaluates to false, the **atomic update** is performed unconditionally. If *use_semantics* is not specified, the effect is as if *use_semantics* evaluates to true.

Cross References

- `atomic` directive, see [Section 16.8.5](#)

16.8.3.3 fail Clause

Name: <code>fail</code>	Properties: innermost-leaf, unique
--------------------------------	---

Arguments

Name	Type	Properties
<i>memorder</i>	Keyword: <code>acquire</code> , <code>relaxed</code> , <code>seq_cst</code>	<i>default</i>

Modifiers

Name	Modifies	Type	Properties
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: <code>directive-name</code>	unique

Directives

`atomic`

Semantics

The `fail` clause extends the semantics of the `atomic` construct 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 an `atomic` construct has `atomic conditional update` semantics and the `fail` clause is not specified, 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. If the `atomic` construct does not have `atomic conditional update` semantics, the `fail` clause has no effect.

Restrictions

Restrictions to the `fail` clause are as follows:

- *memorder* may not be `acq_rel` or `release`.

Cross References

- *memory-order* Clauses, see [Section 16.8.1](#)
- `atomic` directive, see [Section 16.8.5](#)

16.8.3.4 weak Clause

Name: <code>weak</code>	Properties: innermost-leaf, unique
--------------------------------	---

Arguments

Name	Type	Properties
<i>use_semantics</i>	expression of OpenMP logical type	constant, optional

Modifiers

Name	Modifies	Type	Properties
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: directive-name	unique

Directives

atomic

Semantics

If *use_semantics* evaluates to true, the **weak clause** has the same effect as the **compare clause** and, in addition, the **atomic construct** has weak comparison semantics, which mean that the comparison may spuriously fail, evaluating to not equal even when the values are equal. If *use_semantics* evaluates to false, the semantics of the **atomic construct** are not extended. If *use_semantics* is not specified, the effect is as if *use_semantics* evaluates to true.

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.

Cross References

- **atomic** directive, see [Section 16.8.5](#)

16.8.4 memscope Clause

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

Arguments

Name	Type	Properties
<i>scope-specifier</i>	Keyword: all, cgroup, device	<i>default</i>

Modifiers

Name	Modifies	Type	Properties
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: directive-name	unique

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 **flushes** 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 16.8.5](#)
- **flush** directive, see [Section 16.8.6](#)

16.8.5 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 5.3.3](#)).

Binding

The **memscope** clause determines the **binding thread set** for an **atomic** **region**. If the **memscope** **clause** is not present, the behavior is as if the **memscope** **clause** appeared on the **construct** with the **device** *scope-specifier*.

Semantics

The **atomic construct** ensures that a specific **storage location** is accessed atomically so that possible simultaneous reads and writes by multiple **threads** do not result in indeterminate values. An **atomic region** enforces exclusive access with respect to other **atomic regions** that access the same **storage location** x among all **threads** in the **binding thread set** without regard to the **teams** to which the **threads** belong.

An **atomic construct** with the **read clause** results in an **atomic read** of the **storage location** designated by x . An **atomic construct** with the **write clause** results in an **atomic write** of the **storage location** designated by x . An **atomic construct** with the **update clause** results in an **atomic update** of the **storage location** designated by x using the designated operator or intrinsic. Only the read and write of the **storage location** designated by x are performed mutually atomically. The evaluation of $expr$ or $expr$ -list need not be atomic with respect to the read or write of the **storage location** designated by x . No **task scheduling points** are allowed between the read and the write of the **storage location** designated by x .

If the **capture clause** is present, the **atomic update** is an **atomic captured update** — an **atomic update** to the **storage location** designated by x using the designated operator or intrinsic while also capturing the original or final value of the **storage location** designated by x with respect to the **atomic update**. The original or final value of the **storage location** designated by x is written in the **storage location** designated by v based on the **base language** semantics of **structured block** or statements of the **atomic construct**. Only the read and write of the **storage location** designated by x are performed mutually atomically. Neither the evaluation of $expr$ or $expr$ -list, nor the write to the **storage location** designated by v , need be atomic with respect to the read or write of the **storage location** designated by x .

If the **compare clause** is present, the **atomic update** is an **atomic conditional update**. For forms that use an equality comparison, the operation is an atomic compare-and-swap. It atomically compares the value of x to e and writes the value of d into the **storage location** designated by x if they are equal. Based on the **base language** semantics of the associated **structured block**, the original or final value of the **storage location** designated by x is written to the **storage location** designated by v , which is allowed to be the same **storage location** as designated by e , or the result of the comparison is written to the **storage location** designated by r . Only the read and write of the **storage location** designated by x are performed mutually atomically. Neither the evaluation of either e or d nor writes to the **storage locations** designated by v and r need be atomic with respect to the read or write of the **storage location** designated by x .

▼ C / C++ ▼

If the **compare clause** is present, forms that use *ordop* are logically an atomic maximum or minimum, but they may be implemented with a compare-and-swap loop with short-circuiting. For forms where *statement* is *cond-expr-stmt*, if the result of the condition implies that the value of x 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-order clause** is
3 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 **storage location** designated by x is
6 performed on entry to and exit from the **atomic operation**, ensuring that the set of all **atomic**
7 **operations** applied to the same **storage location** in a race-free program has a total completion order.
8 If the **write** or **update clause** is specified, the **atomic operation** is not an **atomic conditional**
9 **update** for which the comparison fails, and the effective memory ordering is **release, acq_rel**,
10 or **seq_cst**, the **strong flush** on entry to the **atomic operation** is also a **release flush**. If the **read**
11 or **update clause** is specified and the effective memory ordering is **acquire, acq_rel**, or
12 **seq_cst** then the **strong flush** on exit from the **atomic operation** is also an **acquire flush**.
13 Therefore, if the effective memory ordering is not **relaxed**, **release flushes** and/or **acquire flushes**
14 are implied and permit synchronization between the **threads** without the use of explicit **flush**
15 **directives**.

16 For all forms of the **atomic construct**, any combination of two or more of these **atomic**
17 **constructs** enforces mutually exclusive access to the **storage locations** designated by x among
18 **threads** in the **binding thread set**. To avoid data races, all accesses of the **storage locations**
19 designated by x that could potentially occur in parallel must be protected with an **atomic**
20 **construct**.

21 **atomic regions** do not guarantee exclusive access with respect to any accesses outside of **atomic**
22 **regions** to the same **storage location** x even if those accesses occur during a **critical** or
23 **ordered region**, while an OpenMP lock is owned by the executing **task**, or during the execution
24 of a **reduction clause**.

25 However, other OpenMP synchronization can ensure the desired exclusive access. For example, a
26 **barrier** that follows a series of **atomic updates** to x guarantees that subsequent accesses do not form
27 a race with the atomic accesses.

28 A **compliant implementation** may enforce exclusive access between **atomic regions** that update
29 different **storage locations**. The circumstances under which this occurs are **implementation defined**.

30 If the **storage location** designated by x is not size-aligned (that is, if the byte alignment of x is not a
31 multiple of the size of x), then the behavior of the **atomic region** is **implementation defined**.

32 Execution Model Events

33 The *atomic-acquiring event* occurs in the **thread** that encounters the **atomic construct** on entry to
34 the **atomic region** before initiating synchronization for the **region**.

35 The *atomic-acquired event* occurs in the **thread** that encounters the **atomic construct** after it
36 enters the **region**, but before it executes the **structured block** of the **atomic region**.

37 The *atomic-released event* occurs in the **thread** that encounters the **atomic construct** after it
38 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 • **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 **OpenMP**
32 **program** are 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 [OpenMP 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 16.1.2](#)
- [memscope](#) clause, see [Section 16.8.4](#)
- [barrier](#) directive, see [Section 16.3.1](#)
- [critical](#) directive, see [Section 16.2](#)
- [flush](#) directive, see [Section 16.8.6](#)
- [requires](#) directive, see [Section 9.5](#)
- Lock Routines, see [Section 19.9](#)
- OpenMP Atomic Structured Blocks, see [Section 5.3.3](#)
- Synchronization Hints, see [Section 16.1](#)
- [ompt_callback_mutex_acquire_t](#), see [Section 20.5.2.14](#)
- [ompt_callback_mutex_t](#), see [Section 20.5.2.15](#)
- [ompt_mutex_t](#), see [Section 20.4.4.17](#)
- [ordered](#) Construct, see [Section 16.10](#)

16.8.6 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 the **temporary**
7 **view of memory** of a **thread** 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 **flush**
11 in order to be guaranteed to observe the effects of the **flush** of the **encountering thread**. See the
12 **memory** model description in **Section 1.4** and the **memscope** clause description in **Section 16.8.4**
13 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 **storage**
17 **locations** that are accessible to the **thread** are flushed by a **strong flush**; that is, the **flush** has the
18 **strong flush** property. A **flush** construct with a **list** applies a **strong flush** to the items in the **list**,
19 and the **flush** does not complete until the operation is complete for all specified **list items**. An
20 implementation may implement a **flush** construct with a **list** by ignoring the **list** and treating it
21 the same as a **flush** construct with the **seq_cst** clause.

22 If no **list items** are specified, the **flush** operation has the **release flush** property and/or the **acquire**
23 **flush** property:

- 24 • If the **memory-order** clause is **seq_cst** or **acq_rel**, the **flush** is both a **release flush** and
25 an **acquire flush**.
- 26 • If the **memory-order** clause is **release**, the **flush** is a **release flush**.
- 27 • If the **memory-order** clause is **acquire**, the **flush** is an **acquire flush**.

▼ C / C++ ▼

28 If a pointer is present in the **list**, the pointer itself is flushed, not the **storage locations** to which the
29 pointer refers.

30 A **flush** construct without a **list** corresponds to a call to **atomic_thread_fence**, where the
31 argument is given by the identifier that results from prefixing **memory_order_** to the
32 **memory-order** clause name.

33 For a **flush** construct without a **list**, the generated **flush** region implicitly performs the
34 corresponding call to **atomic_thread_fence**. The behavior of an explicit call to
35 **atomic_thread_fence** that occurs in an **OpenMP program** and does not have the argument
36 **memory_order_consume** is as if the call is replaced by its corresponding **flush** construct.

▲ C / C++ ▲

Fortran

1 If the [list item](#) or a subobject of the [list item](#) has the **POINTER** attribute, the allocation or
2 association status of the **POINTER** item is flushed, but the pointer target is not. If the [list item](#) is of
3 type **C_PTR**, the [variable](#) is flushed, but the [storage location](#) that corresponds to that address is not
4 flushed. 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.

Fortran

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 16.8.4](#)
- 18 • **ompt_callback_flush_t**, see [Section 20.5.2.17](#)

19 16.8.7 Implicit Flushes

20 [Flushes](#) implied when executing an **atomic** region are described in [Section 16.8.5](#).

21 A [flush region](#) that corresponds to a **flush** directive with the **release** clause present is implied
22 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](#);
- 30 • Immediately before every [task scheduling point](#);
- 31 • At exit from the [task region](#) of each [implicit task](#);

- At exit from an **ordered region**, if a **threads clause** or a **doacross clause** with a **source task-dependence-type** is present, or if no **clauses** are present; and
- During a **cancel region**, if the *cancel-var ICV* is *true*.

For a **target construct**, the **thread-set** of an implicit **release flush** that is performed in a **target task** during the generation of the **target region** and that is performed on exit from the **initial task region** that implicitly encloses the **target region** consists of the **thread** that executes the **target task** and the **initial thread** that executes the **target region**.

A **flush region** that corresponds to a **flush directive** with the **acquire clause** present is implied at the following locations:

- During a **barrier region**;
- At exit from a **teams region**;
- At entry to a **critical region**;
- If the **region** causes the lock to be set, during:
 - an **omp_set_lock region**;
 - an **omp_test_lock region**;
 - an **omp_set_nest_lock region**; and
 - an **omp_test_nest_lock region**;
- Immediately after every **task scheduling point**;
- At entry to the **task region** of each **implicit task**;
- At entry to an **ordered region**, if a **threads clause** or a **doacross clause** with a **sink task-dependence-type** is present, or if no **clauses** are present; and
- Immediately before a **cancellation point**, if the *cancel-var ICV* is *true* and **cancellation** has been activated.

For a **target construct**, the **thread-set** of an implicit **acquire flush** that is performed in a **target task** following the generation of the **target region** or that is performed on entry to the **initial task region** that implicitly encloses the **target region** consists of the **thread** that executes the **target task** and the **initial thread** that executes the **target region**.

Note – A **flush region** is not implied at the following locations:

- At entry to **worksharing regions**; and
- At entry to or exit from **masked regions**.

1 The synchronization behavior of **implicit flushes** is as follows:

- 2 • When a **thread** executes an **atomic region** for which the corresponding **construct** has the
3 **release**, **acq_rel**, or **seq_cst** clause and specifies an **atomic operation** that starts a
4 given **release sequence**, the **release flush** that is performed on entry to the **atomic operation**
5 **synchronizes with** an **acquire flush** that is performed by a different **thread** and has an
6 associated **atomic operation** that reads a value written by a modification in the **release**
7 **sequence**.
- 8 • When a **thread** executes an **atomic region** for which the corresponding **construct** has the
9 **acquire**, **acq_rel**, or **seq_cst** clause and specifies an **atomic operation** that reads a
10 value written by a given modification, a **release flush** that is performed by a different **thread**
11 and has an associated **release sequence** that contains that modification **synchronizes with** the
12 **acquire flush** that is performed on exit from the **atomic operation**.
- 13 • When a **thread** executes a **critical region** that has a given *name*, the behavior is as if the
14 **release flush** performed on exit from the **region synchronizes with** the **acquire flush**
15 performed on entry to the next **critical region** with the same *name* that is performed by a
16 different **thread**, if it exists.
- 17 • When a **thread** team executes a **barrier region**, the behavior is as if the **release flush**
18 performed by each **thread** within the **region**, and the **release flush** performed by any other
19 **thread** upon fulfilling the *allow-completion event* for a **detachable task** bound to the binding
20 **parallel region** of the **region**, **synchronizes with** the **acquire flush** performed by all other
21 **threads** within the **region**.
- 22 • When a **thread** executes a **taskwait region** that does not result in the creation of a
23 **dependent task** and the **task** that encounters the corresponding **taskwait construct** has at
24 least one **child task**, the behavior is as if each **thread** that executes a **child task** that is
25 generated before the **taskwait region** performs a **release flush** upon completion of the
26 associated **structured block** of the **child task** that **synchronizes with** an **acquire flush**
27 performed in the **taskwait region**. If the **child task** is a **detachable task**, the **thread** that
28 fulfills its *allow-completion event* performs a **release flush** upon fulfilling the **event** that
29 **synchronizes with** the **acquire flush** performed in the **taskwait region**.
- 30 • When a **thread** executes a **taskgroup region**, the behavior is as if each **thread** that executes
31 a remaining **descendent task** performs a **release flush** upon completion of the associated
32 **structured block** of the **descendent task** that **synchronizes with** an **acquire flush** performed on
33 exit from the **taskgroup region**. If the **descendent task** is a **detachable task**, the **thread** that
34 fulfills its *allow-completion event* performs a **release flush** upon fulfilling the **event** that
35 **synchronizes with** the **acquire flush** performed in the **taskgroup region**.
- 36 • When a **thread** executes an **ordered region** that does not arise from a stand-alone
37 **ordered directive**, the behavior is as if the **release flush** performed on exit from the **region**
38 **synchronizes with** the **acquire flush** performed on entry to an **ordered region** encountered
39 in the next **collapsed iteration** to be executed by a different **thread**, if it exists.

- 1 • When a **thread** executes an **ordered region** that arises from a stand-alone **ordered**
2 **directive**, the behavior is as if the **release flush** performed in the **ordered region** from a
3 given source **doacross iteration synchronizes with** the **acquire flush** performed in all
4 **ordered regions** executed by a different **thread** that are waiting for dependences on that
5 **doacross iteration** to be satisfied.
- 6 • When a **team** begins execution of a **parallel region**, the behavior is as if the **release flush**
7 performed by the **primary thread** on entry to the **parallel region synchronizes with** the
8 **acquire flush** performed on entry to each **implicit task** that is assigned to a different **thread**.
- 9 • When an **initial thread** begins execution of a **target region** that is generated by a different
10 **thread** from a **target task**, the behavior is as if the **release flush** performed by the generating
11 **thread** in the **target task synchronizes with** the **acquire flush** performed by the **initial thread** on
12 entry to its **initial task region**.
- 13 • When an **initial thread** completes execution of a **target region** that is generated by a
14 different **thread** from a **target task**, the behavior is as if the **release flush** performed by the
15 **initial thread** on exit from its **initial task region synchronizes with** the **acquire flush** performed
16 by the generating **thread** in the **target task**.
- 17 • When a **thread** encounters a **teams construct**, the behavior is as if the **release flush**
18 performed by the **thread** on entry to the **teams region synchronizes with** the **acquire flush**
19 performed on entry to each **initial task** that is executed by a different **initial thread** that
20 participates in the execution of the **teams region**.
- 21 • When a **thread** that encounters a **teams construct** reaches the end of the **teams region**, the
22 behavior is as if the **release flush** performed by each different participating **initial thread** at
23 exit from its **initial task synchronizes with** the **acquire flush** performed by the **thread** at exit
24 from the **teams region**.
- 25 • When a **task** generates an **explicit task** that begins execution on a different **thread**, the
26 behavior is as if the **thread** that is executing the **generating task** performs a **release flush** that
27 **synchronizes with** the **acquire flush** performed by the **thread** that begins to execute the
28 **explicit task**.
- 29 • When an **underrdeferred task** completes execution on a given **thread** that is different from the
30 **thread** on which its **generating task** is suspended, the behavior is as if a **release flush**
31 performed by the **thread** that completes execution of the associated **structured block** of the
32 **underrdeferred task synchronizes with** an **acquire flush** performed by the **thread** that resumes
33 execution of the **generating task**.
- 34 • When a **dependent task** with one or more **predecessor tasks** begins execution on a given
35 **thread**, the behavior is as if each **release flush** performed by a different **thread** on completion
36 of the associated **structured block** of a **predecessor task synchronizes with** the **acquire flush**
37 performed by the **thread** that begins to execute the **dependent task**. If the **predecessor task** is a
38 **detachable task**, the **thread** that fulfills its *allow-completion event* performs a **release flush**
39 upon fulfilling the **event** that **synchronizes with** the **acquire flush** performed when the

- 1 dependent task begins to execute.
- 2 • When a **task** begins execution on a given **thread** and it is mutually exclusive with respect to
- 3 another **sibling task** that is executed by a different **thread**, the behavior is as if each **release**
- 4 **flush** performed on completion of the **sibling task** synchronizes with the **acquire flush**
- 5 performed by the **thread** that begins to execute the **task**.
- 6 • When a **thread** executes a **cancel region**, the *cancel-var ICV* is *true*, and **cancellation** is not
- 7 already activated for the specified **region**, the behavior is as if the **release flush** performed
- 8 during the **cancel region** synchronizes with the **acquire flush** performed by a different
- 9 **thread** immediately before a **cancellation point** in which that **thread** observes **cancellation** was
- 10 activated for the **region**.
- 11 • When a **thread** executes an **omp_unset_lock region** that causes the specified lock to be
- 12 unset, the behavior is as if a **release flush** is performed during the **omp_unset_lock**
- 13 **region** that synchronizes with an **acquire flush** that is performed during the next
- 14 **omp_set_lock** or **omp_test_lock region** to be executed by a different **thread** that
- 15 causes the specified lock to be set.
- 16 • When a **thread** executes an **omp_unset_nest_lock region** that causes the specified
- 17 nested lock to be unset, the behavior is as if a **release flush** is performed during the
- 18 **omp_unset_nest_lock region** that synchronizes with an **acquire flush** that is performed
- 19 during the next **omp_set_nest_lock** or **omp_test_nest_lock region** to be
- 20 executed by a different **thread** that causes the specified nested lock to be set.

21 16.9 OpenMP Dependences

22 This section describes **constructs** and **clauses** in OpenMP that support the specification and

23 enforcement of **dependences**. OpenMP supports two kinds of **dependences**: **task dependences**,

24 which enforce orderings between **tasks**; and **doacross dependences**, which enforce orderings

25 between **doacross iterations** of a loop.

26 16.9.1 *task-dependence-type* Modifier

27 Modifiers

Name	Modifies	Type	Properties
<i>task-dependence-type</i>	<i>locator-list</i>	Keyword: depobj , in , inout , inoutset , mutexinoutset , out	required, ultimate

29 Clauses

30 **depend**, **update**

Semantics

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 16.9.5](#).

Cross References

- [depend](#) clause, see [Section 16.9.5](#)
- [update](#) clause, see [Section 16.9.3](#)

16.9.2 Depend Objects

OpenMP depend objects can be used to supply user-computed [dependences](#) to [depend](#) clauses. Depend objects must be accessed only through the [depobj](#) construct or through the [depend](#) clause; [OpenMP programs](#) that otherwise access depend objects are [non-conforming programs](#).

A depend object can be in one of the following states: *uninitialized* or *initialized*. Initially, depend objects are in the uninitialized state.

16.9.3 update Clause

Name: <code>update</code>	Properties: innermost-leaf, unique
----------------------------------	---

Arguments

Name	Type	Properties
<i>task-dependence-type</i>	Keyword: <code>depobj</code> , <code>in</code> , <code>inout</code> , <code>inoutset</code> , <code>mutexinoutset</code> , <code>out</code>	<i>default</i>

Modifiers

Name	Modifies	Type	Properties
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: <code>directive-name</code>	unique

Directives

[depobj](#)

Semantics

The [update](#) clause sets the [dependence](#) type of a 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 16.9.4](#)
- `task-dependence-type` modifier, see [Section 16.9.1](#)

16.9.4 `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: depend , destroy , update
--	---

Binding

The [binding thread set](#) for a `depobj` region is the [encountering thread](#).

Semantics

The `depobj` construct initializes, updates or destroys a 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 specify a *locator-list* with only one [list item](#).
- 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 16.9.5](#)
- **destroy** clause, see [Section 4.6](#)
- **update** clause, see [Section 16.9.3](#)
- **task-dependence-type** modifier, see [Section 16.9.1](#)

16.9.5 depend Clause

Name: depend	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
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: directive-name	unique

Directives

depobj, **dispatch**, **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**.

1 For the **out** *task-dependence-type* and **inout** *task-dependence-type*, if the **storage location** of at
2 least one of the list items matches the **storage location** of a list item appearing in a **depend** clause
3 with an **in**, **out**, **inout**, **mutexinoutset**, or **inoutset** *task-dependence-type* on a **construct**
4 from which a **sibling task** was previously generated, then the generated **task** will be a **dependent**
5 **task** of that **sibling task**.

6 For the **mutexinoutset** *task-dependence-type*, if the **storage location** of at least one of the list
7 items matches the **storage location** of a list item appearing in a **depend** clause with an **in**, **out**,
8 **inout**, or **inoutset** *task-dependence-type* on a **construct** from which a **sibling task** was
9 previously generated, then the generated **task** will be a **dependent task** of that **sibling task**.

10 If a list item appearing in a **depend** clause with a **mutexinoutset** *task-dependence-type* on a
11 **task-generating construct** matches a list item appearing in a **depend** clause with a
12 **mutexinoutset** *task-dependence-type* on a different **task-generating construct**, and both
13 **constructs** generate **sibling tasks**, the **sibling tasks** will be **mutually exclusive tasks**.

14 For the **inoutset** *task-dependence-type*, if the **storage location** of at least one of the list items
15 matches the **storage location** of a list item appearing in a **depend** clause with an **in**, **out**, **inout**,
16 or **mutexinoutset** *task-dependence-type* on a **construct** from which a **sibling task** was
17 previously generated, then the generated **task** will be a **dependent task** of that **sibling task**.

18 When the *task-dependence-type* is **depobj**, the **task dependences** are derived from the **task**
19 **dependences** represented by the **depend** objects specified in the **depend** clause as if the **depend**
20 **clauses** of the **depobj** **constructs** were specified in the current **construct**.

21 The list items that appear in the **depend** clause may reference any *iterator-identifier* defined in its
22 *iterator* modifier.

23 The list items that appear in the **depend** clause may include **array sections** or the
24 **omp_all_memory** reserved locator.

Fortran

25 If a list item has the **ALLOCATABLE** attribute and its allocation status is unallocated, the behavior
26 is unspecified. If a list item has the **POINTER** attribute and its association status is disassociated or
27 undefined, the behavior is unspecified.

Fortran

C / C++

28 The list items that appear in a **depend** clause may use shape-operators.

C / C++

29

30 **Note** – The enforced **task dependence** establishes a synchronization of **memory** accesses
31 performed by a **dependent task** with respect to accesses performed by the **predecessor tasks**.
32 However, the programmer must properly synchronize with respect to other concurrent accesses that
33 occur outside of those **tasks**.

34

Execution Model Events

The *task-dependences* event occurs in a *thread* that encounters a *task-generating construct* or a *taskwait* construct with a *depend* clause immediately after the *task-create* event for the new *task* or the *taskwait-init* event.

The *task-dependence* event indicates an unfulfilled *dependence* for the generated *task*. This event occurs in a *thread* that observes the unfulfilled *dependence* before it is satisfied.

Tool Callbacks

A *thread* dispatches the `ompt_callback_dependences` callback for each occurrence of the *task-dependences* event to announce its *dependences* with respect to the list items in the *depend* clause. This callback has type signature `ompt_callback_dependences_t`.

A *thread* dispatches the `ompt_callback_task_dependence` callback for a *task-dependence* event to report a *dependence* between a *predecessor task* (*src_task_data*) and a *dependent task* (*sink_task_data*). This callback has type signature `ompt_callback_task_dependence_t`.

Restrictions

Restrictions to the *depend* clause are as follows:

- List items, other than reserved locators, used in *depend* clauses of the same *task* or *sibling tasks* must indicate identical *storage locations* or disjoint *storage locations*.
- List items used in *depend* clauses cannot be zero-length array sections.
- The `omp_all_memory` reserved locator can only be used in a *depend* clause with an *out* or *inout* *task-dependence-type*.
- *Array sections* cannot be specified in *depend* clauses with the `depobj` *task-dependence-type*.
- List items used in *depend* clauses with the `depobj` *task-dependence-type* must be expressions of the OpenMP *depend* type that correspond to depend objects in the initialized state.
- List items that are expressions of the OpenMP *depend* type can only be used in *depend* clauses with the `depobj` *task-dependence-type*.

Fortran

- A common block name cannot appear in a *depend* clause.

Fortran

C / C++

- A bit-field cannot appear in a *depend* clause.

C / C++

Cross References

- `depobj` directive, see [Section 16.9.4](#)
- `dispatch` directive, see [Section 8.6](#)
- `interop` directive, see [Section 15.1](#)
- `target` directive, see [Section 14.8](#)
- `target enter data` directive, see [Section 14.6](#)
- `target exit data` directive, see [Section 14.7](#)
- `target update` directive, see [Section 14.9](#)
- `task` directive, see [Section 13.6](#)
- `taskwait` directive, see [Section 16.5](#)
- Array Sections, see [Section 4.2.5](#)
- Array Shaping, see [Section 4.2.4](#)
- `iterator` modifier, see [Section 4.2.6](#)
- `task-dependence-type` modifier, see [Section 16.9.1](#)
- `ompt_callback_dependences_t`, see [Section 20.5.2.8](#)
- `ompt_callback_task_dependence_t`, see [Section 20.5.2.9](#)

16.9.6 doacross Clause

Name: <code>doacross</code>	Properties: required
------------------------------------	-----------------------------

Arguments

Name	Type	Properties
<i>iteration-specifier</i>	OpenMP iteration specifier	<i>default</i>

Modifiers

Name	Modifies	Type	Properties
<i>dependence-type</i>	<i>iteration-specifier</i>	Keyword: sink, source	required
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: directive-name	unique

Directives

[ordered-standalone](#)

Semantics

The **doacross** clause identifies **doacross dependences** that imply additional constraints on the scheduling of **doacross logical iterations** of a **doacross loop nest**. These constraints establish **dependences** only between **doacross iterations**. The *iteration-specifier* specifies a **doacross iteration** and is either a **loop-iteration vector** or uses the **omp_cur_iteration** keyword (see [Section 5.4.2](#)).

The **source dependence-type** specifies that the current **doacross iteration** is a **source iteration** and, thus, satisfies **doacross dependences** that arise from the current **doacross iteration**. If the **source dependence-type** is specified then the *iteration-specifier* argument is optional; if *iteration-specifier* is omitted, it is assumed to be **omp_cur_iteration**.

The **sink dependence-type** specifies the current **doacross iteration** is a **sink iteration** and, thus, has a **doacross dependence**, where *iteration-specifier* indicates the **doacross iteration** that satisfies the **dependence**. If *iteration-specifier* indicates a **doacross iteration** that does not occur in the **doacross 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 an *iteration-specifier* that does not indicate an earlier iteration of the **doacross iteration space**, deadlock may occur.

Restrictions

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

- If *iteration-specifier* is a **loop-iteration vector** and it has n elements, the innermost **loop-nest-associated construct** that encloses the **construct** on which the **clause** appears must specify an **ordered** clause for which the parameter value equals n .
- If *iteration-specifier* is specified with the **omp_cur_iteration** keyword and with **sink** as the **dependence-type** then it must be **omp_cur_iteration - 1**.
- If *iteration-specifier* is specified with **source** as the **dependence-type** then it must be **omp_cur_iteration**.
- If *iteration-specifier* is a **loop-iteration vector** and the **sink dependence-type** is specified then for each element, 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++

- If *iteration-specifier* is a **loop-iteration vector** and the **sink dependence-type** is specified then for each element, 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 5.4.4](#)
- `ordered` directive, see [Section 16.10.1](#)
- OpenMP Loop-Iteration Spaces and Vectors, see [Section 5.4.2](#)

16.10 `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

- `omp_t_callback_mutex_acquire_t`, see [Section 20.5.2.14](#)
- `omp_t_callback_mutex_t`, see [Section 20.5.2.15](#)

16.10.1 Stand-alone ordered Construct

Name: <code>ordered</code>	Association: none
Category: <code>executable</code>	Properties: <code>default</code>

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 innermost enclosing [worksharing-loop construct](#) of a stand-alone `ordered` construct is associated with a `doacross` loop nest of n associated loops given by the argument in the `ordered` clause of that construct.

The stand-alone `ordered` construct specifies that execution must not violate [doacross dependences](#) as specified in the `doacross` clauses that appear on the construct. When a thread that is executing a `doacross` iteration encounters an `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 `doacross` 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 `doacross` 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_dependencies` [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_dependencies` [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_dependencies_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 16.9.6](#)
- Worksharing-Loop Constructs, see [Section 12.6](#)
- `ompt_callback_dependencies_t`, see [Section 20.5.2.8](#)

16.10.2 Block-associated `ordered` Construct

Name: <code>ordered</code> Category: executable	Association: <code>block</code> Properties: simdizable , thread-limiting , thread-exclusive
--	---

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 region](#), [simd region](#) 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 **collapsed 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 **collapsed iterations**. With either **parallelization-level**, execution of code outside the **region** for different **collapsed iterations** can run in parallel; execution of that code within the same **collapsed iteration** must observe any constraints imposed by the **base language semantics**.

When the **thread** that is executing the first **collapsed iteration** of the loop encounters a block-associated **ordered construct**, it can enter the **ordered region** without waiting. When a **thread** that is executing any subsequent **collapsed 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 **collapsed 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 clause** is specified.
- If the **simd parallelization-level clause** is specified, the **binding region** must be a **simd region** or one that corresponds to a **combined construct** or **composite construct** for which the **simd construct** is a **leaf construct**.
- If the **threads parallelization-level clause** is specified, the **binding region** must be a **worksharing-loop region** or one that corresponds to a **combined construct** or **composite construct** for which a **worksharing-loop construct** is a **leaf construct**.
- If the **threads parallelization-level clause** is specified and the **binding region** corresponds to a **combined construct** or **composite construct** then the **simd construct** must not be a **leaf construct** unless the **simd parallelization-level clause** is also specified.
- During execution of the **collapsed iteration** associated with a **loop-nest-associated directive**, a **thread** must not execute more than one block-associated **ordered region** that binds to the corresponding **region** of the **loop-nest-associated directive**.
- 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 16.10.3](#)
- **ordered clause**, see [Section 5.4.4](#)
- **simd directive**, see [Section 11.5](#)
- **Worksharing-Loop Constructs**, see [Section 12.6](#)

16.10.3 *parallelization-level* Clauses

Clause groups

Properties: unique	Members: Clauses simd , threads
---------------------------	---

Directives

[ordered-blockassoc](#)

Semantics

The *parallelization-level clause group* defines a set of [clauses](#) that indicate the level of parallelization with which to associate a [construct](#).

Cross References

- [ordered](#) directive, see [Section 16.10.2](#)

16.10.3.1 *threads* Clause

Name: threads	Properties: innermost-leaf, unique
--------------------------------------	---

Arguments

Name	Type	Properties
apply-to-threads	expression of OpenMP logical type	constant, optional

Modifiers

Name	Modifies	Type	Properties
directive-name-modifier	<i>all arguments</i>	Keyword: directive-name	unique

Directives

[ordered-blockassoc](#)

Semantics

If [apply_to_threads](#) evaluates to true, the effect is as if the [threads parallelization-level clause](#) is specified. If [apply_to_threads](#) evaluates to false, the effect is as if the [threads clause](#) is not specified. If [apply_to_threads](#) is not specified, the effect is as if [apply_to_threads](#) evaluates to true.

Cross References

- [ordered](#) directive, see [Section 16.10.2](#)

16.10.3.2 `simd` Clause

Name: <code>simd</code>	Properties: innermost-leaf, unique
-------------------------	------------------------------------

Arguments

Name	Type	Properties
<i>apply-to-simd</i>	expression of OpenMP logical type	constant, optional

Modifiers

Name	Modifies	Type	Properties
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: directive-name	unique

Directives

`ordered-blockassoc`

Semantics

If *apply_to_simd* evaluates to true, the effect is as if the `simd parallelization-level` clause is specified. If *apply_to_simd* evaluates to false, the effect is as if the `simd` clause is not specified. If *apply_to_simd* is not specified, the effect is as if *apply_to_simd* evaluates to true.

Cross References

- `ordered` directive, see [Section 16.10.2](#)

17 Cancellation Constructs

This chapter defines constructs related to cancellation of OpenMP regions.

17.1 *cancel-directive-name* Clauses

Clause groups

Properties: required, unique, exclusive	Members: Clauses do, for, parallel, sections, taskgroup
--	--

Modifiers

Name	Modifies	Type	Properties
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: directive-name	unique

Directives

cancel, **cancellation point**

Semantics

For each **directive** that has the **cancellable property** (i.e., the **directive** is a **cancellable construct**), a corresponding **clause** for which *clause-name* is the *directive-name* of that **directive** is a member of the *cancel-directive-name clause group*. Each member of the *cancel-directive-name clause group* takes an optional argument, *apply-to-directive*, that must be a constant expression of logical type. For each member of the **clause group**, if *apply_to_directive* evaluates to true then the semantics of the **construct** on which the **clause** appears are applied for the **directive** with the *directive-name* specified by the **clause**. If *apply_to_directive* evaluates to false, the effect is equivalent to specifying an **if clause** for which *if-expression* evaluates to false. If *apply_to_directive* is not specified, the effect is as if *apply_to_directive* evaluates to true.

Restrictions

Restrictions to any **clauses** in the *cancel-directive-name clause group* are as follows:

- If *apply_to_directive* evaluates to false and an **if clause** is specified for the same constituent **construct**, *if-expression* must evaluate to false.

Cross References

- `cancel` directive, see [Section 17.2](#)
- `cancellation point` directive, see [Section 17.3](#)
- `do` directive, see [Section 12.6.2](#)
- `for` directive, see [Section 12.6.1](#)
- `parallel` directive, see [Section 11.2](#)
- `sections` directive, see [Section 12.3](#)
- `taskgroup` directive, see [Section 16.4](#)

17.2 `cancel` Construct

Name: <code>cancel</code> Category: <code>executable</code>	Association: none Properties: <code>default</code>
--	---

Clause groups

cancel-directive-name

Clauses

if

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 If the canceled **construct** specifies a **reduction scoping clause** or **lastprivate clause**, the final
2 values of the **list items** that appear in those **clauses** are **undefined**.

3 When an **if clause** is present on a **cancel construct** and *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 *if-expression*.

6
7 **Note** – The programmer is responsible for releasing locks and other synchronization data structures
8 that might cause a deadlock when a **cancel construct** is encountered and blocked **threads** cannot
9 be canceled. The programmer is also responsible for ensuring proper synchronizations to avoid
10 deadlocks that might arise from **cancellation** of **regions** that contain synchronization **constructs**.
11

12 Execution Model Events

13 If a **task** encounters a **cancel construct** that will activate **cancellation** then a *cancel event* occurs.

14 A *discarded-task event* occurs for any discarded **tasks**.

15 Tool Callbacks

16 A **thread** dispatches a registered **ompt_callback_cancel callback** for each occurrence of a
17 *cancel event* in the context of the **encountering task**. This **callback** has type signature
18 **ompt_callback_cancel_t; (flags & ompt_cancel_activated)** always evaluates to
19 *true* in the dispatched **callback**; **(flags & ompt_cancel_parallel)** evaluates to *true* in the
20 dispatched **callback** if *cancel-directive-name* is **parallel**;
21 **(flags & ompt_cancel_sections)** evaluates to *true* in the dispatched **callback** if
22 *cancel-directive-name* is **sections**; **(flags & ompt_cancel_loop)** evaluates to *true* in the
23 dispatched **callback** if *cancel-directive-name* is **for** or **do**; and
24 **(flags & ompt_cancel_taskgroup)** evaluates to *true* in the dispatched **callback** if
25 *cancel-directive-name* is **taskgroup**.

26 A **thread** dispatches a registered **ompt_callback_cancel callback** with its *task_data*
27 argument pointing to the **ompt_data_t** object associated with the discarded **task** and with
28 **ompt_cancel_discarded_task** as its *flags* argument for each occurrence of a
29 *discarded-task event*. The **callback** occurs in the context of the **task** that discards the **task** and has
30 type signature **ompt_callback_cancel_t**.

31 Restrictions

32 Restrictions to the **cancel construct** are as follows:

- 33 • The behavior for concurrent **cancellation** of a **region** and a **region** nested within it is
34 unspecified.
- 35 • If *cancel-directive-name* is **taskgroup**, the **cancel construct** must be a **closely nested**
36 **construct** of a **task** or a **taskloop construct** and the **cancel region** must be a **closely**
37 **nested region** of a **taskgroup region**.

- 1 • If *cancel-directive-name* is not **taskgroup**, the **cancel** construct must be a **closely nested**
2 **construct** of a **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 reduction-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 reduction-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 • During execution of a **construct** that may be subject to **cancellation**, a **thread** must not
14 encounter an orphaned **cancellation point**. That is, a **cancellation point** must only be
15 encountered within that **construct** and must not be encountered elsewhere in its **region**.

16 **Cross References**

- 17 • **firstprivate** clause, see [Section 6.4.4](#)
- 18 • **if** clause, see [Section 4.5](#)
- 19 • **nowait** clause, see [Section 16.6](#)
- 20 • **ordered** clause, see [Section 5.4.4](#)
- 21 • **private** clause, see [Section 6.4.3](#)
- 22 • **reduction** clause, see [Section 6.5.9](#)
- 23 • **barrier** directive, see [Section 16.3.1](#)
- 24 • **cancellation point** directive, see [Section 17.3](#)
- 25 • **declare reduction** directive, see [Section 6.5.13](#)
- 26 • **task** directive, see [Section 13.6](#)
- 27 • *cancel-var* ICV, see [Table 2.1](#)
- 28 • **omp_get_cancellation**, see [Section 19.2.8](#)
- 29 • **ompt_callback_cancel_t**, see [Section 20.5.2.18](#)
- 30 • **ompt_cancel_flag_t**, see [Section 20.4.4.26](#)

17.3 cancellation point Construct

Name: <code>cancellation point</code>	Association: none
Category: <code>executable</code>	Properties: <code>default</code>

Clause groups

cancel-directive-name

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 task` 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`. The semantics, including the execution model events and tool callbacks, for when an `implicit task` or `explicit task` reaches a `user-defined cancellation point` are identical to those of any other `cancellation point` and are defined in Section 17.2.

Restrictions

Restrictions to the `cancellation point` construct are as follows:

- A `cancellation point` construct for which *cancel-directive-name* is `taskgroup` must be a `closely nested construct` of a `task` or `taskloop` construct, and the `cancellation point` region must be a `closely nested region` of a `taskgroup` region.
- A `cancellation point` construct for which *cancel-directive-name* is not `taskgroup` must be a `closely nested construct` inside a `construct` that matches *cancel-directive-name*.

Cross References

- *cancel-var* ICV, see Table 2.1
- `omp_get_cancellation`, see Section 19.2.8
- `ompt_callback_cancel_t`, see Section 20.5.2.18

18 Composition of Constructs

This chapter defines rules and mechanisms for nesting [regions](#) and for combining [constructs](#).

18.1 Nesting of Regions

This section describes a set of restrictions on the nesting of regions. The restrictions on nesting are as follows:

- A [team-executed region](#) may not be closely nested inside a [partitioned worksharing region](#), a [region](#) that corresponds to a [thread-exclusive construct](#), or a [region](#) that corresponds to a [task-generating construct](#) that is not to a [team-generating construct](#).
- 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**, **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.
- If a **target update**, **target data**, **target enter data**, or **target exit data** construct is encountered during execution of a **target** region, the behavior is unspecified.

- 1 • If a **target** construct is encountered during execution of a **target** region and a **device**
2 clause in which the **ancestor device-modifier** appears is not present on the construct, the
3 behavior is unspecified.
- 4 • A **teams** region must be strictly nested either within the implicit parallel region that
5 surrounds the whole OpenMP program or within a **target** region. If a **teams** construct is
6 nested within a **target** construct, that **target** construct must contain no statements,
7 declarations or directives outside of the **teams** construct.
- 8 • **distribute** regions, including any **distribute** regions arising from composite
9 constructs, **parallel** regions, including any **parallel** regions arising from combined
10 constructs, **loop** regions, **omp_get_num_teams()** regions, and
11 **omp_get_team_num()** regions are the only OpenMP regions that may be strictly nested
12 inside the **teams** region.
- 13 • A **loop** region that binds to a **teams** region must be strictly nested inside a **teams** region.
- 14 • A **distribute** region must be strictly nested inside a **teams** region.
- 15 • If *cancel-directive-name* is **taskgroup**, the **cancel** construct must be closely nested
16 inside a **task** construct and the **cancel** region must be closely nested inside a
17 **taskgroup** region. Otherwise, the **cancel** construct must be closely nested inside an
18 OpenMP construct for which *directive-name* is *cancel-directive-name*.
- 19 • A **cancellation point** construct for which *cancel-directive-name* is **taskgroup** must
20 be closely nested inside a **task** construct, and the **cancellation point** region must be
21 closely nested inside a **taskgroup** region. Otherwise, a **cancellation point**
22 construct must be closely nested inside an OpenMP construct for which *directive-name* is
23 *cancel-directive-name*.
- 24 • The only constructs that may be encountered inside a region that corresponds to a construct
25 with an **order** clause that specifies **concurrent** are the **loop**, **parallel** and **simd**
26 constructs, and combined constructs for which *directive-name-A* is **parallel**.
- 27 • A region that corresponds to a construct with an **order** clause that specifies **concurrent**
28 may not contain calls to the OpenMP Runtime API or to [procedures](#) that contain OpenMP
29 directives.

30 18.2 Clauses on Combined and Composite 31 Constructs

32 This section specifies the handling of [clauses](#) on [combined constructs](#) or [composite constructs](#) and
33 the handling of implicit [clauses](#) from variables with predetermined data sharing if they are not
34 predetermined only on a particular [construct](#). Some [clauses](#) are permitted only on a single [leaf](#)
35 [construct](#) of the [combined construct](#) or [composite construct](#), in which case the effect is as if the
36 [clause](#) is applied to that specific [construct](#). Other [clauses](#) that are permitted on more than one [leaf](#)

1 **construct** have the effect as if they are applied to a subset of those **construct**, as detailed in this
2 section. Unless otherwise specified, the effect of a **clause** on a **combined directive** or **composite**
3 **directive** is as if it is applied to all **leaf constructs** that permit it (i.e., it has the default
4 all-constituents property).

5 Unless otherwise specified, certain **clause** properties determine how each **clause** with those
6 properties applies to the constituents of **combined directives** and **composite directives**. Regardless
7 of any specified *directive-name-modifier*, the effect of any **clause** with the once-for-all-constituents
8 property on a **combined construct** or **composite construct** is as if it is applied once to the **combined**
9 **construct** or **composite construct** regardless of how many **constituent constructs** to which they may
10 apply. The effect of any **clause** with the all-privatizing property on a **combined directive** or
11 **composite directive** is as if it is applied to all **leaf constructs** that permit the **clause** and to which a
12 data-sharing attribute **clause** that may create a private copy of the same list item is applied. Unless
13 otherwise specified, the effect of any **clause** with the innermost-leaf property on a **combined**
14 **construct** or **composite construct** is as if it is applied only to the innermost **leaf construct** that
15 permits it. Unless otherwise specified, the effect of any **clause** with the outermost-leaf property on a
16 **combined construct** or **composite construct** is as if it is applied only to the outermost **leaf construct**
17 that permits it.

18 The effect of the **firstprivate** **clause** is as if it is applied to one or more **leaf constructs** as
19 follows:

- 20 • To the **distribute** **construct** if it is among the **constituent constructs**;
- 21 • To the **teams** **construct** if it is among the **constituent constructs** and the **distribute**
22 **construct** is not;
- 23 • To a **worksharing** **construct** that accepts the **clause** if one is among the **constituent constructs**;
- 24 • To the **taskloop** **construct** if it is among the **constituent constructs**;
- 25 • To the **parallel** **construct** if it is among the **constituent construct** and neither a
26 **taskloop** **construct** nor a **worksharing** **construct** that accepts the **clause** is among them;
- 27 • To the **target** **construct** if it is among the **constituent constructs** and the same list item
28 neither appears in a **lastprivate** **clause** nor is the **base variable** or **base pointer** of a list
29 item that appears in a **map** **clause**.

30 If the **parallel** **construct** is among the constituent constructs and the effect is not as if the
31 **firstprivate** **clause** is applied to it by the above rules, then the effect is as if the **shared**
32 **clause** with the same list item is applied to the **parallel** **construct**. If the **teams** **construct** is
33 among the constituent constructs and the effect is not as if the **firstprivate** **clause** is applied to
34 it by the above rules, then the effect is as if the **shared** **clause** with the same list item is applied to
35 the **teams** **construct**.

36 The effect of the **lastprivate** **clause** is as if it is applied to all leaf constructs that permit the
37 **clause**. If the **parallel** **construct** is among the constituent constructs and the list item is not also
38 specified in the **firstprivate** **clause**, then the effect of the **lastprivate** **clause** is as if the

1 **shared** clause with the same list item is applied to the **parallel** construct. If the **teams**
2 construct is among the constituent constructs and the list item is not also specified in the
3 **firstprivate** clause, then the effect of the **lastprivate** clause is as if the **shared** clause
4 with the same list item is applied to the **teams** construct. If the **target** construct is among the
5 constituent constructs and the list item is not the base variable or base pointer of a list item that
6 appears in a **map** clause, the effect of the **lastprivate** clause is as if the same list item appears
7 in a **map** clause with a *map-type* of **tofrom**.

8 The effect of the **reduction** clause is as if it is applied to all **leaf constructs** that permit the
9 **clause**, except for the following **constructs**:

- 10 • The **parallel** construct, when combined with the **sections**, worksharing-loop, **loop**,
11 or **taskloop** construct; and
- 12 • The **teams** construct, when combined with the **loop** construct.

13 For the **parallel** and **teams** constructs above, the effect of the **reduction** clause instead is as
14 if each list item or, for any list item that is an **array item**, its corresponding **base array** or
15 corresponding **base pointer** appears in a **shared** clause for the **construct**. If the **task**
16 **reduction-modifier** is specified, the effect is as if it only modifies the behavior of the **reduction**
17 **clause** on the innermost **leaf construct** that accepts the **modifier** (see Section 6.5.9). If the **inscan**
18 **reduction-modifier** is specified, the effect is as if it modifies the behavior of the **reduction** clause
19 on all **constructs** of the **combined construct** to which the **clause** is applied and that accept the
20 **modifier**. If a list item in a **reduction** clause on a **combined target construct** does not have the
21 same **base variable** or **base pointer** as a list item in a **map** clause on the **construct**, then the effect is
22 as if the list item in the **reduction** clause appears as a list item in a **map** clause with a *map-type*
23 of **tofrom**.

24 The effect of the **linear** clause is as if it is applied to the innermost leaf construct. Additionally,
25 if the list item is not the iteration variable of a **simd** or worksharing-loop SIMD construct, the
26 effect on the outer leaf constructs is as if the list item was specified in **firstprivate** and
27 **lastprivate** clauses on the combined or composite construct, with the rules specified above
28 applied. If a list item of the **linear** clause is the iteration variable of a **simd** or worksharing-loop
29 SIMD construct and it is not declared in the construct, the effect on the outer leaf constructs is as if
30 the list item was specified in a **lastprivate** clause on the combined or composite construct with
31 the rules specified above applied.

32 If the clauses have expressions on them, such as for various clauses where the argument of the
33 clause is an expression, or *lower-bound*, *length*, or *stride* expressions inside array sections (or
34 *subscript* and *stride* expressions in *subscript-triplet* for Fortran), or *linear-step* or *alignment*
35 expressions, the expressions are evaluated immediately before the construct to which the clause has
36 been split or duplicated per the above rules (therefore inside of the outer leaf constructs). However,
37 the expressions inside the **num_teams** and **thread_limit** clauses are always evaluated before
38 the outermost leaf construct.

39 The restriction that a list item may not appear in more than one data sharing clause with the
40 exception of specifying a variable in both **firstprivate** and **lastprivate** clauses applies

1 after the clauses are split or duplicated per the above rules.

2 **Restrictions**

3 Restrictions to clauses on combined and composite constructs are as follows:

- 4 • A clause that appears on a combined or composite construct must apply to at least one of the
5 leaf constructs per the rules defined in this section.

6 **18.3 Combined and Composite Directive Names**

7 **Combined directives** are shortcuts for specifying one **directive** immediately nested inside another
8 **directive**. **Composite directives** are also shortcuts for specifying the effect of one **directive**
9 immediately following the effect of another construct. However, **composite directives** define
10 semantics to combine **directive** that cannot otherwise be immediately nested.

11 For all combined and composite constructs, *directive-name* concatenates *directive-name-A*, the
12 directive name of the enclosing construct, with an intervening space followed by *directive-name-B*,
13 the directive name of the **nested construct**. If *directive-name-A* and *directive-name-B* both
14 correspond to loop-associated constructs then *directive-name* is a composite construct. Otherwise
15 *directive-name* is a combined construct.

16 If *directive-name-A* is **taskloop**, **for** or **do** then *directive-name-B* may be **simd**.

17 If *directive-name-A* is **masked** then *directive-name-B* may be **taskloop** or the directive name of
18 a combined or composite construct for which *directive-name-A* is **taskloop**.

19 If *directive-name-A* is **parallel** then *directive-name-B* may be **loop**, **sections**,
20 **workshare**, **masked**, **for**, **do** or the directive name of a combined or composite construct for
21 which *directive-name-A* is **masked**, **for** or **do**.

22 If *directive-name-A* is **distribute** then *directive-name-B* may be **simd** or the directive name of
23 a combined or composite construct for which *directive-name-A* is **parallel** and **for** or **do** is a
24 leaf construct.

25 If *directive-name-A* is **teams** then *directive-name-B* may be **loop**, **coexecute**, **distribute**
26 or the directive name of a combined or composite construct for which *directive-name-A* is
27 **distribute**.

28 If *directive-name-A* is **target** then *directive-name-B* may be **simd**, **parallel**, **teams**, the
29 directive name of a combined or composite construct for which *directive-name-A* is **teams** or the
30 directive name of a combined or composite construct for which *directive-name-A* is **parallel**
31 and **loop**, **for** or **do** is a leaf construct.

32 **Cross References**

- 33 • **coexecute** directive, see [Section 12.5](#)
- 34 • **distribute** directive, see [Section 12.7](#)

- 1 • **do** directive, see [Section 12.6.2](#)
- 2 • **for** directive, see [Section 12.6.1](#)
- 3 • **loop** directive, see [Section 12.8](#)
- 4 • **masked** directive, see [Section 11.6](#)
- 5 • **parallel** directive, see [Section 11.2](#)
- 6 • **sections** directive, see [Section 12.3](#)
- 7 • **target** directive, see [Section 14.8](#)
- 8 • **taskloop** directive, see [Section 13.7](#)
- 9 • **teams** directive, see [Section 11.3](#)
- 10 • **workshare** directive, see [Section 12.4](#)

11 18.4 Combined Construct Semantics

12 The semantics of the combined constructs are identical to that of explicitly specifying the first
13 construct containing one instance of the second construct and no other statements. All combined
14 and composite directives for which a loop-associated construct is a leaf construct are themselves
15 loop-associated constructs. For combined constructs, tool callbacks are invoked as if the constructs
16 were explicitly nested.

17 Restrictions

18 Restrictions to combined constructs are as follows:

- 19 • The restrictions of *directive-name-A* and *directive-name-B* apply.
- 20 • If *directive-name-A* is **parallel**, the **in_reduction** clause must not be specified.
- 21 • If *directive-name-A* is **parallel** and **target** is not among the constituent constructs, the
22 **nowait** clause must not be specified.
- 23 • If *directive-name-A* is **target**, the **copyin** clause must not be specified.

24 Cross References

- 25 • **copyin** clause, see [Section 6.7.1](#)
- 26 • **in_reduction** clause, see [Section 6.5.11](#)
- 27 • **nowait** clause, see [Section 16.6](#)
- 28 • **parallel** directive, see [Section 11.2](#)
- 29 • **target** directive, see [Section 14.8](#)

18.5 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 **collapsed iteration** 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 6.5.11](#)
- **linear** clause, see [Section 6.4.6](#)
- **ordered** clause, see [Section 5.4.4](#)
- **reduction** clause, see [Section 6.5.9](#)
- **distribute** directive, see [Section 12.7](#)
- **simd** directive, see [Section 11.5](#)
- **taskloop** directive, see [Section 13.7](#)

1

Part III

2

Runtime Library Routines

19 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

19.1 Runtime Library Definitions

For each base language, a compliant implementation must supply a set of definitions for the OpenMP API runtime library routines and the special data types of their parameters. The set of definitions must contain a declaration for each OpenMP API runtime library routine and variable and a definition of each required data type listed below. In addition, each set of definitions may specify other implementation specific values.


C / C++

The library routines are external functions with “C” linkage.

Prototypes for the C/C++ runtime library routines described in this chapter shall be provided in a header file named `omp.h`. This file also defines the following:

- The type `omp_allocator_handle_t`, which must be an implementation-defined (for C++ possibly scoped) enum type with at least the `omp_null_allocator` enumerator with the value zero and an enumerator for each predefined memory allocator in Table 7.3;
- `omp_atv_default`, which is an instance of a type compatible with `omp_uintptr_t` with the value -1;
- The type `omp_control_tool_result_t`;
- The type `omp_control_tool_t`;
- The type `omp_depend_t`;
- The type `omp_event_handle_t`, which must be an implementation-defined (for C++ possibly scoped) enum type;
- The enumerator `omp_initial_device` with value -1;
- The type `omp_interop_t`, which must be an implementation-defined integral or pointer type;
- The type `omp_interop_fr_t`, which must be an implementation-defined enum type with enumerators named `omp_ifr_name` where *name* is a foreign runtime name that is defined in the [OpenMP Additional Definitions document](#);
- The type `omp_intptr_t`, which is a signed integer type that is at least the size of a pointer on any device;
- The enumerator `omp_invalid_device` with an implementation-defined value less than -1;
- The type `omp_lock_t`;
- The type `omp_memspace_handle_t`, which must be an implementation-defined (for C++ possibly scoped) enum type with at least the `omp_null_mem_space` enumerator with the value zero and an enumerator for each predefined memory space in Table 7.1;

- 1 • The type `omp_nest_lock_t`;
- 2 • The type `omp_pause_resource_t`;
- 3 • The type `omp_proc_bind_t`;
- 4 • The type `omp_sched_t`;
- 5 • The type `omp_sync_hint_t`; and
- 6 • The type `omp_uintptr_t`, which is an unsigned integer type capable of holding a pointer
- 7 on any device.
- 8 • The enumerator `omp_unassigned_thread` with an implementation-defined value less
- 9 than -1;



10 The OpenMP enumeration types provided in the `omp.h` header file shall not be scoped
 11 enumeration types unless explicitly allowed.

12 The `omp.h` header file also defines a class template that models the **Allocator** concept in the
 13 `omp::allocator` namespace for each predefined memory allocator in Table 7.3 for which the
 14 name includes neither the `omp_` prefix nor the `_alloc` suffix.



15 The OpenMP Fortran API runtime library routines are external procedures. The return values of
 16 these routines are of default kind, unless otherwise specified.

17 Interface declarations for the OpenMP Fortran runtime library routines described in this chapter
 18 shall be provided in the form of a Fortran **module** named `omp_lib` or a Fortran **include** file
 19 named `omp_lib.h`. Whether the `omp_lib.h` file provides derived-type definitions or those
 20 routines that require an explicit interface is implementation defined. Whether the **include** file or
 21 the **module** file (or both) is provided is also implementation defined.

22 These files also define the following:

- 23 • The default integer named constant `omp_allocator_handle_kind`;
- 24 • An integer named constant of kind `omp_allocator_handle_kind` for each predefined
- 25 memory allocator in Table 7.3;
- 26 • The integer named constant `omp_null_allocator` of kind
- 27 `omp_allocator_handle_kind`;
- 28 • The default integer named constant `omp_alloctrait_key_kind`;
- 29 • The default integer named constant `omp_alloctrait_val_kind`;
- 30 • The default integer named constant `omp_control_tool_kind`;

- 1 • The default integer named constant `omp_control_tool_result_kind`;
- 2 • The default integer named constant `omp_depend_kind`;
- 3 • The default integer named constant `omp_event_handle_kind`;
- 4 • The default integer named constant `omp_initial_device` with value -1;
- 5 • The default integer named constant `omp_interop_kind`;
- 6 • The default integer named constant `omp_interop_fr_kind`;
- 7 • An integer named constant `omp_ifr_name` of kind `omp_interop_fr_kind` for each
- 8 *name* that is a foreign runtime name that is defined in the [OpenMP Additional Definitions](#)
- 9 [document](#);
- 10 • The default integer named constant `omp_invalid_device` with an
- 11 implementation-defined value less than -1;
- 12 • The default integer named constant `omp_lock_kind`;
- 13 • The default integer named constant `omp_memspace_handle_kind`;
- 14 • An integer named constant of kind `omp_memspace_handle_kind` for each predefined
- 15 memory space in Table 7.1;
- 16 • The integer named constant `omp_null_mem_space` of kind
- 17 `omp_memspace_handle_kind`;
- 18 • The default integer named constant `omp_nest_lock_kind`;
- 19 • The default integer named constant `omp_pause_resource_kind`;
- 20 • The default integer named constant `omp_proc_bind_kind`;
- 21 • The default integer named constant `omp_sched_kind`;
- 22 • The default integer named constant `omp_sync_hint_kind`; and
- 23 • The default integer named constant `omp_unassigned_thread` with an
- 24 implementation-defined value less than -1;
- 25 • The default integer named constant `openmp_version` with a value *yyyymm* where *yyyy*
- 26 and *mm* are the year and month designations of the version of the OpenMP Fortran API that
- 27 the implementation supports; this value matches that of the C preprocessor macro `_OPENMP`,
- 28 when a macro preprocessor is supported (see [Section 4.3](#)).

29 Whether any of the OpenMP runtime library routines that take an argument are extended with a
30 generic interface so arguments of different **KIND** type can be accommodated is implementation
31 defined.



Fortran

19.2 Thread Team Routines

This section describes routines that affect and monitor thread teams that execute tasks in the current contention group.

19.2.1 `omp_set_num_threads`

Summary

The `omp_set_num_threads` routine affects the number of threads to be used for subsequent `parallel` regions that do not specify a `num_threads` clause, by setting the value of the first element of the `nthreads-var` ICV of the current task.

Format

	C / C++	
<code>void omp_set_num_threads(int num_threads);</code>		
	C / C++	
	Fortran	
<code>subroutine omp_set_num_threads(num_threads)</code> <code>integer num_threads</code>		
	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 11.2.2](#)
- `parallel` directive, see [Section 11.2](#)
- `nthreads-var` ICV, see [Table 2.1](#)
- Determining the Number of Threads for a `parallel` Region, see [Section 11.2.1](#)

19.2.2 `omp_get_num_threads`

Summary

The `omp_get_num_threads` routine returns the number of threads in the current team.

1

Format

C / C++

2

```
int omp_get_num_threads(void);
```

C / C++

Fortran

3

```
integer function omp_get_num_threads()
```

Fortran

4

Binding

5

The binding region for an `omp_get_num_threads` region is the innermost enclosing parallel region.

6

7

Effect

8

The `omp_get_num_threads` routine returns the number of threads in the team that is executing the parallel region to which the routine region binds.

9

10

19.2.3 `omp_get_max_threads`

11

Summary

12

The `omp_get_max_threads` routine returns an upper bound on the number of threads that could be used to form a new team if a `parallel` construct without a `num_threads` clause is encountered after execution returns from this routine.

13

14

15

Format

C / C++

16

```
int omp_get_max_threads(void);
```

C / C++

Fortran

17

```
integer function omp_get_max_threads()
```

Fortran

18

Binding

19

The binding task set for an `omp_get_max_threads` region is the generating task.

20

Effect

21

The value returned by `omp_get_max_threads` is the value of the first element of the `nthreads-var` ICV of the current task. This value is also an upper bound on the number of threads that could be used to form a new team if a parallel region without a `num_threads` clause is encountered after execution returns from this routine.

22

23

24

Cross References









- `num_threads` clause, see [Section 11.2.2](#)
- `parallel` directive, see [Section 11.2](#)
- `nthreads-var` ICV, see [Table 2.1](#)
- Determining the Number of Threads for a `parallel` Region, see [Section 11.2.1](#)

19.2.4 `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++	
11	<code>int omp_get_thread_num(void);</code>		
		C / C++	
		Fortran	
12	<code>integer function omp_get_thread_num()</code>		
		Fortran	

Binding

The binding thread set for an `omp_get_thread_num` region is the current team. The binding region for an `omp_get_thread_num` region is the innermost enclosing parallel region.

Effect

The `omp_get_thread_num` routine returns the [thread number](#) of the calling [thread](#), within the [team](#) that is executing the [parallel region](#) to which the [routine region](#) binds. For [assigned threads](#), 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. For [unassigned threads](#), the [thread number](#) is the value `omp_unassigned_thread`.

Cross References

- `omp_get_num_threads`, see [Section 19.2.2](#)

19.2.5 `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*.

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23

Format

```
int omp_in_parallel(void);
logical function omp_in_parallel()
```

C / C++
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 11.2](#)
- `active-levels-var` ICV, see [Table 2.1](#)

19.2.6 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

```
void omp_set_dynamic(int dynamic_threads);
subroutine omp_set_dynamic(dynamic_threads)
logical dynamic_threads
```

C / C++
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](#)

19.2.7 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++
		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](#)

19.2.8 `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](#)

19.2.9 `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

```
! schedule kinds
integer(kind=omp_sched_kind), &
    parameter :: omp_sched_static = &
        int(Z'1', kind=omp_sched_kind)
integer(kind=omp_sched_kind), &
    parameter :: omp_sched_dynamic = &
        int(Z'2', kind=omp_sched_kind)
integer(kind=omp_sched_kind), &
    parameter :: omp_sched_guided = &
        int(Z'3', kind=omp_sched_kind)
integer(kind=omp_sched_kind), &
    parameter :: omp_sched_auto = &
        int(Z'4', kind=omp_sched_kind)

! schedule modifier
integer(kind=omp_sched_kind), &
    parameter :: omp_sched_monotonic = &
        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](#)

19.2.10 omp_get_schedule

Summary

The **omp_get_schedule** routine returns the schedule that is applied when the runtime schedule is used.

Format

	C / C++	
<pre>void omp_get_schedule(omp_sched_t *kind, int *chunk_size);</pre>		
	C / C++	
	Fortran	
<pre>subroutine omp_get_schedule(kind, chunk_size) integer (kind=omp_sched_kind) kind integer chunk_size</pre>		
	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 19.2.9](#), 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

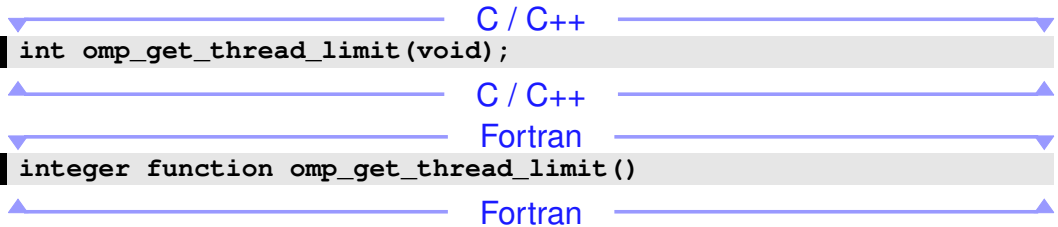
- *run-sched-var* ICV, see [Table 2.1](#)

19.2.11 `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



```
int omp_get_thread_limit(void);
```

```
integer function omp_get_thread_limit()
```

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

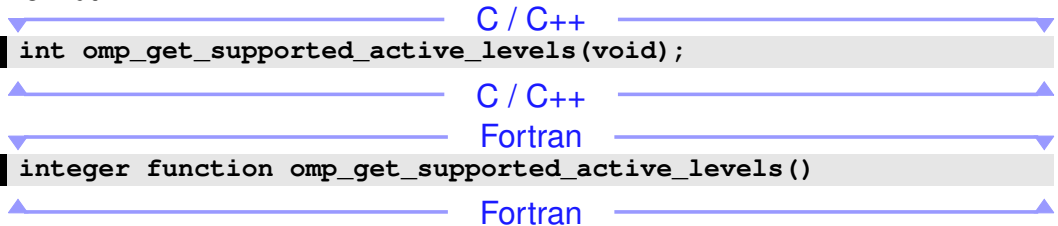
- *thread-limit-var* ICV, see [Table 2.1](#)

19.2.12 `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



```
int omp_get_supported_active_levels(void);
```

```
integer function omp_get_supported_active_levels()
```


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 level](#) 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](#)

19.2.13 `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.

Format

	C / C++	
<pre>void omp_set_max_active_levels(int max_levels);</pre>		
	C / C++	
	Fortran	
<pre>subroutine omp_set_max_active_levels(max_levels) integer max_levels</pre>		
	Fortran	

Constraints on Arguments

The value of the argument passed to this routine must evaluate to a non-negative integer, otherwise the behavior of this routine is [implementation defined](#).

Binding

The [binding task set](#) for an `omp_set_max_active_levels` region is the [generating task](#).

Effect

The effect of this routine is to set the value of the *max-active-levels-var* ICV to the value specified in the argument.

If the number of active levels requested exceeds the number of active levels of parallelism supported by the implementation, the value of the *max-active-levels-var* ICV will be set to the number of active levels supported by the implementation. If the number of active levels requested is less than the value of the *active-levels-var* ICV, the value of the *max-active-levels-var* ICV will be set to an implementation-defined value between the requested number and *active-levels-var*, inclusive.

Cross References

- *max-active-levels-var* ICV, see [Table 2.1](#)

19.2.14 `omp_get_max_active_levels`

Summary

The `omp_get_max_active_levels` routine returns the value of the *max-active-levels-var* ICV, which determines the maximum number of nested active parallel regions when the innermost parallel region is generated by the current task.

Format

	C / C++
<code>int omp_get_max_active_levels(void);</code>	
	C / C++
	Fortran
<code>integer function omp_get_max_active_levels()</code>	
	Fortran

Binding

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](#)

19.2.15 `omp_get_level`

Summary

The `omp_get_level` routine returns the value of the *levels-var* ICV.

Format

	C / C++
<code>int omp_get_level(void);</code>	
	C / C++
	Fortran
<code>integer function omp_get_level()</code>	
	Fortran

1 Binding

2 The binding task set for an `omp_get_level` region is the generating task.

3 Effect

4 The effect of the `omp_get_level` routine is to return the number of nested **parallel** regions
5 (whether active or inactive) that enclose the current task such that all of the **parallel** regions are
6 enclosed by the outermost initial task region on the current device.

7 Cross References

- 8 • `parallel` directive, see [Section 11.2](#)
- 9 • `levels-var` ICV, see [Table 2.1](#)

10 19.2.16 `omp_get_ancestor_thread_num`

11 Summary

12 The `omp_get_ancestor_thread_num` routine returns, for a given nested level of the
13 [encountering thread](#), the [thread number](#) of the [ancestor thread](#) of the [encountering thread](#).

14 Format

15 C / C++
16 Fortran
17 Fortran

18 Binding

19 The [binding thread set](#) for an `omp_get_ancestor_thread_num` [region](#) is the [encountering thread](#).
20 The [binding region](#) for an `omp_get_ancestor_thread_num` [region](#) is the innermost
21 enclosing [parallel region](#).

22 Effect

23 The `omp_get_ancestor_thread_num` routine returns the [thread number](#) of the [ancestor thread](#)
24 at a given nest level of the [encountering thread](#) or the [thread number](#) of the [encountering thread](#).
25 If the requested nest level is outside the range of 0 and the nest level of the [encountering thread](#),
26 as returned by the `omp_get_level` routine, the routine returns -1.

27 Note – When the `omp_get_ancestor_thread_num` routine is called with value of
28 `level=0`, the routine always returns 0. If `level=omp_get_level()`, the routine has the
29 same effect as the `omp_get_thread_num` routine.
30
31

Cross References

- `parallel` directive, see [Section 11.2](#)
- `omp_get_level`, see [Section 19.2.15](#)
- `omp_get_thread_num`, see [Section 19.2.4](#)

19.2.17 `omp_get_team_size`

Summary

The `omp_get_team_size` routine returns, for a given nested level of the [encountering thread](#), the size of the [current team](#) to which the [ancestor thread](#) or the [encountering task](#) belongs.

Format

	C / C++	
<code>int omp_get_team_size(int level);</code>		
	C / C++	
	Fortran	
<code>integer function omp_get_team_size(level)</code>		
<code>integer level</code>		
	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 [current team](#) to which the [ancestor thread](#) or the [encountering task](#) belongs. If the requested nested level is outside the range of 0 and the nested level of the [encountering 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 11.2](#)
- `omp_get_level`, see [Section 19.2.15](#)
- `omp_get_num_threads`, see [Section 19.2.2](#)

19.2.18 `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 11.2](#)
- *active-levels-var* ICV, see [Table 2.1](#)

19.3 Thread Affinity Routines

This section describes routines that affect and access [thread affinity](#) policies that are in effect.

19.3.1 `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

```
integer (kind=omp_proc_bind_kind), &
    parameter :: omp_proc_bind_false = 0
integer (kind=omp_proc_bind_kind), &
    parameter :: omp_proc_bind_true = 1
integer (kind=omp_proc_bind_kind), &
    parameter :: omp_proc_bind_primary = 2
integer (kind=omp_proc_bind_kind), &
    parameter :: omp_proc_bind_close = 3
integer (kind=omp_proc_bind_kind), &
    parameter :: omp_proc_bind_spread = 4
```

Fortran

Binding

The binding task set for an `omp_get_proc_bind` region is the generating task.

Effect

The effect of this routine is to return the value of the first element of the *bind-var* ICV of the current task. See [Section 11.2.3](#) for the rules that govern the thread affinity policy.

Cross References

- `parallel` directive, see [Section 11.2](#)
- Controlling OpenMP Thread Affinity, see [Section 11.2.3](#)
- *bind-var* ICV, see [Table 2.1](#)

19.3.2 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

Fortran
`integer function omp_get_num_places()`

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](#)

19.3.3 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++
`int omp_get_place_num_procs(int place_num);`

C / C++
Fortran

Fortran
`integer function omp_get_place_num_procs(place_num)
integer place_num`

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 19.3.2](#)

19.3.4 `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

```
void omp_get_place_proc_ids(int place_num, int *ids);
```

C / C++

```
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 3.1.5](#)
- `omp_get_num_places`, see [Section 19.3.2](#)
- `omp_get_place_num_procs`, see [Section 19.3.3](#)

19.3.5 `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.

Format

	C / C++	
<code>int omp_get_place_num(void);</code>		
	C / C++	
	Fortran	
<code>integer function omp_get_place_num()</code>		
	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 19.3.2](#)

19.3.6 `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.

Format

	C / C++	
<code>int omp_get_partition_num_places(void);</code>		
	C / C++	

Fortran

```
integer function omp_get_partition_num_places()
```

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](#)

19.3.7 omp_get_partition_place_nums

Summary

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](#).

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 19.3.6](#)

19.3.8 `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
`subroutine omp_set_affinity_format(format)
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 3.2.5](#)
- `OMP_DISPLAY_AFFINITY`, see [Section 3.2.4](#)
- Controlling OpenMP Thread Affinity, see [Section 11.2.3](#)
- `omp_capture_affinity`, see [Section 19.3.11](#)
- `omp_display_affinity`, see [Section 19.3.10](#)
- `omp_get_affinity_format`, see [Section 19.3.9](#)

19.3.9 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++
`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 11.2](#)
- `teams` directive, see [Section 11.3](#)
- `affinity-format-var` ICV, see [Table 2.1](#)

19.3.10 `omp_display_affinity`

Summary

The `omp_display_affinity` routine prints the OpenMP thread affinity information using the format specification provided.

Format

	C / C++	
<pre>void omp_display_affinity(const char *format);</pre>		
	C / C++	
	Fortran	
<pre>subroutine omp_display_affinity (format) character(len=*), intent(in) :: format</pre>		
	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](#)

19.3.11 `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

```
integer function omp_capture_affinity(buffer,format)  
character(len=*), intent(out) :: buffer  
character(len=*), intent(in)  :: format
```

Fortran

Binding

The binding thread set for an `omp_capture_affinity` region is the encountering thread.

Effect

C / C++

The `omp_capture_affinity` routine returns the number of characters in the entire thread affinity information string excluding the terminating null byte ('`\0`'). 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](#)

19.4 Teams Region Routines

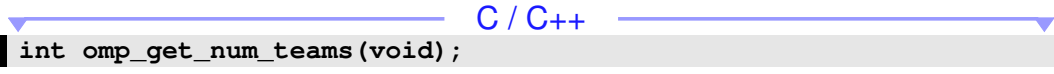


This section describes routines that affect and monitor the league of teams that may execute a **teams** region.

19.4.1 `omp_get_num_teams`

Summary

The `omp_get_num_teams` routine returns the number of initial teams in the current **teams** region.

Format


`int omp_get_num_teams(void);`

`integer function omp_get_num_teams()`


Binding

The binding task set for an `omp_get_num_teams` region is the generating task

Effect

The effect of this routine is to return the number of initial teams in the current **teams** region. The routine returns 1 if it is called from outside of a **teams** region.

Cross References



- **teams** directive, see [Section 11.3](#)

19.4.2 `omp_get_team_num`

Summary

The `omp_get_team_num` routine returns the initial team number of the calling thread.

Format


`int omp_get_team_num(void);`


Fortran

1 | integer function omp_get_team_num()

Fortran

2 Binding

3 The binding task set for an `omp_get_team_num` region is the generating task.

4 Effect

5 The `omp_get_team_num` routine returns the initial team number of the calling thread. The
6 initial team number is an integer between 0 and one less than the value returned by
7 `omp_get_num_teams()`, inclusive. The routine returns 0 if it is called outside of a `teams`
8 region.

9 Cross References

- 10 • `teams` directive, see [Section 11.3](#)
- 11 • `omp_get_num_teams`, see [Section 19.4.1](#)

12 19.4.3 omp_set_num_teams

13 Summary

14 The `omp_set_num_teams` routine affects the number of threads to be used for subsequent
15 `teams` regions that do not specify a `num_teams` clause, by setting the value of the *nteam*s-var
16 ICV of the current device.

17 Format

C / C++

18 | void omp_set_num_teams(int num_teams);

C / C++

Fortran

19 | subroutine omp_set_num_teams(num_teams)

20 | integer num_teams

Fortran

21 Constraints on Arguments

22 The value of the argument passed to this routine must evaluate to a positive integer, or else the
23 behavior of this routine is implementation defined.

24 Binding

25 The binding task set for an `omp_set_num_teams` region is the generating task.

26 Effect

27 The effect of this routine is to set the value of the *nteam*s-var ICV of the current device to the value
28 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


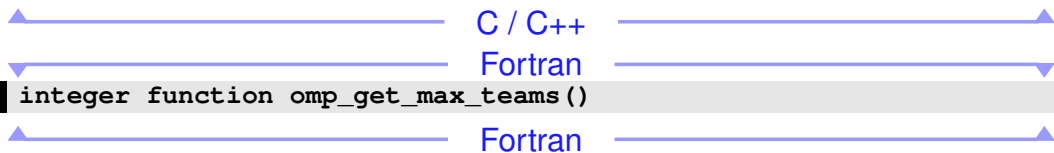
- `num_teams` clause, see [Section 11.3.1](#)
- `teams` directive, see [Section 11.3](#)
- `ntteams-var` ICV, see [Table 2.1](#)

19.4.4 `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

 <pre>int omp_get_max_teams(void);</pre>	C / C++
 <pre>integer function omp_get_max_teams()</pre>	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 `ntteams-var` 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 11.3.1](#)
- `teams` directive, see [Section 11.3](#)
- `ntteams-var` ICV, see [Table 2.1](#)

19.4.5 `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

C / C++

```
void omp_set_teams_thread_limit(int thread_limit);
```

C / C++

Fortran

```
subroutine omp_set_teams_thread_limit(thread_limit)  
integer thread_limit
```

Fortran

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 14.3](#)
- `teams` directive, see [Section 11.3](#)
- `teams-thread-limit-var` ICV, see [Table 2.1](#)

19.4.6 `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++
`int omp_get_teams_thread_limit(void);`

C / C++

Fortran
`integer function omp_get_teams_thread_limit()`

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 11.3](#)
- `teams-thread-limit-var` ICV, see [Table 2.1](#)

19.5 Tasking Routines

This section describes routines that pertain to OpenMP explicit tasks.

19.5.1 `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++
`int omp_get_max_task_priority(void);`

C / C++

Fortran
`integer function omp_get_max_task_priority()`

Fortran

Binding

The binding thread set for an `omp_get_max_task_priority` region is all threads on the device. The effect of executing this routine is not related to any specific region 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 13.5](#)
- *max-task-priority-var* ICV, see [Table 2.1](#)

19.5.2 `omp_in_explicit_task`

Summary

The `omp_in_explicit_task` routine returns the value of the *explicit-task-var* ICV.

Format

	C / C++	
<code>int omp_in_explicit_task(void);</code>		
	C / C++	
	Fortran	
<code>logical function omp_in_explicit_task()</code>		
	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 13.6](#)
- *explicit-task-var* ICV, see [Table 2.1](#)

19.5.3 `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*.

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22

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*.

19.5.4 `omp_is_free_agent`

Summary

The `omp_is_free_agent` routine returns *true* if the encountering thread is a free-agent thread; otherwise, it returns *false*.

Format

```

C / C++
int omp_is_free_agent(void);
C / C++
Fortran
logical function omp_is_free_agent()
Fortran

```

Binding

The binding task set for an `omp_is_free_agent` region is the generating task.

Effect

The `omp_is_free_agent` routine returns *true* if a free-agent thread is executing the enclosing task region at the time the routine is called. Otherwise, it returns *false*.

Cross References

- `threadset` clause, see [Section 13.4](#)
- `task` directive, see [Section 13.6](#)

19.5.5 omp_ancestor_is_free_agent

Summary

The `omp_ancestor_is_free_agent` routine returns *true* if the `ancestor thread` of the `encountering thread` is a `free-agent thread`, for a given nested level of the `encountering thread`; otherwise, it returns *false*.

Format

C / C++

```
int omp_ancestor_is_free_agent(int level);
```

C / C++

Fortran

```
logical function omp_ancestor_is_free_agent(level)  
integer level
```

Fortran

Binding

The `binding task set` for an `omp_ancestor_is_free_agent` region is the `generating task`.

Effect

The `omp_ancestor_is_free_agent` routine returns *true* if the `ancestor thread` of the `encountering thread` is a `free-agent thread`, for a given nested level of the `encountering thread`; otherwise, it returns *false*. If the requested nesting level is outside the range of 0 and the nesting level of the `current task`, as returned by the `omp_get_level` routine, the routine returns *false*.

Note – When the `omp_ancestor_is_free_agent` routine is called with a value of `level = omp_get_level()`, the routine has the same effect as the `omp_is_free_agent` routine.

Cross References

- `threadset` clause, see [Section 13.4](#)
- `task` directive, see [Section 13.6](#)
- `omp_get_level`, see [Section 19.2.15](#)

19.6 Resource Relinquishing Routines

This section describes routines that relinquish resources used by the OpenMP runtime.

19.6.1 `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
```

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_stop_tool = 3
} 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
integer (kind=omp_pause_resource_kind), parameter :: &
    omp_pause_stop_tool = 3

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 The **omp_pause_resource** routine implies a barrier.

7 If successful, the **omp_pause_hard** value results in a hard pause for which the OpenMP state is
8 not guaranteed to persist across the **omp_pause_resource** call. A hard pause may relinquish
9 any data allocated by OpenMP on a given device, including data allocated by memory routines for
10 that device as well as data present on the device as a result of a declare target directive or
11 **target data** construct. A hard pause may also relinquish any data associated with a
12 **threadprivate** directive. When relinquished and when applicable, base language appropriate
13 deallocation/finalization is performed. When relinquished and when applicable, mapped data on a
14 device will not be copied back from the device to the host.

15 If successful, the **omp_pause_soft** value results in a soft pause for which the OpenMP state is
16 guaranteed to persist across the call, with the exception of any data associated with a
17 **threadprivate** directive, which may be relinquished across the call. When relinquished and
18 when applicable, base language appropriate deallocation/finalization is performed.

19 
20 **Note** – A hard pause may relinquish more resources, but may resume processing OpenMP regions
21 more slowly. A soft pause allows OpenMP regions to restart more quickly, but may relinquish fewer
22 resources. An OpenMP implementation will reclaim resources as needed for OpenMP regions
23 encountered after the **omp_pause_resource** region. Since a hard pause may unmap data on the
24 specified device, appropriate data mapping is required before using data on the specified device
25 after the **omp_pause_region** region.
26 

27 The routine returns zero in case of success, and non-zero otherwise.

28 Tool Callbacks

29 If the tool is not allowed to interact with the specified device after encountering this call, then the
30 runtime must call the tool finalizer for that device.

31 Restrictions

32 Restrictions to the **omp_pause_resource** routine are as follows:

- 33 • The **omp_pause_resource** region may not be nested in any explicit OpenMP region.
- 34 • The routine may only be called when all explicit tasks that do not bind to the implicit parallel
35 region to which the encountering thread binds have finalized execution.
- 36 • The **omp_pause_stop_tool** value must not be specified.

Cross References

- `target data` directive, see [Section 14.5](#)
- `threadprivate` directive, see [Section 6.2](#)
- Declare Target Directives, see [Section 8.8](#)

19.6.2 `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 `omp_pause_resource_all` routine implies a barrier.

The argument `kind` passed to this routine can be one of the valid OpenMP pause kind as defined in [Section 19.6.1](#), or any implementation-specific pause kind.

If successful, the `omp_pause_stop_tool` value results in a hard pause for which the OpenMP state is not guaranteed to persist across the `omp_pause_resource` call. In addition to the effects described above, the implementation will shutdown the OMPT interface as if the program execution was ending.

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 that do not bind to the implicit parallel region to which the encountering thread binds have finalized execution.

Cross References

- `omp_pause_resource`, see [Section 19.6.1](#)

19.7 Device Information Routines

This section describes routines that pertain to the set of devices that are available to an OpenMP program.

19.7.1 `omp_get_num_procs`

Summary

The `omp_get_num_procs` routine returns the number of processors available to the device.

Format

<code>int omp_get_num_procs(void);</code>	C / C++
<code>integer function omp_get_num_procs()</code>	Fortran

Binding

The binding thread set for an `omp_get_num_procs` region is all threads on a device. The effect of executing this routine is not related to any specific region corresponding to any construct or API routine.

Effect

The `omp_get_num_procs` routine returns the number of processors that are available to the device at the time the routine is called. This value may change between the time that it is determined by the `omp_get_num_procs` routine and the time that it is read in the calling context due to system actions outside the control of the OpenMP implementation.

19.7.2 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_max_progress_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 11.2](#)

19.7.3 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

C / C++

```
void omp_set_default_device(int device_num);
```

C / C++

Fortran

```
subroutine omp_set_default_device(device_num)  
integer device_num
```

Fortran

Binding

The binding task set for an `omp_set_default_device` region is the generating task.

Effect

The effect of this routine is to set the value of the *default-device-var* ICV of the current task to the value specified in the argument. When called from within a **target** region the effect of this routine is unspecified.

Cross References

- **target** directive, see [Section 14.8](#)
- *default-device-var* ICV, see [Table 2.1](#)

19.7.4 omp_get_default_device

Summary

The `omp_get_default_device` routine returns the default target device.

Format

	C / C++	
<code>int omp_get_default_device(void);</code>		
	C / C++	
	Fortran	
<code>integer function omp_get_default_device()</code>		
	Fortran	

Binding

The binding task set for an `omp_get_default_device` region is the generating task.

Effect

The `omp_get_default_device` routine returns the value of the *default-device-var* ICV of the current task. When called from within a **target** region the effect of this routine is unspecified.

Cross References

- **target** directive, see [Section 14.8](#)
- *default-device-var* ICV, see [Table 2.1](#)

19.7.5 omp_get_num_devices

Summary

The `omp_get_num_devices` routine returns the number of non-host devices available for offloading code or data.

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_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 14.8](#)

19.7.6 `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

```
int omp_get_device_num(void);  
integer function omp_get_device_num()
```

Binding

The binding task set for an `omp_get_device_num` region is the generating task.

Effect

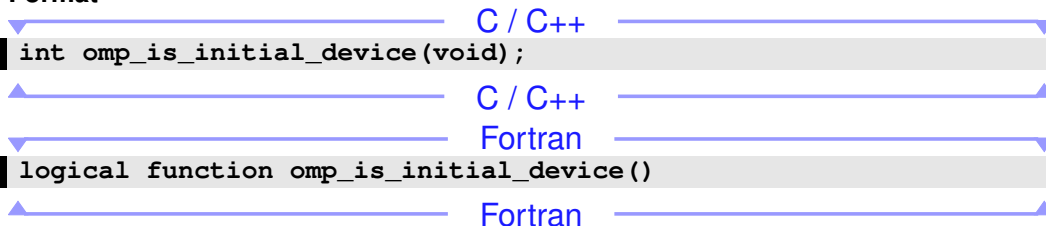
The `omp_get_device_num` routine returns the device number of the device on which the calling thread is executing. When called on the host device, it will return the same value as the `omp_get_initial_device` routine.

19.7.7 `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



```
int omp_is_initial_device(void);
```

```
logical function omp_is_initial_device()
```

Binding

The binding task set for an `omp_is_initial_device` region is the generating task.

Effect

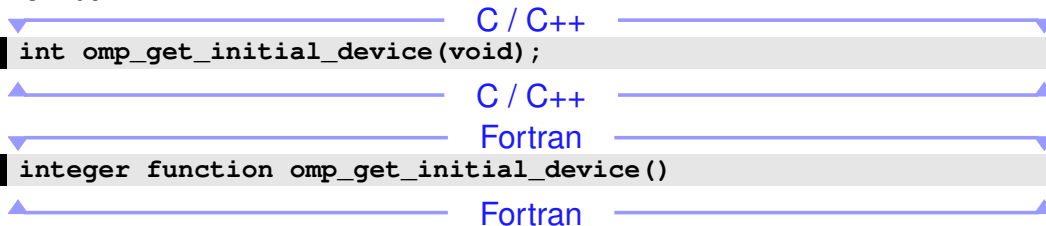
The effect of this routine is to return *true* if the current task is executing on the host device; otherwise, it returns *false*.

19.7.8 `omp_get_initial_device`

Summary

The `omp_get_initial_device` routine returns a device number that represents the host device.

Format



```
int omp_get_initial_device(void);
```

```
integer function omp_get_initial_device()
```

Binding

The binding task set for an `omp_get_initial_device` region is the generating task.

Effect

The effect of this routine is to return the device number of the host device. The value of the device number is the value returned by the `omp_get_num_devices` routine. When called from within a `target` region the effect of this routine is unspecified.

Cross References

- `target` directive, see [Section 14.8](#)

19.8 Device Memory Routines

This section describes routines that support allocation of memory and management of pointers in the data environments of target devices.

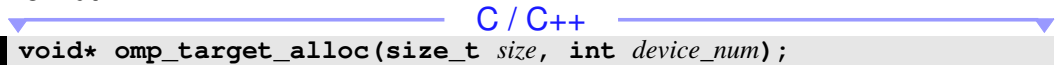
If the `device_num`, `src_device_num`, or `dst_device_num` argument of a [device memory](#) routine has the value `omp_invalid_device`, [runtime error termination](#) is performed.

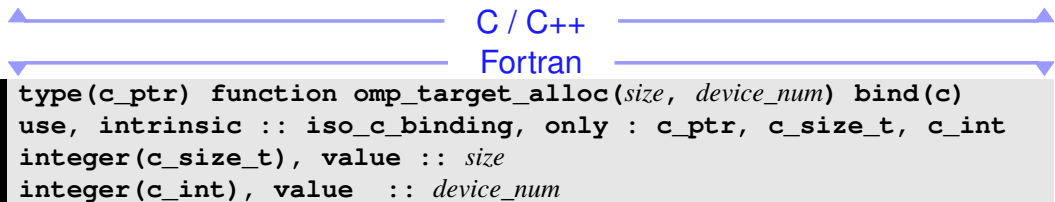
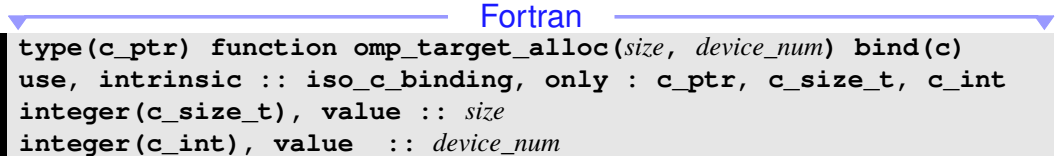
19.8.1 `omp_target_alloc`

Summary

The `omp_target_alloc` routine allocates memory in a device data environment and returns a device pointer to that memory.

Format

 C / C++
`void* omp_target_alloc(size_t size, int device_num);`

 C / C++
 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 or if `size` is 0. The device pointer returned by `omp_target_alloc` can be used in an `is_device_ptr` clause (see [Section 6.4.7](#)).

Fortran

1 The `omp_target_alloc` routine requires an explicit interface and so might not be provided in
2 `omp_lib.h`.

Fortran

Execution Model Events

3 The *target-data-allocation-begin* event occurs before a thread initiates a data allocation on a target
4 device.
5

6 The *target-data-allocation-end* event occurs after a thread initiates a data allocation on a target
7 device.

Tool Callbacks

8 A thread dispatches a registered `ompt_callback_target_data_op_emi` callback with
9 `ompt_scope_begin` as its endpoint argument for each occurrence of a
10 *target-data-allocation-begin* event in that thread. Similarly, a thread dispatches a registered
11 `ompt_callback_target_data_op_emi` callback with `ompt_scope_end` as its endpoint
12 argument for each occurrence of a *target-data-allocation-end* event in that thread. These callbacks
13 have type signature `ompt_callback_target_data_op_emi_t`.
14

15 A thread dispatches a registered `ompt_callback_target_data_op` callback for each
16 occurrence of a *target-data-allocation-end* event in that thread. The callback occurs in the context
17 of the target task and has type signature `ompt_callback_target_data_op_t`.

Restrictions

18 Restrictions to the `omp_target_alloc` routine are as follows.

- 19 • Freeing the storage returned by `omp_target_alloc` with any routine other than
20 `omp_target_free` results in unspecified behavior.
21
- 22 • When called from within a `target` region the effect is unspecified.

C / C++

- 23 • Unless the `unified_address` clause appears on a `requires` directive in the
24 compilation unit, pointer arithmetic is not supported on the device pointer returned by
25 `omp_target_alloc`.

C / C++

Cross References

- 26 • `is_device_ptr` clause, see [Section 6.4.7](#)
- 27
- 28 • `target` directive, see [Section 14.8](#)
- 29 • `omp_target_free`, see [Section 19.8.2](#)
- 30 • `ompt_callback_target_data_op_emi_t` and
31 `ompt_callback_target_data_op_t`, see [Section 20.5.2.25](#)

19.8.2 `omp_target_free`

Summary

The `omp_target_free` routine frees the device memory allocated by the `omp_target_alloc` routine.

Format

C / C++

```
void omp_target_free(void *device_ptr, int device_num);
```

C / C++

Fortran

```
subroutine omp_target_free(device_ptr, device_num) bind(c)  
use, intrinsic :: iso_c_binding, only : c_ptr, c_int  
type(c_ptr), value :: device_ptr  
integer(c_int), value :: device_num
```

Fortran

Constraints on Arguments

An [OpenMP program](#) that calls `omp_target_free` with a [non-null pointer](#) that does not have a value returned from `omp_target_alloc` is a [non-conforming program](#). The `device_num` argument must be a 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.

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-free-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-free-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-free-begin* 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_free` routine are as follows.

- 13 • When called from within a **target** region the effect is unspecified.

14 Cross References

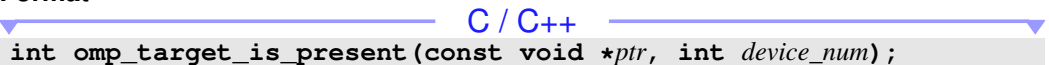
- 15 • `target` directive, see [Section 14.8](#)
- 16 • `omp_target_alloc`, see [Section 19.8.1](#)
- 17 • `ompt_callback_target_data_op_emi_t` and
18 `ompt_callback_target_data_op_t`, see [Section 20.5.2.25](#)

19 19.8.3 omp_target_is_present

20 Summary

21 The `omp_target_is_present` routine tests whether a host pointer refers to storage that is
22 mapped to a given device.

23 Format

24  `int omp_target_is_present(const void *ptr, int device_num);`

 `C / C++`
 `Fortran`

25 `integer(c_int) function omp_target_is_present(ptr, device_num) &`
26 `bind(c)`
27 `use, intrinsic :: iso_c_binding, only : c_ptr, c_int`
28 `type(c_ptr), value :: ptr`
29 `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 14.8](#)

19.8.4 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 that corresponds to the [address range](#) 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 14.8](#)

19.8.5 `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

```
1 integer(c_int) function omp_target_memcpy(dst, src, length, &  
2     dst_offset, src_offset, dst_device_num, src_device_num) bind(c)  
3 use, intrinsic :: iso_c_binding, only : c_ptr, c_int, c_size_t  
4 type(c_ptr), value :: dst, src  
5 integer(c_size_t), value :: length, dst_offset, src_offset  
6 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 14.8](#)
- `ompt_callback_target_data_op_emi_t` and `ompt_callback_target_data_op_t`, see [Section 20.5.2.25](#)

19.8.6 `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,  
    const size_t *src_offsets,  
    const size_t *dst_dimensions,  
    const size_t *src_dimensions,  
    int dst_device_num,  
    int src_device_num  
);
```

C / C++

Fortran

```
integer(c_int) function omp_target_memcpy_rect(dst,src,element_size, &  
    num_dims, volume, dst_offsets, src_offsets, dst_dimensions, &  
    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 :: element_size  
integer(c_int), value :: num_dims, dst_device_num, src_device_num  
integer(c_size_t), intent(in) :: volume(*), dst_offsets(*), &  
    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*).

The **omp_target_memcpy_rect** 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. The routine contains a task scheduling point.

An application can determine the inclusive number of 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** 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** region.

The *target-data-op-end* event occurs after a thread initiates a data transfer in the **omp_target_memcpy_rect** 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` routine are as follows.

- 13 • When called from within a **target** region the effect is unspecified.

14 Cross References

- 15 • `target` directive, see [Section 14.8](#)
- 16 • `ompt_callback_target_data_op_emi_t` and
17 `ompt_callback_target_data_op_t`, see [Section 20.5.2.25](#)

18 19.8.7 omp_target_memcpy_async

19 Summary

20 The `omp_target_memcpy_async` routine asynchronously performs a copy between any
21 combination of host and device pointers.

22 Format

```
23 int omp_target_memcpy_async(  
24     void *dst,  
25     const void *src,  
26     size_t length,  
27     size_t dst_offset,  
28     size_t src_offset,  
29     int dst_device_num,  
30     int src_device_num,  
31     int depobj_count,  
32     omp_depend_t *depobj_list  
33 );
```

C / C++

C / C++

Fortran

```
1 integer(c_int) function omp_target_memcpy_async(dst, src, length, &  
2 dst_offset, src_offset, dst_device_num, src_device_num, &  
3 depobj_count, depobj_list) bind(c)  
4 use, intrinsic :: iso_c_binding, only : c_ptr, c_int, c_size_t  
5 type(c_ptr), value :: dst, src  
6 integer(c_size_t), value :: length, dst_offset, src_offset  
7 integer(c_int), value :: dst_device_num, src_device_num, depobj_count  
8 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.

The routine returns zero if successful. Otherwise, it returns a non-zero value. The routine contains a task scheduling point.

Fortran

The **omp_target_memcpy_async** routine requires an explicit interface and so might not be provided in **omp_lib.h**.

Fortran

Execution Model Events

Events associated with a *target task* are the same as for the **task** construct defined in [Section 13.6](#). Events associated with task dependences that result from *depobj_list* are the same as for a **depend** clause with the **debobj** *task-dependence-type* defined in [Section 16.9.5](#).

The *target-data-op-begin* event occurs before a thread initiates a data transfer in the **omp_target_memcpy_async** region.

The *target-data-op-end* event occurs after a thread initiates a data transfer in the **omp_target_memcpy_async** region.

Tool Callbacks

Callbacks associated with events for *target tasks* are the same as for the **task** construct defined in [Section 13.6](#); (*flags & omp_target_target*) always evaluates to *true* in the dispatched callback.

Callbacks associated with events for task dependences are the same as for the **depend** clause defined in [Section 16.9.5](#).

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_async** routine are as follows.

- When called from within a **target** region the effect is unspecified.

Cross References

- **target** directive, see [Section 14.8](#)
- Depend Objects, see [Section 16.9.2](#)
- **ompt_callback_target_data_op_emi_t** and **ompt_callback_target_data_op_t**, see [Section 20.5.2.25](#)

19.8.8 omp_target_memcpy_rect_async

Summary

The **omp_target_memcpy_rect_async** routine asynchronously performs a copy between 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_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. 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 *depobj_count*. *depobj_list* is ignored if the value of *depobj_count* is zero.

The routine returns zero if successful. Otherwise, it returns a non-zero value. The routine contains a task scheduling point.

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

Events associated with a *target task* are the same as for the `task` construct defined in [Section 13.6](#). Events associated with task dependences that result from *depobj_list* are the same as for a `depend` clause with the `deboj task-dependence-type` defined in [Section 16.9.5](#).

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 Callbacks associated with events for *target tasks* are the same as for the **task** construct defined in
3 [Section 13.6](#); (*flags & omp_task_target*) always evaluates to *true* in the dispatched callback.

4 Callbacks associated with events for task dependences are the same as for the **depend** clause
5 defined in [Section 16.9.5](#).

6 A thread dispatches a registered **omp_callback_target_data_op_emi** callback with
7 **omp_scope_begin** as its endpoint argument for each occurrence of a *target-data-op-begin*
8 event in that thread. Similarly, a thread dispatches a registered
9 **omp_callback_target_data_op_emi** callback with **omp_scope_end** as its endpoint
10 argument for each occurrence of a *target-data-op-end* event in that thread. These callbacks have
11 type signature **omp_callback_target_data_op_emi_t**.

12 A thread dispatches a registered **omp_callback_target_data_op** callback for each
13 occurrence of a *target-data-op-end* event in that thread. The callback occurs in the context of the
14 target task and has type signature **omp_callback_target_data_op_t**.

15 Restrictions

16 Restrictions to the **omp_target_memcpy_rect_async** routine are as follows.

- 17 • When called from within a **target** region the effect is unspecified.

18 Cross References

- 19 • **target** directive, see [Section 14.8](#)
- 20 • Depend Objects, see [Section 16.9.2](#)
- 21 • **omp_callback_target_data_op_emi_t** and
22 **omp_callback_target_data_op_t**, see [Section 20.5.2.25](#)

23 19.8.9 omp_target_memset

24 Summary

25 The **omp_target_memset** routine fills [memory](#) in a [device data environment](#) with a given value.

26 Format

```
27 void* omp_target_memset(void *ptr, int val, size_t count,  
28 int device_num);
```

29 C / C++

Fortran

```
1 type(c_ptr) function omp_target_memset(ptr, val, count, device_num) &  
2   bind(c)  
3 use, intrinsic :: iso_c_binding, only : c_ptr, c_int, c_size_t  
4 type(c_ptr), value :: ptr  
5 integer(c_int), value :: val  
6 integer(c_size_t), value :: count  
7 integer(c_int), value :: device_num
```

Fortran

Constraints on Arguments

The value of *ptr* 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_memset region](#) is the [generating task](#), which is the [target task](#) generated by the call to the [omp_target_memset routine](#).

Effect

The [omp_target_memset routine](#) fills the first *count* bytes pointed to by *ptr* with the value *val* (converted to [unsigned char](#)) in the [device data environment](#) associated with [device device_num](#). If *count* is zero, the [routine](#) has no effect. If *ptr* is [NULL](#), the effect is unspecified. The [omp_target_memset routine](#) returns *ptr*.

The [omp_target_memset 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_memset routine](#) contains a [task scheduling point](#).

Fortran

The [omp_target_memset 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 filling the [memory](#) in the [omp_target_memset region](#).

The [target-data-op-end event](#) occurs after a [thread](#) initiates filling the [memory](#) in the [omp_target_memset 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

The restrictions to the [omp_target_memset routine](#) are as follows:

- When called from within a [target region](#) the effect is unspecified.

Cross References

- [omp_target_alloc](#), see [Section 19.8.1](#)
- [omp_target_free](#), see [Section 19.8.2](#)
- [ompt_callback_target_data_op_emi_t](#) and [ompt_callback_target_data_op_t](#), see [Section 20.5.2.25](#)

19.8.10 omp_target_memset_async

Summary

The [omp_target_memset_async routine](#) fills [memory](#) in the [device data environment](#) with a given value.

Format

```
void* omp_target_memset_async(void *ptr, int val, size_t count,  
                             int device_num,  
                             int depobj_count,  
                             omp_depend_t *depobj_list);
```

C / C++

Fortran

```
1 type(c_ptr) function omp_target_memset_async(ptr, val, count, &
2                                     device_num, &
3                                     depobj_count, depobj_list) &
4     bind(c)
5     use, intrinsic :: iso_c_binding, only : c_ptr, c_int, c_size_t
6     type(c_ptr), value :: ptr
7     integer(c_int), value :: val
8     integer(c_size_t), value :: count
9     integer(c_int), value :: device_num
10    integer(c_int), value :: depobj_count
11    integer(omp_depend_kind), optional :: depobj_list(*)
```

Fortran

Constraints on Arguments

The value of *ptr* 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_memset_async` region is the [generating task](#), which is the [target task](#) generated by the call to the `omp_target_memset_async` routine.

Effect

The `omp_target_memset_async` routine fills the first *count* bytes pointed to by *ptr* with the value *val* (converted to `unsigned char`) in the [device data environment](#) associated with [device device_num](#). If *count* is zero, the routine has no effect. If *ptr* is `NULL`, the effect is unspecified. The `omp_target_memset_async` routine returns *ptr*.

The `omp_target_memset_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*. The *depobj_list* argument is ignored if the value of *depobj_count* is zero.

The routine contains a [task scheduling point](#).

Fortran

The `omp_target_memset_async` routine requires an explicit interface and so might not be provided in `omp_lib.h`.

Fortran

Execution Model Events

Events associated with a **target task** are the same as for the **task construct** defined in [Section 13.6](#). Events associated with **task dependences** that result from *depobj_list* are the same as for a **depend clause** with the **depobj task-dependence-type** defined in [Section 16.9.5](#).

The *target-data-op-begin* and *target-data-op-end* events in the **omp_target_memset_async region** are the same as those in the **omp_target_memset region**.

Tool Callbacks

Callbacks associated with events for *target tasks* are the same as for the **task** construct defined in [Section 13.6](#); (*flags & omp_target_target*) always evaluates to *true* in the dispatched callback.

Callbacks associated with events for task dependences are the same as for the **depend** clause defined in [Section 16.9.5](#).

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

The restrictions to the **omp_target_memset_async** routine are as follows:

- When called from within a **target** region the effect is unspecified.

Cross References

- Depend Objects, see [Section 16.9.2](#)
- **omp_target_alloc**, see [Section 19.8.1](#)
- **omp_target_free**, see [Section 19.8.2](#)
- **ompt_callback_target_data_op_emi_t** and **ompt_callback_target_data_op_t**, see [Section 20.5.2.25](#)

19.8.11 omp_target_associate_ptr

Summary

The **omp_target_associate_ptr** routine maps a device pointer, which may be returned from **omp_target_alloc** or implementation-defined runtime routines, to a host pointer.

Format

C / C++

```
1 int omp_target_associate_ptr(  
2     const void *host_ptr,  
3     const void *device_ptr,  
4     size_t size,  
5     size_t device_offset,  
6     int device_num  
7 );  
8
```

C / C++

Fortran

```
9 integer(c_int) function omp_target_associate_ptr(host_ptr, &  
10     device_ptr, size, device_offset, device_num) bind(c)  
11 use, intrinsic :: iso_c_binding, only : c_ptr, c_size_t, c_int  
12 type(c_ptr), value :: host_ptr, device_ptr  
13 integer(c_size_t), value :: size, device_offset  
14 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

1 on the same device with the same offset has no effect and returns zero. Associating pointers that
2 share underlying storage will result in unspecified behavior. The `omp_target_is_present`
3 function can be used to test whether a given host pointer has a corresponding variable in the device
4 data environment.

Fortran

5 The `omp_target_associate_ptr` routine requires an explicit interface and so might not be
6 provided in `omp_lib.h`.

Fortran

Execution Model Events

7 The *target-data-associate* event occurs before a thread initiates a device pointer association on a
8 target device.
9

Tool Callbacks

10 A thread dispatches a registered `ompt_callback_target_data_op` callback, or a registered
11 `ompt_callback_target_data_op_emi` callback with `ompt_scope_beginend` as its
12 endpoint argument for each occurrence of a *target-data-associate* event in that thread. These
13 callbacks have type signature `ompt_callback_target_data_op_t` or
14 `ompt_callback_target_data_op_emi_t`, respectively.
15

Restrictions

16 Restrictions to the `omp_target_associate_ptr` routine are as follows.
17

- 18 • When called from within a `target` region the effect is unspecified.

Cross References

- 19 • `target` directive, see [Section 14.8](#)
- 20 • `omp_target_alloc`, see [Section 19.8.1](#)
- 21 • `omp_target_disassociate_ptr`, see [Section 19.8.12](#)
- 22 • `omp_target_is_present`, see [Section 19.8.3](#)
- 23 • `ompt_callback_target_data_op_emi_t` and
24 `ompt_callback_target_data_op_t`, see [Section 20.5.2.25](#)
25

19.8.12 `omp_target_disassociate_ptr`

Summary

27 The `omp_target_disassociate_ptr` removes the associated pointer for a given device
28 from a host pointer.
29

Format

C / C++

```
30 | int omp_target_disassociate_ptr(const void *ptr, int device_num);
```

C / C++

Fortran

```
1 integer(c_int) function omp_target_disassociate_ptr(ptr, &  
2 device_num) bind(c)  
3 use, intrinsic :: iso_c_binding, only : c_ptr, c_int  
4 type(c_ptr), value :: ptr  
5 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 14.8](#)
- `ompt_callback_target_data_op_emi_t` and `ompt_callback_target_data_op_t`, see [Section 20.5.2.25](#)

19.8.13 `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 19.7.8](#)

19.9 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

C / C++

The type `omp_lock_t` represents a simple lock. For the following routines, a simple lock variable must be of `omp_lock_t` type. All simple lock routines require an argument that is a pointer to a variable of type `omp_lock_t`.

C / C++

Fortran

For the following routines, a simple lock variable must be an integer variable of `kind=omp_lock_kind`.

Fortran

The simple lock routines are as follows:

- The `omp_init_lock` routine initializes a simple lock;
- The `omp_init_lock_with_hint` routine initializes a simple lock and attaches a hint to it;
- The `omp_destroy_lock` routine uninitialized a simple lock;
- The `omp_set_lock` routine waits until a simple lock is available and then sets it;
- The `omp_unset_lock` routine unsets a simple lock; and
- The `omp_test_lock` routine tests a simple lock and sets it if it is available.

Nestable Lock Routines

C / C++

The type `omp_nest_lock_t` represents a nestable lock. For the following routines, a nestable lock variable must be of `omp_nest_lock_t` type. All nestable lock routines require an argument that is a pointer to a variable of type `omp_nest_lock_t`.

C / C++

Fortran

For the following routines, a nestable lock variable must be an integer variable of `kind=omp_nest_lock_kind`.

Fortran

The nestable lock routines are as follows:

- The `omp_init_nest_lock` routine initializes a nestable lock;
- The `omp_init_nest_lock_with_hint` routine initializes a nestable lock and attaches a hint to it;
- The `omp_destroy_nest_lock` routine uninitialized a nestable lock;

- The `omp_set_nest_lock` routine waits until a nestable lock is available and then sets it;
- The `omp_unset_nest_lock` routine unsets a nestable lock; and
- The `omp_test_nest_lock` routine tests a nestable lock and sets it if it is available.

Restrictions

Restrictions to OpenMP lock routines are as follows:

- The use of the same OpenMP lock in different contention groups results in unspecified behavior.

19.9.1 `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.

1 Tool Callbacks

2 A thread dispatches a registered `ompt_callback_lock_init` callback with
3 `omp_sync_hint_none` as the *hint* argument and `ompt_mutex_lock` as the *kind* argument
4 for each occurrence of a *lock-init* event in that thread. Similarly, a thread dispatches a registered
5 `ompt_callback_lock_init` callback with `omp_sync_hint_none` as the *hint* argument
6 and `ompt_mutex_nest_lock` as the *kind* argument for each occurrence of a *nest-lock-init*
7 event in that thread. These callbacks have the type signature
8 `ompt_callback_mutex_acquire_t` and occur in the task that encounters the routine.

9 Cross References

- 10 • `ompt_callback_mutex_acquire_t`, see [Section 20.5.2.14](#)

11 19.9.2 `omp_init_lock_with_hint` and 12 `omp_init_nest_lock_with_hint`

13 Summary

14 These routines initialize an OpenMP lock with a hint. The effect of the hint is
15 implementation-defined. The OpenMP implementation can ignore the hint without changing
16 program semantics.

17 Format

C / C++

```
18 void omp_init_lock_with_hint(  
19     omp_lock_t *lock,  
20     omp_sync_hint_t hint  
21 );  
22 void omp_init_nest_lock_with_hint(  
23     omp_nest_lock_t *lock,  
24     omp_sync_hint_t hint  
25 );
```

C / C++

Fortran

```
26 subroutine omp_init_lock_with_hint(svar, hint)  
27     integer (kind=omp_lock_kind) svar  
28     integer (kind=omp_sync_hint_kind) hint  
29  
30 subroutine omp_init_nest_lock_with_hint(nvar, hint)  
31     integer (kind=omp_nest_lock_kind) nvar  
32     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 16.1](#).

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.

Tool Callbacks

A thread dispatches a registered `ompt_callback_lock_init` callback with the same value for its *hint* argument as the *hint* argument of the call to `omp_init_lock_with_hint` and `ompt_mutex_lock` as the *kind* argument for each occurrence of a *lock-init-with-hint* event in that thread. Similarly, a thread dispatches a registered `ompt_callback_lock_init` callback with the same value for its *hint* argument as the *hint* argument of the call to `omp_init_nest_lock_with_hint` and `ompt_mutex_nest_lock` as the *kind* argument for each occurrence of a *nest-lock-init-with-hint* event in that thread. These callbacks have the type signature `ompt_callback_mutex_acquire_t` and occur in the task that encounters the routine.

Cross References

- Synchronization Hints, see [Section 16.1](#)
- `ompt_callback_mutex_acquire_t`, see [Section 20.5.2.14](#)

19.9.3 `omp_destroy_lock` and `omp_destroy_nest_lock`

Summary

These routines ensure that the OpenMP lock is uninitialized.

Format

```
void omp_destroy_lock(omp_lock_t *lock);
void omp_destroy_nest_lock(omp_nest_lock_t *lock);
```

Fortran

```
1  subroutine omp_destroy_lock(svar)
2  integer (kind=omp_lock_kind) svar
3
4  subroutine omp_destroy_nest_lock(nvar)
5  integer (kind=omp_nest_lock_kind) nvar
```

Fortran

Constraints on Arguments

A program that accesses a lock that is not in the unlocked state through either routine is non-conforming.

Effect

The effect of these routines is to change the state of the lock to uninitialized.

Execution Model Events

The *lock-destroy* event occurs in a thread that executes an `omp_destroy_lock` region before it finishes the region. The *nest-lock-destroy* 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 20.5.2.15](#)

19.9.4 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

```
1  subroutine omp_set_lock(svar)
2  integer (kind=omp_lock_kind) svar
3
4  subroutine omp_set_nest_lock(nvar)
5  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.

1 Tool Callbacks

2 A thread dispatches a registered `ompt_callback_mutex_acquire` callback for each
3 occurrence of a *lock-acquire* or *nest-lock-acquire* event in that thread. This callback has the type
4 signature `ompt_callback_mutex_acquire_t`.

5 A thread dispatches a registered `ompt_callback_mutex_acquired` callback for each
6 occurrence of a *lock-acquired* or *nest-lock-acquired* event in that thread. This callback has the type
7 signature `ompt_callback_mutex_t`.

8 A thread dispatches a registered `ompt_callback_nest_lock` callback with
9 `ompt_scope_begin` as its *endpoint* argument for each occurrence of a *nest-lock-owned* event in
10 that thread. This callback has the type signature `ompt_callback_nest_lock_t`.

11 The above callbacks occur in the task that encounters the lock function. The *kind* argument of these
12 callbacks is `ompt_mutex_lock` when the events arise from an `omp_set_lock` region while it
13 is `ompt_mutex_nest_lock` when the events arise from an `omp_set_nest_lock` region.

14 Cross References

- 15 • `ompt_callback_mutex_acquire_t`, see [Section 20.5.2.14](#)
- 16 • `ompt_callback_mutex_t`, see [Section 20.5.2.15](#)
- 17 • `ompt_callback_nest_lock_t`, see [Section 20.5.2.16](#)

18 19.9.5 omp_unset_lock and omp_unset_nest_lock

19 Summary

20 These routines provide the means of unsetting an OpenMP lock.

21 Format

C / C++

```
22 void omp_unset_lock(omp_lock_t *lock);  
23 void omp_unset_nest_lock(omp_nest_lock_t *lock);
```

C / C++

Fortran

```
24 subroutine omp_unset_lock(svar)  
25 integer (kind=omp_lock_kind) svar  
26  
27 subroutine omp_unset_nest_lock(nvar)  
28 integer (kind=omp_nest_lock_kind) nvar
```

Fortran

29 Constraints on Arguments

30 A program that accesses a lock that is not in the locked state or that is not owned by the task that
31 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.

Tool Callbacks

A thread dispatches a registered `ompt_callback_mutex_released` callback with `ompt_mutex_lock` as the *kind* argument for each occurrence of a *lock-release* event in that thread. Similarly, a thread dispatches a registered `ompt_callback_mutex_released` callback with `ompt_mutex_nest_lock` as the *kind* argument for each occurrence of a *nest-lock-release* event in that thread. These callbacks have the type signature `ompt_callback_mutex_t` and occur in the task that encounters the routine.

A thread dispatches a registered `ompt_callback_nest_lock` callback with `ompt_scope_end` as its *endpoint* argument for each occurrence of a *nest-lock-held* event in that thread. This callback has the type signature `ompt_callback_nest_lock_t`.

Cross References

- `ompt_callback_mutex_t`, see [Section 20.5.2.15](#)
- `ompt_callback_nest_lock_t`, see [Section 20.5.2.16](#)

19.9.6 `omp_test_lock` and `omp_test_nest_lock`

Summary

These routines attempt to set an OpenMP lock but do not suspend execution of the task that executes the routine.

Format

```
int omp_test_lock(omp_lock_t *lock);
int omp_test_nest_lock(omp_nest_lock_t *lock);
```

C / C++

Fortran

```
1 logical function omp_test_lock(svar)
2 integer (kind=omp_lock_kind) svar
3
4 integer function omp_test_nest_lock(nvar)
5 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. The behavior is unspecified if a simple lock accessed by `omp_test_lock` is in the locked state and is owned by the task that contains the call.

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`.

1 The above callbacks occur in the task that encounters the lock function. The *kind* argument of these
2 callbacks is `ompt_mutex_test_lock` when the events arise from an `omp_test_lock`
3 region while it is `ompt_mutex_test_nest_lock` when the events arise from an
4 `omp_test_nest_lock` region.

5 Cross References

- 6 • `ompt_callback_mutex_acquire_t`, see [Section 20.5.2.14](#)
- 7 • `ompt_callback_mutex_t`, see [Section 20.5.2.15](#)
- 8 • `ompt_callback_nest_lock_t`, see [Section 20.5.2.16](#)

9 19.10 Timing Routines

10 This section describes routines that support a portable wall clock timer.

11 19.10.1 `omp_get_wtime`

12 Summary

13 The `omp_get_wtime` routine returns elapsed wall clock time in seconds.

14 Format

15	<code>double omp_get_wtime(void);</code>	C / C++
		Fortran
16	<code>double precision function omp_get_wtime()</code>	Fortran

17 Binding

18 The binding thread set for an `omp_get_wtime` region is the encountering thread. The routine's
19 return value is not guaranteed to be consistent across any set of threads.

20 Effect

21 The `omp_get_wtime` routine returns a value equal to the elapsed wall clock time in seconds
22 since some *time-in-the-past*. The actual *time-in-the-past* is arbitrary, but it is guaranteed not to
23 change during the execution of the application program. The time returned is a *per-thread time*, so
24 it is not required to be globally consistent across all threads that participate in an application.

19.10.2 `omp_get_wtick`

Summary

The `omp_get_wtick` routine returns the precision of the timer used by `omp_get_wtime`.

Format

C / C++
`double omp_get_wtick(void);`

C / C++
Fortran
`double precision function omp_get_wtick()`

Fortran

Binding

The binding thread set for an `omp_get_wtick` region is the encountering thread. The routine's return value is not guaranteed to be consistent across any set of threads.

Effect

The `omp_get_wtick` routine returns a value equal to the number of seconds between successive clock ticks of the timer used by `omp_get_wtime`.

19.11 Event Routine

This section describes a routine that supports OpenMP event objects.

Binding

The binding thread set for all event routine regions is the encountering thread.

19.11.1 `omp_fulfill_event`

Summary

This routine fulfills and destroys an OpenMP event.

Format

C / C++
`void omp_fulfill_event(omp_event_handle_t event);`

C / C++
Fortran
`subroutine omp_fulfill_event(event)
integer (kind=omp_event_handle_kind) event`

Fortran

Constraints on Arguments

A program that calls this routine on an event that was already fulfilled is non-conforming. A program that calls this routine with an event handle that was not created by the **detach** clause is non-conforming.

Effect

The effect of this routine is to fulfill the event associated with the event handle argument. The effect of fulfilling the event will depend on how the event was created. The event is destroyed and cannot be accessed after calling this routine, and the event handle becomes unassociated with any event.

Execution Model Events

The *task-fulfill* event occurs in a thread that executes an **omp_fulfill_event** region before the event is fulfilled if the OpenMP event object was created by a **detach** clause on a task.

Tool Callbacks

A thread dispatches a registered **ompt_callback_task_schedule** callback with **NULL** as its *next_task_data* argument while the argument *prior_task_data* binds to the detachable task for each occurrence of a *task-fulfill* event. If the *task-fulfill* event occurs before the detachable task finished the execution of the associated *structured-block*, the callback has **ompt_task_early_fulfill** as its *prior_task_status* argument; otherwise the callback has **ompt_task_late_fulfill** as its *prior_task_status* argument. This callback has type signature **ompt_callback_task_schedule_t**.

Restrictions

Restrictions to the **omp_fulfill_event** routine are as follows:

- The event handler passed to the routine must have been created by a thread in the same device as the thread that invoked the routine.

Cross References

- **detach** clause, see [Section 13.6.2](#)
- **ompt_callback_task_schedule_t**, see [Section 20.5.2.10](#)

▼ C / C++ ▼

19.12 Interoperability Routines

The interoperability routines provide mechanisms to inspect the properties associated with an **omp_interop_t** object. Such objects may be initialized, destroyed or otherwise used by an **interop** construct. Additionally, an **omp_interop_t** object can be initialized to **omp_interop_none**, which is defined to be zero. An **omp_interop_t** object may only be accessed or modified through OpenMP directives and API routines.

An **omp_interop_t** object can be copied without affecting, or copying, the underlying state. Destruction of an **omp_interop_t** object destroys the state to which all copies of the object refer.

TABLE 19.1: 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>			

1 OpenMP reserves all negative values for properties, as listed in Table 19.1; implementation-defined
 2 properties may use zero and positive values. The special property, `omp_ipr_first`, will always
 3 have the lowest property value, which may change in future versions of this specification. Valid
 4 values and types for the properties that Table 19.1 lists are specified in the *OpenMP Additional*
 5 *Definitions* document or are implementation defined unless otherwise specified.

6 Table 19.2 lists the return codes used by routines that take an `int* ret_code` argument.

7 Binding

8 The binding task set for all interoperability routine regions is the generating task.



9 19.12.1 omp_get_num_interop_properties

10 Summary

11 The `omp_get_num_interop_properties` routine retrieves the number of
 12 implementation-defined properties available for an `omp_interop_t` object.

TABLE 19.2: 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 19.1
<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>

1 **Format**

2 `int omp_get_num_interop_properties(const omp_interop_t interop);`

3 **Effect**

4 The `omp_get_num_interop_properties` routine returns the number of
 5 implementation-defined properties available for *interop*. The total number of properties available
 6 for *interop* is the returned value minus `omp_ipr_first`.



7 **19.12.2 omp_get_interop_int**

8 **Summary**

9 The `omp_get_interop_int` routine retrieves an integer property from an `omp_interop_t`
 10 object.

11 **Format**

12 `omp_intptr_t omp_get_interop_int(const omp_interop_t interop,`
 13 `omp_interop_property_t property_id,`
 14 `int *ret_code);`

15 **Effect**

16 The `omp_get_interop_int` routine returns the requested integer property, if available, and
 17 zero if an error occurs or no value is available. If the *interop* is `omp_interop_none`, an empty
 18 error occurs. If the *property_id* is less than `omp_ipr_first` or greater than or equal to
 19 `omp_get_num_interop_properties(interop)`, an out of range error occurs. If the
 20 requested property value is not convertible into an integer value, a type error occurs.

21 If a [non-null pointer](#) is passed to *ret_code*, an `omp_interop_rc_t` value that indicates the
 22 return code is stored in the object to which *ret_code* points. If an error occurred, the stored value
 23 will be negative and it will match the error as defined in Table 19.2. On success, zero will be stored.

1 If no error occurred but no meaningful value can be returned, `omp_irc_no_value`, which is
2 one, will be stored.

3 **Restrictions**

4 Restrictions to the `omp_get_interop_int` routine are as follows:

- 5 • The behavior of the routine is unspecified if an invalid `omp_interop_t` object is provided.

6 **Cross References**

- 7 • `omp_get_num_interop_properties`, see [Section 19.12.1](#)

▲ C / C++ ▲

▼ C / C++ ▼

8 **19.12.3 omp_get_interop_ptr**

9 **Summary**

10 The `omp_get_interop_ptr` routine retrieves a pointer property from an `omp_interop_t`
11 object.

12 **Format**

```
13 void* omp_get_interop_ptr(const omp_interop_t interop,  
14                          omp_interop_property_t property_id,  
15                          int *ret_code);
```

16 **Effect**

17 The `omp_get_interop_ptr` routine returns the requested pointer property, if available, and
18 `NULL` if an error occurs or no value is available. If the `interop` is `omp_interop_none`, an
19 empty error occurs. If the `property_id` is less than `omp_ipr_first` or greater than or equal to
20 `omp_get_num_interop_properties(interop)`, an out of range error occurs. If the
21 requested property value is not convertible into a pointer value, a type error occurs.

22 If a [non-null pointer](#) is passed to `ret_code`, an `omp_interop_rc_t` value that indicates the
23 return code is stored in the object to which the `ret_code` points. If an error occurred, the stored
24 value will be negative and it will match the error as defined in [Table 19.2](#). On success, zero will be
25 stored. If no error occurred but no meaningful value can be returned, `omp_irc_no_value`,
26 which is one, will be stored.

27 **Restrictions**

28 Restrictions to the `omp_get_interop_ptr` routine are as follows:

- 29 • The behavior of the routine is unspecified if an invalid `omp_interop_t` object is provided.
- 30 • Memory referenced by the pointer returned by the `omp_get_interop_ptr` routine is
31 managed by the OpenMP implementation and should not be freed or modified.

32 **Cross References**

- 33 • `omp_get_num_interop_properties`, see [Section 19.12.1](#)

▲ C / C++ ▲

19.12.4 `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 19.2](#). 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 19.12.1](#)

19.12.5 `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 19.1. 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 19.12.1](#)

← C / C++ →

→ C / C++ ←

19.12.6 `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 19.12.1](#)

▲ C / C++ ▲

▼ C / C++ ▼

19.12.7 `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++ ▲

19.13 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++

19.13.1 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,
```

```

1      omp_atv_thread = 8,
2      omp_atv_pteam = 9,
3      omp_atv_cgroup = 10,
4      omp_atv_default_mem_fb = 11,
5      omp_atv_null_fb = 12,
6      omp_atv_abort_fb = 13,
7      omp_atv_allocator_fb = 14,
8      omp_atv_environment = 15,
9      omp_atv_nearest = 16,
10     omp_atv_blocked = 17,
11     omp_atv_interleaved = 18,
12     omp_atv_all = 19,
13     omp_atv_single = 20,
14     omp_atv_multiple = 21,
15     omp_atv_memspace = 22
16 } omp_alloctrail_value_t;
17
18 typedef struct omp_alloctrail_t {
19     omp_alloctrail_key_t key;
20     omp_uintptr_t value;
21 } omp_alloctrail_t;

```

C / C++

Fortran

```

22 integer(kind=omp_alloctrail_key_kind), &
23     parameter :: omp_atk_sync_hint = 1
24 integer(kind=omp_alloctrail_key_kind), &
25     parameter :: omp_atk_alignment = 2
26 integer(kind=omp_alloctrail_key_kind), &
27     parameter :: omp_atk_access = 3
28 integer(kind=omp_alloctrail_key_kind), &
29     parameter :: omp_atk_pool_size = 4
30 integer(kind=omp_alloctrail_key_kind), &
31     parameter :: omp_atk_fallback = 5
32 integer(kind=omp_alloctrail_key_kind), &
33     parameter :: omp_atk_fb_data = 6
34 integer(kind=omp_alloctrail_key_kind), &
35     parameter :: omp_atk_pinned = 7
36 integer(kind=omp_alloctrail_key_kind), &
37     parameter :: omp_atk_partition = 8
38 integer(kind=omp_alloctrail_key_kind), &
39     parameter :: omp_atk_pin_device = 9
40 integer(kind=omp_alloctrail_key_kind), &
41     parameter :: omp_atk_preferred_device = 10

```

```

1 integer(kind=omp_alloctrail_key_kind), &
2   parameter :: omp_atk_device_access = 11
3 integer(kind=omp_alloctrail_key_kind), &
4   parameter :: omp_atk_target_access = 12
5 integer(kind=omp_alloctrail_key_kind), &
6   parameter :: omp_atk_atomic_scope = 13
7 integer(kind=omp_alloctrail_key_kind), &
8   parameter :: omp_atk_part_size = 14
9
10 integer(kind=omp_alloctrail_val_kind), &
11   parameter :: omp_atv_default = -1
12 integer(kind=omp_alloctrail_val_kind), &
13   parameter :: omp_atv_false = 0
14 integer(kind=omp_alloctrail_val_kind), &
15   parameter :: omp_atv_true = 1
16 integer(kind=omp_alloctrail_val_kind), &
17   parameter :: omp_atv_contended = 3
18 integer(kind=omp_alloctrail_val_kind), &
19   parameter :: omp_atv_uncontended = 4
20 integer(kind=omp_alloctrail_val_kind), &
21   parameter :: omp_atv_serialized = 5
22 integer(kind=omp_alloctrail_val_kind), &
23   parameter :: omp_atv_private = 6
24 integer(kind=omp_alloctrail_val_kind), &
25   parameter :: omp_atv_device = 7
26 integer(kind=omp_alloctrail_val_kind), &
27   parameter :: omp_atv_thread = 8
28 integer(kind=omp_alloctrail_val_kind), &
29   parameter :: omp_atv_pteam = 9
30 integer(kind=omp_alloctrail_val_kind), &
31   parameter :: omp_atv_cgroup = 10
32 integer(kind=omp_alloctrail_val_kind), &
33   parameter :: omp_atv_default_mem_fb = 11
34 integer(kind=omp_alloctrail_val_kind), &
35   parameter :: omp_atv_null_fb = 12
36 integer(kind=omp_alloctrail_val_kind), &
37   parameter :: omp_atv_abort_fb = 13
38 integer(kind=omp_alloctrail_val_kind), &
39   parameter :: omp_atv_allocator_fb = 14
40 integer(kind=omp_alloctrail_val_kind), &
41   parameter :: omp_atv_environment = 15
42 integer(kind=omp_alloctrail_val_kind), &
43   parameter :: omp_atv_nearest = 16

```

```

1  integer(kind=omp_alloctrail_val_kind), &
2     parameter :: omp_atv_blocked = 17
3  integer(kind=omp_alloctrail_val_kind), &
4     parameter :: omp_atv_interleaved = 18
5  integer(kind=omp_alloctrail_val_kind), &
6     parameter :: omp_atv_all = 19
7  integer(kind=omp_alloctrail_val_kind), &
8     parameter :: omp_atv_single = 20
9  integer(kind=omp_alloctrail_val_kind), &
10     parameter :: omp_atv_multiple = 21
11 integer(kind=omp_alloctrail_val_kind), &
12     parameter :: omp_atv_memspace = 22
13
14 ! omp_alloctrail might not be provided in omp_lib.h.
15 type omp_alloctrail
16     integer(kind=omp_alloctrail_key_kind) key
17     integer(kind=omp_alloctrail_val_kind) value
18 end type omp_alloctrail
19
20 integer(kind=omp_memspace_handle_kind), &
21     parameter :: omp_null_mem_space = 0
22
23 integer(kind=omp_allocator_handle_kind), &
24     parameter :: omp_null_allocator = 0

```

Fortran

25 19.13.2 Memory Space Routines

26 Summary

27 The following routines return a memory space that represents a set of resources accessible by one
28 or more devices.

29 Format

C / C++

```

30 omp_memspace_handle_t omp_get_devices_memspace (
31     int ndevs,
32     const int *devs,
33     omp_memspace_handle_t memspace
34 );

```

```

1  omp_memspace_handle_t omp_get_device_memspace(
2      int dev,
3      omp_memspace_handle_t memspace
4  );
5
6  omp_memspace_handle_t omp_get_devices_and_host_memspace(
7      int ndevs,
8      const int *devs,
9      omp_memspace_handle_t memspace
10 );
11
12 omp_memspace_handle_t omp_get_device_and_host_memspace(
13     int dev,
14     omp_memspace_handle_t memspace
15 );
16
17 omp_memspace_handle_t omp_get_devices_all_memspace(
18     omp_memspace_handle_t memspace
19 );

```

↔ C / C++ ↔

↔ Fortran ↔

```

20 integer(kind=omp_memspace_handle_kind) &
21 function omp_get_devices_memspace(ndevs, devs, memspace)
22 integer, intent(in) :: ndevs
23 integer, intent(in) :: devs(*)
24 integer(kind=omp_memspace_handle_kind), intent(in) :: memspace
25
26 integer(kind=omp_memspace_handle_kind) &
27 function omp_get_device_memspace(dev, memspace)
28 integer, intent(in) :: dev
29 integer(kind=omp_memspace_handle_kind), intent(in) :: memspace
30
31 integer(kind=omp_memspace_handle_kind) &
32 function omp_get_devices_and_host_memspace(ndevs, devs, memspace)
33 integer, intent(in) :: ndevs
34 integer, intent(in) :: devs(*)
35 integer(kind=omp_memspace_handle_kind), intent(in) :: memspace
36
37 integer(kind=omp_memspace_handle_kind) &
38 function omp_get_device_and_host_memspace(dev, memspace)
39 integer, intent(in) :: dev
40 integer(kind=omp_memspace_handle_kind), intent(in) :: memspace
41

```

```
1 integer(kind=omp_memspace_handle_kind) &  
2 function omp_get_devices_all_memspace(memspace)  
3 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.

The *dev* argument to `omp_get_device_memspace` and `omp_get_device_and_host_memspace` must be a conforming device number.

Binding

The binding thread set for these routines region is all threads on the device.

Effect

The effect of these routines is to return a handle to a memory space that represents a set of storage resources such that for each storage resource the following requirements are true:

- The storage resource is accessible by each of the devices selected by the routine; and
- The storage resource is part of *memspace* in each of the devices selected by the routine.

If no set of storage resources matches the above requirements, then the special value `omp_null_mem_space` is returned.

The devices selected by `omp_get_devices_memspace` are those specified in the *devs* argument.

The device selected by `omp_get_device_memspace` is the device specified in the *dev* argument.

The devices selected by `omp_get_devices_and_host_memspace` are those specified in the *devs* argument and the initial device.

The device selected by `omp_get_device_and_host_memspace` are the device specified in the *dev* argument and the initial device.

The devices selected by `omp_get_devices_all_memspace` are all available devices.

The [memory spaces](#) returned by these routine are [target memory spaces](#) if any of the selected [devices](#) is not the current [device](#).

Restrictions

The restrictions to these routines are as follows:

- These routines must only be invoked on the initial device.

Cross References

- **requires** directive, see [Section 9.5](#)
- **target** directive, see [Section 14.8](#)
- Memory Spaces, see [Section 7.1](#)

19.13.3 omp_init_allocator

Summary

The `omp_init_allocator` routine initializes an allocator and associates it with a memory 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 7.2 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 9.5](#)
- `target` directive, see [Section 14.8](#)
- Memory Allocators, see [Section 7.2](#)
- Memory Spaces, see [Section 7.1](#)

19.13.4 Memory Allocator Routines

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
);
```



```

1  omp_allocator_handle_t omp_get_device_allocator(
2      int dev,
3      omp_memspace_handle_t memspace
4  );
5
6  omp_allocator_handle_t omp_get_devices_and_host_allocator(
7      int ndevs,
8      const int *devs,
9      omp_memspace_handle_t memspace
10 );
11
12 omp_allocator_handle_t omp_get_device_and_host_allocator(
13     int dev,
14     omp_memspace_handle_t memspace
15 );
16
17 omp_allocator_handle_t omp_get_devices_all_allocator(
18     omp_memspace_handle_t memspace
19 );

```

▲ C / C++ ▲

▼ Fortran ▼

```

20 integer(kind=omp_allocator_handle_kind) &
21 function omp_get_devices_allocator(ndevs, devs, memspace)
22 integer, intent(in) :: ndevs
23 integer, intent(in) :: devs(*)
24 integer(kind=omp_memspace_handle_kind), intent(in) :: memspace
25
26 integer(kind=omp_allocator_handle_kind) &
27 function omp_get_device_allocator(dev, memspace)
28 integer, intent(in) :: dev
29 integer(kind=omp_memspace_handle_kind), intent(in) :: memspace
30
31 integer(kind=omp_allocator_handle_kind) &
32 function omp_get_devices_and_host_allocator(ndevs, devs, memspace)
33 integer, intent(in) :: ndevs
34 integer, intent(in) :: devs(*)
35 integer(kind=omp_memspace_handle_kind), intent(in) :: memspace
36
37 integer(kind=omp_allocator_handle_kind) &
38 function omp_get_device_and_host_allocator(dev, memspace)
39 integer, intent(in) :: dev
40 integer(kind=omp_memspace_handle_kind), intent(in) :: memspace
41

```

```
1 integer(kind=omp_allocator_handle_kind) &  
2 function omp_get_devices_all_allocator(memspace)  
3 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_allocator` and `omp_get_devices_and_host_allocator` must be greater than zero. The *devs* argument to `omp_get_devices_allocator` and `omp_get_devices_and_host_allocator` 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.

The *dev* argument to `omp_get_device_allocator` and `omp_get_device_and_host_allocator` must be a conforming device number.

Binding

The binding thread set for these routines 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 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 7.3. 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 9.5](#)
- **target** directive, see [Section 14.8](#)
- Memory Allocators, see [Section 7.2](#)
- Memory Spaces, see [Section 7.1](#)

19.13.5 `omp_destroy_allocator`

Summary

The `omp_destroy_allocator` routine releases all resources used by the allocator handle.

Format

	C / C++	
<pre>void omp_destroy_allocator(omp_allocator_handle_t allocator);</pre>		
	C / C++	
	Fortran	
<pre>subroutine omp_destroy_allocator(allocator) integer(kind=omp_allocator_handle_kind), intent(in) :: allocator</pre>		
	Fortran	

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 9.5](#)
- **target** directive, see [Section 14.8](#)
- Memory Allocators, see [Section 7.2](#)

19.13.6 `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 7.6](#)
- `allocate` directive, see [Section 7.5](#)
- `allocators` directive, see [Section 7.7](#)
- Memory Allocators, see [Section 7.2](#)
- *def-allocator-var* ICV, see [Table 2.1](#)

19.13.7 `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++
Fortran
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 7.6](#)
- `allocate` directive, see [Section 7.5](#)
- `allocators` directive, see [Section 7.7](#)
- Memory Allocators, see [Section 7.2](#)
- *def-allocator-var* ICV, see [Table 2.1](#)

19.13.8 `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++

```
1 void *omp_alloc(  
2     size_t size,  
3     omp_allocator_handle_t allocator=omp_null_allocator  
4 );  
5 void *omp_aligned_alloc(  
6     size_t alignment,  
7     size_t size,  
8     omp_allocator_handle_t allocator=omp_null_allocator  
9 );
```

C++

Fortran

```
10 type(c_ptr) function omp_alloc(size, allocator) bind(c)  
11 use, intrinsic :: iso_c_binding, only : c_ptr, c_size_t  
12 integer(c_size_t), value :: size  
13 integer(omp_allocator_handle_kind), value :: allocator  
14  
15 type(c_ptr) function omp_aligned_alloc(alignment, &  
16     size, allocator) bind(c)  
17 use, intrinsic :: iso_c_binding, only : c_ptr, c_size_t  
18 integer(c_size_t), value :: alignment, size  
19 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`.

1 Memory allocated by `omp_alloc` will be byte-aligned to at least the maximum of the alignment
2 required by `malloc` and the `alignment` trait of the allocator. Memory allocated by
3 `omp_aligned_alloc` will be byte-aligned to at least the maximum of the alignment required by
4 `malloc`, the `alignment` trait of the allocator and the `alignment` argument value.

5 Pointers returned by these routines are considered device pointers if at least one of the devices
6 associated with the allocator is not the current device.

Fortran

7 The `omp_alloc` and `omp_aligned_alloc` routines require an explicit interface and so might
8 not be provided in `omp_lib.h`.

Fortran

Cross References

- `requires` directive, see [Section 9.5](#)
- `target` directive, see [Section 14.8](#)
- Memory Allocators, see [Section 7.2](#)
- `def-allocator-var` ICV, see [Table 2.1](#)

19.13.9 `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 7.2](#)
- `omp_destroy_allocator`, see [Section 19.13.5](#)

19.13.10 `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

```
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

C++

```
1 void *omp_malloc(  
2     size_t nmemb,  
3     size_t size,  
4     omp_allocator_handle_t allocator=omp_null_allocator  
5 );  
6 void *omp_aligned_malloc(  
7     size_t alignment,  
8     size_t nmemb,  
9     size_t size,  
10    omp_allocator_handle_t allocator=omp_null_allocator  
11 );
```

C++

Fortran

```
12 type(c_ptr) function omp_malloc(nmemb, size, allocator) bind(c)  
13 use, intrinsic :: iso_c_binding, only : c_ptr, c_size_t  
14 integer(c_size_t), value :: nmemb, size  
15 integer(omp_allocator_handle_kind), value :: allocator  
16  
17 type(c_ptr) function omp_aligned_malloc(alignment, nmemb, size, &  
18     allocator) bind(c)  
19 use, intrinsic :: iso_c_binding, only : c_ptr, c_size_t  
20 integer(c_size_t), value :: alignment, nmemb, size  
21 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` and `omp_aligned_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 9.5](#)
- `target` directive, see [Section 14.8](#)
- Memory Allocators, see [Section 7.2](#)
- `def-allocator-var` ICV, see [Table 2.1](#)

19.13.11 `omp_realloc`

Summary

The `omp_realloc` routine deallocates previously allocated memory and requests a memory allocation from a memory allocator.

Format

```
void *omp_realloc(  
    void *ptr,  
    size_t size,  
    omp_allocator_handle_t allocator,  
    omp_allocator_handle_t free_allocator  
);
```

C

C

C++

```
1 void *omp_realloc(  
2     void *ptr,  
3     size_t size,  
4     omp_allocator_handle_t allocator=omp_null_allocator,  
5     omp_allocator_handle_t free_allocator=omp_null_allocator  
6 );
```

C++

Fortran

```
7 type(c_ptr) &  
8 function omp_realloc(ptr, size, allocator, free_allocator) bind(c)  
9 use, intrinsic :: iso_c_binding, only : c_ptr, c_size_t  
10 type(c_ptr), value :: ptr  
11 integer(c_size_t), value :: size  
12 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 9.5](#)
- `target` directive, see [Section 14.8](#)
- Memory Allocators, see [Section 7.2](#)
- `omp_alloc` and `omp_aligned_alloc`, see [Section 19.13.8](#)
- `omp_destroy_allocator`, see [Section 19.13.5](#)

19.13.12 `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 7.1](#)

19.13.13 `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 7.1](#)

19.14 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
```

Fortran

Constraints on Arguments

The following enumeration type defines four standard commands. Table 19.3 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

```
1 integer (kind=omp_control_tool_kind), &  
2     parameter :: omp_control_tool_start = 1  
3 integer (kind=omp_control_tool_kind), &  
4     parameter :: omp_control_tool_pause = 2  
5 integer (kind=omp_control_tool_kind), &  
6     parameter :: omp_control_tool_flush = 3  
7 integer (kind=omp_control_tool_kind), &  
8     parameter :: omp_control_tool_end = 4
```

Fortran

9 Tool-specific values for *command* must be greater or equal to 64. Tools must ignore *command*
10 values that they are not explicitly designed to handle. Other values accepted by a tool for *command*,
11 and any values for *modifier* and *arg* are tool-defined.

TABLE 19.3: 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.

Binding

12 The binding task set for an `omp_control_tool` region is the generating task.

Effect

13
14 An OpenMP program may use `omp_control_tool` to pass commands to a tool. An application
15 can use `omp_control_tool` to request that a tool starts or restarts data collection when a code
16 region of interest is encountered, that a tool pauses data collection when leaving the region of
17 interest, that a tool flushes any data that it has collected so far, or that a tool ends data collection.
18 Additionally, `omp_control_tool` can be used to pass tool-specific commands to a particular
19 tool. The following types correspond to return values from `omp_control_tool`:
20

C / C++

```
1 typedef enum omp_control_tool_result_t {  
2     omp_control_tool_notool = -2,  
3     omp_control_tool_nocallback = -1,  
4     omp_control_tool_success = 0,  
5     omp_control_tool_ignored = 1  
6 } omp_control_tool_result_t;
```

C / C++

Fortran

```
7 integer (kind=omp_control_tool_result_kind), &  
8     parameter :: omp_control_tool_notool = -2  
9 integer (kind=omp_control_tool_result_kind), &  
10     parameter :: omp_control_tool_nocallback = -1  
11 integer (kind=omp_control_tool_result_kind), &  
12     parameter :: omp_control_tool_success = 0  
13 integer (kind=omp_control_tool_result_kind), &  
14     parameter :: omp_control_tool_ignored = 1
```

Fortran

15 If the **OMPT interface state** is **OMPT inactive**, the OpenMP implementation returns
16 **omp_control_tool_notool**. If the **OMPT interface state** is **OMPT active**, but no callback is
17 registered for the *tool-control* event, the OpenMP implementation returns
18 **omp_control_tool_nocallback**. An OpenMP implementation may return other
19 **implementation defined** negative values strictly smaller than -64; an **OpenMP program** may assume
20 that any negative return value indicates that a **tool** has not received the command. A return value of
21 **omp_control_tool_success** indicates that the **tool** has performed the specified command. A
22 return value of **omp_control_tool_ignored** indicates that the **tool** has ignored the specified
23 command. A **tool** may return other positive values strictly greater than 64 that are tool-defined.

24 Execution Model Events

25 The *tool-control* event occurs in the **thread** that encounters a call to **omp_control_tool** at a
26 point inside its corresponding **region**.

27 Tool Callbacks

28 A **thread** dispatches a registered **ompt_callback_control_tool** callback for each
29 occurrence of a *tool-control* event. The **callback** executes in the context of the call that occurs in the
30 user program and has type signature **ompt_callback_control_tool_t**. The **callback** may
31 return any non-negative value, which will be returned to the **OpenMP program** by the OpenMP
32 implementation as the return value of the **omp_control_tool** call that triggered the **callback**.

33 Arguments passed to the **callback** are those passed by the user to **omp_control_tool**. If the
34 call is made in Fortran, the **tool** will be passed **NULL** as the third argument to the **callback**. If any
35 of the four standard commands is presented to a **tool**, the **tool** will ignore the *modifier* and *arg*
36 argument values.

Restrictions

Restrictions on access to the state of an OpenMP [first-party tool](#) are as follows:

- An [OpenMP program](#) may access the [tool](#) state modified by an OMPT [callback](#) only by using `omp_control_tool`.

Cross References

- OMPT Interface, see [Chapter 20](#)
- `ompt_callback_control_tool_t`, see [Section 20.5.2.29](#)

19.15 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 3](#).

Format

C / C++	▼
<code>void omp_display_env(int verbose);</code>	
C / C++	▲
Fortran	▼
<code>subroutine omp_display_env(verbose) logical,intent(in) :: verbose</code>	
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 3](#). 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".

1 If the *verbose* argument evaluates to *false*, the runtime displays the OpenMP version number
2 defined by the `_OPENMP` version macro (or the `openmp_version` named constant for Fortran)
3 value and the initial ICV values for the environment variables listed in [Chapter 3](#). If the *verbose*
4 argument evaluates to *true*, the runtime may also display the values of vendor-specific ICVs that
5 may be modified by vendor-specific environment variables.

6 Example output:

```
7 OPENMP DISPLAY ENVIRONMENT BEGIN  
8   _OPENMP=' 202111'  
9   [host] OMP_SCHEDULE=' GUIDED, 4'  
10  [host] OMP_NUM_THREADS=' 4, 3, 2'  
11  [device] OMP_NUM_THREADS=' 2'  
12  [host,device] OMP_DYNAMIC=' TRUE'  
13  [host] OMP_PLACES=' {0:4}, {4:4}, {8:4}, {12:4}'  
14  ...  
15 OPENMP DISPLAY ENVIRONMENT END
```

16 Restrictions

17 Restrictions to the `omp_display_env` routine are as follows.

- 18 • When called from within a `target` region the effect is unspecified.

19 Cross References

- 20 • `OMP_DISPLAY_ENV`, see [Section 3.7](#)

1

Part IV

2

Tool Interfaces

20 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 **thread state** associated with a **thread**, to interpret the call stack of a **thread**, to receive notification about **events**, to trace activity on **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 program**.

20.1 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++

20.2 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.

20.2.1 **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 **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 20.4.1](#)




FIGURE 20.1: First-Party Tool Activation Flow Chart

20.2.2 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 OpenMP program executes. In this case, the OMPT interface state remains OMPT inactive.

Otherwise, the OMPT interface state changes to OMPT 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 program;
- By introducing a dynamically-linked library that includes its definition of *ompt_start_tool* into the program's address space; or
- By providing, in the *tool-libraries-var* ICV, the name of a dynamically-linked library that is appropriate for the OpenMP architecture and operating system used by the OpenMP program

1 and that 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 [OpenMP program](#) or in a dynamically-linked library loaded in the [address
6 space](#). If multiple implementations of `ompt_start_tool` are available, the OpenMP
7 implementation will use the first tool-provided implementation of `ompt_start_tool` that it
8 finds.

9 If the implementation does not find a tool-provided implementation of `ompt_start_tool` in the
10 [address space](#), it consults the `tool-libraries-var` *ICV*, which contains a (possibly empty) list of
11 dynamically-linked libraries. As described in detail in [Section 3.3.2](#), the libraries in
12 `tool-libraries-var` are then searched for the first usable implementation of `ompt_start_tool`
13 that one of the libraries in the list provides.

14 If the implementation finds a tool-provided definition of `ompt_start_tool`, it invokes that
15 method; if a `NULL` pointer is returned, the [OMPT interface state](#) remains [OMPT pending](#) and the
16 implementation continues to look for implementations of `ompt_start_tool`; otherwise a
17 [non-null pointer](#) to an `ompt_start_tool_result_t` structure is returned, the [OMPT
18 interface state](#) changes to [OMPT active](#) and the OpenMP implementation makes the [OMPT
19 interface](#) available as the program executes. In this case, as the OpenMP implementation completes
20 its initialization, it initializes the [OMPT interface](#).

21 If no `tool` can be found, the [OMPT interface state](#) changes to [OMPT inactive](#).

22 **Cross References**

- 23 • Tool Initialization and Finalization, see [Section 20.4.1](#)
- 24 • `tool-libraries-var` *ICV*, see [Table 2.1](#)
- 25 • `tool-var` *ICV*, see [Table 2.1](#)
- 26 • `ompt_start_tool`, see [Section 20.2.1](#)

27 **20.2.3 Initializing a First-Party Tool**

28 To initialize the [OMPT interface](#), the OpenMP implementation invokes the `tool` initializer that is
29 specified in the `ompt_start_tool_result_t` [structure](#) that is indicated by the [non-null
30 pointer](#) that `ompt_start_tool` returns. The initializer is invoked prior to the occurrence of any
31 OpenMP [event](#).

32 A tool initializer, described in [Section 20.5.1.1](#), uses the function specified in its *lookup* argument
33 to look up pointers to OMPT interface runtime entry points that the OpenMP implementation
34 provides; this process is described in [Section 20.2.3.1](#). Typically, a tool initializer obtains a pointer
35 to the `ompt_set_callback` runtime entry point with type signature

1 `ompt_set_callback_t` and then uses this runtime entry point to perform [callback registration](#)
2 for [events](#), as described in [Section 20.2.4](#).

3 A tool initializer may use the `ompt_enumerate_states` runtime entry point, which has type
4 signature `ompt_enumerate_states_t`, to determine the thread states that an OpenMP
5 implementation employs. Similarly, it may use the `ompt_enumerate_mutex_impls` runtime
6 entry point, which has type signature `ompt_enumerate_mutex_impls_t`, to determine the
7 mutual exclusion implementations that the OpenMP implementation employs.

8 If a tool initializer returns a non-zero value, the OMPT interface state remains *active* for the
9 execution; otherwise, the OMPT interface state changes to *inactive*.

10 Cross References

- 11 • Tool Initialization and Finalization, see [Section 20.4.1](#)
- 12 • `ompt_enumerate_mutex_impls_t`, see [Section 20.6.1.2](#)
- 13 • `ompt_enumerate_states_t`, see [Section 20.6.1.1](#)
- 14 • `ompt_set_callback_t`, see [Section 20.6.1.3](#)
- 15 • `ompt_start_tool`, see [Section 20.2.1](#)

16 20.2.3.1 Binding Entry Points in the OMPT Callback Interface

17 Functions that an OpenMP implementation provides to support the OMPT interface are not defined
18 as global function symbols. Instead, they are defined as runtime entry points that a tool can only
19 identify through the *lookup* function that is provided as an argument with type signature
20 `ompt_function_lookup_t` to the tool initializer. A tool can use this function to obtain a
21 pointer to each of the runtime entry points that an OpenMP implementation provides to support the
22 OMPT interface. Once a tool has obtained a *lookup* function, it may employ it at any point in the
23 future.

24 For each runtime entry point in the OMPT interface for the host device, [Table 20.1](#) provides the
25 string name by which it is known and its associated type signature. Implementations can provide
26 additional implementation-specific names and corresponding entry points. Any names that begin
27 with `ompt_` are reserved names.

28 During initialization, a tool should look up each runtime entry point in the OMPT interface by
29 name and bind a pointer maintained by the tool that can later be used to invoke the entry point. The
30 entry points described in [Table 20.1](#) enable a tool to assess the thread states and mutual exclusion
31 implementations that an OpenMP implementation supports for [callback registration](#), to inspect
32 [registered callbacks](#), to introspect OpenMP state associated with threads, and to use tracing to
33 monitor computations that execute on target devices.

34 Detailed information about each runtime entry point listed in [Table 20.1](#) is included as part of the
35 description of its type signature.

TABLE 20.1: 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

- ompt_enumerate_mutex_impls_t, see [Section 20.6.1.2](#)
- ompt_enumerate_states_t, see [Section 20.6.1.1](#)
- Lookup Entry Points: ompt_function_lookup_t, see [Section 20.6.3](#)
- ompt_get_callback_t, see [Section 20.6.1.4](#)
- ompt_get_num_devices_t, see [Section 20.6.1.17](#)
- ompt_get_num_places_t, see [Section 20.6.1.7](#)
- ompt_get_num_procs_t, see [Section 20.6.1.6](#)
- ompt_get_parallel_info_t, see [Section 20.6.1.13](#)
- ompt_get_partition_place_nums_t, see [Section 20.6.1.10](#)
- ompt_get_place_num_t, see [Section 20.6.1.9](#)
- ompt_get_place_proc_ids_t, see [Section 20.6.1.8](#)
- ompt_get_proc_id_t, see [Section 20.6.1.11](#)
- ompt_get_state_t, see [Section 20.6.1.12](#)

- 1 • `ompt_get_target_info_t`, see [Section 20.6.1.16](#)
- 2 • `ompt_get_task_info_t`, see [Section 20.6.1.14](#)
- 3 • `ompt_get_task_memory_t`, see [Section 20.6.1.15](#)
- 4 • `ompt_get_thread_data_t`, see [Section 20.6.1.5](#)
- 5 • `ompt_get_unique_id_t`, see [Section 20.6.1.18](#)
- 6 • `ompt_set_callback_t`, see [Section 20.6.1.3](#)

7 **20.2.4 Monitoring Activity on the Host with OMPT**

8 To monitor the execution of an OpenMP program on the host device, a tool initializer must register
9 to receive notification of events that occur as an OpenMP program executes. A tool can use the
10 **`ompt_set_callback`** runtime entry point to perform [callback registrations](#) for [events](#). The
11 return codes for **`ompt_set_callback`** use the **`ompt_set_result_t`** enumeration type. If
12 the **`ompt_set_callback`** runtime entry point is called outside a tool initializer, [callback](#)
13 [registration](#) may fail for supported [callbacks](#) with a return value of **`ompt_set_error`**.

14 All [registered callbacks](#) and all [callbacks](#) returned by **`ompt_get_callback`** use the dummy type
15 signature **`ompt_callback_t`**.

16 For [callbacks](#) listed in [Table 20.2](#), **`ompt_set_always`** is the only registration return code that is
17 allowed. An OpenMP implementation must guarantee that the callback will be invoked every time
18 that a runtime [event](#) that is associated with it occurs. Support for such callbacks is required in a
19 minimal implementation of the [OMPT](#) interface.

20 For any other [callbacks](#) not listed in [Table 20.2](#), the **`ompt_set_callback`** runtime entry may
21 return any non-error code. Whether an OpenMP implementation invokes a registered [callback](#)
22 never, sometimes, or always is [implementation defined](#). If registration for a [callback](#) allows a return
23 code of **`ompt_set_never`**, support for invoking such a [callback](#) may not be present in a minimal
24 implementation of the [OMPT](#) interface. The return code from [callback registration](#) indicates the
25 [implementation defined](#) level of support for the [callback](#).

26 Two techniques reduce the size of the OMPT interface. First, in cases where events are naturally
27 paired, for example, the beginning and end of a region, and the arguments needed by the callback at
28 each endpoint are identical, a tool registers a single callback for the pair of events, with
29 **`ompt_scope_begin`** or **`ompt_scope_end`** provided as an argument to identify for which
30 endpoint the callback is invoked. Second, when a class of events is amenable to uniform treatment,
31 OMPT provides a single callback for that class of events, for example, an
32 **`ompt_callback_sync_region_wait`** callback is used for multiple kinds of synchronization
33 regions, such as barrier, taskwait, and taskgroup regions. Some events, for example,
34 **`ompt_callback_sync_region_wait`**, use both techniques.

35 **Cross References**

- 36 • `ompt_get_callback_t`, see [Section 20.6.1.4](#)

TABLE 20.2: 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>
<code>ompt_callback_error</code>

- 1 • `ompt_set_callback_t`, see [Section 20.6.1.3](#)
- 2 • `ompt_set_result_t`, see [Section 20.4.4.2](#)

3 20.2.5 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

TABLE 20.3: OMPT Tracing Interface Runtime Entry Point Names and Their Type Signatures

Entry Point String Name	Type Signature
"ompt_get_device_num_procs"	ompt_get_device_num_procs_t
"ompt_get_device_time"	ompt_get_device_time_t
"ompt_translate_time"	ompt_translate_time_t
"ompt_set_trace_ompt"	ompt_set_trace_ompt_t
"ompt_set_trace_native"	ompt_set_trace_native_t
"ompt_start_trace"	ompt_start_trace_t
"ompt_pause_trace"	ompt_pause_trace_t
"ompt_flush_trace"	ompt_flush_trace_t
"ompt_stop_trace"	ompt_stop_trace_t
"ompt_advance_buffer_cursor"	ompt_advance_buffer_cursor_t
"ompt_get_record_type"	ompt_get_record_type_t
"ompt_get_record_ompt"	ompt_get_record_ompt_t
"ompt_get_record_native"	ompt_get_record_native_t
"ompt_get_record_abstract"	ompt_get_record_abstract_t

- 1 implementation or target device does not support tracing, the OpenMP implementation
2 passes `NULL` to the device initializer of the tool for its `lookup` argument; otherwise, the
3 OpenMP implementation passes a pointer to a device-specific runtime entry point with type
4 signature `ompt_function_lookup_t` to the device initializer of the tool.
- 5 • If the `lookup` argument of the `device` initializer of the `tool` is a `non-null pointer`, the `tool` may
6 use it to determine the `runtime entry points` in the tracing interface that are available for the
7 `device` and may bind the returned function pointers to `tool variables`. Table 20.3 indicates the
8 names of `runtime entry points` that may be available for a `device`; an implementation may
9 provide additional `implementation defined` names and corresponding entry points. The driver
10 for the `device` provides the `runtime entry points` that enable a `tool` to control the trace
11 collection interface of the `device`. The `native` trace format that the interface uses may be
12 `device` specific and the available kinds of `trace records` are `implementation defined`. Some
13 `devices` may allow a `tool` to collect traces of records in a standard format known as OMPT
14 `trace records`. Each OMPT `trace record` serves as a substitute for an OMPT `callback` that is
15 not appropriate to be dispatched on the `device`. The fields in each `trace record` type are
16 defined in the description of the `callback` that the record represents. If this type of record is
17 provided then the `lookup` function returns values for the `runtime entry points`
18 `ompt_set_trace_ompt` and `ompt_get_record_ompt`, which support collecting
19 and decoding OMPT traces. If the native tracing format for a device is the OMPT format then
20 tracing can be controlled using the `runtime entry points` for native or OMPT tracing.
 - 21 • The tool uses the `ompt_set_trace_native` and/or the `ompt_set_trace_ompt`
22 runtime entry point to specify what types of events or activities to monitor on the device. The
23 return codes for `ompt_set_trace_ompt` and `ompt_set_trace_native` use the
24 `ompt_set_result_t` enumeration type. If the `ompt_set_trace_native` or the
25 `ompt_set_trace_ompt` runtime entry point is called outside a device initializer,

- 1 registration of supported callbacks may fail with a return code of `ompt_set_error`.
- 2
- 3
- 4
- 5
- 6
- 7
- 8
- 9
- 10
- 11
- 12
- 13
- 14
- 15
- 16
- 17
- 18
- 19
- 20
- 21
- 22
- 23
- 24
- 25
- 26
- 27
- 28
- 29
- 30
- 31
- 32
- 33
- 34
- 35
- 36
- 37
- 38
- 39
- 40
- The tool initiates tracing of device activity by invoking `ompt_start_trace`. Arguments to `ompt_start_trace` include two tool callbacks through which the OpenMP implementation can manage traces associated with the device. One callback allocates a buffer in which device activity can be deposited. The second callback processes a buffer of trace events from the device.
 - If the OpenMP implementation requires a trace buffer for device activity, the OpenMP implementation invokes the tool-supplied callback function on the host device to request a new buffer.
 - The OpenMP implementation monitors the execution of OpenMP constructs on the device and records a trace of events or activities into a trace buffer. If possible, device [trace records](#) are marked with a `host_op_id`—an identifier that associates device activities with the target operation that the host initiated to cause these activities. To correlate activities on the host with activities on a device, a tool can register a `ompt_callback_target_submit_emi` callback. Before and after the host initiates creation of an initial task on a device associated with a structured block for a `target` construct, the OpenMP implementation dispatches the `ompt_callback_target_submit_emi` callback on the host in the thread that is executing the task that encounters the `target` construct. This callback provides the tool with a pair of identifiers: one that identifies the `target` region and a second that uniquely identifies the initial task associated with that region. These identifiers help the tool correlate activities on the target device with their `target` region.
 - When appropriate, for example, when a trace buffer fills or needs to be flushed, the OpenMP implementation invokes the tool-supplied buffer completion callback to process a non-empty sequence of records in a trace buffer that is associated with the device.
 - The tool-supplied buffer completion callback may return immediately, ignoring records in the trace buffer, or it may iterate through them using the `ompt_advance_buffer_cursor` entry point to inspect each record. A tool may use the `ompt_get_record_type` runtime entry point to inspect the type of the record at the current cursor position. Three runtime entry points (`ompt_get_record_ompt`, `ompt_get_record_native`, and `ompt_get_record_abstract`) allow tools to inspect the contents of some or all records in a trace buffer. The `ompt_get_record_native` runtime entry point uses the native trace format of the device. The `ompt_get_record_abstract` runtime entry point decodes the contents of a [native trace record](#) and summarizes them as an `ompt_record_abstract_t` record. The [ompt_get_record_ompt runtime entry point](#) can only be used to retrieve records in OMPT format.
 - Once device tracing has been started, a tool may pause or resume device tracing at any time by invoking `ompt_pause_trace` with an appropriate flag value as an argument.
 - A tool may invoke the `ompt_flush_trace` runtime entry point for a device at any time between device initialization and finalization to cause the pending trace records for that

1 device to be flushed.

- 2 • At any time, a tool may use the `ompt_start_trace` runtime entry point to start or the
3 `ompt_stop_trace` runtime entry point to stop device tracing. When device tracing is
4 stopped, the OpenMP implementation eventually gathers all trace records already collected
5 from device tracing and presents them to the tool using the buffer completion callback.
- 6 • An OpenMP implementation can be shut down while device tracing is in progress.
- 7 • When an OpenMP implementation is shut down, it finalizes each device. Device finalization
8 occurs in three steps. First, the OpenMP implementation halts any tracing in progress for the
9 device. Second, the OpenMP implementation flushes all trace records collected for the
10 device and uses the buffer completion callback associated with that device to present them to
11 the tool. Finally, the OpenMP implementation dispatches any
12 `ompt_callback_device_finalize` callback registered for the device.

13 Restrictions

14 Restrictions on tracing activity on devices are as follows:

- 15 • Implementation-defined names must not start with the prefix `ompt_`, which is reserved for
16 the OpenMP specification.

17 Cross References

- 18 • `ompt_advance_buffer_cursor_t`, see [Section 20.6.2.11](#)
- 19 • `ompt_callback_device_finalize_t`, see [Section 20.5.2.20](#)
- 20 • `ompt_callback_device_initialize_t`, see [Section 20.5.2.19](#)
- 21 • `ompt_flush_trace_t`, see [Section 20.6.2.9](#)
- 22 • `ompt_get_device_num_procs_t`, see [Section 20.6.2.1](#)
- 23 • `ompt_get_device_time_t`, see [Section 20.6.2.2](#)
- 24 • `ompt_get_record_abstract_t`, see [Section 20.6.2.15](#)
- 25 • `ompt_get_record_native_t`, see [Section 20.6.2.14](#)
- 26 • `ompt_get_record_ompt_t`, see [Section 20.6.2.13](#)
- 27 • `ompt_get_record_type_t`, see [Section 20.6.2.12](#)
- 28 • `ompt_pause_trace_t`, see [Section 20.6.2.8](#)
- 29 • `ompt_set_trace_native_t`, see [Section 20.6.2.5](#)
- 30 • `ompt_set_trace_ompt_t`, see [Section 20.6.2.4](#)
- 31 • `ompt_start_trace_t`, see [Section 20.6.2.7](#)
- 32 • `ompt_stop_trace_t`, see [Section 20.6.2.10](#)
- 33 • `ompt_translate_time_t`, see [Section 20.6.2.3](#)

20.3 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 20.5.1.2](#)

20.4 OMPT Data Types

The C/C++ header file (`omp-tools.h`) provides the definitions of the types that are specified throughout this subsection.

20.4.1 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 values](#).

Cross References

- `ompt_data_t`, see [Section 20.4.4.4](#)
- `ompt_finalize_t`, see [Section 20.5.1.2](#)
- `ompt_initialize_t`, see [Section 20.5.1.1](#)
- `ompt_start_tool`, see [Section 20.2.1](#)

20.4.2 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++

```
typedef enum ompt_callbacks_t {
    ompt_callback_thread_begin          = 1,
    ompt_callback_thread_end            = 2,
    ompt_callback_parallel_begin        = 3,
    ompt_callback_parallel_end          = 4,
    ompt_callback_task_create           = 5,
    ompt_callback_task_schedule         = 6,
    ompt_callback_implicit_task        = 7,
    ompt_callback_target                = 8,
    ompt_callback_target_data_op       = 9,
    ompt_callback_target_submit        = 10,
    ompt_callback_control_tool         = 11,
    ompt_callback_device_initialize     = 12,
    ompt_callback_device_finalize      = 13,
    ompt_callback_device_load          = 14,
    ompt_callback_device_unload        = 15,
    ompt_callback_sync_region_wait     = 16,
    ompt_callback_mutex_released       = 17,
    ompt_callback_dependences          = 18,
    ompt_callback_task_dependence      = 19,
    ompt_callback_work                 = 20,
    ompt_callback_masked               = 21,
    ompt_callback_target_map           = 22,
    ompt_callback_sync_region          = 23,
    ompt_callback_lock_init            = 24,
    ompt_callback_lock_destroy         = 25,
    ompt_callback_mutex_acquire        = 26,
    ompt_callback_mutex_acquired       = 27,
    ompt_callback_nest_lock            = 28,
    ompt_callback_flush                = 29,
    ompt_callback_cancel               = 30,
    ompt_callback_reduction            = 31,
    ompt_callback_dispatch              = 32,
    ompt_callback_target_emi           = 33,
    ompt_callback_target_data_op_emi   = 34,
    ompt_callback_target_submit_emi    = 35,
```



```
1  ompt_callback_target_map_emi          = 36,  
2  ompt_callback_error                  = 37  
3  } ompt_callbacks_t;
```

C / C++

20.4.3 Tracing

OpenMP provides type definitions that support tracing with OMPT.

20.4.3.1 Record Type

Summary

The `ompt_record_t` enumeration type indicates the integer codes used to identify OpenMP trace record formats.

Format

C / C++

```
11 typedef enum ompt_record_t {  
12     ompt_record_ompt          = 1,  
13     ompt_record_native       = 2,  
14     ompt_record_invalid      = 3  
15 } ompt_record_t;
```

C / C++

20.4.3.2 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++

```
21 typedef enum ompt_record_native_t {  
22     ompt_record_native_info = 1,  
23     ompt_record_native_event = 2  
24 } ompt_record_native_t;
```

C / C++

20.4.3.3 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`.

20.4.3.4 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.

20.4.4 Miscellaneous Type Definitions

This section describes miscellaneous types and enumerations used by the tool interface.

20.4.4.1 ompt_callback_t

Summary

Pointers to [tool callback](#) functions with different type signatures are passed to the `ompt_set_callback` [runtime entry point](#) and returned by the `ompt_get_callback` [runtime entry point](#). For convenience, these [runtime entry points](#) expect all type signatures to be cast to a dummy type `ompt_callback_t`.

Format

```

typedef void (*ompt_callback_t) (void);

```

C / C++

20.4.4.2 `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

C / C++

```
typedef enum ompt_set_result_t {
    ompt_set_error          = 0,
    ompt_set_never          = 1,
    ompt_set_impossible     = 2,
    ompt_set_sometimes      = 3,
    ompt_set_sometimes_paired = 4,
    ompt_set_always         = 5
} ompt_set_result_t;
```

C / C++

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, [callback dispatch](#) 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 20.6.1.3](#)
- `ompt_set_trace_native_t`, see [Section 20.6.2.5](#)
- `ompt_set_trace_ompt_t`, see [Section 20.6.2.4](#)

20.4.4.3 `ompt_id_t`

Summary

The `ompt_id_t` type is used to provide various identifiers to tools.

Format

C / C++

```
typedef uint64_t ompt_id_t;
```

C / C++

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.

20.4.4.4 `ompt_data_t`

Summary

The `ompt_data_t` type represents data associated with threads and with parallel and task regions.

Format

C / C++

```
typedef union ompt_data_t {  
    uint64_t value;  
    void *ptr;  
} ompt_data_t;
```

C / C++

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.

20.4.4.5 `ompt_device_t`

Summary

The `ompt_device_t` opaque object type represents a device.

```
1 Format
2 | typedef void ompt_device_t;
3 C / C++
```

20.4.4.6 ompt_device_time_t

Summary

The `ompt_device_time_t` type represents raw device time values.

```
6 Format
7 | typedef uint64_t ompt_device_time_t;
8 C / C++
```

Semantics

The `ompt_device_time_t` opaque object type represents raw device time values.

`ompt_time_none` refers to an unknown or unspecified time and is defined as an instance of type `ompt_device_time_t` with the value 0.

20.4.4.7 ompt_buffer_t

Summary

The `ompt_buffer_t` opaque object type is a handle for a target buffer.

```
15 Format
16 | typedef void ompt_buffer_t;
17 C / C++
```

20.4.4.8 ompt_buffer_cursor_t

Summary

The `ompt_buffer_cursor_t` opaque type is a handle for a position in a target buffer.

```
20 Format
21 | typedef uint64_t ompt_buffer_cursor_t;
22 C / C++
```

20.4.4.9 `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** or **doacross** clause. For task dependences, the `variable.ptr` field points to the storage location of the dependence. For `doacross` dependences, the `variable.value` field contains the value of a vector element that describes the dependence. The `dependence_type` field indicates the type of the dependence. For task dependences with the reserved locator **omp_all_memory**, the value of `variable` is undefined and the `dependence_type` field contains the value of an enumerator that has the `_all_memory` suffix.

Cross References

- `ompt_dependence_type_t`, see [Section 20.4.4.24](#)

20.4.4.10 `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 **threads** that are **thread-pool-worker threads** have **thread** type `ompt_thread_worker`. A **native thread** that an OpenMP implementation uses but that does not execute user code has **thread** type `ompt_thread_other`. Any **native thread** that is created outside an OpenMP implementation and that is not an *initial thread* has **thread** type `ompt_thread_unknown`.

20.4.4.11 `ompt_scope_endpoint_t`

Summary

The `ompt_scope_endpoint_t` enumeration type defines valid scope endpoint values.

Format

C / C++

```
typedef enum ompt_scope_endpoint_t {
    ompt_scope_begin           = 1,
    ompt_scope_end             = 2,
    ompt_scope_beginend        = 3
} ompt_scope_endpoint_t;
```

C / C++

20.4.4.12 `ompt_dispatch_t`

Summary

The `ompt_dispatch_t` enumeration type defines the valid dispatch kind values.

Format

C / C++

```
typedef enum ompt_dispatch_t {
    ompt_dispatch_iteration     = 1,
    ompt_dispatch_section       = 2,
    ompt_dispatch_ws_loop_chunk = 3,
    ompt_dispatch_taskloop_chunk = 4,
    ompt_dispatch_distribute_chunk = 5
} ompt_dispatch_t;
```

C / C++

20.4.4.13 `ompt_dispatch_chunk_t`

Summary

The `ompt_dispatch_chunk_t` type represents a the chunk information for a dispatched chunk.

Format

C / C++

```
typedef struct ompt_dispatch_chunk_t {  
    uint64_t start;  
    uint64_t iterations;  
} ompt_dispatch_chunk_t;
```

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.

20.4.4.14 `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++

20.4.4.15 `ompt_target_data_op_t`

Summary

The `ompt_target_data_op_t` enumeration type defines the valid target data operation values.

Format

C/C++

```
1
2 typedef enum ompt_target_data_op_t {
3     ompt_target_data_alloc                = 1,
4     ompt_target_data_transfer_to_device  = 2, // deprecated
5     ompt_target_data_transfer_from_device = 3, // deprecated
6     ompt_target_data_delete              = 4,
7     ompt_target_data_associate           = 5,
8     ompt_target_data_disassociate        = 6,
9     ompt_target_data_transfer            = 7,
10    ompt_target_data_memset               = 8,
11    ompt_target_data_alloc_async          = 17,
12    ompt_target_data_transfer_to_device_async = 18, //
13    deprecated
14    ompt_target_data_transfer_from_device_async = 19, //
15    deprecated
16    ompt_target_data_delete_async         = 20,
17    ompt_target_data_transfer_async       = 23,
18    ompt_target_data_memset_async         = 24
19 } ompt_target_data_op_t;
```

C/C++

Semantics

The `ompt_target_data_op_t` enumeration type indicates the kind of target data operation for `ompt_callback_target_data_op_emi_t` which can be *alloc*, *delete*, *associate*, *disassociate*, or *transfer*. For asynchronous data operations the corresponding value with `_async` suffix is used.

20.4.4.16 ompt_work_t

Summary

The `ompt_work_t` enumeration type defines the valid work type values.

Format

C/C++

```
29 typedef enum ompt_work_t {
30     ompt_work_loop                = 1,
31     ompt_work_sections            = 2,
32     ompt_work_single_executor     = 3,
33     ompt_work_single_other        = 4,
34     ompt_work_workshare           = 5,
35     ompt_work_distribute          = 6,
36     ompt_work_taskloop            = 7,
37     ompt_work_scope               = 8,
```

```

1  ompt_work_loop_static      = 10,
2  ompt_work_loop_dynamic    = 11,
3  ompt_work_loop_guided     = 12,
4  ompt_work_loop_other      = 13,
5  ompt_work_coexecute       = 14
6  } ompt_work_t;

```

C / C++

20.4.4.17 ompt_mutex_t

Summary

The `ompt_mutex_t` enumeration type defines the valid mutex kind values.

Format

C / C++

```

11 typedef enum ompt_mutex_t {
12     ompt_mutex_lock          = 1,
13     ompt_mutex_test_lock     = 2,
14     ompt_mutex_nest_lock     = 3,
15     ompt_mutex_test_nest_lock = 4,
16     ompt_mutex_critical      = 5,
17     ompt_mutex_atomic        = 6,
18     ompt_mutex_ordered       = 7
19 } ompt_mutex_t;

```

C / C++

20.4.4.18 ompt_native_mon_flag_t

Summary

The `ompt_native_mon_flag_t` enumeration type defines the valid native monitoring flag values.

Format

C / C++

```

25 typedef enum ompt_native_mon_flag_t {
26     ompt_native_data_motion_explicit = 0x01,
27     ompt_native_data_motion_implicit = 0x02,
28     ompt_native_kernel_invocation    = 0x04,
29     ompt_native_kernel_execution     = 0x08,
30     ompt_native_driver                = 0x10,
31     ompt_native_runtime               = 0x20,
32     ompt_native_overhead              = 0x40,
33     ompt_native_idleness              = 0x80
34 } ompt_native_mon_flag_t;

```

C / C++

20.4.4.19 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_undelayed       = 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.

20.4.4.20 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.

20.4.4.21 `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++

20.4.4.22 `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 the code that implements the associated block of the region is invoked or encountered.

The value `ompt_parallel_invoker_program` indicates that the encountering thread for a `parallel` or `teams` region will execute the code that implements the associated block of the region as if directly invoked or encountered from application code. The value `ompt_parallel_invoker_runtime` indicates that the encountering thread for a `parallel` or `teams` region invokes the code that implements the associated block of the region from the runtime.

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.

20.4.4.23 `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.

20.4.4.24 `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_out_all_memory = 34,
    ompt_dependence_type_inout_all_memory = 35
} ompt_dependence_type_t;
```

C / C++

Semantics

The `ompt_dependence_type_`*dependence-type* value represents the *task-dependence-type* present in a `depend` clause or the *dependence-type* present in a `doacross` clause. If *dependence-type* is *task-dependence-type* `_all_memory`, then it represents a dependence for the `omp_all_memory` reserved locator.

20.4.4.25 `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++

20.4.4.26 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++

20.4.4.27 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 20.4.3.3](#)

20.4.4.28 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_work_free_agent      = 0x003,

    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,

    ompt_state_wait_mutex            = 0x040,
    ompt_state_wait_lock             = 0x041,
    ompt_state_wait_critical         = 0x042,
    ompt_state_wait_atomic           = 0x043,
    ompt_state_wait_ordered          = 0x044,

    ompt_state_wait_target           = 0x080,
    ompt_state_wait_target_map       = 0x081,
    ompt_state_wait_target_update    = 0x082,

    ompt_state_idle                  = 0x100,
    ompt_state_overhead               = 0x101,
    ompt_state_undefined             = 0x102
} ompt_state_t;
```

C / C++

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_work_free_agent` indicates that the thread is executing code within the scope of a `task` while not being assigned of its `current team`. 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 specification but is 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`. The value `ompt_state_wait_ordered` indicates that the `thread` is waiting to enter an `ordered region`. The value `ompt_state_wait_target` indicates that the thread is waiting for a `target region` to complete. The value `ompt_state_wait_target_map` indicates that the `thread` is waiting for a target data mapping operation to complete. An implementation may report `ompt_state_wait_target` for `target data constructs`. The value `ompt_state_wait_target_update` indicates that the `thread` is waiting for a `target update` operation to complete. An implementation may report `ompt_state_wait_target` for `target update constructs`. The value `ompt_state_idle` indicates that the `native thread` is an `idle thread`, that is, it is an `unassigned thread`. The value `ompt_state_overhead` indicates that the `thread` is in the overhead state at any point while executing within the OpenMP runtime, except while waiting at a synchronization point. The value `ompt_state_undefined` indicates that the `native thread` is not created by the OpenMP implementation.

20.4.4.29 `ompt_frame_t`

Summary

The `ompt_frame_t` type describes [procedure](#) frame information for an OpenMP task.

Format

```
typedef struct ompt_frame_t {  
    ompt_data_t exit_frame;  
    ompt_data_t enter_frame;  
    int exit_frame_flags;  
    int enter_frame_flags;  
} ompt_frame_t;
```

C / C++

Semantics

Each `ompt_frame_t` object is associated with the task to which the [procedure](#) frames belong. Each non-merged initial, implicit, explicit, or target task with one or more frames on the stack of a [native thread](#) has an associated `ompt_frame_t` object.

The `exit_frame` field of an `ompt_frame_t` object contains information to identify the first [procedure](#) frame executing the task region. The `exit_frame` for the `ompt_frame_t` object associated with the *initial task* that is not nested inside any OpenMP construct is `ompt_data_none`.

The `enter_frame` field of an `ompt_frame_t` object contains information to identify the latest still active [procedure frame](#) executing the [task region](#) before entering the OpenMP runtime implementation or before executing a different [task](#). If a [task](#) with [frames](#) on the stack is not executing [implementation code](#) in the OpenMP runtime, the value of `enter_frame` for the `ompt_frame_t` object associated with the [task](#) will be `ompt_data_none`.

For `exit_frame`, the `exit_frame_flags` and, for `enter_frame`, the `enter_frame_flags` field indicates that the provided [frame](#) information points to a runtime or an [OpenMP program frame](#) address. The same fields also specify the kind of information that is provided to identify the [frame](#). These fields are a disjunction of values in the `ompt_frame_flag_t` enumeration type.

The lifetime of an `ompt_frame_t` object begins when a task is created and ends when the task is destroyed. Tools should not assume that a frame structure remains at a constant location in memory throughout the lifetime of the task. A pointer to an `ompt_frame_t` object is passed to some callbacks; a pointer to the `ompt_frame_t` object of a task can also be retrieved by a tool at any time, including in a [signal handler](#), by invoking the `ompt_get_task_info` [runtime entry point](#) (described in [Section 20.6.1.14](#)). A pointer to an `ompt_frame_t` object that a tool retrieved is valid as long as the tool does not pass back control to the OpenMP implementation.

Note – A monitoring tool that uses asynchronous sampling can observe values of *exit_frame* and *enter_frame* at inconvenient times. Tools must be prepared to handle `ompt_frame_t` objects observed just prior to when their field values will be set or cleared.

20.4.4.30 `ompt_frame_flag_t`

Summary

The `ompt_frame_flag_t` enumeration type defines valid frame information flags.

Format

```
C / C++  
typedef enum ompt_frame_flag_t {  
    ompt_frame_runtime      = 0x00,  
    ompt_frame_application  = 0x01,  
    ompt_frame_cfa         = 0x10,  
    ompt_frame_framepointer = 0x20,  
    ompt_frame_stackaddress = 0x30  
} ompt_frame_flag_t;
```

Semantics

The value `ompt_frame_runtime` of the `ompt_frame_flag_t` type indicates that a [frame](#) address is a [procedure frame](#) in the OpenMP runtime implementation. The value `ompt_frame_application` of the `ompt_frame_flag_t` type indicates that a [frame](#) address is a [procedure frame](#) in the [OpenMP program](#)

Higher order bits indicate the kind of provided information that is unique for the particular [frame](#) pointer. The value `ompt_frame_cfa` indicates that a [frame](#) address specifies a [canonical frame address](#). The value `ompt_frame_framepointer` indicates that a [frame](#) address provides the value of the [frame](#) pointer register. The value `ompt_frame_stackaddress` indicates that a [frame](#) address specifies a pointer address that is contained in the current stack frame.

20.4.4.31 `ompt_wait_id_t`

Summary

The `ompt_wait_id_t` type describes [wait identifiers](#) for a [thread](#).

Format

```
C / C++  
typedef uint64_t ompt_wait_id_t;
```

Semantics

Each **thread** maintains a **wait identifier** of type `ompt_wait_id_t`. When a **task** that a **thread** executes is waiting for mutual exclusion, the **wait identifier** of the **thread** indicates the reason that the **thread** is waiting. A **wait identifier** may represent a critical section *name*, a lock, a **variable** accessed in an **atomic region**, or a synchronization object that is internal to an OpenMP implementation. When a **thread** is not in a wait state then the value of the **wait identifier** of the **thread** is undefined. `ompt_wait_id_none` is defined as an instance of type `ompt_wait_id_t` with the value 0.

20.5 OMPT Tool Callback Signatures and Trace Records

The C/C++ header file (`ompt-tools.h`) provides the definitions of the types that are specified throughout this subsection. Restrictions to the OpenMP tool callbacks are as follows:

Restrictions

- Tool callbacks may not use OpenMP directives or call any runtime library routines described in [Chapter 19](#).
- Tool callbacks must exit by either returning to the caller or aborting.

20.5.1 Initialization and Finalization Callback Signature

20.5.1.1 `ompt_initialize_t`

Summary

A callback with type signature `ompt_initialize_t` initializes the use of the OMPT interface.

Format

```
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 **OMPT inactive** as described in [Section 20.2.3](#).

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 20.4.1](#)
- `omp_get_initial_device`, see [Section 19.7.8](#)
- `ompt_data_t`, see [Section 20.4.4.4](#)
- `ompt_start_tool`, see [Section 20.2.1](#)

20.5.1.2 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

```
typedef void (*ompt_finalize_t) (  
    ompt_data_t *tool_data  
);
```

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 20.4.1](#)
- `ompt_data_t`, see [Section 20.4.4.4](#)
- `ompt_start_tool`, see [Section 20.2.1](#)

20.5.2 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 20.4.4.4](#)
- `ompt_id_t`, see [Section 20.4.4.3](#)

20.5.2.1 `ompt_callback_thread_begin_t`

Summary

The `ompt_callback_thread_begin_t` type is used for callbacks that are dispatched when [native threads](#) are created.

Format

```
typedef void (*ompt_callback_thread_begin_t) (  
    ompt_thread_t thread_type,  
    ompt_data_t *thread_data  
);
```

Trace Record

```
typedef struct ompt_record_thread_begin_t {  
    ompt_thread_t thread_type;  
} ompt_record_thread_begin_t;
```

Description of Arguments

The *thread_type* argument indicates the type of the new thread: initial, worker, or other. The binding of the *thread_data* argument is the new thread.

Cross References

- `parallel` directive, see [Section 11.2](#)
- `teams` directive, see [Section 11.3](#)
- Initial Task, see [Section 13.9](#)
- `ompt_data_t`, see [Section 20.4.4.4](#)
- `ompt_thread_t`, see [Section 20.4.4.10](#)

20.5.2.2 `ompt_callback_thread_end_t`

Summary

The `ompt_callback_thread_end_t` type is used for callbacks that are dispatched when [native threads](#) are destroyed.

Format

C / C++

```
typedef void (*ompt_callback_thread_end_t) (  
    ompt_data_t *thread_data  
);
```

C / C++

Description of Arguments

The binding of the `thread_data` argument is the thread that will be destroyed.

Cross References

- `parallel` directive, see [Section 11.2](#)
- `teams` directive, see [Section 11.3](#)
- Initial Task, see [Section 13.9](#)
- Standard Trace Record Type, see [Section 20.4.3.4](#)
- `ompt_data_t`, see [Section 20.4.4.4](#)

20.5.2.3 `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 **ompt_callback_parallel_begin_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 11.2](#)
- `teams` directive, see [Section 11.3](#)
- `ompt_data_t`, see [Section 20.4.4.4](#)
- `ompt_frame_t`, see [Section 20.4.4.29](#)
- `ompt_parallel_flag_t`, see [Section 20.4.4.22](#)

20.5.2.4 `ompt_callback_parallel_end_t`

Summary

The `ompt_callback_parallel_end_t` type is used for callbacks that are dispatched when a `parallel` or `teams` region ends.

Format

```
C / C++
typedef void (*ompt_callback_parallel_end_t) (
    ompt_data_t *parallel_data,
    ompt_data_t *encountering_task_data,
    int flags,
    const void *codeptr_ra
);
```

Trace Record

```
C / C++
typedef struct ompt_record_parallel_end_t {
    ompt_id_t parallel_id;
    ompt_id_t encountering_task_id;
    int flags;
    const void *codeptr_ra;
} ompt_record_parallel_end_t;
```

Description of Arguments

The binding of the `parallel_data` argument is the `parallel` or `teams` region that is ending.

The binding of the `encountering_task_data` argument is the encountering task.

The `flags` argument indicates whether the execution of the region is inlined into the application or invoked by the runtime and also whether it is a `parallel` or `teams` region. Values for `flags` are a disjunction of elements in the enum `ompt_parallel_flag_t`.

The *codeptr_ra* argument relates the implementation of an OpenMP region to its source code. If a runtime routine implements the region associated with a callback that has type signature **ompt_callback_parallel_end_t** then *codeptr_ra* contains the return address of the call to that runtime routine. If the implementation of the region is inlined then *codeptr_ra* contains the return address of the callback invocation. If attribution to source code is impossible or inappropriate, *codeptr_ra* may be **NULL**.

Cross References

- **parallel** directive, see [Section 11.2](#)
- **teams** directive, see [Section 11.3](#)
- **ompt_data_t**, see [Section 20.4.4.4](#)
- **ompt_parallel_flag_t**, see [Section 20.4.4.22](#)

20.5.2.5 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

```
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

```
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** or **coexecute** construct, *count* represents the **units of work**, as defined by the **workshare** or **coexecute** 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 13.7](#)
- Work-Distribution Constructs, see [Chapter 12](#)
- **ompt_data_t**, see [Section 20.4.4.4](#)
- **ompt_scope_endpoint_t**, see [Section 20.4.4.11](#)
- **ompt_work_t**, see [Section 20.4.4.16](#)

20.5.2.6 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 12.3](#)
- `taskloop` directive, see [Section 13.7](#)
- Worksharing-Loop Constructs, see [Section 12.6](#)
- `ompt_data_t`, see [Section 20.4.4.4](#)
- `ompt_dispatch_chunk_t`, see [Section 20.4.4.13](#)
- `ompt_dispatch_t`, see [Section 20.4.4.12](#)

20.5.2.7 `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
);
```

Trace Record

```
C / C++
typedef struct ompt_record_task_create_t {
    ompt_id_t encountering_task_id;
    ompt_id_t new_task_id;
    int flags;
    int has_dependences;
    const void *codeptr_ra;
} ompt_record_task_create_t;
```

Description of Arguments

The binding of the *encountering_task_data* argument is the encountering task.

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 13.6](#)
- Initial Task, see [Section 13.9](#)
- `ompt_data_t`, see [Section 20.4.4.4](#)
- `ompt_frame_t`, see [Section 20.4.4.29](#)
- `ompt_task_flag_t`, see [Section 20.4.4.19](#)

20.5.2.8 `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  
);
```

C / C++

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 16.9.5](#)
- **ordered** directive, see [Section 16.10.1](#)
- **ompt_data_t**, see [Section 20.4.4.4](#)
- **ompt_dependence_t**, see [Section 20.4.4.9](#)

20.5.2.9 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 omp_t_record_task_dependence_t {
    omp_t_id_t src_task_id;
    omp_t_id_t sink_task_id;
} omp_t_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 16.9.5](#)
- `omp_data_t`, see [Section 20.4.4.4](#)

20.5.2.10 `omp_callback_task_schedule_t`

Summary

The `omp_callback_task_schedule_t` type is used for callbacks that are dispatched when task scheduling decisions are made.

Format

C / C++

```
typedef void (*omp_callback_task_schedule_t) (
    omp_data_t *prior_task_data,
    omp_task_status_t prior_task_status,
    omp_data_t *next_task_data
);
```

C / C++

Trace Record

C / C++

```
typedef struct omp_t_record_task_schedule_t {
    omp_t_id_t prior_task_id;
    omp_task_status_t prior_task_status;
    omp_t_id_t next_task_id;
} omp_t_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.

1 The binding of the *prior_task_data* argument is the task that arrived at the scheduling point. This
2 argument can be `NULL` if no task was active when the next task is scheduled.

3 The binding of the *next_task_data* argument is the task that is resumed at the scheduling point.
4 This argument is `NULL` if the callback is dispatched for a *task-fulfill* event or if the callback signals
5 completion of a `taskwait` construct. This argument can be `NULL` if no task was active when the
6 prior task was scheduled.

7 Cross References

- 8 • Task Scheduling, see [Section 13.10](#)
- 9 • `ompt_data_t`, see [Section 20.4.4.4](#)
- 10 • `ompt_task_status_t`, see [Section 20.4.4.20](#)

11 20.5.2.11 `ompt_callback_implicit_task_t`

12 Summary

13 The `ompt_callback_implicit_task_t` type is used for callbacks that are dispatched when
14 initial tasks and implicit tasks are generated and completed.

15 Format

```
16 typedef void (*ompt_callback_implicit_task_t) (  
17     ompt_scope_endpoint_t endpoint,  
18     ompt_data_t *parallel_data,  
19     ompt_data_t *task_data,  
20     unsigned int actual_parallelism,  
21     unsigned int index,  
22     int flags  
23 );
```

24 Trace Record

```
25 typedef struct ompt_record_implicit_task_t {  
26     ompt_scope_endpoint_t endpoint;  
27     ompt_id_t parallel_id;  
28     ompt_id_t task_id;  
29     unsigned int actual_parallelism;  
30     unsigned int index;  
31     int flags;  
32 } ompt_record_implicit_task_t;
```

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 11.2](#)
- **teams** directive, see [Section 11.3](#)
- **ompt_data_t**, see [Section 20.4.4.4](#)
- **ompt_scope_endpoint_t**, see [Section 20.4.4.11](#)

20.5.2.12 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 11.6](#)
- **ompt_data_t**, see [Section 20.4.4.4](#)
- **ompt_scope_endpoint_t**, see [Section 20.4.4.11](#)

20.5.2.13 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 16.3.1](#)
- `taskgroup` directive, see [Section 16.4](#)
- `taskwait` directive, see [Section 16.5](#)
- Implicit Barriers, see [Section 16.3.2](#)
- Properties Common to All Reduction Clauses, see [Section 6.5.6](#)
- `ompt_data_t`, see [Section 20.4.4.4](#)
- `ompt_scope_endpoint_t`, see [Section 20.4.4.11](#)
- `ompt_sync_region_t`, see [Section 20.4.4.14](#)

20.5.2.14 `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 16.8.5](#)
- `critical` directive, see [Section 16.2](#)
- `ompt_wait_id_t`, see [Section 20.4.4.31](#)
- `omp_init_lock` and `omp_init_nest_lock`, see [Section 19.9.1](#)
- `ompt_mutex_t`, see [Section 20.4.4.17](#)
- `ordered` Construct, see [Section 16.10](#)

20.5.2.15 `ompt_callback_mutex_t`

Summary

The `ompt_callback_mutex_t` type is used for callbacks that indicate important synchronization events.

Format

C / C++

```
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;
```

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 16.8.5](#)
- `critical` directive, see [Section 16.2](#)
- `omp_destroy_lock` and `omp_destroy_nest_lock`, see [Section 19.9.3](#)
- `ompt_wait_id_t`, see [Section 20.4.4.31](#)
- `omp_set_lock` and `omp_set_nest_lock`, see [Section 19.9.4](#)
- `omp_test_lock` and `omp_test_nest_lock`, see [Section 19.9.6](#)
- `omp_unset_lock` and `omp_unset_nest_lock`, see [Section 19.9.5](#)
- `ompt_mutex_t`, see [Section 20.4.4.17](#)
- `ordered` Construct, see [Section 16.10](#)

20.5.2.16 `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
);
```

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;
```

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 20.4.4.31](#)
- **omp_set_lock** and **omp_set_nest_lock**, see [Section 19.9.4](#)
- **omp_test_lock** and **omp_test_nest_lock**, see [Section 19.9.6](#)
- **omp_unset_lock** and **omp_unset_nest_lock**, see [Section 19.9.5](#)
- **ompt_scope_endpoint_t**, see [Section 20.4.4.11](#)

20.5.2.17 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

1 runtime routine. If the implementation of the region is inlined then *codeptr_ra* contains the return
2 address of the callback invocation. If attribution to source code is impossible or inappropriate,
3 *codeptr_ra* may be `NULL`.

4 **Cross References**

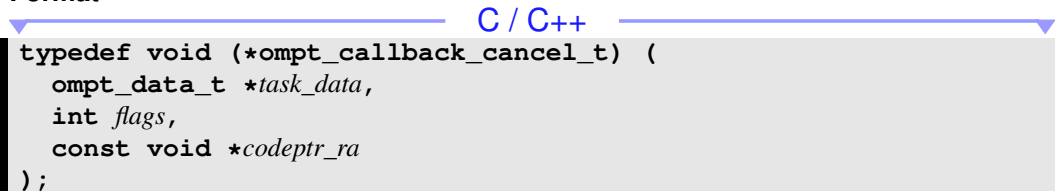
- 5 • `flush` directive, see [Section 16.8.6](#)
- 6 • `ompt_data_t`, see [Section 20.4.4.4](#)

7 **20.5.2.18 ompt_callback_cancel_t**

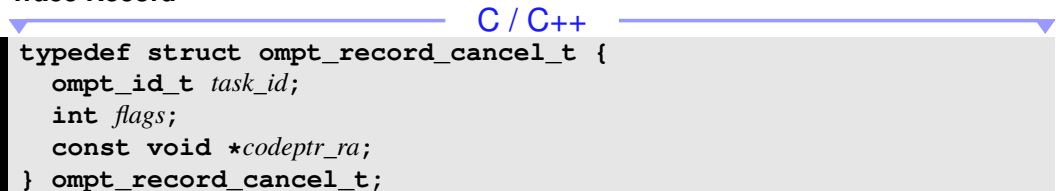
8 **Summary**

9 The `ompt_callback_cancel_t` type is used for callbacks that are dispatched for *cancellation*,
10 *cancel* and *discarded-task* events.

11 **Format**

12  C / C++
13 `typedef void (*ompt_callback_cancel_t) (
14 ompt_data_t *task_data,
15 int flags,
16 const void *codeptr_ra
17);`

17 **Trace Record**

18  C / C++
19 `typedef struct ompt_record_cancel_t {
20 ompt_id_t task_id;
21 int flags;
22 const void *codeptr_ra;
23 } ompt_record_cancel_t;`

23 **Description of Arguments**

24 The binding of the *task_data* argument is the task that encounters a **cancel** construct, a
25 **cancellation point** construct, or a construct defined as having an implicit cancellation
26 point.

27 The *flags* argument, defined by the `ompt_cancel_flag_t` enumeration type, indicates whether
28 cancellation is activated by the current task or detected as being activated by another task. The
29 construct that is being canceled is also described in the *flags* argument. When several constructs are
30 detected as being concurrently canceled, each corresponding bit in the argument will be set.

1 The *codeptr_ra* argument relates the implementation of an OpenMP region to its source code. If a
2 runtime routine implements the region associated with a callback that has type signature
3 **ompt_callback_cancel_t** then *codeptr_ra* contains the return address of the call to that
4 runtime routine. If the implementation of the region is inlined then *codeptr_ra* contains the return
5 address of the callback invocation. If attribution to source code is impossible or inappropriate,
6 *codeptr_ra* may be **NULL**.

7 Cross References

- 8 • **ompt_cancel_flag_t**, see [Section 20.4.4.26](#)

9 20.5.2.19 ompt_callback_device_initialize_t

10 Summary

11 The **ompt_callback_device_initialize_t** type is used for callbacks that initialize
12 device tracing interfaces.

13 Format

```
14 typedef void (*ompt_callback_device_initialize_t) (  
15     int device_num,  
16     const char *type,  
17     ompt_device_t *device,  
18     ompt_function_lookup_t lookup,  
19     const char *documentation  
20 );
```

C / C++

21 Semantics

22 Registration of a callback with type signature **ompt_callback_device_initialize_t** for
23 the **ompt_callback_device_initialize** event enables asynchronous collection of a trace
24 for a device. The OpenMP implementation invokes this callback after OpenMP is initialized for the
25 device but before execution of any OpenMP construct is started on the device.

26 Description of Arguments

27 The *device_num* argument identifies the logical device that is being initialized.

28 The *type* argument is a C string that indicates the type of the device. A device type string is a
29 semicolon-separated character string that includes, at a minimum, the vendor and model name of
30 the device. These names may be followed by a semicolon-separated sequence of properties that
31 describe the hardware or software of the device.

32 The *device* argument is a pointer to an opaque object that represents the target device instance.
33 Functions in the device tracing interface use this pointer to identify the device that is being
34 addressed.

1 The *lookup* argument points to a runtime callback that a tool must use to obtain pointers to runtime
2 entry points in the device's OMPT tracing interface. If a device does not support tracing then
3 *lookup* is [NULL](#).

4 The *documentation* argument is a C string that describes how to use any device-specific runtime
5 entry points that can be obtained through the *lookup* argument. This documentation string may be a
6 pointer to external documentation, or it may be inline descriptions that include names and type
7 signatures for any device-specific interfaces that are available through the *lookup* argument along
8 with descriptions of how to use these interface functions to control monitoring and analysis of
9 device traces.

10 **Constraints on Arguments**

11 The *type* and *documentation* arguments must be immutable strings that are defined for the lifetime
12 of program execution.

13 **Effect**

14 A device initializer must fulfill several duties. First, the *type* argument should be used to determine
15 if any special knowledge about the hardware and/or software of a device is employed. Second, the
16 *lookup* argument should be used to look up pointers to runtime entry points in the OMPT tracing
17 interface for the device. Finally, these runtime entry points should be used to set up tracing for the
18 device. Initialization of tracing for a target device is described in [Section 20.2.5](#).

19 **Cross References**

- 20 • Lookup Entry Points: `ompt_function_lookup_t`, see [Section 20.6.3](#)

21 **20.5.2.20 ompt_callback_device_finalize_t**

22 **Summary**

23 The `ompt_callback_device_initialize_t` type is used for callbacks that finalize device
24 tracing interfaces.

25 **Format**

```
26 typedef void (*ompt_callback_device_finalize_t) (  
27     int device_num  
28 );
```

C / C++

29 **Description of Arguments**

30 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 20.5.2.24](#)

20.5.2.21 `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.

1 The *bytes* argument indicates the size of the device code object in bytes.
2 The *host_addr* argument indicates the address at which a copy of the device code is available in
3 host memory. A value of `ompt_addr_none` indicates that a host code address is not available.
4 The *device_addr* argument indicates the address at which the device code has been loaded in device
5 memory. A value of `ompt_addr_none` indicates that a device code address is not available.
6 The *module_id* argument is an identifier that is associated with the device code object.

7 **Cross References**

- 8 • Device Directives and Clauses, see [Chapter 14](#)

9 **20.5.2.22 ompt_callback_device_unload_t**

10 **Summary**

11 The `ompt_callback_device_unload_t` type is used for callbacks that the OpenMP
12 runtime invokes to indicate that it is about to unload code from the specified device.

13 **Format**

```
14 typedef void (*ompt_callback_device_unload_t) (  
15     int device_num,  
16     uint64_t module_id  
17 );
```

C / C++

18 **Description of Arguments**

19 The *device_num* argument specifies the device.

20 The *module_id* argument is an identifier that is associated with the device code object.

21 **Cross References**

- 22 • Device Directives and Clauses, see [Chapter 14](#)

23 **20.5.2.23 ompt_callback_buffer_request_t**

24 **Summary**

25 The `ompt_callback_buffer_request_t` type is used for callbacks that are dispatched
26 when a buffer to store event records for a device is requested.

27 **Format**

```
28 typedef void (*ompt_callback_buffer_request_t) (  
29     int device_num,  
30     ompt_buffer_t **buffer,  
31     size_t *bytes  
32 );
```

C / C++

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 20.4.4.7](#)
- `ompt_get_buffer_limits_t`, see [Section 20.6.2.6](#)

20.5.2.24 `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

```
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  
);
```

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 20.4.4.8](#)
- `ompt_buffer_t`, see [Section 20.4.4.7](#)

20.5.2.25 `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++

```
typedef void (*ompt_callback_target_data_op_emi_t) (  
    ompt_scope_endpoint_t endpoint,  
    ompt_data_t *target_task_data,  
    ompt_data_t *target_data,  
    ompt_id_t *host_op_id,  
    ompt_target_data_op_t otype,  
    void *dev1_addr,  
    int dev1_device_num,  
    void *dev2_addr,  
    int dev2_device_num,  
    size_t bytes,  
    const void *codeptr_ra  
);
```



```

1  typedef void (*ompt_callback_target_data_op_t) (
2      ompt_id_t target_id,
3      ompt_id_t host_op_id,
4      ompt_target_data_op_t optype,
5      void *dev1_addr,
6      int dev1_device_num,
7      void *dev2_addr,
8      int dev2_device_num,
9      size_t bytes,
10     const void *codeptr_ra
11 );

```

C / C++

12 Trace Record

```

13 typedef struct ompt_record_target_data_op_t {
14     ompt_id_t host_op_id;
15     ompt_target_data_op_t optype;
16     void *dev1_addr;
17     int dev1_device_num;
18     void *dev2_addr;
19     int dev2_device_num;
20     size_t bytes;
21     ompt_device_time_t end_time;
22     const void *codeptr_ra;
23 } ompt_record_target_data_op_t;

```

C / C++

24 Semantics

25 A thread dispatches a registered `ompt_callback_target_data_op_emi` or
26 `ompt_callback_target_data_op` callback when device memory is allocated or freed, as
27 well as when data is copied to or from a device.

28

29 Note – An OpenMP implementation may aggregate program variables and data operations upon
30 them. For instance, an OpenMP implementation may synthesize a composite to represent multiple
31 scalars and then allocate, free, or copy this composite as a whole rather than performing data
32 operations on each scalar individually. Thus, callbacks may not be dispatched as separate data
33 operations on each variable.

34

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 *dev1_addr* argument indicates the data address on the *device* given by Table 20.4 or `NULL` for `omp_target_alloc` and `omp_target_free`.

The *dev1_device_num* argument indicates the *device number* on the *device* given by Table 20.4.

The *dev2_addr* argument indicates the data address on the *device* given by Table 20.4.

The *dev2_device_num* argument indicates the *device number* on the *device* given by Table 20.4.

Whether in some operations *dev1_addr* or *dev2_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`.

If `ompt_set_trace_ompt` has configured the implementation to trace data operations to device memory then the implementation will log an `ompt_record_target_data_op_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_data_op_t` record with its associated `ompt_callback_target_data_op_emi` or `ompt_callback_target_data_op` callback on the host;
- The *src_addr*, *src_device_num*, *dest_addr*, *dest_device_num*, *bytes*, and *codeptr_ra* fields contain the values described above for the associated callback;
- The time when the data operation began execution for the device is recorded in the *time* field of an enclosing `ompt_record_t` structure; and
- The time when the data operation completed execution for the device is recorded in the *end_time* field.

TABLE 20.4: Association of dev1 and dev2 arguments for target data operations

Data op	dev1	dev2
alloc	host	device
transfer	<i>from</i> device	<i>to</i> device
delete	host	device
associate	host	device
disassociate	host	device

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 6.8.3](#)
- `ompt_data_t`, see [Section 20.4.4.4](#)
- `ompt_id_t`, see [Section 20.4.4.3](#)
- `ompt_scope_endpoint_t`, see [Section 20.4.4.11](#)
- `ompt_target_data_op_t`, see [Section 20.4.4.15](#)

20.5.2.26 `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

```

17
18
19
20
21
22
23
24
25

```

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  
);
```

```

1  typedef void (*ompt_callback_target_t) (
2      ompt_target_t kind,
3      ompt_scope_endpoint_t endpoint,
4      int device_num,
5      ompt_data_t *task_data,
6      ompt_id_t target_id,
7      const void *codeptr_ra
8  );

```

C / C++

Trace Record

```

9
10 typedef struct ompt_record_target_t {
11     ompt_target_t kind;
12     ompt_scope_endpoint_t endpoint;
13     int device_num;
14     ompt_id_t task_id;
15     ompt_id_t target_id;
16     const void *codeptr_ra;
17 } 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 14.8](#)
- `target data` directive, see [Section 14.5](#)
- `target enter data` directive, see [Section 14.6](#)
- `target exit data` directive, see [Section 14.7](#)
- `target update` directive, see [Section 14.9](#)
- `ompt_data_t`, see [Section 20.4.4.4](#)
- `ompt_id_t`, see [Section 20.4.4.3](#)
- `ompt_scope_endpoint_t`, see [Section 20.4.4.11](#)
- `ompt_target_t`, see [Section 20.4.4.21](#)

20.5.2.27 `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  
);
```

```

1  typedef void (*ompt_callback_target_map_t) (
2      ompt_id_t target_id,
3      unsigned int nitems,
4      void **host_addr,
5      void **device_addr,
6      size_t *bytes,
7      unsigned int *mapping_flags,
8      const void *codeptr_ra
9  );

```

▲ C / C++ ▲

Trace Record

```

10
11  typedef struct ompt_record_target_map_t {
12      ompt_id_t target_id;
13      unsigned int nitems;
14      void **host_addr;
15      void **device_addr;
16      size_t *bytes;
17      unsigned int *mapping_flags;
18      const void *codeptr_ra;
19  } 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.

1 The *bytes* argument indicates an array of sizes of data.

2 The *mapping_flags* argument indicates the kind of mapping operations, which may result from
3 explicit **map** clauses or the implicit data-mapping rules defined in [Section 6.8](#). Flags for the
4 mapping operations include one or more values specified by the **ompt_target_map_flag_t**
5 type.

6 The *codeptr_ra* argument relates the implementation of an OpenMP region to its source code. If a
7 runtime routine implements the region associated with a callback that has type signature
8 **ompt_callback_target_map_t** or **ompt_callback_target_map_emi_t** then
9 *codeptr_ra* contains the return address of the call to that runtime routine. If the implementation of
10 the region is inlined then *codeptr_ra* contains the return address of the callback invocation. If
11 attribution to source code is impossible or inappropriate, *codeptr_ra* may be **NULL**.

12 Restrictions

13 Restrictions to the **ompt_callback_target_data_map_emi** and
14 **ompt_callback_target_data_map** callbacks are as follows:

- These callbacks must not be registered at the same time.

16 Cross References

- **target** directive, see [Section 14.8](#)
- **target data** directive, see [Section 14.5](#)
- **target enter data** directive, see [Section 14.6](#)
- **target exit data** directive, see [Section 14.7](#)
- **ompt_callback_target_data_op_emi_t** and
22 **ompt_callback_target_data_op_t**, see [Section 20.5.2.25](#)
- **ompt_data_t**, see [Section 20.4.4.4](#)
- **ompt_id_t**, see [Section 20.4.4.3](#)
- **ompt_target_map_flag_t**, see [Section 20.4.4.23](#)

26 20.5.2.28 **ompt_callback_target_submit_emi_t** and 27 **ompt_callback_target_submit_t**

28 Summary

29 The **ompt_callback_target_submit_emi_t** and
30 **ompt_callback_target_submit_t** types are used for callbacks that are dispatched before
31 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 implementation to trace kernel execution for a device then the implementation will log an `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 14.8](#)
- `ompt_data_t`, see [Section 20.4.4.4](#)
- `ompt_id_t`, see [Section 20.4.4.3](#)
- `ompt_scope_endpoint_t`, see [Section 20.4.4.11](#)

20.5.2.29 `ompt_callback_control_tool_t`

Summary

The `ompt_callback_control_tool_t` type is used for callbacks that dispatch *tool-control* events.

Format

```

C / C++
typedef int (*ompt_callback_control_tool_t) (
    uint64_t command,
    uint64_t modifier,
    void *arg,
    const void *codeptr_ra
);
C / C++

```

Trace Record

C / C++

```
typedef struct ompt_record_control_tool_t {  
    uint64_t command;  
    uint64_t modifier;  
    const void *codeptr_ra;  
} ompt_record_control_tool_t;
```

C / C++

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 19.14](#).

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 `ompt_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 19.14](#)

20.5.2.30 `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 9.1](#)
- `ompt_severity_t`, see [Section 20.4.4.25](#)

20.6 OMPT Runtime Entry Points for Tools

OMPT supports two principal sets of runtime entry points for tools. One set of runtime entry points enables a tool to register callbacks for OpenMP events and to inspect the state of an OpenMP thread while executing in a tool callback or a [signal handler](#). The second set of [runtime entry points](#) enables a tool to trace activities on a [device](#). When directed by the tracing interface, an OpenMP implementation will trace activities on a [device](#), collect buffers of [trace records](#), and invoke callbacks on the host to process these records. OMPT runtime entry points should not be global symbols since tools cannot rely on the visibility of such symbols.

OMPT also supports runtime entry points for two classes of lookup routines. The first class of lookup routines contains a single member: a routine that returns runtime entry points in the OMPT callback interface. The second class of lookup routines includes a unique lookup routine for each kind of device that can return runtime entry points in a device's OMPT tracing interface.

The `omp-tools.h` C/C++ header file provides the definitions of the types that are specified throughout this subsection.

Binding

The binding thread set for each of the entry points in this section is the encountering thread unless otherwise specified. The binding task set is the task executing on the encountering thread.

Restrictions

Restrictions on OMPT runtime entry points are as follows:

- OMPT runtime entry points must not be called from a [signal handler](#) on a [native thread](#) before a *native-thread-begin* or after a *native-thread-end* event.
- OMPT device runtime entry points must not be called after a *device-finalize* event for that device.

20.6.1 Entry Points in the OMPT Callback Interface

Entry points in the OMPT callback interface enable a tool to register callbacks for OpenMP events and to inspect the state of an OpenMP thread while executing in a tool callback or a [signal handler](#). Pointers to these [runtime entry points](#) are obtained through the lookup function that is provided through the OMPT initializer.

20.6.1.1 `omp_enumerate_states_t`

Summary

The `omp_enumerate_states_t` type is the type signature of the `omp_enumerate_states` runtime entry point, which enumerates the thread states that an OpenMP implementation supports.

Format

C / C++

```
typedef int (*ompt_enumerate_states_t) (  
    int current_state,  
    int *next_state,  
    const char **next_state_name  
);
```

C / C++

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 20.4.4.28](#)

20.6.1.2 `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
);
```

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.

20.6.1.3 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 20.4.2](#)
- Monitoring Activity on the Host with OMPT, see [Section 20.2.4](#)
- `ompt_callback_t`, see [Section 20.4.4.1](#)
- `ompt_get_callback_t`, see [Section 20.6.1.4](#)
- `ompt_set_result_t`, see [Section 20.4.4.2](#)

20.6.1.4 `ompt_get_callback_t`

Summary

The `ompt_get_callback_t` type is the type signature of the `ompt_get_callback` runtime entry point, which retrieves a pointer to a registered tool callback routine (if any) that an OpenMP implementation invokes when a host OpenMP event occurs.

Format

```

C / C++
typedef int (*ompt_get_callback_t) (
    ompt_callbacks_t event,
    ompt_callback_t *callback
);
C / C++
```

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 the [tool callback](#) that is registered for the specified [event](#) is not `NULL`, 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 20.4.2](#)
- `ompt_callback_t`, see [Section 20.4.4.1](#)
- `ompt_set_callback_t`, see [Section 20.6.1.3](#)

20.6.1.5 `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 20.4.4.4](#)

20.6.1.6 `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](#).

20.6.1.7 `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

```
typedef int (*ompt_get_num_places_t) (void);
```

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 3.1.5](#)
- *place-partition-var* ICV, see [Table 2.1](#)

20.6.1.8 `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.

Description of Arguments

The `place_num` argument specifies the place that is being queried.

The `ids` argument is an array in which the routine can return a vector of processor identifiers in the specified place.

The `ids_size` argument indicates the size of the result array that is specified by `ids`.

Effect

If the `ids` array of size `ids_size` is large enough to contain all identifiers then they are returned in `ids` and their order in the array is implementation defined. Otherwise, if the `ids` array is too small, the values in `ids` when the function returns are unspecified. The routine always returns the number of numerical identifiers of the processors that are available to the execution environment in the specified place.

20.6.1.9 `ompt_get_place_num_t`

Summary

The `ompt_get_place_num_t` type is the type signature of the `ompt_get_place_num` runtime entry point, which returns the place number of the place to which the current thread is bound.

Format

C / C++

```
typedef int (*ompt_get_place_num_t) (void);
```

C / C++

Semantics

When the current thread is bound to a place, `ompt_get_place_num` returns the place number associated with the thread. The returned value is between 0 and one less than the value returned by `ompt_get_num_places`, inclusive. When the current thread is not bound to a place, the routine returns -1. This runtime entry point is [async signal safe](#).

20.6.1.10 `ompt_get_partition_place_nums_t`

Summary

The `ompt_get_partition_place_nums_t` type is the type signature of the `ompt_get_partition_place_nums` runtime entry point, which returns a list of place numbers that correspond to the places in the *place-partition-var* ICV of the innermost implicit task.

Format

```
C / C++
typedef int (*ompt_get_partition_place_nums_t) (
    int place_nums_size,
    int *place_nums
);
C / C++
```

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 3.1.5](#)
- *place-partition-var* ICV, see [Table 2.1](#)

20.6.1.11 `ompt_get_proc_id_t`

Summary

The `ompt_get_proc_id_t` type is the type signature of the `ompt_get_proc_id` runtime entry point, which returns the numerical identifier of the processor of the current thread.

Format

```
typedef int (*ompt_get_proc_id_t) (void);
```

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](#).

20.6.1.12 `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  
);
```

Semantics

Each [thread](#) has an associated state and a [wait identifier](#). If the [thread state](#) indicates that the [thread](#) is waiting for mutual exclusion then its [wait identifier](#) contains a [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](#) and the [wait identifier](#) passed as the `wait_id` argument is not `NULL` then the value of the [wait identifier](#) is assigned to that 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 [runtime entry point](#) must be a reference to a [variable](#) of the specified type or `NULL`.

Cross References

- `ompt_wait_id_t`, see [Section 20.4.4.31](#)
- `ompt_enumerate_states_t`, see [Section 20.6.1.1](#)
- `ompt_state_t`, see [Section 20.4.4.28](#)

20.6.1.13 `ompt_get_parallel_info_t`

Summary

The `ompt_get_parallel_info_t` type is the type signature of the `ompt_get_parallel_info` runtime entry point, which returns information about the parallel region, if any, at the specified ancestor level for the current execution context.

Format

```

C / C++
typedef int (*ompt_get_parallel_info_t) (
    int ancestor_level,
    ompt_data_t **parallel_data,
    int *team_size
);
```

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.

1 A tool may use the pointer to the data object of a parallel region that it obtains from this runtime
2 entry point to inspect or to modify the value of the data object. When a parallel region is created, its
3 data object will be initialized with the value `ompt_data_none`.

4 This runtime entry point is *async signal safe*.

5 Between a *parallel-begin* event and an *implicit-task-begin* event, a call to
6 `ompt_get_parallel_info(0, ...)` may return information about the outer parallel team or
7 the new parallel team.

8 If a thread is in the state `ompt_state_wait_barrier_implicit_parallel` then a call to
9 `ompt_get_parallel_info` may return a pointer to a copy of the specified parallel region's
10 *parallel_data* rather than a pointer to the data word for the region itself. This convention enables
11 the primary thread for a parallel region to free storage for the region immediately after the region
12 ends, yet avoid having some other thread in the team that is executing the region potentially
13 reference the *parallel_data* object for the region after it has been freed.

14 Description of Arguments

15 The *ancestor_level* argument specifies the parallel region of interest by its ancestor level. Ancestor
16 level 0 refers to the innermost parallel region; information about enclosing parallel regions may be
17 obtained using larger values for *ancestor_level*.

18 The *parallel_data* argument returns the parallel data if the argument is not `NULL`.

19 The *team_size* argument returns the team size if the argument is not `NULL`.

20 Effect

21 If the runtime entry point returns 0 or 1, no argument is modified. Otherwise,
22 `ompt_get_parallel_info` has the following effects:

- 23 • If a [non-null value](#) was passed for *parallel_data*, the value returned in *parallel_data* is a
24 pointer to a data word that is associated with the parallel [region](#) at the specified level; and
- 25 • If a [non-null value](#) was passed for *team_size*, the value returned in the integer to which
26 *team_size* point is the number of [threads](#) in the [team](#) that is associated with the parallel [region](#).

27 Constraints on Arguments

28 While argument *ancestor_level* is passed by value, all other arguments to the entry point must be
29 pointers to variables of the specified types or `NULL`.

30 Cross References

- 31 • `ompt_data_t`, see [Section 20.4.4.4](#)

32 20.6.1.14 `ompt_get_task_info_t`

33 Summary

34 The `ompt_get_task_info_t` type is the type signature of the `ompt_get_task_info`
35 runtime entry point, which returns information about the task, if any, at the specified ancestor level
36 in the current execution context.

Format

C / C++

```
typedef int (*ompt_get_task_info_t) (  
    int ancestor_level,  
    int *flags,  
    ompt_data_t **task_data,  
    ompt_frame_t **task_frame,  
    ompt_data_t **parallel_data,  
    int *thread_num  
);
```

C / C++

Semantics

During execution, a [thread](#) may be executing a [task](#). Additionally, the stack of the [thread](#) may contain [procedure frames](#) that are associated with suspended [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](#).

1 The *task_frame* argument returns the task frame pointer if the argument is not [NULL](#).

2 The *parallel_data* argument returns the parallel data if the argument is not [NULL](#).

3 The *thread_num* argument returns the thread number if the argument is not [NULL](#).

4 **Effect**

5 If the runtime entry point returns 0 or 1, no argument is modified. Otherwise,
6 **ompt_get_task_info** has the following effects:

- 7 • If a [non-null value](#) was passed for *flags* then the value returned in the integer to which *flags*
8 points represents the type of the task at the specified level; possible task types include initial,
9 implicit, explicit, and target tasks;
- 10 • If a [non-null value](#) was passed for *task_data* then the value that is returned in the object to
11 which it points is a pointer to a data word that is associated with the task at the specified level;
- 12 • If a [non-null value](#) was passed for *task_frame* then the value that is returned in the object to
13 which *task_frame* points is a pointer to the **ompt_frame_t** structure that is associated with
14 the task at the specified level;
- 15 • If a [non-null value](#) was passed for *parallel_data* then the value that is returned in the object to
16 which *parallel_data* points is a pointer to a data word that is associated with the parallel
17 region that contains the task at the specified level or, if the task at the specified level is an
18 initial task, [NULL](#); and
- 19 • If a [non-null value](#) was passed for *thread_num*, then the value that is returned in the object to
20 which *thread_num* points indicates the number of the thread in the parallel region that is
21 executing the task at the specified level.

22 **Constraints on Arguments**

23 While argument *ancestor_level* is passed by value, all other arguments to
24 **ompt_get_task_info** must be pointers to variables of the specified types or [NULL](#).

25 **Cross References**

- 26 • **ompt_data_t**, see [Section 20.4.4.4](#)
- 27 • **ompt_frame_t**, see [Section 20.4.4.29](#)
- 28 • **ompt_task_flag_t**, see [Section 20.4.4.19](#)

29 **20.6.1.15 ompt_get_task_memory_t**

30 **Summary**

31 The **ompt_get_task_memory_t** type is the type signature of the
32 **ompt_get_task_memory** runtime entry point, which returns information about memory ranges
33 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.

20.6.1.16 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

C / C++

```
typedef int (*ompt_get_target_info_t) (  
    uint64_t *device_num,  
    ompt_id_t *target_id,  
    ompt_id_t *host_op_id  
);
```

C / C++

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 20.4.4.3](#)

20.6.1.17 `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

```
typedef int (*ompt_get_num_devices_t) (void);
```

C / C++

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](#).

20.6.1.18 `ompt_get_unique_id_t`

Summary

The `ompt_get_unique_id_t` type is the type signature of the `ompt_get_unique_id` runtime entry point, which returns a unique number.

Format

```
typedef uint64_t (*ompt_get_unique_id_t) (void);
```

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](#).

20.6.1.19 `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 will dispatch all callbacks as if the program were exiting.

Restrictions

Restrictions to the `ompt_finalize_tool` routine are as follows:

- The `ompt_finalize_tool` routine must not be called from inside an [explicit region](#).
- As the `ompt_finalize_tool` routine should only be called when a tool detects that the execution of an OpenMP program is ending, a [thread](#) encountering an [explicit region](#) after the `ompt_finalize_tool` routine has completed results in unspecified behavior.

20.6.2 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.

20.6.2.1 `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 20.4.4.5](#)

20.6.2.2 `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 20.4.4.5](#)
- `ompt_device_time_t`, see [Section 20.4.4.6](#)

20.6.2.3 `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  
);
```

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 19.10.1](#)
- `ompt_device_t`, see [Section 20.4.4.5](#)
- `ompt_device_time_t`, see [Section 20.4.4.6](#)

20.6.2.4 `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.

If any of the events that correspond to the `ompt_callback_target_data_op`, `ompt_callback_data_op_emi`, `ompt_callback_target_submit` or `ompt_callback_target_submit_emi` callbacks are specified by *etype* then tracing, if supported, is enabled or disabled for those events when they occur on the host device. If any other event corresponds to the callback specified by *etype* then tracing, if supported, is enabled or disabled for the specified events when they occur on a target device.

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 20.4.2](#)
- Tracing Activity on Target Devices with OMPT, see [Section 20.2.5](#)
- `ompt_device_t`, see [Section 20.4.4.5](#)
- `ompt_set_result_t`, see [Section 20.4.4.2](#)

20.6.2.5 `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

```
typedef ompt_set_result_t (*ompt_set_trace_native_t) (  
    ompt_device_t *device,  
    int enable,  
    int flags  
);
```

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 20.2.5](#)
- `ompt_device_t`, see [Section 20.4.4.5](#)
- `ompt_native_mon_flag_t`, see [Section 20.4.4.18](#)
- `ompt_set_result_t`, see [Section 20.4.4.2](#)

20.6.2.6 `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

```
typedef void (*ompt_get_buffer_limits_t) (  
    ompt_device_t *device,  
    int *max_concurrent_allocs,  
    size_t *recommended_bytes  
);
```

C / C++

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 performing [callback dispatch](#) 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 20.5.2.24](#)
- `ompt_callback_buffer_request_t`, see [Section 20.5.2.23](#)
- `ompt_device_t`, see [Section 20.4.4.5](#)
- `ompt_start_trace_t`, see [Section 20.6.2.7](#)
- `ompt_stop_trace_t`, see [Section 20.6.2.10](#)

20.6.2.7 `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

```
                                     C / C++
typedef int (*ompt_start_trace_t) (
    ompt_device_t *device,
    ompt_callback_buffer_request_t request,
    ompt_callback_buffer_complete_t complete
);
```

C / C++

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 20.5.2.24](#)
- `ompt_callback_buffer_request_t`, see [Section 20.5.2.23](#)
- `ompt_device_t`, see [Section 20.4.4.5](#)

20.6.2.8 `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

```
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 20.4.4.5](#)

20.6.2.9 `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

```
typedef int (*ompt_flush_trace_t) (  
    ompt_device_t *device  
);
```

C / C++

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 20.4.4.5](#)

20.6.2.10 `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++

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 20.4.4.5](#)

20.6.2.11 `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 20.4.4.8](#)
- `ompt_device_t`, see [Section 20.4.4.5](#)

20.6.2.12 `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
);
```

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 20.4.3.1](#)
- `ompt_buffer_cursor_t`, see [Section 20.4.4.8](#)
- `ompt_buffer_t`, see [Section 20.4.4.7](#)

20.6.2.13 `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

```
C / C++
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_get_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 20.4.3.4](#)
- `ompt_buffer_cursor_t`, see [Section 20.4.4.8](#)
- `ompt_device_t`, see [Section 20.4.4.5](#)

20.6.2.14 `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 value](#) 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 20.4.4.8](#)
- `ompt_buffer_t`, see [Section 20.4.4.7](#)
- `ompt_id_t`, see [Section 20.4.4.3](#)

20.6.2.15 `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

```
typedef ompt_record_abstract_t *(*ompt_get_record_abstract_t) (  
    void *native_record  
);
```

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 20.4.3.3](#)

20.6.3 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 20.1.

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 20.3, except for `ompt_set_trace_ompt` and `ompt_get_record_ompt`, as described in Section 20.2.5.

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 20.6.1
- Entry Points in the OMPT Device Tracing Interface, see Section 20.6.2
- Tracing Activity on Target Devices with OMPT, see Section 20.2.5
- `ompt_initialize_t`, see Section 20.5.1.1

21 OMPD Interface

This chapter describes **OMPD**, which is an interface for **third-party tool**. **third-party tool** 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.

OMPD allows **third-party tools** such as debuggers to inspect the OpenMP state of a live **OpenMP program** or core file in an implementation-agnostic manner. That is, a **third-party tool** that uses **OMPD** should work with any **compliant 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 **OpenMP 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 21.2.2** and **Section 21.2.3**. 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 21.6**. 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.

1 If the [third-party tool](#) requires notification of an [event](#), it can plant a breakpoint at the matching
2 location. The location can, but may not, be a function. It can, for example, simply be a label.
3 However, the names of the locations must have external **C** linkage.

4 21.1 OMPD Interfaces Definitions

C / C++

5 A compliant implementation must supply a set of definitions for the OMPD runtime entry points,
6 OMPD third-party tool callback signatures, third-party tool interface functions and the special data
7 types of their parameters and return values. These definitions, which are listed throughout this
8 chapter, and their associated declarations shall be provided in a header file named `omp-tools.h`.
9 In addition, the set of definitions may specify other implementation-specific values.

10 The `ompd_dll_locations` variable, all OMPD third-party tool interface functions, and all
11 OMPD runtime entry points are external symbols with **C** linkage.

C / C++

12 21.2 Activating a Third-Party Tool

13 The third-party tool and the OpenMP program exist as separate processes. Thus, coordination is
14 required between the OpenMP runtime and the third-party tool for OMPD.

15 21.2.1 Enabling Runtime Support for OMPD

16 In order to support third-party tools, the OpenMP runtime may need to collect and to store
17 information that it may not otherwise maintain. The OpenMP runtime collects whatever
18 information is necessary to support OMPD if the environment variable `OMP_DEBUG` is set to
19 *enabled*.

20 Cross References

- 21 • `OMP_DEBUG`, see [Section 3.4.1](#)

22 21.2.2 `ompd_dll_locations`

23 Summary

24 The `ompd_dll_locations` global variable points to the locations of OMPD libraries that are
25 compatible with the OpenMP implementation.

26 Format

```
27 | extern const char **ompd_dll_locations;
```

C

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 a `non-null value`, the vector of strings to which it points shall be valid and complete.

Cross References

- `ompd_dll_locations_valid`, see [Section 21.2.3](#)

21.2.3 `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);
```

C

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.

21.3 OMPD Data Types

This section defines OMPD data types.

21.3.1 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++

21.3.2 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 20.4.4.31](#)

21.3.3 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.

21.3.4 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 21.3.3](#)

21.3.5 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 21.3.4](#)
- Basic Value Types, see [Section 21.3.3](#)
- `ompt_frame_flag_t`, see [Section 20.4.4.30](#)

21.3.6 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/>.

21.3.7 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/>.

21.3.8 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 21.6](#)) 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.

21.3.9 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;
```

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 21.1](#).

TABLE 21.1: 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 native 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

21.3.10 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;

```

C / C++

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.

21.3.11 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.

21.3.12 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 [OpenMP process](#) that it is monitoring. Similarly, it uniquely defines a [native thread context](#) to identify a [native thread](#) of the [OpenMP process](#) that it is monitoring. These [tool contexts](#) are opaque to the [OMP](#) library.

21.3.13 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,  
};
```

```
1  ompd_rc_needs_state_tracking = 6,  
2  ompd_rc_incompatible        = 7,  
3  ompd_rc_device_read_error   = 8,  
4  ompd_rc_device_write_error  = 9,  
5  ompd_rc_nomem               = 10,  
6  ompd_rc_incomplete          = 11,  
7  ompd_rc_callback_error      = 12,  
8  ompd_rc_incompatible_handle = 13  
9  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:

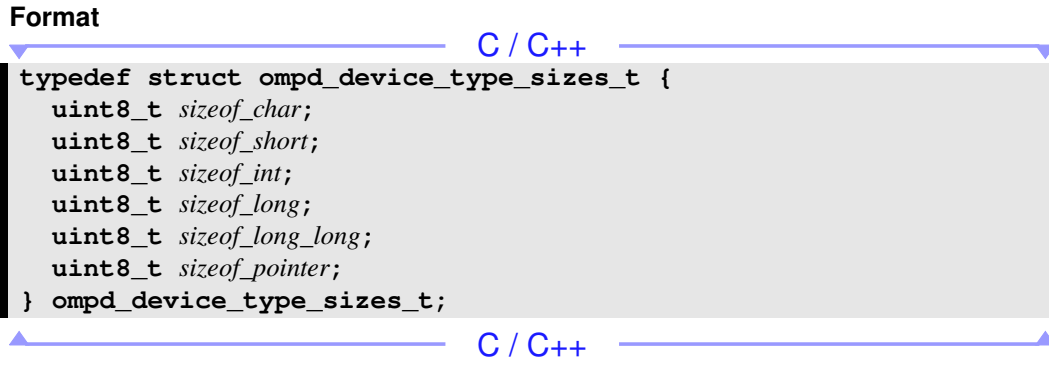
- `ompd_rc_ok` is returned when the operation is successful;
- `ompd_rc_unavailable` is returned when information is not available for the specified context;
- `ompd_rc_stale_handle` is returned when the specified handle is no longer valid;
- `ompd_rc_incompatible_handle` is returned when the specified 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;
- `ompd_rc_device_read_error` is returned when a read operation failed on the device;
- `ompd_rc_device_write_error` is returned when a write operation failed on the device;
- `ompd_rc_incompatible` is returned when this OMPD library is incompatible with the OpenMP program or is not capable of handling it;
- `ompd_rc_nomem` is returned when a memory allocation fails;
- `ompd_rc_incomplete` is returned when the information provided on return is incomplete, while the arguments are still set to valid values; and
- `ompd_rc_callback_error` is returned when the callback interface or any one of the required callback routines provided by the third-party tool is invalid.

21.3.14 Primitive Type Sizes

Summary

The `ompd_device_type_sizes_t` type provides the size of primitive types in the OpenMP architecture address space.

Format



```
typedef struct ompd_device_type_sizes_t {
    uint8_t sizeof_char;
    uint8_t sizeof_short;
    uint8_t sizeof_int;
    uint8_t sizeof_long;
    uint8_t sizeof_long_long;
    uint8_t sizeof_pointer;
} ompd_device_type_sizes_t;
```

Semantics

The `ompd_device_type_sizes_t` type is used in operations through which the OMPD library can interrogate the third-party tool about the size of primitive types for the target architecture of the OpenMP runtime, as returned by the `sizeof` operator. The fields of `ompd_device_type_sizes_t` give the sizes of the eponymous basic types used by the OpenMP runtime. As the third-party tool and the OMPD library, by definition, execute on the same architecture, the size of the fields can be given as `uint8_t`.

Cross References

- `ompd_callback_sizeof_fn_t`, see [Section 21.4.2.2](#)

21.4 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.

21.4.1 Memory Management of OMPD Library

`ompd_callback_memory_alloc_fn_t` (see [Section 21.4.1.1](#)) and `ompd_callback_memory_free_fn_t` (see [Section 21.4.1.2](#)) 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.

A third-party tool creates contexts and passes them to the OMPD library. The OMPD library does not release contexts; instead the third-party tool releases them after it releases any handles that may reference the contexts.

21.4.1.1 `ompd_callback_memory_alloc_fn_t`

Summary

The `ompd_callback_memory_alloc_fn_t` type is the type signature of the callback routine that the third-party tool provides to the OMPD library to allocate memory.

Format

```
typedef ompd_rc_t (*ompd_callback_memory_alloc_fn_t) (  
    ompd_size_t nbytes,  
    void **ptr  
);
```

Semantics

The `ompd_callback_memory_alloc_fn_t` type is the type signature of the memory allocation callback routine that the third-party tool provides. The OMPD library may call the `ompd_callback_memory_alloc_fn_t` callback function to allocate memory.

Description of Arguments

The *nbytes* argument is the size in bytes of the block of memory to allocate.

The address of the newly allocated block of memory is returned in the location to which the *ptr* argument points. The newly allocated block is suitably aligned for any type of variable and is not guaranteed to be set to zero.

Description of Return Codes

Routines that use the `ompd_callback_memory_alloc_fn_t` type may return the general return codes listed at the beginning of [Section 21.4](#).

Cross References

- Return Code Types, see [Section 21.3.13](#)
- Size Type, see [Section 21.3.1](#)
- The Callback Interface, see [Section 21.4.6](#)

21.4.1.2 `ompd_callback_memory_free_fn_t`

Summary

The `ompd_callback_memory_free_fn_t` type is the type signature of the callback routine that the 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 21.4](#).

Cross References

- Return Code Types, see [Section 21.3.13](#)
- The Callback Interface, see [Section 21.4.6](#)
- `ompd_callback_memory_alloc_fn_t`, see [Section 21.4.1.1](#)

21.4.2 Context Management and Navigation

Summary

The third-party tool provides the OMPD library with callbacks to manage and to navigate context relationships.

21.4.2.1 `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 [native 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 [tool context](#) that maps a [callback](#) that the [third-party tool](#) provides. This [callback](#) maps a [native thread identifier](#) to a [third-party tool native thread context](#). The [native thread identifier](#) is within the [address space](#) that `address_space_context` identifies. The OMPD library can use the [native 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 native thread context](#).

Description of Return Codes

In addition to the general return codes listed at the beginning of [Section 21.4](#), 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 [OMP](#) library returns from the [OMP third-party tool](#) interface routine.

Cross References

- Native Thread Identifiers, see [Section 21.3.7](#)
- Return Code Types, see [Section 21.3.13](#)
- Size Type, see [Section 21.3.1](#)
- The Callback Interface, see [Section 21.4.6](#)
- Tool Context Types, see [Section 21.3.12](#)

21.4.2.2 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 [OMP](#) 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 21.4](#).

Cross References

- Primitive Type Sizes, see [Section 21.3.14](#)
- Return Code Types, see [Section 21.3.13](#)
- The Callback Interface, see [Section 21.4.6](#)
- Tool Context Types, see [Section 21.3.12](#)

21.4.3 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.

21.4.3.1 `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 [native thread context](#) for the symbol lookup for the purpose of calculating thread local storage addresses. In this case, the [native thread](#) to which `thread_context` refers must be associated with either the [OpenMP 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, thread local storage variable, function, or untyped label. The symbol can have local, global, or weak binding.

The `file_name` argument is an optional input parameter that indicates the name of the shared library in which the symbol is defined, and it is intended to help the [third-party tool](#) disambiguate symbols that are defined multiple times across the executable or shared library files. The shared library name may not be an exact match for the name seen by the [third-party tool](#). If `file_name` is `NULL` then the [third-party tool](#) first tries to find the symbol in the executable file, and, if the symbol is not found, the [third-party tool](#) tries to find the symbol in the shared libraries in the order in which the shared libraries are loaded into the [address space](#). If `file_name` is a [non-null value](#) then the [third-party tool](#) first tries to find the symbol in the libraries that match the name in the `file_name` argument, and, if the symbol is not found, the [third-party tool](#) then uses the same [procedure](#) as when `file_name` is `NULL`.

The callback does not support finding either symbols that are dynamically allocated on the call stack or statically allocated symbols that are defined within the scope of a function or subroutine.

The callback returns the address of the symbol in the location to which `symbol_addr` points.

Description of Return Codes

In addition to the general return codes listed at the beginning of [Section 21.4](#), routines that use the `ompd_callback_symbol_addr_fn_t` type may also return the following return codes:

- `ompd_rc_error` if the requested symbol is not found; or
- `ompd_rc_bad_input` if no symbol name is provided.

Restrictions

Restrictions on routines that use the `ompd_callback_symbol_addr_fn_t` type are as follows:

- The `address_space_context` argument must be a [non-null value](#).
- The symbol that the `symbol_name` argument specifies must be defined.

Cross References

- Address Type, see [Section 21.3.4](#)
- Return Code Types, see [Section 21.3.13](#)
- The Callback Interface, see [Section 21.4.6](#)
- Tool Context Types, see [Section 21.3.12](#)

21.4.3.2 `ompd_callback_memory_read_fn_t`

Summary

The `ompd_callback_memory_read_fn_t` type is the type signature of the callback that the third-party tool provides to read data (`read_memory`) or a string (`read_string`) from an OpenMP program.

Format

```
typedef ompd_rc_t (*ompd_callback_memory_read_fn_t) (  
    ompd_address_space_context_t *address_space_context,  
    ompd_thread_context_t *thread_context,  
    const ompd_address_t *addr,  
    ompd_size_t nbytes,  
    void *buffer  
);
```

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 a non-null value, *thread_context* identifies the native thread 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 21.4.4) to interpret the data.

Description of Return Codes

In addition to the general return codes listed at the beginning of Section 21.4, 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 21.3.4
- Return Code Types, see Section 21.3.13
- Size Type, see Section 21.3.1
- The Callback Interface, see Section 21.4.6
- Tool Context Types, see Section 21.3.12
- Data Format Conversion: `ompd_callback_device_host_fn_t`, see Section 21.4.4

21.4.3.3 `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 a [non-null value](#), then *thread_context* identifies the [native thread 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 21.4.4](#)) 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 21.4](#).

Cross References

- Address Type, see [Section 21.3.4](#)
- Return Code Types, see [Section 21.3.13](#)
- Size Type, see [Section 21.3.1](#)
- The Callback Interface, see [Section 21.4.6](#)
- Tool Context Types, see [Section 21.3.12](#)
- Data Format Conversion: `ompd_callback_device_host_fn_t`, see [Section 21.4.4](#)

21.4.4 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 21.4](#).

Cross References

- Return Code Types, see [Section 21.3.13](#)
- Size Type, see [Section 21.3.1](#)
- The Callback Interface, see [Section 21.4.6](#)
- Tool Context Types, see [Section 21.3.12](#)

21.4.5 `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 21.4](#).

Cross References

- Return Code Types, see [Section 21.3.13](#)
- The Callback Interface, see [Section 21.4.6](#)

21.4.6 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;
```



```

1  ompd_callback_memory_write_fn_t write_memory;
2  ompd_callback_memory_read_fn_t read_string;
3  ompd_callback_device_host_fn_t device_to_host;
4  ompd_callback_device_host_fn_t host_to_device;
5  ompd_callback_get_thread_context_for_thread_id_fn_t
6      get_thread_context_for_thread_id;
7  ompd_callbacks_t;

```

C

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 21.5.1.1](#)). Each field points to a function that the OMPD library must use either to interact with the OpenMP program or for memory operations.

The `alloc_memory` and `free_memory` fields are pointers to functions the OMPD library uses to allocate and to release dynamic memory.

The `print_string` field points to a function that prints a string.

The architecture on which the OMPD library and third-party tool execute may be different from the architecture on which the OpenMP program that is being examined executes. The `sizeof_type` field points to a function that allows the OMPD library to determine the sizes of the basic integer and pointer types that the OpenMP program uses. Because of the potential differences in the targeted architectures, the conventions for representing data in the OMPD library and the OpenMP program may be different. The `device_to_host` field points to a function that translates data from the conventions that the OpenMP program uses to those that the third-party tool and OMPD library use. The reverse operation is performed by the function to which the `host_to_device` field points.

The `symbol_addr_lookup` field points to a callback that the OMPD library can use to find the address of a global or thread local storage symbol. The `read_memory`, `read_string` and `write_memory` fields are pointers to functions for reading from and writing to global memory or thread local storage in the OpenMP program.

The `get_thread_context_for_thread_id` field is a pointer to a function that the OMPD library can use to obtain a [native thread context](#) that corresponds to a [native thread identifier](#).

Cross References

- Data Format Conversion: `ompd_callback_device_host_fn_t`, see [Section 21.4.4](#)
- `ompd_callback_get_thread_context_for_thread_id_fn_t`, see [Section 21.4.2.1](#)
- `ompd_callback_memory_alloc_fn_t`, see [Section 21.4.1.1](#)
- `ompd_callback_memory_free_fn_t`, see [Section 21.4.1.2](#)
- `ompd_callback_memory_read_fn_t`, see [Section 21.4.3.2](#)

- `ompd_callback_memory_write_fn_t`, see [Section 21.4.3.3](#)
- `ompd_callback_print_string_fn_t`, see [Section 21.4.5](#)
- `ompd_callback_sizeof_fn_t`, see [Section 21.4.2.2](#)
- `ompd_callback_symbol_addr_fn_t`, see [Section 21.4.3.1](#)

21.5 OMPD Tool Interface Routines

This section defines the interface provided by the OMPD library to be used by the third-party tool. Some interface routines require one or more specified threads to be *stopped* for the returned values to be meaningful. In this context, a stopped thread is a thread that is not modifying the observable 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.

21.5.1 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 21.5.1.2](#)). If the tool supports the version that `ompd_get_api_version` returns, the tool starts the initialization by calling `ompd_initialize` (see [Section 21.5.1.1](#)) 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.

21.5.1.1 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 21.5](#) 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 21.3.13](#)
- The Callback Interface, see [Section 21.4.6](#)
- `ompd_get_api_version`, see [Section 21.5.1.2](#)

1 21.5.1.2 ompd_get_api_version

2 Summary

3 The `ompd_get_api_version` function returns the OMPD API version.

4 Format

5  C

```
6        | ompd_rc_t ompd_get_api_version(ompd_word_t *version);
```

6 Semantics

7 The tool may call the `ompd_get_api_version` function to obtain the latest OMPD API
8 version number of the OMPD library. The OMPD API version number is equal to the value of the
9 `_OPENMP` macro defined in the associated OpenMP implementation, if the C preprocessor is
10 supported. If the associated OpenMP implementation compiles Fortran codes without the use of a
11 C preprocessor, the OMPD API version number is equal to the value of the Fortran integer
12 parameter `openmp_version`.

13 Description of Arguments

14 The latest version number is returned into the location to which the `version` argument points.

15 Description of Return Codes

16 This routine must return any of the general return codes listed at the beginning of [Section 21.5](#).

17 Cross References

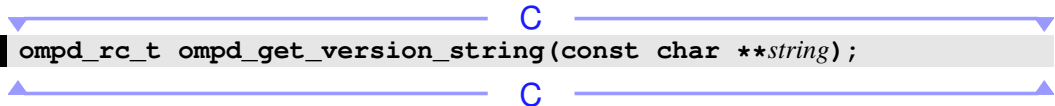
- 18 • Return Code Types, see [Section 21.3.13](#)

19 21.5.1.3 ompd_get_version_string

20 Summary

21 The `ompd_get_version_string` function returns a descriptive string for the OMPD library
22 version.

23 Format

24  C

```
25       | ompd_rc_t ompd_get_version_string(const char **string);
```

25 Semantics

26 The tool may call this function to obtain a pointer to a descriptive version string of the OMPD
27 library vendor, implementation, internal version, date, or any other information that may be useful
28 to a tool user or vendor. An implementation should provide a different string for every change to its
29 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 21.5.1.1](#)). Accordingly, the OMPD library must not use heap or stack memory for the string.

The signatures of `ompd_get_api_version` (see [Section 21.5.1.2](#)) 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 21.5](#).

Cross References

- Return Code Types, see [Section 21.3.13](#)

21.5.1.4 `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 21.5](#) or the following return code:

- `ompd_rc_unsupported` if the OMPD library is not initialized.

Cross References

- Return Code Types, see [Section 21.3.13](#)

21.5.2 Per OpenMP Process Initialization and Finalization

21.5.2.1 `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 21.5](#) 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 21.3.8](#)
- Return Code Types, see [Section 21.3.13](#)
- Tool Context Types, see [Section 21.3.12](#)
- `ompd_rel_address_space_handle`, see [Section 21.5.2.3](#)

21.5.2.2 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 21.5](#) 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 21.3.8](#)
- Return Code Types, see [Section 21.3.13](#)
- Size Type, see [Section 21.3.1](#)
- System Device Identifiers, see [Section 21.3.6](#)
- Tool Context Types, see [Section 21.3.12](#)

21.5.2.3 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 21.5](#).

Cross References

- OMPD Handle Types, see [Section 21.3.8](#)
- Return Code Types, see [Section 21.3.13](#)

21.5.2.4 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 21.5](#).

Cross References

- Native Thread Identifiers, see [Section 21.3.7](#)
- OMPD Handle Types, see [Section 21.3.8](#)
- Return Code Types, see [Section 21.3.13](#)
- Size Type, see [Section 21.3.1](#)

21.5.3 Thread and Signal Safety

The [OMP library](#) does not need to be reentrant. The [tool](#) must ensure that only one [native thread](#) enters the [OMP library](#) at a time. The [OMP library](#) must not install [signal handlers](#) or otherwise interfere with the [signal](#) configuration of the [tool](#).

21.5.4 Address Space Information

21.5.4.1 `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 21.5](#).

Cross References

- OMPD Handle Types, see [Section 21.3.8](#)
- Return Code Types, see [Section 21.3.13](#)

21.5.4.2 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 21.5](#).

Cross References

- OMPD Handle Types, see [Section 21.3.8](#)
- Return Code Types, see [Section 21.3.13](#)

21.5.5 Thread Handles

21.5.5.1 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 [native 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 [thread number](#) and selects the [thread](#), the [handle](#) for which is to be returned. On return, the *thread_handle* argument is a [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 21.5](#) 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 21.3.8](#)
- Return Code Types, see [Section 21.3.13](#)
- `ompd_get_icv_from_scope`, see [Section 21.5.10.2](#)

21.5.5.2 `ompd_get_thread_handle`

Summary

The `ompd_get_thread_handle` function maps a [native thread](#) to a [native 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 [native thread handle](#) for the [native thread](#) to which the [OpenMP thread](#) is mapped.

Description of Arguments

The *handle* argument is a [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 a [handle](#) to the [native 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 21.5](#) 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 [native thread](#) is not an [OpenMP thread](#).

Cross References

- Native Thread Identifiers, see [Section 21.3.7](#)
- OMPD Handle Types, see [Section 21.3.8](#)
- Return Code Types, see [Section 21.3.13](#)
- Size Type, see [Section 21.3.1](#)

21.5.5.3 ompd_rel_thread_handle

Summary

The `ompd_rel_thread_handle` function releases a [native 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 [native 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 21.5](#).

Cross References

- OMPD Handle Types, see [Section 21.3.8](#)
- Return Code Types, see [Section 21.3.13](#)

21.5.5.4 ompd_thread_handle_compare

Summary

The `ompd_thread_handle_compare` function allows tools to compare two [native 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 [native thread handles](#) is opaque to a tool. While the tool can easily compare pointers to [native thread handles](#), it cannot determine whether [handles](#) of two different addresses refer to the same underlying [native thread](#). The `ompd_thread_handle_compare` function compares [native 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 [native threads](#) compare: a value less than, equal to, or greater than 0 indicates that the [native 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 [handles](#) for [native 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 21.5](#).

Cross References

- OMPD Handle Types, see [Section 21.3.8](#)
- Return Code Types, see [Section 21.3.13](#)

21.5.5.5 ompd_get_thread_id

Summary

The `ompd_get_thread_id` function maps a [native 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 a [native 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 a [native 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 21.5](#) 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 [native thread](#) is not supported.

Cross References

- Native Thread Identifiers, see [Section 21.3.7](#)
- OMPD Handle Types, see [Section 21.3.8](#)
- Return Code Types, see [Section 21.3.13](#)
- Size Type, see [Section 21.3.1](#)

21.5.5.6 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 a [native thread handle](#) that represents a [native thread](#) to which an [OpenMP thread](#) is mapped. On return, the `device` argument is the address of a pointer to an [address space handle](#).

Description of Return Codes

This routine must return any of the general return codes listed at the beginning of [Section 21.5](#).

Cross References

- OMPD Handle Types, see [Section 21.3.8](#)
- Return Code Types, see [Section 21.3.13](#)

21.5.6 Parallel Region Handles

21.5.6.1 `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 21.5](#) 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 21.3.8](#)
- Return Code Types, see [Section 21.3.13](#)
- `ompd_rel_parallel_handle`, see [Section 21.5.6.4](#)

21.5.6.2 `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 21.5](#) or the following return code:

- `ompd_rc_unavailable` if no enclosing parallel region exists.

Cross References

- OMPD Handle Types, see [Section 21.3.8](#)
- Return Code Types, see [Section 21.3.13](#)
- `ompd_rel_parallel_handle`, see [Section 21.5.6.4](#)

21.5.6.3 `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 21.5](#).

Cross References

- OMPD Handle Types, see [Section 21.3.8](#)
- Return Code Types, see [Section 21.3.13](#)
- `ompd_rel_parallel_handle`, see [Section 21.5.6.4](#)

21.5.6.4 `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 21.5](#).

Cross References

- OMPD Handle Types, see [Section 21.3.8](#)
- Return Code Types, see [Section 21.3.13](#)

21.5.6.5 `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 21.5](#).

Cross References

- OMPD Handle Types, see [Section 21.3.8](#)
- Return Code Types, see [Section 21.3.13](#)

21.5.7 Task Handles

21.5.7.1 `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 21.5](#) or the following return code:

- **ompd_rc_unavailable** if the thread is currently not executing a task.

Cross References

- OMPD Handle Types, see [Section 21.3.8](#)
- Return Code Types, see [Section 21.3.13](#)
- **ompd_rel_task_handle**, see [Section 21.5.7.5](#)

21.5.7.2 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 [task construct](#) that generated the [task](#) represented by *task_handle*. The [generating task](#) is the [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 21.5](#) or the following return code:

- `ompd_rc_unavailable` if no [generating task region](#) exists.

Cross References

- OMPD Handle Types, see [Section 21.3.8](#)
- Return Code Types, see [Section 21.3.13](#)
- `ompd_rel_task_handle`, see [Section 21.5.7.5](#)

21.5.7.3 `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 21.5](#) or the following return code:

- `ompd_rc_unavailable` if no scheduling task exists.

Cross References

- OMPD Handle Types, see [Section 21.3.8](#)
- Return Code Types, see [Section 21.3.13](#)
- `ompd_rel_task_handle`, see [Section 21.5.7.5](#)

21.5.7.4 `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 21.5](#) 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 21.3.8](#)
- Return Code Types, see [Section 21.3.13](#)
- `ompd_get_icv_from_scope`, see [Section 21.5.10.2](#)

21.5.7.5 `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 21.5](#).

Cross References

- OMPD Handle Types, see [Section 21.3.8](#)
- Return Code Types, see [Section 21.3.13](#)

21.5.7.6 `ompd_task_handle_compare`

Summary

The `ompd_task_handle_compare` function compares task handles.

Format

```
ompd_rc_t ompd_task_handle_compare(  
    ompd_task_handle_t *task_handle_1,  
    ompd_task_handle_t *task_handle_2,  
    int *cmp_value  
);
```


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.

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.

Description of Return Codes

This routine must return any of the general return codes listed at the beginning of [Section 21.5](#).

Cross References

- OMPD Handle Types, see [Section 21.3.8](#)
- Return Code Types, see [Section 21.3.13](#)

21.5.7.7 ompd_get_task_function

Summary

This `ompd_get_task_function` function returns the entry point of the code that corresponds to the body of a task.

Format

```
ompd_rc_t ompd_get_task_function (  
    ompd_task_handle_t *task_handle,  
    ompd_address_t *entry_point  
);
```

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 21.5](#).

Cross References

- Address Type, see [Section 21.3.4](#)
- OMPD Handle Types, see [Section 21.3.8](#)
- Return Code Types, see [Section 21.3.13](#)

21.5.7.8 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 21.5](#).

Cross References

- Address Type, see [Section 21.3.4](#)
- Frame Information Type, see [Section 21.3.5](#)
- OMPD Handle Types, see [Section 21.3.8](#)
- Return Code Types, see [Section 21.3.13](#)
- `ompt_frame_t`, see [Section 20.4.4.29](#)

21.5.8 Querying Thread States

21.5.8.1 `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 21.5](#) 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 21.3.8](#)
- Return Code Types, see [Section 21.3.13](#)
- **ompt_state_t**, see [Section 20.4.4.28](#)

21.5.8.2 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 a **non-null value** 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 21.5](#).

Cross References

- OMPD Handle Types, see [Section 21.3.8](#)
- Return Code Types, see [Section 21.3.13](#)
- Wait ID Type, see [Section 21.3.2](#)
- `ompd_enumerate_states`, see [Section 21.5.8.1](#)

21.5.9 Display Control Variables

21.5.9.1 `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 21.5](#).

Cross References

- OMPD Handle Types, see [Section 21.3.8](#)
- Return Code Types, see [Section 21.3.13](#)
- `ompd_initialize`, see [Section 21.5.1.1](#)
- `ompd_rel_display_control_vars`, see [Section 21.5.9.2](#)

21.5.9.2 `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 21.5](#).

Cross References

- Return Code Types, see [Section 21.3.13](#)
- `ompd_get_display_control_vars`, see [Section 21.5.9.1](#)

21.5.10 Accessing Scope-Specific Information

21.5.10.1 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.

When the *current* argument is set to the identifier of a supported ICV, `ompd_enumerate_icvs` assigns the value, string name, and scope of the next ICV in the enumeration to the locations to which the *next_id*, *next_icv_name*, and *next_scope* arguments point. On return, the third-party tool owns the *next_icv_name* string. The OMPD library uses the memory allocation callback that the tool provides to allocate the string storage; the tool is responsible for releasing the memory.

On return, the location to which the *more* argument points has the value of 1 whenever one or more ICV are left in the enumeration. On return, that location has the value 0 when *current* is the last ICV in the enumeration.

Description of Arguments

The *address_space_handle* argument identifies the address space. The *current* argument must be an ICV that the OpenMP implementation supports. To begin enumerating the ICVs, a tool should pass `ompd_icv_undefined` as the value of *current*. Subsequent calls to `ompd_enumerate_icvs` should pass the value returned by the call in the *next_id* output argument. On return, the *next_id* argument points to an integer with the value of the ID of the next ICV in the enumeration. On return, the *next_icv_name* argument points to a character string with the name of the next ICV. On return, the *next_scope* argument points to the scope enum value of the scope of the next ICV. On return, the *more_enums* argument points to an integer with the value of 1 when more ICVs are left to enumerate and the value of 0 when no more ICVs are left.

Description of Return Codes

This routine must return any of the general return codes listed at the beginning of [Section 21.5](#) or the following return code:

- `ompd_rc_bad_input` if an unknown value is provided in *current*.

Cross References

- ICV ID Type, see [Section 21.3.11](#)
- OMPD Handle Types, see [Section 21.3.8](#)
- OMPD Scope Types, see [Section 21.3.9](#)
- Return Code Types, see [Section 21.3.13](#)

21.5.10.2 `ompd_get_icv_from_scope`

Summary

The `ompd_get_icv_from_scope` function returns the value of an ICV.

Format

```
ompd_rc_t ompd_get_icv_from_scope (
    void *handle,
    ompd_scope_t scope,
    ompd_icv_id_t icv_id,
    ompd_word_t *icv_value
);
```

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 21.3.11](#).

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 21.5](#) 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 21.3.11](#)
- OMPD Handle Types, see [Section 21.3.8](#)
- OMPD Scope Types, see [Section 21.3.9](#)
- Return Code Types, see [Section 21.3.13](#)
- `ompd_enumerate_icvs`, see [Section 21.5.10.1](#)

21.5.10.3 `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 21.3.11](#).

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 21.5](#) 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 21.3.11](#)
- OMPD Handle Types, see [Section 21.3.8](#)
- OMPD Scope Types, see [Section 21.3.9](#)
- Return Code Types, see [Section 21.3.13](#)
- `ompd_enumerate_icvs`, see [Section 21.5.10.1](#)

21.5.10.4 `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 21.5](#) or the following return code:

- `ompd_rc_unsupported` if the runtime library does not support OMPT.

Cross References

- OMPD Handle Types, see [Section 21.3.8](#)
- OMPD Scope Types, see [Section 21.3.9](#)
- Return Code Types, see [Section 21.3.13](#)
- `ompt_data_t`, see [Section 20.4.4.4](#)

21.6 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.

21.6.1 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 11.2](#)
- `ompd_get_curr_parallel_handle`, see [Section 21.5.6.1](#)
- `ompd_get_curr_task_handle`, see [Section 21.5.7.1](#)

21.6.2 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 11.2](#)
- `ompd_get_curr_parallel_handle`, see [Section 21.5.6.1](#)
- `ompd_get_curr_task_handle`, see [Section 21.5.7.1](#)
- `ompd_rel_parallel_handle`, see [Section 21.5.6.4](#)

21.6.3 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 11.3](#)
- `ompd_get_curr_parallel_handle`, see [Section 21.5.6.1](#)
- `ompd_get_curr_task_handle`, see [Section 21.5.7.1](#)

21.6.4 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 11.3](#)
- `ompd_get_curr_parallel_handle`, see [Section 21.5.6.1](#)
- `ompd_get_curr_task_handle`, see [Section 21.5.7.1](#)
- `ompd_rel_parallel_handle`, see [Section 21.5.6.4](#)

21.6.5 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 21.5.7.1](#)

21.6.6 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 21.5.7.1](#)
- `ompd_rel_task_handle`, see [Section 21.5.7.5](#)

21.6.7 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 11.2](#)
- Initial Task, see [Section 13.9](#)

21.6.8 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 11.2](#)
- Initial Task, see [Section 13.9](#)
- `ompd_rel_thread_handle`, see [Section 21.5.5.3](#)

21.6.9 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 14.8](#)
- `ompd_get_curr_parallel_handle`, see [Section 21.5.6.1](#)
- `ompd_get_curr_task_handle`, see [Section 21.5.7.1](#)

21.6.10 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 14.8](#)
- `ompd_get_curr_parallel_handle`, see [Section 21.5.6.1](#)
- `ompd_get_curr_task_handle`, see [Section 21.5.7.1](#)
- `ompd_rel_parallel_handle`, see [Section 21.5.6.4](#)

21.6.11 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 14.4](#)

21.6.12 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 14.4](#)

8
 - **ompd_rel_address_space_handle**, see [Section 21.5.2.3](#)

Part V

Appendices

1

2

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](#)).
- **Device:** An implementation-defined logical execution engine (see [Section 1.2](#)).
- **Device pointer:** An *implementation-defined handle* that refers to a device address (see [Section 1.2](#)).
- **Supported active levels of parallelism:** The maximum number of [active parallel regions](#) that may enclose any [region](#) of code in an [OpenMP program](#) is [implementation defined](#) (see [Section 1.2](#)).
- **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](#)).
- **Memory model:** The minimum size at which a memory update may also read and write back adjacent [variables](#) that are part of an [aggregate variable](#) 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](#)).
- **Device Data Environments:** Whether a [variable](#) with [static storage duration](#) that is accessible on a [device](#) and is not a [device local variable](#) is mapped with a [persistent self map](#) at the beginning of the program is [implementation defined](#) (see [Section 1.4.2](#)).

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:

- **OMP_DYNAMIC environment variable:** If the value is neither **true** nor **false**, the behavior of the program is [implementation defined](#) (see [Section 3.1.1](#)).
- **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 3.1.2](#)).
- **OMP_THREAD_LIMIT environment variable:** If the requested value is greater than the number of [threads](#) that an implementation can support, or if the value is not a positive integer, the behavior of the program is [implementation defined](#) (see [Section 3.1.3](#)).
- **OMP_MAX_ACTIVE_LEVELS environment variable:** If the value is a negative integer or is greater than the maximum number of nested [active levels](#) that an implementation can support then the behavior of the program is [implementation defined](#) (see [Section 3.1.4](#)).
- **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 3.1.5](#)).
- **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 3.1.6](#)).
- **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 3.2.1](#)).
- **OMP_STACKSIZE environment variable:** If the value does not conform to the specified format or the implementation cannot provide a stack of the specified size then the behavior is [implementation defined](#) (see [Section 3.2.2](#)).
- **OMP_WAIT_POLICY environment variable:** The details of the **active** and **passive** behaviors are [implementation defined](#) (see [Section 3.2.3](#)).
- **OMP_DISPLAY_AFFINITY environment variable:** For all values of the environment variables other than **true** or **false**, the display action is [implementation defined](#) (see

1 Section 3.2.4).

- 2 • **OMP_AFFINITY_FORMAT environment variable:** Additional [implementation defined](#)
3 field types can be added (see [Section 3.2.5](#)).
- 4 • **OMP_CANCELLATION environment variable:** If the value is set to neither **true** nor
5 **false**, the behavior of the program is [implementation defined](#) (see [Section 3.2.6](#)).
- 6 • **OMP_TARGET_OFFLOAD environment variable:** The support of **disabled** is
7 [implementation defined](#) (see [Section 3.2.9](#)).
- 8 • **OMP_THREADS_RESERVE environment variable:** If the requested values are greater than
9 **OMP_THREAD_LIMIT**, the behavior of the program is [implementation defined](#) (see
10 [Section 3.2.10](#)).
- 11 • **OMP_TOOL_LIBRARIES environment variable:** Whether the value of the environment
12 variable is case sensitive is [implementation defined](#) (see [Section 3.3.2](#)).
- 13 • **OMP_TOOL_VERBOSE_INIT environment variable:** Support for logging to **stdout** or
14 **stderr** is [implementation defined](#). Whether the value of the environment variable is case
15 sensitive when it is treated as a filename is [implementation defined](#). The format and detail of
16 the log is [implementation defined](#) (see [Section 3.3.3](#)).
- 17 • **OMP_DEBUG environment variable:** If the value is neither **disabled** nor **enabled**, the
18 behavior is [implementation defined](#) (see [Section 3.4.1](#)).
- 19 • **OMP_NUM_TEAMS environment variable:** If the value is not a positive integer or is greater
20 than the number of **teams** that an implementation can support, the behavior of the program is
21 [implementation defined](#) (see [Section 3.6.1](#)).
- 22 • **OMP_TEAMS_THREAD_LIMIT environment variable:** If the value is not a positive integer
23 or is greater than the number of **threads** that an implementation can support, the behavior of
24 the program is [implementation defined](#) (see [Section 3.6.2](#)).

25 Chapter 4:

C / C++

- 26 • A pragma directive that uses **omp_x** as the first processing token is [implementation defined](#)
27 (see [Section 4.1](#)).
- 28 • The attribute namespace of an attribute specifier or the optional namespace qualifier within a
29 **sequence** attribute that uses **omp_x** is [implementation defined](#) (see [Section 4.1](#)).

C / C++

C++

- 30 • Whether a **throw** executed inside a **region** that arises from an [exception-aborting directive](#)
31 results in [runtime error termination](#) is [implementation defined](#) (see [Section 4.1](#)).

C++

Fortran

- Any directive that uses **omx** or **omp_x** in the sentinel is implementation defined (see [Section 4.1](#)).

Fortran

Chapter 5:

- **Loop-iteration spaces and vectors:** The particular integer type used to compute the iteration count for the collapsed loop is implementation defined (see [Section 5.4.2](#)).

Chapter 6:

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 6.1.2](#)).
- **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 6.2](#)).

Fortran

- **is_device_ptr clause:** Support for pointers created outside of the OpenMP device data management routines is implementation defined (see [Section 6.4.7](#)).

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 6.4.9](#) and [Section 6.4.10](#)).

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 6.11](#)).

Chapter 7:

- **Memory spaces:** The actual storage resources that each memory space defined in [Table 7.1](#) 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 7.1](#)).
- **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 7.2](#)) 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 7.3](#)) are implementation defined (see [Section 7.2](#)).

Chapter 8:

- **OpenMP context:** The accepted *isa-name* values for the *isa trait*, the accepted *arch-name* values for the *arch trait* and the accepted *extension-name* values for the *extension trait* are [implementation defined](#) (see [Section 8.1](#)).
- **Metadirectives:** The number of times that each expression of the context selector of a **when** clause is evaluated is [implementation defined](#) (see [Section 8.4.1](#)).
- **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 8.5](#)).
- **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 8.7](#)).
- **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 8.8](#)).

Chapter 9:

- **requires directive:** Support for any feature specified by a requirement clause on a **requires** directive is [implementation defined](#) (see [Section 9.5](#)).

Chapter 10:

- **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 10.2](#)).

Chapter 11:

- **Dynamic adjustment of threads:** Providing the ability to adjust the number of threads dynamically is [implementation defined](#) (see [Section 11.2.1](#)).
- **Compile-time message:** If the implementation determines that the requested number of threads can never be provided and therefore performs compile-time error termination, the effect of any **message** clause associated with the directive is [implementation defined](#) (see [Section 11.2.2](#)).
- **Thread affinity:** If another [OpenMP thread](#) is bound to the [place](#) associated with its position, the [place](#) to which a [free-agent thread](#) is bound is [implementation defined](#). For the **spread thread affinity**, if $T \leq P$ and T does not divide P evenly, which subpartitions contain $\lceil P/T \rceil$ [places](#) is [implementation defined](#). For the **close** and **spread thread affinity** policies, if

1 *ET* is not zero, which sets have *AT* positions and which sets have *BT* positions is
2 [implementation defined](#). Further, the positions assigned to the groups that are assigned sets
3 with *BT* positions to make the number of positions assigned to each group *AT* is
4 [implementation defined](#). The determination of whether the [thread affinity](#) request can be
5 fulfilled is [implementation defined](#). If the [thread affinity](#) request cannot be fulfilled, then the
6 [thread affinity](#) of [threads](#) in the [team](#) is [implementation defined](#) (see [Section 11.2.3](#)).

- 7 • **teams construct**: The number of teams that are created is implementation defined, but it is
8 greater than or equal to the lower bound and less than or equal to the upper bound values of
9 the `num_teams` clause if specified. If the `num_teams` clause is not specified, the number
10 of teams is less than or equal to the value of the `nteams-var` ICV if its value is greater than
11 zero. Otherwise it is an implementation defined value greater than or equal to one (see
12 [Section 11.3](#)).
- 13 • **simd construct**: The number of iterations that are executed concurrently at any given time
14 is implementation defined (see [Section 11.5](#)).

Chapter 12:

- 15 • **single construct**: The method of choosing a thread to execute the structured block each
16 time the team encounters the construct is implementation defined (see [Section 12.1](#)).
- 17 • **sections construct**: The method of scheduling the structured block sequences among
18 threads in the team is implementation defined (see [Section 12.3](#)).
- 19 • **Worksharing-loop directive**: The schedule that is used is implementation defined if the
20 `schedule` clause is not specified or if the specified schedule has the kind `auto`. The value
21 of `simd_width` for the `simd` schedule modifier is implementation defined (see [Section 12.6](#)).
- 22 • **distribute construct**: If no `dist_schedule` clause is specified then the schedule for
23 the `distribute` construct is implementation defined (see [Section 12.7](#)).

Chapter 13:

- 24 • **taskloop construct**: The number of loop iterations assigned to a task created from a
25 `taskloop` construct is implementation defined, unless the `grainsize` or `num_tasks`
26 clause is specified (see [Section 13.7](#)).

27 C++

- 28 • **taskloop construct**: For `firstprivate` variables of class type, the number of
29 invocations of copy constructors to perform the initialization is implementation defined (see
30 [Section 13.7](#)).

31 C++

Chapter 14:

- 32 • **thread_limit clause**: The maximum number of threads that participate in executing
33 tasks in the contention group that each team initiates is implementation defined if no
34 `thread_limit` clause is specified on the construct. Otherwise, it has the implementation
35 defined upper bound of the `teams-thread-limit-var` ICV, if the value of this ICV is greater
36 than zero (see [Section 14.3](#)).

1 **Chapter 15:**

- 2 • **interop Construct:** The *foreign-runtime-id* values for the **prefer_type** clause that the
3 implementation supports, including non-standard names compatible with this clause, and the
4 default choice when the implementation supports multiple values are implementation defined
5 (see [Section 15.1](#)).

6 **Chapter 16:**

- 7 • **atomic construct:** A compliant implementation may enforce exclusive access between
8 **atomic** regions that update different storage locations. The circumstances under which this
9 occurs are implementation defined. If the storage location designated by *x* is not size-aligned
10 (that is, if the byte alignment of *x* is not a multiple of the size of *x*), then the behavior of the
11 atomic region is implementation defined (see [Section 16.8.5](#)).

12 **Chapter 17:**

- 13 • None.

14 **Chapter 18:**

- 15 • None.

16 **Chapter 19:**

- 17 • Runtime Routine names that begin with the **omp_x** prefix are implementation-defined
18 extensions to the OpenMP Runtime API (see [Chapter 19](#)).

19 **C / C++**

- 20 • **Runtime library definitions:** The enum types for **omp_allocator_handle_t**,
21 **omp_event_handle_t**, **omp_interop_fr_t** and **omp_memspace_handle_t** are
22 implementation defined. The integral or pointer type for **omp_interop_t** is
23 implementation defined. The value of the **omp_invalid_device** enumerator is
24 implementation defined. The value of the **omp_unknown_thread** enumerator is
implementation defined (see [Section 19.1](#)).

25 **C / C++**

26 **Fortran**

- 27 • **Runtime library definitions:** Whether the include file **omp_lib.h** or the module
28 **omp_lib** (or both) is provided is implementation defined. Whether the **omp_lib.h** file
29 provides derived-type definitions or those routines that require an explicit interface is
30 implementation defined. Whether any of the OpenMP runtime library routines that take an
31 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 19.1](#)).

Fortran

- 1 • **omp_set_num_threads routine:** If the argument is not a positive integer, the behavior is
2 implementation defined (see [Section 19.2.1](#)).
- 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 19.2.9](#)).
- 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 19.2.10](#)).
- 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 19.2.12](#)).
- 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 19.2.13](#)).
- 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 19.3.4](#)).
- 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 19.3.8](#)).
- 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 19.3.9](#)).
- 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 19.3.10](#)).
- 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 19.3.11](#)).
- 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 19.4.3](#)).
- 32 • **omp_set_teams_thread_limit routine:** If the argument is not a positive integer, the
33 behavior is implementation defined (see [Section 19.4.5](#)).
- 34 • **omp_pause_resource_all routine:** The behavior of this routine is implementation
35 defined if the argument kind is not listed in [Section 19.6.1](#) (see [Section 19.6.2](#)).
- 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 19.8.6](#) and [Section 19.8.8](#)).

- 2 • **Lock routines:** If a lock contains a synchronization hint, the effect of the hint is
3 implementation defined (see [Section 19.9](#)).
- 4 • **Interoperability routines:** Implementation-defined properties may use zero and positive
5 values for properties associated with an `omp_interop_t` object (see [Section 19.12](#)).

6 **Chapter 20:**

- 7 • **Tool callbacks:** If a tool attempts to register a callback not listed in [Table 20.2](#), whether the
8 registered callback may never, sometimes or always invoke this callback for the associated
9 events is implementation defined (see [Section 20.2.4](#)).
- 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 20.2.5](#)).
- 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 20.2.5](#)).
- 17 • **Native record abstract type:** The meaning of a *hwid* value for a device is implementation
18 defined (see [Section 20.4.3.3](#)).
- 19 • **ompt_dispatch_chunk_t type:** Whether the chunk of a taskloop is contiguous is
20 implementation defined (see [Section 20.4.4.13](#)).
- 21 • **ompt_record_abstract_t type:** The set of OMPT thread states supported is
22 implementation defined (see [Section 20.4.4.28](#)).
- 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 20.5.2.13](#)).
- 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 20.5.2.25](#)).
- 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 20.6.1.8](#)).
- 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 20.6.1.10](#)).
- 36 • **ompt_get_proc_id_t entry point type:** The meaning of the numerical identifier
37 returned is implementation defined (see [Section 20.6.1.11](#)).

1
2
3
4
5
6
7

Chapter 21:

- **ompd_callback_print_string_fn_t callback type:** The value of *category* is implementation defined (see [Section 21.4.5](#)).
- **ompd_parallel_handle_compare operation:** The means by which parallel region handles are ordered is implementation defined (see [Section 21.5.6.5](#)).
- **ompd_task_handle_compare operation:** The means by which task handles are ordered is implementation defined (see [Section 21.5.7.6](#)).

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:

- The syntax of the **declare reduction** directive that specifies the combiner expression in the directive argument was deprecated.
- The `ompt_target_data_transfer_to_device`, `ompt_target_data_transfer_from_device`, `ompt_target_data_transfer_to_device_async`, and `ompt_target_data_transfer_from_device_async` values in `ompt_target_data_op_t` enum were [deprecated](#) (see [Section 20.4.4.15](#)).

B.2 Version 5.2 to 6.0 Differences

- All features [deprecated](#) in versions 5.2, 5.1 and 5.0 were removed.
- Full support for C23 was added (see [Section 1.7](#)).
- Full support for C++23 was added (see [Section 1.7](#)).
- 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 3](#)).
- The handling of the *nthreads-var* ICV was updated (see [Section 2.4](#)) and the *nthreads* argument of the **num_threads** clause was changed to a list (see [Section 11.2.2](#)) to support context-specific reservation of inner parallelism.
- 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 3.1.5](#)).
- The environment variable **OMP_AVAILABLE_DEVICES** was added and the environment variable **OMP_DEFAULT_DEVICE** was extended to support [device](#) selection by [traits](#) (see [Section 3.2.7](#) and [Section 3.2.8](#)).

- 1 • The environment variable **OMP_THREADS_RESERVE** was added to reserve a number of
2 structured threads and free-agent threads (see [Section 3.2.10](#)).

▼ C++ ▼

- 3 • The **decl** attribute was added to improve the attribute syntax for declarative directives (see
4 [Section 4.1](#)).

▲ C++ ▲

▼ C ▼

- 5 • The OpenMP directive syntax was extended to include C attribute specifiers (see
6 [Section 4.1](#)).

▲ C ▲

- 7 • To improve consistency in clause format, all inarguable clauses were extended to take an
8 optional argument for which the default value yields equivalent semantics to the existing
9 inarguable semantics (see [Section 4.2](#)).

▼ Fortran ▼

- 10 • The definitions of locator list items and assignable OpenMP types were extended to include
11 function references that have data pointer results (see [Section 4.2.1](#)).

▲ Fortran ▲

▼ C / C++ ▼

- 12 • Array section definition was extended to permit, where explicitly allowed, omission of length
13 when the size of the array dimension is not known (see [Section 4.2.5](#)).

▲ C / C++ ▲

- 14 • To support greater specificity on combined and composite constructs, all clauses were
15 extended to accept the *directive-name-modifier*, which identifies the constituent directives to
16 which the clause applies (see [Section 4.4](#)).

▼ Fortran ▼

- 17 • OpenMP atomic structured blocks were extended to allow **BLOCK** constructs (see
18 [Section 5.3.3](#)).
- 19 • *conditional-update-statement* was extended to allow more forms and comparisons (see
20 [Section 5.3.3](#)).

▲ Fortran ▲

- 21 • The concept of [canonical loop sequences](#) and the **looprange** clause were defined (see
22 [Section 5.4.6](#) and [Section 5.4.7](#)).

- 23 • The semantics of the **use_device_ptr** and **use_device_addr** clauses on a
24 **target data** construct were altered to imply a reference count update on entry and exit
25 from the region for the corresponding objects that they reference in the device data
26 environment (see [Section 6.4.8](#) and [Section 6.4.10](#)).

- Support for induction operations was added (see [Section 6.5](#)) through the **induction** clause (see [Section 6.5.12](#)) and the **declare induction** directive (see [Section 6.5.16](#)), which supports user-defined induction operators.

C++

- The circumstances under which implicitly declared reduction identifiers are supported for variables of class type were clarified (see [Section 6.5.3](#) and [Section 6.5.6](#)).

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 6.8.3](#)).
- The **self map-type-modifier** was added to the **map** clause and the **self implicit-behavior** was added to the **defaultmap** clause to explicitly request that the **corresponding list item** refer to the same object as the **original list item** (see [Section 6.8.3](#) and [Section 6.8.6](#)).
- The **map** clause was extended to permit mapping of assumed-size arrays (see [Section 6.8.3](#)).
- The **groupprivate** directive was added to specify that variables should be privatized with respect to a contention group (see [Section 6.12](#)).
- The **local** clause was added to the **declare target** directive to specify that variables should be replicated locally for each device (see [Section 6.13](#)).
- The allocator trait **part_size** was added to specify the size of the **interleaved** allocator partitions (see [Section 7.2](#)).
- 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 7.2](#)).
- 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 7.2](#)).
- The **interop** operation of the **append_args** clause was extended to allow specification of all modifiers of the **init** clause (see [Section 8.5.3](#) and [Section 15.1.2](#)).
- The **dispatch** construct was extended with the **interop** clause to support appending arguments specific to a call site (see [Section 8.6](#) and [Section 8.6.1](#)).
- The **message** and **severity** clauses were added to the **parallel** directive to support customization of any **error termination** associated with the **directive** (see [Section 9.3](#), [Section 9.4](#), and [Section 11.2](#)).
- The **self_maps requirement** clause was added to require that all **mapping operations** are **self maps** (see [Section 9.5.1.6](#)).

- The *assumption* clause group was extended with the **no_omp_constructs** clause to support identification of **regions** in which no **constructs** will be encountered (see [Section 9.6.1](#) and [Section 9.6.1.5](#)).
- The **reverse** construct was added to reverse the iteration order of a loop (see [Section 10.3](#)).
- The **interchange** construct was added to permute the order of loops in a loop nest (see [Section 10.4](#)).
- The **fuse** construct was added to fuse two or more loops in a **canonical loop sequences** (see [Section 10.5](#)).
- The **apply** clause was added to enable more flexible composition of loop-transforming constructs (see [Section 10.6](#)).
- The **omp_curr_progress_width** identifier (see [Section 11.1](#)), **safesync** clause on the **parallel** construct (see [Section 11.2.5](#)) and the **omp_get_max_progress_width** runtime routine (see [Section 19.7.2](#)) 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 11.2.2](#)).
- To support a wider range of synchronization choices, the **atomic construct** was added to the **constructs** that may be encountered inside a **region** that corresponds to a **construct** with an **order** clause that specifies **concurrent** (see [Section 11.4](#)).
- The **coexecute directive** was added to support Fortran array expressions in **teams constructs** (see [Section 12.5](#)).

Fortran

- The **loop construct** was extended to allow **DO CONCURRENT** loops as the associated loops (see [Section 12.8](#)).

Fortran

- The **threadset** clause was added to **task-generating constructs** to specify the **binding thread set** of the generated task (see [Section 13.4](#)).
- The **nowait** clause was added to the clauses that may appear on the **target** construct when the **device** clause is specified with the **ancestor device-modifier** (see [Section 14.8](#)).
- The *do_not_synchronize* argument for the **nowait** clause (see [Section 16.6](#)) and **nogroup** clause (see [Section 16.7](#)) was updated to permit non-constant expressions.
- The **memscope** clause was added to the **atomic** and **flush** constructs to allow the binding thread set to span multiple devices (see [Section 16.8.4](#)).

- 1 • The `omp_is_free_agent` and `omp_ancestor_is_free_agent` routines were
2 added to test whether the `encountering thread`, or the `ancestor thread`, is a `free-agent thread`
3 (see [Section 19.5.4](#) and [Section 19.5.5](#)).
- 4 • The `omp_target_memset` and `omp_target_memset_rect_async` routine were
5 added to fill memory in a `device data environment` of a `device` (see [Section 19.8.9](#) and
6 [Section 19.8.10](#)).
- 7 • New routines were added to obtain `memory spaces` and `memory allocators` to allocate remote
8 and shared memory (see [Section 19.13](#)).
- 9 • The `omp_get_memspace_num_resources` routine was added to be able to query the
10 number of available resources of a memory space (see [Section 19.13.12](#)).
- 11 • The `omp_get_submemspace` routine was added to obtain a memory space with a subset
12 of the original memory space resources (see [Section 19.13.13](#)).
- 13 • The more general values `ompt_target_data_transfer` and
14 **`ompt_target_data_transfer_async`** were added to the
15 **`ompt_target_data_op_t`** enum and supersede the values
16 **`ompt_target_data_transfer_to_device`**,
17 **`ompt_target_data_transfer_from_device`**,
18 **`ompt_target_data_transfer_to_device_async`**, and
19 **`ompt_target_data_transfer_from_device_async`** (see [Section 20.4.4.15](#)).
20 The superseded values were [deprecated](#).
- 21 • The `ompt_get_buffer_limits` runtime entry point was added to the OMPT device
22 tracing interface so that a first party tool can obtain an upper limit on the sizes of the trace
23 buffers that it should make available to the implementation (see [Section 20.5.2.23](#) and
24 [Section 20.6.2.6](#)).

25 B.3 Version 5.1 to 5.2 Differences

- 26 • The *explicit-task-var* ICV has replaced the *implicit-task-var* ICV and has the opposite
27 meaning and semantics (see [Chapter 2](#)). The `omp_in_explicit_task` routine was
28 added to query if a code region is executed from an explicit task region (see [Section 19.5.2](#)).
- 29 • Major reorganization and numerous changes were made to improve the quality of the
30 specification of OpenMP syntax and to increase consistency of restrictions and their wording.
31 These changes frequently result in the possible perception of differences to preceding versions
32 of the OpenMP specification. However, those differences almost always resolve ambiguities,
33 which may nonetheless have implications for existing implementations and programs.
- 34 • For OpenMP directives, reserved the `omp` sentinel (see [Section 4.1](#), [Section 4.1.1](#) and
35 [Section 4.1.2](#)) and, for implementation-defined directives that extend the OpenMP directives
36 reserved the `omp_x` sentinel for C/C++ and free source form Fortran (see [Section 4.1](#) and

1 Section 4.1.2) and the **omx** sentinel for fixed source form Fortran to accommodate character
2 position requirements (see Section 4.1.1). Reserved clause names that begin with the **omp_**
3 prefix for implementation-defined clauses on OpenMP directives (see Section 4.2). Reserved
4 names in the base language that start with the **omp_** and **omp_** prefix and reserved the **omp**
5 and **omp_** namespaces (see Chapter 5) for the OpenMP runtime API and for
6 implementation-defined extensions to that API (see Chapter 19).

- 7 • Allowed any clause that can be specified on a paired **end** directive to be specified on the
8 directive (see Section 4.1), including the **copyprivate** clause (see Section 6.7.2) and the
9 **nowait** clause in Fortran (see Section 16.6).
- 10 • Allowed **if** clause on **teams** construct (see Section 4.5 and Section 11.3).
- 11 • For consistency with the syntax of other definitions of the **clause**, the syntax of the **destroy**
12 **clause** on the **depobj** construct with no argument was **deprecated** (see Section 4.6).
- 13 • For consistency with the syntax of other **clauses**, the syntax of the **linear** clause that
14 specifies its argument and *linear-modifier* as *linear-modifier (list)* was **deprecated** and the
15 *step modifier* was added for specifying the linear step (see Section 6.4.6).
- 16 • The *minus* (–) operator for reductions was **deprecated** (see Section 6.5.6).
- 17 • The syntax of **modifiers** without comma separators in the **map** clause was **deprecated** (see
18 Section 6.8.3).
- 19 • To support the complete range of user-defined mappers and to improve consistency of **map**
20 clause usage, the **declare mapper** directive was extended to accept *iterator-modifier* and
21 the **present** *map-type-modifier* (see Section 6.8.3 and Section 6.8.7).
- 22 • Mapping of a pointer list item was updated such that if a matched candidate is not found in
23 the data environment, firstprivate semantics apply and the pointer retains its original value
24 (see Section 6.8.3).
- 25 • The **enter** clause was added as a synonym for the **to** clause on the declare target **directive**,
26 and the corresponding **to** clause was **deprecated** to reduce parsing ambiguity (see
27 Section 6.8.4 and Section 8.8).

Fortran

- 28 • Metadirectives (see Section 8.4), assumption directives (see Section 9.6), **nothing**
29 directives (see Section 9.7), **error** directives (see Section 9.1) and loop transformation
30 constructs (see Chapter 10) were added to the list of directives that are allowed in a pure
31 procedure (see Chapter 4).
- 32 • The **allocators** construct was added to support the use of OpenMP allocators for
33 **variables** that are allocated by a Fortran **ALLOCATE** statement, and the application of
34 **allocate** directives to an **ALLOCATE** statement was **deprecated** (see Section 7.7).

- 1 • For consistency with other constructs with associated base language code, the **dispatch**
2 construct was extended to allow an optional paired **end** directive to be specified (see
3 [Section 8.6](#)).

Fortran

- 4 • To support the full range of allocators and to improve consistency with the syntax of other
5 **clauses**, the argument that specified the arguments of the **uses_allocators** clause as a
6 comma-separated list in which each list item is a *clause-argument-specification* of the form
7 `allocator[(traits)]` was **deprecated** (see [Section 7.8](#)).
- 8 • To improve code clarity and to reduce ambiguity in this specification, the **otherwise**
9 **clause** was added as a synonym for the **default** clause on **metadirectives** and the
10 corresponding **default** clause syntax was **deprecated** (see [Section 8.4.2](#)).

C / C++

- 11 • To improve overall syntax consistency and to reduce redundancy, the delimited form of the
12 **declare target** directive was **deprecated** (see [Section 8.8.2](#)).

C / C++

- 13 • The behavior of the **order** clause with the **concurrent** parameter was changed so that it
14 only affects whether a loop schedule is reproducible if a modifier is explicitly specified (see
15 [Section 11.4](#)).
- 16 • Support for the **allocate** and **firstprivate** clauses on the **scope** directive was
17 added (see [Section 12.2](#)).
- 18 • The **ompt_callback_work** callback work types for worksharing loop were added (see
19 [Section 12.6](#)).
- 20 • To simplify usage, the **map** clause on a **target enter data** or **target exit data**
21 construct now has a default map type that provides the same behavior as the **to** or **from** map
22 types, respectively (see [Section 14.6](#) and [Section 14.7](#)).
- 23 • The **interop** construct was updated to allow the **init** clause to accept an *interop_type* in
24 any position of the modifier list (see [Section 15.1](#)).
- 25 • The **doacross** clause was added as a synonym for the **depend** clause with the keywords
26 **source** and **sink** as *dependence-type modifiers* and the corresponding **depend** clause
27 syntax was **deprecated** to improve code clarity and to reduce parsing ambiguity. Also, the
28 **omp_cur_iteration** keyword was added to represent an iteration vector that refers to
29 the current **logical iteration** (see [Section 16.9.6](#)).

B.4 Version 5.0 to 5.1 Differences

- Full support of C11, C++11, C++14, C++17, C++20 and Fortran 2008 was completed (see [Section 1.7](#)).
- Various changes throughout the specification were made to provide initial support of Fortran 2018 (see [Section 1.7](#)).
- To support device-specific ICV settings the environment variable syntax was extended to support device-specific variables (see [Section 2.2](#) and [Chapter 3](#)).
- The OpenMP directive syntax was extended to include C++ attribute specifiers (see [Section 4.1](#)).
- The **omp_all_memory** reserved locator was added (see [Section 4.1](#)), and the **depend** clause was extended to allow its use (see [Section 16.9.5](#)).
- Support for **private** and **firstprivate** as an argument to the **default** clause in C and C++ was added (see [Section 6.4.1](#)).
- Support was added so that iterators may be defined and used in a **map** clause (see [Section 6.8.3](#)) or in data-motion clause on a **target update** directive (see [Section 14.9](#)).
- The **present** argument was added to the **defaultmap** clause (see [Section 6.8.6](#)).
- Support for the **align** clause on the **allocate** directive and **allocator** and **align** modifiers on the **allocate** clause was added (see [Chapter 7](#)).
- The *target_device* trait set was added to the OpenMP context (see [Section 8.1](#)), and the **target_device** selector set was added to context selectors (see [Section 8.2](#)).
- For C/C++, the declare variant **directive** was extended to support elision of [preprocessed code](#) and to allow enclosed function definitions to be interpreted as variant functions (see [Section 8.5](#)).
- The **declare variant** directive was extended with new clauses (**adjust_args** and **append_args**) that support adjustment of the interface between the original function and its variants (see [Section 8.5](#)).
- The **dispatch** construct was added to allow users to control when variant substitution happens and to define additional information that can be passed as arguments to the function variants (see [Section 8.6](#)).
- Support was added for indirect calls to the device version of a [procedure](#) in **target** regions (see [Section 8.8](#)).
- Assumption directives were added to allow users to specify invariants (see [Section 9.6](#)).
- To support clarity in metadirectives, the **nothing** directive was added (see [Section 9.7](#)).

- 1 • To allow users to control the compilation process and runtime error actions, the **error**
2 directive was added (see [Section 9.1](#)).
- 3 • Loop transformation constructs were added (see [Chapter 8](#)).
- 4 • The **masked construct** was added to support restricting execution to a specific **thread** to
5 replace the **deprecated master construct** (see [Section 11.6](#)).
- 6 • The **scope** directive was added to support reductions without requiring a **parallel** or
7 worksharing region (see [Section 12.2](#)).
- 8 • The **grainsize** and **num_tasks** clauses for the **taskloop** construct were extended
9 with a **strict** modifier to ensure a deterministic distribution of logical iterations to tasks
10 (see [Section 13.7](#)).
- 11 • The **thread_limit** clause was added to the **target** construct to control the upper bound
12 on the number of threads in the created contention group (see [Section 14.8](#)).
- 13 • The **has_device_addr** clause was added to the **target** construct to allow access to
14 variables or array sections that already have a device address (see [Section 14.8](#)).
- 15 • The **interop directive** was added to enable portable interoperability with **foreign execution**
16 **contexts** used to implement OpenMP (see [Section 15.1](#)). Runtime routines that facilitate use
17 of **omp_interop_t** objects were also added (see [Section 19.12](#)).
- 18 • The **nowait** clause was added to the **taskwait** directive to support insertion of
19 non-blocking join operations in a task dependence graph (see [Section 16.5](#)).
- 20 • Support was added for compare-and-swap and (for C and C++) minimum and maximum
21 atomic operations through the **compare** clause. Support was also added for the specification
22 of the memory order to apply to a failed comparing atomic operation with the **fail** clause
23 (see [Section 16.8.5](#)).
- 24 • Specification of the **seq_cst** clause on a **flush** construct was allowed, with the same
25 meaning as a **flush** construct without a list and without a clause (see [Section 16.8.6](#)).
- 26 • To support inout sets, the **inoutset** argument was added to the **depend** clause (see
27 [Section 16.9.5](#)).
- 28 • The **omp_set_num_teams** and **omp_set_teams_thread_limit** runtime routines
29 were added to control the number of teams and the size of those teams on the **teams**
30 construct (see [Section 19.4.3](#) and [Section 19.4.5](#)). Additionally, the **omp_get_max_teams**
31 and **omp_get_teams_thread_limit** runtime routines were added to retrieve the
32 values that will be used in the next **teams** construct (see [Section 19.4.4](#) and [Section 19.4.6](#)).
- 33 • The **omp_target_is_accessible** runtime routine was added to test whether host
34 memory is accessible from a given device (see [Section 19.8.4](#)).
- 35 • To support asynchronous device memory management, **omp_target_memcpy_async**
36 and **omp_target_memcpy_rect_async** runtime routines were added (see

1 [Section 19.8.7](#) and [Section 19.8.8](#)).

- 2 • The `omp_get_mapped_ptr` runtime routine was added to support obtaining the device
3 pointer that is associated with a host pointer for a given device (see [Section 19.8.13](#)).
- 4 • The `omp_calloc`, `omp_realloc`, `omp_aligned_alloc` and
5 `omp_aligned_calloc` API routines were added (see [Section 19.13](#)).
- 6 • For the `omp_alloctrail_key_t` enum, the `omp_atv_serialized` value was added
7 and the `omp_atv_default` value was changed (see [Section 19.13.1](#)).
- 8 • The `omp_display_env` runtime routine was added to provide information about ICVs
9 and settings of environment variables (see [Section 19.15](#)).
- 10 • The `ompt_scope_beginend` value was added to the `ompt_scope_endpoint_t`
11 enum to indicate the coincident beginning and end of a scope (see [Section 20.4.4.11](#)).
- 12 • The `ompt_sync_region_barrier_implicit_workshare`,
13 `ompt_sync_region_barrier_implicit_parallel`, and
14 `ompt_sync_region_barrier_teams` values were added to the
15 `ompt_sync_region_t` enum (see [Section 20.4.4.14](#)).
- 16 • Values for asynchronous data transfers were added to the `ompt_target_data_op_t`
17 enum (see [Section 20.4.4.15](#)).
- 18 • The `ompt_state_wait_barrier_implementation` and
19 `ompt_state_wait_barrier_teams` values were added to the `ompt_state_t`
20 enum (see [Section 20.4.4.28](#)).
- 21 • The `ompt_callback_target_data_op_emi_t`,
22 `ompt_callback_target_emi_t`, `ompt_callback_target_map_emi_t`, and
23 `ompt_callback_target_submit_emi_t` callbacks were added to support external
24 monitoring interfaces (see [Section 20.5.2.25](#), [Section 20.5.2.26](#), [Section 20.5.2.27](#) and
25 [Section 20.5.2.28](#)).
- 26 • The `ompt_callback_error_t` type was added (see [Section 20.5.2.30](#)).
- 27 • The `OMP_PLACES` syntax was extended (see [Section 3.1.5](#)).
- 28 • The `OMP_NUM_TEAMS` and `OMP_TEAMS_THREAD_LIMIT` environment variables were
29 added to control the number and size of teams on the `teams` construct (see [Section 3.6.1](#) and
30 [Section 3.6.2](#)).

31 **B.5 Version 4.5 to 5.0 Differences**

- 32 • The memory model was extended to distinguish different types of flush operations according
33 to specified flush properties (see [Section 1.4.4](#)) and to define a happens before order based on
34 synchronizing flush operations (see [Section 1.4.5](#)).

- 1 • Various changes throughout the specification were made to provide initial support of C11,
2 C++11, C++14, C++17 and Fortran 2008 (see [Section 1.7](#)).
- 3 • Full support of Fortran 2003 was completed (see [Section 1.7](#)).
- 4 • The *target-offload-var* internal control variable (see [Chapter 2](#)) and the
5 **OMP_TARGET_OFFLOAD** environment variable (see [Section 3.2.9](#)) were added to support
6 runtime control of the execution of device constructs.
- 7 • Control over whether nested parallelism is enabled or disabled was integrated into the
8 *max-active-levels-var* internal control variable (see [Section 2.2](#)), the default value of which is
9 now implementation defined, unless determined according to the values of the
10 **OMP_NUM_THREADS** (see [Section 3.1.2](#)) or **OMP_PROC_BIND** (see [Section 3.1.6](#))
11 environment variables.
- 12 • Support for array shaping (see [Section 4.2.4](#)) and for array sections with non-unit strides in C
13 and C++ (see [Section 4.2.5](#)) was added to facilitate specification of discontinuous storage,
14 and the **target update** construct (see [Section 14.9](#)) and the **depend** clause (see
15 [Section 16.9.5](#)) were extended to allow the use of shape-operators (see [Section 4.2.4](#)).
- 16 • Iterators (see [Section 4.2.6](#)) were added to support expressions in a list that expand to
17 multiple expressions.
- 18 • The canonical loop form was defined for Fortran and, for all base languages, extended to
19 permit non-rectangular loop nests (see [Section 5.4.1](#)).
- 20 • The *relational-op* in the *canonical loop form* for C/C++ was extended to include **!=** (see
21 [Section 5.4.1](#)).
- 22 • To support conditional assignment to lastprivate variables, the **conditional** modifier was
23 added to the **lastprivate** clause (see [Section 6.4.5](#)).
- 24 • The **inscan** modifier for the **reduction** clause (see [Section 6.5.9](#)) and the **scan**
25 directive (see [Section 6.6](#)) were added to support inclusive and exclusive scan computations.
- 26 • To support task reductions, the **task** modifier was added to the **reduction** clause (see
27 [Section 6.5.9](#)), the **task_reduction** clause (see [Section 6.5.10](#)) was added to the
28 **taskgroup** construct (see [Section 16.4](#)), and the **in_reduction** clause (see
29 [Section 6.5.11](#)) was added to the **task** (see [Section 13.6](#)) and **target** (see [Section 14.8](#))
30 constructs.
- 31 • To support taskloop reductions, the **reduction** (see [Section 6.5.9](#)) and **in_reduction**
32 (see [Section 6.5.11](#)) clauses were added to the **taskloop** construct (see [Section 13.7](#)).
- 33 • The description of the **map** clause was modified to clarify the mapping order when multiple
34 *map-types* are specified for a variable or structure members of a variable on the same
35 construct. The *close map-type-modifier* was added as a hint for the runtime to allocate
36 memory close to the target device (see [Section 6.8.3](#)).

- 1 • The capability to map C/C++ pointer variables and to assign the address of device memory
2 that is mapped by an array section to them was added. Support for mapping of Fortran
3 pointer and allocatable variables, including pointer and allocatable components of variables,
4 was added (see [Section 6.8.3](#)).
- 5 • The **defaultmap** clause (see [Section 6.8.6](#)) was extended to allow selecting the
6 data-mapping or data-sharing attributes for any of the scalar, aggregate, pointer, or
7 allocatable classes on a per-region basis. Additionally it accepts the **none** parameter to
8 support the requirement that all variables referenced in the construct must be explicitly
9 mapped or privatized.
- 10 • The **declare mapper** directive was added to support mapping of data types with direct
11 and indirect members (see [Section 6.8.7](#)).
- 12 • Predefined memory spaces (see [Section 7.1](#)), predefined memory allocators and allocator
13 traits (see [Section 7.2](#)) and directives, clauses and API routines (see [Chapter 7](#) and
14 [Section 19.13](#)) to use them were added to support different kinds of memories.
- 15 • Metadirectives (see [Section 8.4](#)) and declare variant directives (see [Section 8.5](#)) were added
16 to support selection of directive variants and declared function variants at a call site,
17 respectively, based on compile-time traits of the enclosing context.
- 18 • Support for nested **declare target** directives was added (see [Section 8.8](#)).
- 19 • The **requires** directive (see [Section 9.5](#)) was added to support applications that require
20 implementation-specific features.
- 21 • The **teams** construct (see [Section 11.3](#)) was extended to support execution on the host
22 device without an enclosing **target** construct (see [Section 14.8](#)).
- 23 • The **loop** construct and the **order (concurrent)** clause were added to support
24 compiler optimization and parallelization of loops for which iterations may execute in any
25 order, including concurrently (see [Section 11.4](#) and [Section 12.8](#)).
- 26 • The collapse of associated loops that are imperfectly nested loops was defined for the **simd**
27 (see [Section 11.5](#)), worksharing-loop (see [Section 12.6](#)), **distribute** (see [Section 12.7](#))
28 and **taskloop** (see [Section 13.7](#)) constructs.
- 29 • The **simd** construct (see [Section 11.5](#)) was extended to accept the **if**, **nontemporal**, and
30 **order (concurrent)** clauses and to allow the use of **atomic** constructs within it.
- 31 • The default loop schedule modifier for worksharing-loop constructs without the **static**
32 schedule and the **ordered** clause was changed to **nonmonotonic** (see [Section 12.6](#)).
- 33 • The **affinity** clause was added to the **task** construct (see [Section 13.6](#)) to support hints
34 that indicate data affinity of explicit tasks.
- 35 • The **detach** clause for the **task** construct (see [Section 13.6](#)) and the
36 **omp_fulfill_event** runtime routine (see [Section 19.11.1](#)) were added to support
37 execution of detachable tasks.

- 1 • The **taskloop** construct (see [Section 13.7](#)) was added to the list of constructs that can be
2 canceled by the **cancel** construct (see [Section 17.2](#)).
- 3 • To support mutually exclusive inout sets, a **mutexinoutset** *dependence-type* was added
4 to the **depend** clause (see [Section 13.10](#) and [Section 16.9.5](#)).
- 5 • The semantics of the **use_device_ptr** clause for pointer variables was clarified and the
6 **use_device_addr** clause for using the device address of non-pointer variables inside the
7 **target data** construct was added (see [Section 14.5](#)).
- 8 • To support reverse offload, the **ancestor** modifier was added to the **device** clause for the
9 **target** construct (see [Section 14.8](#)).
- 10 • To reduce programmer effort, implicit declare target directives for some functions (C, C++,
11 Fortran) and subroutines (Fortran) were added (see [Section 14.8](#) and [Section 8.8](#)).
- 12 • The **target update** construct (see [Section 14.9](#)) was modified to allow array sections that
13 specify discontinuous storage.
- 14 • The **to** and **from** clauses on the **target update** construct (see [Section 14.9](#)), the
15 **depend** clause on task generating constructs (see [Section 16.9.5](#)), and the **map** clause (see
16 [Section 6.8.3](#)) were extended to allow any lvalue expression as a list item for C/C++.
- 17 • Lock hints were renamed to synchronization hints, and the old names were **deprecated** (see
18 [Section 16.1](#)).
- 19 • The **depend** clause was added to the **taskwait** construct (see [Section 16.5](#)).
- 20 • To support acquire and release semantics with weak memory ordering, the **acq_rel**,
21 **acquire**, and **release** clauses were added to the **atomic** construct (see [Section 16.8.5](#))
22 and **flush** construct (see [Section 16.8.6](#)), and the memory ordering semantics of implicit
23 flushes on various constructs and runtime routines were clarified (see [Section 16.8.7](#)).
- 24 • The **atomic** construct was extended with the **hint** clause (see [Section 16.8.5](#)).
- 25 • The **depend** clause (see [Section 16.9.5](#)) was extended to support iterators and to support
26 depend objects that can be created with the new **depobj** construct.
- 27 • New combined constructs **master taskloop**, **parallel master**,
28 **parallel master taskloop**, **master taskloop simd**
29 **parallel master taskloop simd** (see [Section 18.3](#)) were added.
- 30 • The **omp_set_nested** and **omp_get_nested** routines and the **OMP_NESTED**
31 environment variable were **deprecated**.
- 32 • The **omp_get_supported_active_levels** routine was added to query the number of
33 active levels of parallelism supported by the implementation (see [Section 19.2.12](#)).
- 34 • Runtime routines **omp_set_affinity_format** (see [Section 19.3.8](#)),
35 **omp_get_affinity_format** (see [Section 19.3.9](#)), **omp_set_affinity** (see
36 [Section 19.3.10](#)), and **omp_capture_affinity** (see [Section 19.3.11](#)) and environment

1 variables `OMP_DISPLAY_AFFINITY` (see [Section 3.2.4](#)) and `OMP_AFFINITY_FORMAT`
2 (see [Section 3.2.5](#)) were added to provide OpenMP runtime thread affinity information.

- 3 • The `omp_pause_resource` and `omp_pause_resource_all` runtime routines were
4 added to allow the runtime to relinquish resources used by OpenMP (see [Section 19.6.1](#) and
5 [Section 19.6.2](#)).
- 6 • The `omp_get_device_num` runtime routine (see [Section 19.7.6](#)) was added to support
7 determination of the device on which a thread is executing.
- 8 • Support for a first-party tool interface (see [Chapter 20](#)) was added.
- 9 • Support for a third-party tool interface (see [Chapter 21](#)) was added.
- 10 • Support for controlling offloading behavior with the `OMP_TARGET_OFFLOAD` environment
11 variable was added (see [Section 3.2.9](#)).
- 12 • Stubs for Runtime Library Routines (previously Appendix A) were moved to a separate
13 document.
- 14 • Interface Declarations (previously Appendix B) were moved to a separate document.

15 B.6 Version 4.0 to 4.5 Differences

- 16 • Support for several features of Fortran 2003 was added (see [Section 1.7](#)).
- 17 • The `if` clause was extended to take a *directive-name-modifier* that allows it to apply to
18 combined constructs (see [Section 4.5](#)).
- 19 • The implicit data-sharing attribute for scalar variables in `target` regions was changed to
20 `firstprivate` (see [Section 6.1.1](#)).
- 21 • Use of some C++ reference types was allowed in some data sharing attribute clauses (see
22 [Section 6.4](#)).
- 23 • The `ref`, `val`, and `uval` modifiers were added to the `linear` clause (see [Section 6.4.6](#)).
- 24 • Semantics for reductions on C/C++ array sections were added and restrictions on the use of
25 arrays and pointers in reductions were removed (see [Section 6.5.9](#)).
- 26 • Support was added to the `map` clauses to handle structure elements (see [Section 6.8.3](#)).
- 27 • To support unstructured data mapping for devices, the `map` clause (see [Section 6.8.3](#)) was
28 updated and the `target enter data` (see [Section 14.6](#)) and `target exit data` (see
29 [Section 14.7](#)) constructs were added.
- 30 • The `declare target` directive was extended to allow mapping of global variables to be
31 deferred to specific device executions and to allow an *extended-list* to be specified in C/C++
32 (see [Section 8.8](#)).

- 1 • The **simdlen** clause was added to the **simd** construct (see [Section 11.5](#)) to support
2 specification of the exact number of iterations desired per **SIMD chunk**.
- 3 • A parameter was added to the **ordered** clause of the worksharing-loop construct (see
4 [Section 12.6](#)) and **clauses** were added to the **ordered** construct (see [Section 16.10](#)) to
5 support **doacross loop nests** and use of the **simd** construct on loops with loop-carried
6 backward dependences.
- 7 • The **linear** clause was added to the worksharing-loop construct (see [Section 12.6](#)).
- 8 • The **priority** clause was added to the **task** construct (see [Section 13.6](#)) to support hints
9 that specify the relative execution priority of explicit tasks. The
10 **omp_get_max_task_priority** routine was added to return the maximum supported
11 priority value (see [Section 19.5.1](#)) and the **OMP_MAX_TASK_PRIORITY** environment
12 variable was added to control the maximum priority value allowed (see [Section 3.2.11](#)).
- 13 • The **taskloop** construct (see [Section 13.7](#)) was added to support nestable parallel loops
14 that create OpenMP tasks.
- 15 • To support interaction with native device implementations, the **use_device_ptr** clause
16 was added to the **target data** construct (see [Section 14.5](#)) and the **is_device_ptr**
17 clause was added to the **target** construct (see [Section 14.8](#)).
- 18 • The **nowait** and **depend** clauses were added to the **target** construct (see [Section 14.8](#))
19 to improve support for asynchronous execution of **target** regions.
- 20 • The **private**, **firstprivate** and **defaultmap** clauses were added to the **target**
21 construct (see [Section 14.8](#)).
- 22 • The **hint** clause was added to the **critical** construct (see [Section 16.2](#)).
- 23 • The **source** and **sink** dependence types were added to the **depend** clause (see
24 [Section 16.9.5](#)) to support **doacross loop nests**.
- 25 • To support a more complete set of device construct shortcuts, the **target parallel**,
26 **target parallel worksharing-loop**, **target parallel worksharing-loop SIMD**, and
27 **target simd** (see [Section 18.3](#)) combined constructs were added.
- 28 • Query functions for OpenMP thread affinity were added (see [Section 19.3.2](#) to
29 [Section 19.3.7](#)).
- 30 • Device memory routines were added to allow explicit allocation, deallocation, memory
31 transfers, and memory associations (see [Section 19.8](#)).
- 32 • The lock API was extended with lock routines that support storing a hint with a lock to select
33 a desired lock implementation for a lock's intended usage by the application code (see
34 [Section 19.9.2](#)).
- 35 • C/C++ Grammar (previously Appendix B) was moved to a separate document.

B.7 Version 3.1 to 4.0 Differences

- Various changes throughout the specification were made to provide initial support of Fortran 2003 (see [Section 1.7](#)).
- C/C++ array syntax was extended to support array sections (see [Section 4.2.5](#)).
- The **reduction** clause (see [Section 6.5.9](#)) was extended and the **declare reduction** construct (see [Section 6.5.13](#)) was added to support user defined reductions.
- The **proc_bind** clause (see [Section 11.2.3](#)), the **OMP_PLACES** environment variable (see [Section 3.1.5](#)), and the **omp_get_proc_bind** runtime routine (see [Section 19.3.1](#)) were added to support thread affinity policies.
- SIMD directives were added to support SIMD parallelism (see [Section 11.5](#)).
- Implementation defined task scheduling points for untied tasks were removed (see [Section 13.10](#)).
- Device directives (see [Chapter 14](#)), the **OMP_DEFAULT_DEVICE** environment variable (see [Section 3.2.8](#)), and the **omp_set_default_device**, **omp_get_default_device**, **omp_get_num_devices**, **omp_get_num_teams**, **omp_get_team_num**, and **omp_is_initial_device** routines were added to support execution on devices.
- The **taskgroup** construct (see [Section 16.4](#)) was added to support deep task synchronization.
- The **atomic** construct (see [Section 16.8.5](#)) was extended to support atomic swap with the **capture** clause, to allow new atomic update and capture forms, and to support sequentially consistent atomic operations with a new **seq_cst** clause.
- The **depend** clause (see [Section 16.9.5](#)) was added to support task dependences.
- The **cancel** construct (see [Section 17.2](#)), the **cancellation point** construct (see [Section 17.3](#)), the **omp_get_cancellation** runtime routine (see [Section 19.2.8](#)), and the **OMP_CANCELLATION** environment variable (see [Section 3.2.6](#)) were added to support the concept of cancellation.
- The **OMP_DISPLAY_ENV** environment variable (see [Section 3.7](#)) was added to display the value of ICVs associated with the OpenMP environment variables.
- Examples (previously Appendix A) were moved to a separate document.

B.8 Version 3.0 to 3.1 Differences

- The *bind-var* ICV (see [Section 2.1](#)) and the **OMP_PROC_BIND** environment variable (see [Section 3.1.6](#)) were added to support control of whether threads are bound to processors.

- 1 • Data environment restrictions were changed to allow **intent (in)** and **const**-qualified
2 types for the **firstprivate** clause (see [Section 6.4.4](#)).
- 3 • Data environment restrictions were changed to allow Fortran pointers in **firstprivate**
4 (see [Section 6.4.4](#)) and **lastprivate** (see [Section 6.4.5](#)) clauses.
- 5 • New reduction operators **min** and **max** were added for C and C++ (see [Section 6.5](#)).
- 6 • The *nthreads-var* ICV was modified to be a list of the number of threads to use at each nested
7 parallel region level, and the algorithm for determining the number of threads used in a
8 parallel region was modified to handle a list (see [Section 11.2.1](#)).
- 9 • The **final** and **mergeable** clauses (see [Section 13.6](#)) were added to the **task** construct
10 to support optimization of task data environments.
- 11 • The **taskyield** construct (see [Section 13.8](#)) was added to allow user-defined task
12 scheduling points.
- 13 • The **atomic** construct (see [Section 16.8.5](#)) was extended to include **read**, **write**, and
14 **capture** forms, and an **update** clause was added to apply the already existing form of the
15 **atomic** construct.
- 16 • The nesting restrictions in [Section 18.1](#) were clarified to disallow closely-nested OpenMP
17 regions within an **atomic** region so that an **atomic** region can be consistently defined with
18 other OpenMP regions to include all code in the **atomic** construct.
- 19 • The **omp_in_final** runtime library routine (see [Section 19.5.3](#)) was added to support
20 specialization of final task regions.
- 21 • Descriptions of examples (previously Appendix A) were expanded and clarified.
- 22 • Incorrect use of **omp_integer_kind** in Fortran interfaces was replaced with
23 **selected_int_kind(8)**.

24 **B.9 Version 2.5 to 3.0 Differences**

- 25 • The definition of active **parallel** region was changed so that a **parallel** region is
26 active if it is executed by a team that consists of more than one thread (see [Section 1.2](#)).
- 27 • The concept of tasks was added to the execution model (see [Section 1.2](#) and [Section 1.3](#)).
- 28 • The OpenMP memory model was extended to cover atomicity of memory accesses (see
29 [Section 1.4.1](#)). The description of the behavior of **volatile** in terms of **flush** was
30 removed.
- 31 • The definition of the *nest-var*, *dyn-var*, *nthreads-var* and *run-sched-var* internal control
32 variables (ICVs) were modified to provide one copy of these ICVs per task instead of one
33 copy for the whole program (see [Chapter 2](#)). The **omp_set_num_threads** and

- 1 **omp_set_dynamic** runtime library routines were specified to support their use from
2 inside a **parallel** region (see [Section 19.2.1](#) and [Section 19.2.6](#)).
- 3 • The *thread-limit-var* ICV, the **omp_get_thread_limit** runtime library routine and the
4 **OMP_THREAD_LIMIT** environment variable were added to support control of the maximum
5 number of threads (see [Section 2.1](#), [Section 19.2.11](#) and [Section 3.1.3](#)).
 - 6 • The *max-active-levels-var* ICV, **omp_set_max_active_levels** and
7 **omp_get_max_active_levels** runtime library routines, and
8 **OMP_MAX_ACTIVE_LEVELS** environment variable were added to support control of the
9 number of nested active **parallel** regions (see [Section 2.1](#), [Section 19.2.13](#),
10 [Section 19.2.14](#) and [Section 3.1.4](#)).
 - 11 • The *stacksize-var* ICV and the **OMP_STACKSIZE** environment variable were added to
12 support control of thread stack sizes (see [Section 2.1](#) and [Section 3.2.2](#)).
 - 13 • The *wait-policy-var* ICV and the **OMP_WAIT_POLICY** environment variable were added to
14 control the desired behavior of waiting threads (see [Section 2.1](#) and [Section 3.2.3](#)).
 - 15 • Predetermined data-sharing attributes were defined for Fortran assumed-size arrays (see
16 [Section 6.1.1](#)).
 - 17 • Static class members variables were allowed in **threadprivate** directives (see
18 [Section 6.2](#)).
 - 19 • Invocations of constructors and destructors for private and threadprivate class type variables
20 was clarified (see [Section 6.2](#), [Section 6.4.3](#), [Section 6.4.4](#), [Section 6.7.1](#) and [Section 6.7.2](#)).
 - 21 • The use of Fortran allocatable arrays was allowed in **private**, **firstprivate**,
22 **lastprivate**, **reduction**, **copyin** and **copyprivate** clauses (see [Section 6.2](#),
23 [Section 6.4.3](#), [Section 6.4.4](#), [Section 6.4.5](#), [Section 6.5.9](#), [Section 6.7.1](#) and [Section 6.7.2](#)).
 - 24 • Support for **firstprivate** was added to the **default** clause in Fortran (see
25 [Section 6.4.1](#)).
 - 26 • Implementations were precluded from using the storage of the original list item to hold the
27 new list item on the primary thread for list items in the **private** clause, and the value was
28 made well defined on exit from the **parallel** region if no attempt is made to reference the
29 original list item inside the **parallel** region (see [Section 6.4.3](#)).
 - 30 • Data environment restrictions were changed to allow **intent(in)** and **const**-qualified
31 types for the **firstprivate** clause (see [Section 6.4.4](#)).
 - 32 • Data environment restrictions were changed to allow Fortran pointers in **firstprivate**
33 (see [Section 6.4.4](#)) and **lastprivate** (see [Section 6.4.5](#)).
 - 34 • Determination of the number of threads in **parallel** regions was updated (see
35 [Section 11.2.1](#)).

- 1 • The assignment of iterations to threads in a loop construct with a **static** schedule kind was
2 made deterministic (see [Section 12.6](#)).
- 3 • The worksharing-loop construct was extended to support association with more than one
4 perfectly nested loop through the **collapse** clause (see [Section 12.6](#)).
- 5 • Iteration variables for worksharing-loops were allowed to be random access iterators or of
6 unsigned integer type (see [Section 12.6](#)).
- 7 • The schedule kind **auto** was added to allow the implementation to choose any possible
8 mapping of iterations in a loop construct to threads in the team (see [Section 12.6](#)).
- 9 • The **task** construct (see [Chapter 13](#)) was added to support explicit tasks.
- 10 • The **taskwait** construct (see [Section 16.5](#)) was added to support task synchronization.
- 11 • The runtime library routines **omp_set_schedule** and **omp_get_schedule** were
12 added to set and to retrieve the value of the *run-sched-var* ICV (see [Section 19.2.9](#) and
13 [Section 19.2.10](#)).
- 14 • The **omp_get_level** runtime library routine was added to return the number of nested
15 **parallel** regions that enclose the task that contains the call (see [Section 19.2.15](#)).
- 16 • The **omp_get_ancestor_thread_num** runtime library routine was added to return the
17 thread number of the ancestor of the current thread (see [Section 19.2.16](#)).
- 18 • The **omp_get_team_size** runtime library routine was added to return the size of the
19 thread team to which the ancestor of the current thread belongs (see [Section 19.2.17](#)).
- 20 • The **omp_get_active_level** runtime library routine was added to return the number of
21 active **parallel** regions that enclose the task that contains the call (see [Section 19.2.18](#)).
- 22 • Lock ownership was defined in terms of tasks instead of threads (see [Section 19.9](#)).

Index

Symbols

`_OPENMP` macro, 38, 39, 49, 73

A

absent, 253
acq_rel, 366
acquire, 366
acquire flush, 11
adjust_args, 222
affinity, 276
affinity, 319
align, 201
aligned, 192
allocate, 203, 205
allocator, 202
allocators, 207
append_args, 223
apply Clause, 268
array sections, 68
array shaping, 67
assumes, 258, 259
assumption clauses, 253
assumption directives, 252
at, 243
atomic, 375
atomic, 369
atomic construct, 699
atomic_default_mem_order, 247
attribute clauses, 122
attributes, data-mapping, 170
attributes, data-sharing, 109
auto, 305

B

barrier, 356
barrier, implicit, 357

base language format, 83
begin declare target, 239
begin declare variant, 226
begin metadirective, 219
begin assumes, 259
bind, 311
branch, 233

C

cancel, 400
cancel-directive-name, 399
cancellation constructs, 399
 cancel, 400
 cancellation point, 404
cancellation point, 404
canonical loop nest form, 95
canonical loop sequence form, 106
capture, 372
capture, atomic, 375
clause format, 60
clauses
 absent, 253
 acq_rel, 366
 acquire, 366
 adjust_args, 222
 affinity, 319
 align, 201
 aligned, 192
 allocate, 205
 allocator, 202
 append_args, 223
 apply Clause, 268
 assumption, 253
 at, 243
 atomic, 369

atomic_default_mem_order,
 247
 attribute data-sharing, 122
bind, 311
branch, 233
cancel-directive-name, 399
capture, 372
collapse, 103
collector, 162
combiner, 159
compare, 372
contains, 254
copyin, 166
copyprivate, 168
 data copying, 166
 data-sharing, 122
default, 122
defaultmap, 183
depend, 388
destroy, 81
detach, 320
device, 331
device_type, 330
dist_schedule, 308
doacross, 391
dynamic_allocators, 247
enter, 182
exclusive, 166
extended-atomic, 371
fail, 373
filter, 289
final, 314
firstprivate, 125
from, 191
full, 264
grainsize, 324
has_device_addr, 136
hint, 351, 353
holds, 255
if Clause, 80
in_reduction, 153
inbranch, 233
inclusive, 165
indirect, 240
induction, 155
inductor, 162
init, 348
initializer, 159
interop, 229
is_device_ptr, 134
lastprivate, 128
linear, 131
link, 183
local, 196
map, 173
match, 221
memory-order, 365
memscope, 374
mergeable, 314
message, 244
no_openmp, 255
no_openmp_constructs, 256
no_openmp_routines, 257
no_parallelism, 257
nocontext, 230
nogroup, 364
nontemporal, 286
notinbranch, 234
novariants, 230
nowait, 363
num_tasks, 325
num_teams, 282
num_threads, 275
order, 283
ordered, 104
otherwise, 218
parallelization-level, 397
partial, 265
permutation, 267
priority, 316
private, 124
proc_bind, 278
read, 369
reduction, 150
relaxed, 367
release, 368

- requirement*, 246
- reverse_offload**, 248
- safelen**, 287
- safesync**, 279
- schedule**, 304
- self_maps**, 252
- seq_cst**, 368
- severity**, 244
- shared**, 123
- simd**, 398
- simdlen**, 287
- sizes**, 263
- looprange*, 107
- task_reduction**, 153
- thread_limit**, 332
- threads**, 397
- threadset**, 315
- to**, 190
- unified_address**, 249
- unified_shared_memory**, 250
- uniform**, 192
- untied**, 313
- update**, 370, 386
- use**, 349
- use_device_addr**, 137
- use_device_ptr**, 135
- uses_allocators**, 207
- weak**, 374
- when**, 217
- write**, 370
- coexecute**, 298
- collapse**, 103
- combined and composite directive
 - names, 409
- combined construct semantics, 410
- compare**, 372
- compare**, **atomic**, 375
- compilation sentinels, 74, 75
- compliance, 15
- composite constructs, 411
- composition of constructs, 405
- conditional compilation, 73
- consistent loop schedules, 105
- construct syntax, 51
- constructs
 - allocators**, 207
 - atomic**, 375
 - barrier**, 356
 - cancel**, 400
 - cancellation constructs, 399
 - cancellation point**, 404
 - coexecute**, 298
 - combined constructs, 410
 - composite constructs, 411
 - critical**, 354
 - depobj**, 387
 - device constructs, 330
 - dispatch**, 228
 - distribute**, 306
 - do**, 303
 - flush**, 379
 - for**, 302
 - fuse**, 267
 - interop**, 346
 - loop**, 309
 - masked**, 288
 - ordered**, 393–395
 - parallel**, 270
 - reverse**, 265
 - scope**, 292
 - sections**, 293
 - simd**, 284
 - single**, 291
 - target**, 338
 - target data**, 334
 - target enter data**, 335
 - target exit data**, 337
 - target update**, 343
 - task**, 316
 - taskgroup**, 359
 - tasking constructs, 313
 - taskloop**, 321
 - taskwait**, 361
 - taskyield**, 325
 - teams**, 279
 - interchange**, 266

- tile**, 262
- unroll**, 263
- work-distribution, 290
- workshare**, 295
- worksharing, 290
- worksharing-loop construct, 300
- contains**, 254
- controlling OpenMP thread affinity, 276
- copyin**, 166
- copyprivate**, 168
- critical**, 354

D

- data copying clauses, 166
- data environment, 109
- data-mapping control, 170
- data-motion clauses, 188
- data-sharing attribute clauses, 122
- data-sharing attribute rules, 109
- declare induction**, 160
- declare mapper**, 185
- declare reduction**, 157
- declare simd**, 231
- Declare Target, 234
- declare target**, 236
- declare variant**, 225
- declare variant, 220
- default**, 122
- defaultmap**, 183
- depend**, 388
- depend object, 386
- dependences, 385
- depobj**, 387
- deprecated features, 703
- destroy**, 81
- detach**, 320
- device**, 331
- device constructs
 - device constructs, 330
 - target**, 338
 - target update**, 343
- device data environments, 9, 335, 337
- device directives, 330
- device information routines, 451

- device memory routines, 456
- device_type**, 330
- directive format, 52
- directive syntax, 51
- directive-name-modifier**, 75
- directives
 - allocate**, 203
 - assumes**, 258, 259
 - assumptions, 252
 - begin assumes**, 259
 - begin declare target**, 239
 - begin declare variant**, 226
 - begin metadirective**, 219
 - declare induction**, 160
 - declare mapper**, 185
 - declare reduction**, 157
 - declare simd**, 231
 - Declare Target, 234
 - declare target**, 236
 - declare variant**, 225
 - declare variant, 220
 - error**, 242
 - groupprivate**, 193
 - memory management directives, 197
 - metadirective**, 216, 219
 - nothing**, 259
 - requires**, 245
 - scan Directive**, 163
 - section**, 295
 - threadprivate**, 114
 - variant directives, 210
- dispatch**, 228
- dist_schedule**, 308
- distribute**, 306
- do**, 303
- doacross**, 391
- dynamic**, 305
- dynamic thread adjustment, 697
- dynamic_allocators**, 247

E

- enter**, 182
- environment display routine, 522
- environment variables, 30

OMP_AFFINITY_FORMAT, 39
OMP_ALLOCATOR, 48
OMP_AVAILABLE_DEVICES, 41
OMP_CANCELLATION, 41
OMP_DEBUG, 47
OMP_DEFAULT_DEVICE, 42
OMP_DISPLAY_AFFINITY, 38
OMP_DISPLAY_ENV, 49
OMP_DYNAMIC, 31
OMP_MAX_ACTIVE_LEVELS, 32
OMP_MAX_TASK_PRIORITY, 45
OMP_NUM_TEAMS, 49
OMP_NUM_THREADS, 31
OMP_PLACES, 33
OMP_PROC_BIND, 35
OMP_SCHEDULE, 36
OMP_STACKSIZE, 37
OMP_TARGET_OFFLOAD, 42
OMP_TEAMS_THREAD_LIMIT, 49
OMP_THREAD_LIMIT, 32
OMP_THREADS_RESERVE, 43
OMP_TOOL, 45
OMP_TOOL_LIBRARIES, 45
OMP_TOOL_VERBOSE_INIT, 46
OMP_WAIT_POLICY, 37

event, 490
 event callback registration, 531
 event callback signatures, 560
 event routines, 490
exclusive, 166
 execution model, 3
extended-atomic, 371

F
fail, 373
 features history, 703
filter, 289
final, 314
firstprivate, 125
 fixed source form conditional compilation
 sentinels, 74
 fixed source form directives, 58
flush, 379
 flush operation, 10
 flush synchronization, 11
 flush-set, 10
for, 302
 frames, 556
 free source form conditional compilation
 sentinel, 75
 free source form directives, 59
from, 191
full, 264
fuse, 267

G
 glossary, 2
grainsize, 324
groupprivate, 193
guided, 305

H
 happens before, 11
has_device_addr, 136
 header files, 414
hint, 353
 history of features, 703
holds, 255

I
 ICVs (internal control variables), 19
if Clause, 80
 implementation, 693
 implicit barrier, 357
 implicit data-mapping attribute rules, 170
 implicit flushes, 381
in_reduction, 153
inbranch, 233
 include files, 414
inclusive, 165
indirect, 240
induction, 155
inductor, 162
 informational and utility directives, 242
init, 348
 internal control variables, 693
 internal control variables (ICVs), 19
interop, 229

interoperability, 346
 Interoperability routines, 491
 introduction, 2
is_device_ptr, 134
 iterators, 71

L

lastprivate, 128
linear, 131
link, 183
 list item privatization, 119
local, 196
 lock routines, 479
loop, 309
 loop concepts, 95
 loop iteration spaces, 101
 loop iteration vectors, 101
 loop-transforming constructs, 261

M

map, 173
mapper, 172
 mapper identifiers, 172
masked, 288
match, 221
 memory allocators, 198
 memory management, 197
 memory management directives
 memory management directives, 197
 memory management routines, 498
 memory model, 7
 memory spaces, 197
memory-order, 365
memscope, 374
mergeable, 314
message, 244
 metadirective, 216
metadirective, 219
 modifier
 directive-name-modifier*directive-name-*
 modifier, 75
 task-dependence-type*task-dependence-*
 type, 385
 modifying and retrieving ICV values, 25

modifying ICVs, 21

N

nesting of regions, 405
no_omp, 255
no_omp_constructs, 256
no_omp_routines, 257
no_parallelism, 257
nocontext, 230
nogroup, 364
nontemporal, 286
 normative references, 16
nothing, 259
notinbranch, 234
novariants, 230
nowait, 363
num_tasks, 325
num_teams, 282
num_threads, 275

O

OMP_AFFINITY_FORMAT, 39
omp_aligned_alloc, 510
omp_aligned_calloc, 513
omp_alloc, 510
OMP_ALLOCATOR, 48
omp_ancestor_is_free_agent, 447
OMP_AVAILABLE_DEVICES, 41
omp_calloc, 513
OMP_CANCELLATION, 41
omp_capture_affinity, 438
omp_curr_progress_width, 270
OMP_DEBUG, 47
OMP_DEFAULT_DEVICE, 42
omp_destroy_allocator, 508
omp_destroy_lock, 483
omp_destroy_nest_lock, 483
OMP_DISPLAY_AFFINITY, 38
omp_display_affinity, 438
OMP_DISPLAY_ENV, 49
omp_display_env, 522
OMP_DYNAMIC, 31
omp_free, 512
omp_fulfill_event, 490

omp_get_active_level, 430
 omp_get_affinity_format, 437
 omp_get_ancestor_thread_num, 428
 omp_get_cancellation, 422
 omp_get_default_allocator, 510
 omp_get_default_device, 453
 omp_get_device_allocator, 505
 omp_get_device_and_host_allocator, 505
 omp_get_device_and_host_memspace, 501
 omp_get_device_memspace, 501
 omp_get_device_num, 454
 omp_get_devices_all_allocator, 505
 omp_get_devices_all_memspace, 501
 omp_get_devices_allocator, 505
 omp_get_devices_and_host_allocator, 505
 omp_get_devices_and_host_memspace, 501
 omp_get_devices_memspace, 501
 omp_get_dynamic, 421
 omp_get_initial_device, 455
 omp_get_interop_int, 493
 omp_get_interop_name, 495
 omp_get_interop_ptr, 494
 omp_get_interop_rc_desc, 497
 omp_get_interop_str, 495
 omp_get_interop_type_desc, 496
 omp_get_level, 427
 omp_get_mapped_ptr, 478
 omp_get_max_active_levels, 427
 omp_get_max_progress_width, 452
 omp_get_max_task_priority, 444
 omp_get_max_teams, 442
 omp_get_max_threads, 418
 omp_get_memspaces_num_resources, 517
 omp_get_num_devices, 453
 omp_get_num_interop_properties, 492
 omp_get_num_places, 432
 omp_get_num_procs, 451
 omp_get_num_teams, 440
 omp_get_num_threads, 417
 omp_get_partition_num_places, 434
 omp_get_partition_place_nums, 435
 omp_get_place_num, 434
 omp_get_place_num_procs, 432
 omp_get_place_proc_ids, 433
 omp_get_proc_bind, 430
 omp_get_schedule, 424
 omp_get_submemspace, 518
 omp_get_supported_active_levels, 425
 omp_get_team_num, 440
 omp_get_team_size, 429
 omp_get_teams_thread_limit, 444
 omp_get_thread_limit, 425
 omp_get_thread_num, 419
 omp_get_wtick, 490
 omp_get_wtime, 489
 omp_in_explicit_task, 445
 omp_in_final, 445
 omp_in_parallel, 419
 omp_init_allocator, 504
 omp_init_lock, 481, 482
 omp_init_nest_lock, 481, 482
 omp_is_free_agent, 446
 omp_is_initial_device, 455
 OMP_MAX_ACTIVE_LEVELS, 32
 OMP_MAX_TASK_PRIORITY, 45
 OMP_NUM_TEAMS, 49
 OMP_NUM_THREADS, 31
 omp_pause_resource, 448
 omp_pause_resource_all, 450
 OMP_PLACES, 33
 omp_pool, 315
 OMP_PROC_BIND, 35
 omp_realloc, 515
 OMP_SCHEDULE, 36
 omp_set_affinity_format, 436

- omp_set_default_allocator, 509
- omp_set_default_device, 452
- omp_set_dynamic, 420
- omp_set_lock, 484
- omp_set_max_active_levels, 426
- omp_set_nest_lock, 484
- omp_set_num_teams, 441
- omp_set_num_threads, 417
- omp_set_schedule, 422
- omp_set_teams_thread_limit, 443
- OMP_STACKSIZE, 37
- omp_target_alloc, 456
- omp_target_associate_ptr, 474
- omp_target_disassociate_ptr, 476
- omp_target_free, 458
- omp_target_is_accessible, 460
- omp_target_is_present, 459
- omp_target_memcpy, 461
- omp_target_memcpy_async, 465
- omp_target_memcpy_rect, 463
- omp_target_memcpy_rect_async, 467
- omp_target_memset, 470
- omp_target_memset_async, 472
- OMP_TARGET_OFFLOAD, 42
- omp_team, 315
- OMP_TEAMS_THREAD_LIMIT, 49
- omp_test_lock, 487
- omp_test_nest_lock, 487
- OMP_THREAD_LIMIT, 32
- OMP_THREADS_RESERVE, 43
- OMP_TOOL, 45
- OMP_TOOL_LIBRARIES, 45
- OMP_TOOL_VERBOSE_INIT, 46
- omp_unset_lock, 486
- omp_unset_nest_lock, 486
- OMP_WAIT_POLICY, 37
- ompd_bp_device_begin, 690
- ompd_bp_device_end, 690
- ompd_bp_parallel_begin, 684
- ompd_bp_parallel_end, 685
- ompd_bp_target_begin, 689
- ompd_bp_target_end, 689
- ompd_bp_task_begin, 687
- ompd_bp_task_end, 687
- ompd_bp_teams_begin, 686
- ompd_bp_teams_end, 686
- ompd_bp_thread_begin, 688
- ompd_bp_thread_end, 688
- ompd_callback_device_host_fn_t, 647
- ompd_callback_get_thread_context_for_thread_id_fn_t, 641
- ompd_callback_memory_alloc_fn_t, 639
- ompd_callback_memory_free_fn_t, 640
- ompd_callback_memory_read_fn_t, 645
- ompd_callback_memory_write_fn_t, 646
- ompd_callback_print_string_fn_t, 648
- ompd_callback_sizeof_fn_t, 642
- ompd_callback_symbol_addr_fn_t, 643
- ompd_callbacks_t, 649
- ompd_dll_locations_valid, 629
- ompd_dll_locations, 628
- ompt_callback_buffer_complete_t, 584
- ompt_callback_buffer_request_t, 583
- ompt_callback_cancel_t, 579
- ompt_callback_control_tool_t, 594
- ompt_callback_dependences_t, 568
- ompt_callback_dispatch_t, 565
- ompt_callback_error_t, 595
- ompt_callback_device_finalize_t, 581
- ompt_callback_device_initialize_t, 580
- ompt_callback_flush_t, 578
- ompt_callback_implicit

- `_task_t`, 571
- `ompt_callback_masked_t`, 572
- `ompt_callback_mutex`
 - `_acquire_t`, 575
- `ompt_callback_mutex_t`, 576
- `ompt_callback_nest_lock_t`, 577
- `ompt_callback_parallel`
 - `_begin_t`, 561
- `ompt_callback_parallel`
 - `_end_t`, 563
- `ompt_callback_sync`
 - `_region_t`, 573
- `ompt_callback_device`
 - `_load_t`, 582
- `ompt_callback_device`
 - `_unload_t`, 583
- `ompt_callback_target_data`
 - `_emi_op_t`, 585
- `ompt_callback_target_data`
 - `_op_t`, 585
- `ompt_callback_target_emi_t`, 588
- `ompt_callback_target`
 - `_map_emi_t`, 590
- `ompt_callback_target_map_t`, 590
- `ompt_callback_target`
 - `_submit_emi_t`, 592
- `ompt_callback_target`
 - `_submit_t`, 592
- `ompt_callback_target_t`, 588
- `ompt_callback_task`
 - `_create_t`, 567
- `ompt_callback_task`
 - `_dependence_t`, 569
- `ompt_callback_task`
 - `_schedule_t`, 570
- `ompt_callback_thread`
 - `_begin_t`, 560
- `ompt_callback_thread_end_t`, 561
- `ompt_callback_work_t`, 564
- OpenMP allocator structured blocks, 86
- OpenMP argument lists, 64
- OpenMP atomic structured blocks, 88
- OpenMP compliance, 15

- OpenMP context-specific structured blocks, 85
- OpenMP function dispatch structured blocks, 87
- OpenMP operations, 67
- OpenMP stylized expressions, 85
- OpenMP types, 83
- order**, 283
- ordered**, 104, 393–395
- otherwise**, 218

P

- parallel**, 270
- parallelism generating constructs, 270
- parallelization-level*, 397
- partial*, 265
- permutation**, 267
- priority**, 316
- private**, 124
- proc_bind**, 278

R

- read**, 369
- read, atomic**, 375
- collector**, 162
- combiner**, 159
- initializer**, 159
- reduction**, 150
- reduction clauses, 138
- relaxed**, 367
- release**, 368
- release flush, 11
- requirement*, 246
- requires**, 245
- reserved locators, 66
- resource relinquishing routines, 448
- reverse**, 265
- reverse_offload**, 248
- runtime**, 305
- runtime library definitions, 414
- runtime library routines, 413

S

- safelen**, 287

safesync, 279
scan Directive, 163
schedule, 304
scheduling, 327
scope, 292
section, 295
sections, 293
self_maps, 252
seq_cst, 368
severity, 244
shared, 123
simd, 284, 398
simdlen, 287
Simple Lock Routines, 480
single, 291
sizes, 263
looprange, 107
stand-alone directives, 57
static, 304
strong flush, 10
structured blocks, 85
synchronization constructs, 351
synchronization constructs and clauses, 351
synchronization hint type, 351
synchronization hints, 351

T

target, 338
target data, 334
target memory routines, 456
target update, 343
task, 316
task scheduling, 327
task-dependence-type, 385
task_reduction, 153
taskgroup, 359
tasking constructs, 313
tasking routines, 444
taskloop, 321
taskwait, 361
taskyield, 325
teams, 279
teams region routines, 440
thread affinity, 276
thread affinity routines, 430
thread team routines, 417
thread_limit, 332
threadprivate, 114
threads, 397
threadset, 315
interchange, 266
tile, 262
timer, 489
timing routines, 489
to, 190
tool control, 519
tool initialization, 528
tool interfaces definitions, 525, 628
tools header files, 525, 628
tracing device activity, 532
types
 sync_hint, 351

U

unified_address, 249
unified_shared_memory, 250
uniform, 192
unroll, 263
untied, 313
update, 370, 386
update, atomic, 375
use, 349
use_device_addr, 137
use_device_ptr, 135
uses_allocators, 207

V

variables, environment, 30
variant directives, 210

W

wait identifier, 557
wall clock timer, 489
error, 242
weak, 374
when, 217
work-distribution
 constructs, 290

work-distribution constructs, [290](#)
workshare, [295](#)
worksharing
 constructs, [290](#)
worksharing constructs, [290](#)
worksharing-loop construct, [300](#)
write, [370](#)
write, atomic, [375](#)