



OpenMP Technical Report 14: Version 6.1 Preview

This Technical Report is a first preview of the OpenMP Application Programming Interface Specification version 6.1. This version adds new features to the loop transformation directives, such as the **flatten** directive, the **depth** clause for the **fuse** directive, and the sizes selector for the **sizes** clause to specify the number of tiles. For more control over pointer attachment in map clauses, the **attach** modifier was added. See Appendix B.2 for the complete list of changes relative to version 6.0.

EDITORS

Bronis R. de Supinski
Michael Klemm

November 13, 2025
Expires June 30, 2026

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 Application Programming Interface (API) specification.

The goal of this technical report is to build more widespread existing practice for an expanded OpenMP API. It gives advice on extensions or future directions to those vendors who wish to provide them for trial implementation, allows the OpenMP Architecture Review Board to gather early feedback, supports timing and scheduling differences between official OpenMP releases, and offers a preview to users of the future directions of the OpenMP API 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 the OpenMP API, but they are not currently part of any OpenMP API 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 6.1 Preview November 2025

Copyright ©1997-2025 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 6.0 LaTeX sources: 2489, 3211, 3225, 3284, 3288, 3682, 3833, 3877, 4143, 4371, 4398, 4432-4433, 4443-4444, 4448-4450, 4453, 4455, 4457, 4470, 4476, 4483, 4485, 4491-4492, 4502, 4504-4505, 4509, 4512, 4514-4515, 4521, 4523, 4525, 4530-4531, 4534-4535, 4557-4558, 4561, 4567, 4569, 4571-4572, 4576, 4579-4580, 4586, 4594, 4597-4598, 4600, 4610, 4628

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	Execution Model	2
1.3	Memory Model	7
1.3.1	Structure of the OpenMP Memory Model	7
1.3.2	Device Data Environments Overview	8
1.3.3	Memory Management	9
1.3.4	The Flush Operation	10
1.3.5	Flush Synchronization and Happens-Before Order	11
1.3.6	OpenMP Memory Consistency	13
1.4	Tool Interfaces	14
1.4.1	OMPT	14
1.4.2	OMPD	15
1.5	OpenMP Compliance	15
1.6	Normative References	16
1.7	Organization of this Document	17
2	Glossary	19
3	Internal Control Variables	118
3.1	ICV Descriptions	118
3.2	ICV Initialization	121
3.3	Modifying and Retrieving ICV Values	124
3.4	How the Per-Data Environment ICVs Work	127
3.5	ICV Override Relationships	128

4	Environment Variables	130
4.1	Parallel Region Environment Variables	131
4.1.1	Abstract Name Values	131
4.1.2	<code>OMP_DYNAMIC</code>	132
4.1.3	<code>OMP_NUM_THREADS</code>	132
4.1.4	<code>OMP_THREAD_LIMIT</code>	133
4.1.5	<code>OMP_MAX_ACTIVE_LEVELS</code>	133
4.1.6	<code>OMP_PLACES</code>	134
4.1.7	<code>OMP_PROC_BIND</code>	135
4.2	Teams Environment Variables	136
4.2.1	<code>OMP_NUM_TEAMS</code>	137
4.2.2	<code>OMP_TEAMS_THREAD_LIMIT</code>	137
4.3	Program Execution Environment Variables	137
4.3.1	<code>OMP_SCHEDULE</code>	137
4.3.2	<code>OMP_STACKSIZE</code>	138
4.3.3	<code>OMP_WAIT_POLICY</code>	139
4.3.4	<code>OMP_DISPLAY_AFFINITY</code>	139
4.3.5	<code>OMP_AFFINITY_FORMAT</code>	140
4.3.6	<code>OMP_CANCELLATION</code>	142
4.3.7	<code>OMP_AVAILABLE_DEVICES</code>	142
4.3.8	<code>OMP_DEFAULT_DEVICE</code>	143
4.3.9	<code>OMP_TARGET_OFFLOAD</code>	144
4.3.10	<code>OMP_THREADS_RESERVE</code>	145
4.3.11	<code>OMP_MAX_TASK_PRIORITY</code>	146
4.4	Memory Allocation Environment Variables	147
4.4.1	<code>OMP_ALLOCATOR</code>	147
4.5	OMPT Environment Variables	148
4.5.1	<code>OMP_TOOL</code>	148
4.5.2	<code>OMP_TOOL_LIBRARIES</code>	148
4.5.3	<code>OMP_TOOL_VERBOSE_INIT</code>	149
4.6	OMPD Environment Variables	150
4.6.1	<code>OMP_DEBUG</code>	150
4.7	<code>OMP_DISPLAY_ENV</code>	150

5	Directive and Construct Syntax	151
5.1	Directive Format	153
5.1.1	Free Source Form Directives	160
5.1.2	Fixed Source Form Directives	160
5.2	Clause Format	161
5.2.1	OpenMP Argument Lists	165
5.2.2	Reserved Locators	168
5.2.3	OpenMP Operations	168
5.2.4	Array Shaping	168
5.2.5	Array Sections	169
5.2.6	iterator Modifier	172
5.3	Conditional Compilation	175
5.3.1	Free Source Form Conditional Compilation Sentinel	176
5.3.2	Fixed Source Form Conditional Compilation Sentinels	177
5.4	Intrinsic Identifiers	178
5.5	<i>directive-name-modifier</i> Modifier	178
5.6	if Clause	184
5.7	init Clause	185
5.8	destroy Clause	186
6	Base Language Formats and Restrictions	188
6.1	OpenMP Types and Identifiers	188
6.2	OpenMP Stylized Expressions	190
6.3	Structured Blocks	191
6.3.1	OpenMP Allocator Structured Blocks	192
6.3.2	OpenMP Function Dispatch Structured Blocks	192
6.3.3	OpenMP Atomic Structured Blocks	193
6.4	Loop Concepts	200
6.4.1	Canonical Loop Nest Form	201
6.4.2	Canonical Loop Sequence Form	207
6.4.3	OpenMP Loop-Iteration Spaces and Vectors	208
6.4.4	Consistent Loop Schedules	210
6.4.5	collapse Clause	210
6.4.6	ordered Clause	211

6.4.7	depth Clause	212
6.4.8	looprange Clause	213
II	Data Control	214
7	Data-Sharing Control	215
7.1	Data-Sharing Attribute Rules	215
7.1.1	Variables Referenced in a Construct	215
7.1.2	Variables Referenced in a Region but not in a Construct	219
7.2	List Item Privatization	220
7.3	Data-Sharing Attribute Clauses	223
7.3.1	default Clause	223
7.3.2	shared Clause	225
7.3.3	private Clause	226
7.3.4	firstprivate Clause	227
7.3.5	lastprivate Clause	230
7.3.6	is_device_ptr Clause	233
7.3.7	use_device_ptr Clause	234
7.3.8	has_device_addr Clause	235
7.3.9	use_device_addr Clause	236
7.4	saved Modifier	237
8	Reduction and Induction Data Control	238
8.1	OpenMP Reduction and Induction Identifiers	238
8.2	OpenMP Reduction and Induction Expressions	239
8.2.1	OpenMP Combiner Expressions	239
8.2.2	OpenMP Initializer Expressions	240
8.2.3	OpenMP Inductor Expressions	242
8.2.4	OpenMP Collector Expressions	243
8.3	Implicitly Declared OpenMP Reduction Identifiers	243
8.4	Implicitly Declared OpenMP Induction Identifiers	245
8.5	Properties Common to Reduction and induction Clauses	246
8.6	Properties Common to All Reduction Clauses	248
8.7	Reduction Scoping Clauses	250

8.8	Reduction Participating Clauses	250
8.9	<i>reduction-identifier</i> Modifier	250
8.10	reduction Clause	251
8.11	task_reduction Clause	255
8.12	in_reduction Clause	255
8.13	induction Clause	257
8.14	linear Clause	259
8.15	declare_reduction Directive	262
8.15.1	combiner Clause	265
8.15.2	initializer Clause	265
8.16	declare_induction Directive	266
8.16.1	inductor Clause	268
8.16.2	collector Clause	268
8.17	scan Directive	269
8.17.1	inclusive Clause	272
8.17.2	exclusive Clause	272
8.17.3	init_complete Clause	273
9	Static-Lifetime Data Control	274
9.1	threadprivate Directive	274
9.2	groupprivate Directive	278
9.3	local Clause	280
9.4	enter Clause	281
9.5	link Clause	283
10	Data-Copying Control	285
10.1	copyin Clause	285
10.2	copyprivate Clause	286
11	Data-Mapping Control	289
11.1	Implicit Data-Mapping Attribute Rules	289
11.2	Mapper Identifiers and mapper Modifiers	291
11.3	map Clause	292
11.3.1	<i>map-type</i> Modifier	296
11.3.2	Map Type Decay	297

11.3.3	ref-modifier Modifier	297
11.3.4	attach-modifier Modifier	299
11.3.5	Assumed-Size Array Mapping	300
11.3.6	Array and Structure Mapping	300
11.3.7	Storage Block Mapping	302
11.3.8	Implicit Data-Mapping	305
11.4	defaultmap Clause	305
11.5	declare_mapper Directive	307
12	Data-Motion Control	310
12.1	to Clause	311
12.2	from Clause	313
13	Memory Management	315
13.1	Memory Spaces	315
13.2	Memory Allocators	316
13.3	align Clause	320
13.4	allocator Clause	321
13.5	allocate Directive	321
13.6	allocate Clause	323
13.7	allocators Construct	326
13.8	uses_allocators Clause	326
13.9	dyn_groupprivate Clause	328
14	SIMD Data Control	331
14.1	uniform Clause	331
14.2	aligned Clause	331
III	General Directives and Clauses	334
15	Variant Directives	335
15.1	OpenMP Contexts	335
15.2	Context Selectors	337
15.3	Matching and Scoring Context Selectors	340

15.4	Metadirectives	341
15.4.1	when Clause	342
15.4.2	otherwise Clause	343
15.4.3	metadirective	344
15.4.4	begin metadirective	344
15.5	Semantic Requirement Set	345
15.6	Declare Variant Directives	346
15.6.1	match Clause	347
15.6.2	adjust_args Clause	348
15.6.3	append_args Clause	350
15.6.4	declare_variant Directive	351
15.6.5	begin declare_variant Directive	353
15.7	dispatch Construct	355
15.7.1	interop Clause	356
15.7.2	novariants Clause	357
15.7.3	nocontext Clause	358
15.8	declare_simd Directive	358
15.8.1	<i>branch</i> Clauses	360
15.9	Declare Target Directives	362
15.9.1	declare_target Directive	364
15.9.2	begin declare_target Directive	366
15.9.3	indirect Clause	368
16	Informational and Utility Directives	369
16.1	error Directive	369
16.2	at Clause	370
16.3	message Clause	370
16.4	severity Clause	371
16.5	requires Directive	372
16.5.1	<i>requirement</i> Clauses	373
16.6	Assumption Directives	379
16.6.1	<i>assumption</i> Clauses	380
16.6.2	assumes Directive	385
16.6.3	assume Directive	385

16.6.4	begin assumes Directive	386
16.7	nothing Directive	386
17	Loop-Transforming Constructs	388
17.1	apply Clause	389
17.2	sizes Clause	391
17.3	flatten Construct	392
17.4	fuse Construct	392
17.5	interchange Construct	393
17.5.1	permutation Clause	394
17.6	reverse Construct	395
17.7	split Construct	395
17.7.1	counts Clause	396
17.8	stripe Construct	397
17.9	tile Construct	398
17.10	unroll Construct	400
17.10.1	full Clause	400
17.10.2	partial Clause	401
18	Parallelism Generation and Control	402
18.1	parallel Construct	402
18.1.1	Determining the Number of Threads for a parallel Region	406
18.1.2	num_threads Clause	406
18.1.3	Controlling OpenMP Thread Affinity	407
18.1.4	proc_bind Clause	410
18.1.5	safesync Clause	411
18.2	teams Construct	412
18.2.1	num_teams Clause	415
18.3	order Clause	416
18.4	simd Construct	417
18.4.1	nontemporal Clause	418
18.4.2	safelen Clause	419
18.4.3	simdlen Clause	420

18.5	masked Construct	421
18.5.1	filter Clause	422
19	Work-Distribution Constructs	423
19.1	single Construct	424
19.2	scope Construct	425
19.3	sections Construct	426
19.3.1	section Directive	428
19.4	workshare Construct	428
19.5	workdistribute Construct	431
19.6	Worksharing-Loop Constructs	433
19.6.1	for Construct	435
19.6.2	do Construct	436
19.6.3	schedule Clause	437
19.7	distribute Construct	439
19.7.1	dist_schedule Clause	441
19.8	loop Construct	442
19.8.1	bind Clause	444
20	Tasking Constructs	446
20.1	task Construct	446
20.2	taskloop Construct	449
20.2.1	grainsize Clause	452
20.2.2	num_tasks Clause	453
20.2.3	task_iteration Directive	454
20.3	taskgraph Construct	455
20.3.1	graph_id Clause	458
20.3.2	graph_reset Clause	458
20.4	untied Clause	459
20.5	mergeable Clause	460
20.6	replayable Clause	460
20.7	final Clause	461
20.8	threadset Clause	462
20.9	priority Clause	463

20.10	affinity Clause	464
20.11	detach Clause	465
20.12	taskyield Construct	466
20.13	Initial Task	466
20.14	Task Scheduling	467
21	Device Directives and Clauses	470
21.1	device_type Clause	470
21.2	device Clause	471
21.3	thread_limit Clause	472
21.4	Device Initialization	473
21.5	target_enter_data Construct	474
21.6	target_exit_data Construct	476
21.7	target_data Construct	478
21.8	target Construct	481
21.9	target_update Construct	486
22	Interoperability	488
22.1	interop Construct	488
22.1.1	OpenMP Foreign Runtime Identifiers	489
22.1.2	use Clause	489
22.1.3	prefer-type Modifier	490
23	Synchronization Constructs and Clauses	492
23.1	hint Clause	492
23.2	critical Construct	493
23.3	Barriers	495
23.3.1	barrier Construct	495
23.3.2	Implicit Barriers	496
23.3.3	Implementation-Specific Barriers	497
23.4	taskgroup Construct	498
23.5	taskwait Construct	499
23.6	nowait Clause	501
23.7	nogroup Clause	503

23.8	OpenMP Memory Ordering	504
23.8.1	<i>memory-order</i> Clauses	504
23.8.2	<i>atomic</i> Clauses	508
23.8.3	<i>extended-atomic</i> Clauses	510
23.8.4	memscope Clause	513
23.8.5	atomic Construct	514
23.8.6	flush Construct	518
23.8.7	Implicit Flushes	520
23.9	OpenMP Dependences	524
23.9.1	<i>task-dependence-type</i> Modifier	525
23.9.2	Depend Objects	525
23.9.3	depobj Construct	525
23.9.4	update Clause	526
23.9.5	depend Clause	527
23.9.6	transparent Clause	531
23.9.7	doacross Clause	532
23.10	ordered Construct	533
23.10.1	Stand-alone ordered Construct	534
23.10.2	Block-associated ordered Construct	536
23.10.3	<i>parallelization-level</i> Clauses	537
24	Cancellation Constructs	539
24.1	<i>cancel-directive-name</i> Clauses	539
24.2	cancel Construct	540
24.3	cancellation_point Construct	544
25	Composition of Constructs	545
25.1	Compound Directive Names	545
25.2	Clauses on Compound Constructs	548
25.3	Compound Construct Semantics	551

IV Runtime Library Routines 552

26 Runtime Library Definitions	553
26.1 Predefined Identifiers	554
26.2 Routine Bindings	554
26.3 Routine Argument Properties	555
26.4 General OpenMP Types	556
26.4.1 OpenMP intptr Type	556
26.4.2 OpenMP uintptr Type	556
26.5 OpenMP Parallel Region Support Types	556
26.5.1 OpenMP sched Type	556
26.6 OpenMP Tasking Support Types	558
26.6.1 OpenMP event_handle Type	558
26.7 OpenMP Interoperability Support Types	558
26.7.1 OpenMP interop Type	558
26.7.2 OpenMP interop_fr Type	559
26.7.3 OpenMP interop_property Type	560
26.7.4 OpenMP interop_rc Type	561
26.8 OpenMP Memory Management Types	564
26.8.1 OpenMP access Type	564
26.8.2 OpenMP allocator_handle Type	565
26.8.3 OpenMP alloctrail Type	566
26.8.4 OpenMP alloctrail_key Type	568
26.8.5 OpenMP alloctrail_value Type	570
26.8.6 OpenMP alloctrail_val Type	573
26.8.7 OpenMP mempartition Type	573
26.8.8 OpenMP mempartitioner Type	573
26.8.9 OpenMP mempartitioner_lifetime Type	574
26.8.10 OpenMP mempartitioner_compute_proc Type	575
26.8.11 OpenMP mempartitioner_release_proc Type	576
26.8.12 OpenMP memspace_handle Type	577
26.9 OpenMP Synchronization Types	578
26.9.1 OpenMP depend Type	578
26.9.2 OpenMP impex Type	578

26.9.3	OpenMP lock Type	579
26.9.4	OpenMP nest_lock Type	580
26.9.5	OpenMP sync_hint Type	580
26.10	OpenMP Affinity Support Types	582
26.10.1	OpenMP proc_bind Type	582
26.11	OpenMP Resource Relinquishing Types	584
26.11.1	OpenMP pause_resource Type	584
26.12	OpenMP Tool Types	585
26.12.1	OpenMP control_tool Type	585
26.12.2	OpenMP control_tool_result Type	587
27	Parallel Region Support Routines	588
27.1	omp_set_num_threads Routine	588
27.2	omp_get_num_threads Routine	589
27.3	omp_get_thread_num Routine	589
27.4	omp_get_max_threads Routine	590
27.5	omp_get_thread_limit Routine	591
27.6	omp_in_parallel Routine	591
27.7	omp_set_dynamic Routine	592
27.8	omp_get_dynamic Routine	593
27.9	omp_set_schedule Routine	593
27.10	omp_get_schedule Routine	594
27.11	omp_get_supported_active_levels Routine	595
27.12	omp_set_max_active_levels Routine	596
27.13	omp_get_max_active_levels Routine	597
27.14	omp_get_level Routine	597
27.15	omp_get_ancestor_thread_num Routine	598
27.16	omp_get_team_size Routine	599
27.17	omp_get_active_level Routine	600
28	Teams Region Routines	601
28.1	omp_get_num_teams Routine	601
28.2	omp_set_num_teams Routine	601
28.3	omp_get_team_num Routine	602

28.4	<code>omp_get_max_teams</code> Routine	603
28.5	<code>omp_get_teams_thread_limit</code> Routine	604
28.6	<code>omp_set_teams_thread_limit</code> Routine	604
29	Tasking Support Routines	606
29.1	Tasking Routines	606
29.1.1	<code>omp_get_max_task_priority</code> Routine	606
29.1.2	<code>omp_in_explicit_task</code> Routine	607
29.1.3	<code>omp_in_final</code> Routine	607
29.1.4	<code>omp_is_free_agent</code> Routine	608
29.1.5	<code>omp_ancestor_is_free_agent</code> Routine	608
29.2	Event Routine	609
29.2.1	<code>omp_fulfill_event</code> Routine	609
30	Device Information Routines	612
30.1	<code>omp_set_default_device</code> Routine	612
30.2	<code>omp_get_default_device</code> Routine	613
30.3	<code>omp_get_num_devices</code> Routine	613
30.4	<code>omp_get_device_num</code> Routine	614
30.5	<code>omp_get_num_procs</code> Routine	615
30.6	<code>omp_get_max_progress_width</code> Routine	615
30.7	<code>omp_get_device_from_uid</code> Routine	616
30.8	<code>omp_get_uid_from_device</code> Routine	617
30.9	<code>omp_is_initial_device</code> Routine	618
30.10	<code>omp_get_initial_device</code> Routine	618
30.11	<code>omp_get_device_num_teams</code> Routine	619
30.12	<code>omp_set_device_num_teams</code> Routine	620
30.13	<code>omp_get_device_teams_thread_limit</code> Routine	621
30.14	<code>omp_set_device_teams_thread_limit</code> Routine	621
31	Device Memory Routines	623
31.1	Asynchronous Device Memory Routines	624
31.2	Device Memory Information Routines	624
31.2.1	<code>omp_target_is_present</code> Routine	624
31.2.2	<code>omp_target_is_accessible</code> Routine	625

31.2.3	omp_get_mapped_ptr Routine	626
31.3	omp_target_alloc Routine	627
31.4	omp_target_free Routine	629
31.5	omp_target_associate_ptr Routine	630
31.6	omp_target_disassociate_ptr Routine	631
31.7	Memory Copying Routines	633
31.7.1	omp_target_memcpy Routine	634
31.7.2	omp_target_memcpy_rect Routine	635
31.7.3	omp_target_memcpy_async Routine	637
31.7.4	omp_target_memcpy_rect_async Routine	638
31.8	Memory Setting Routines	640
31.8.1	omp_target_memset Routine	640
31.8.2	omp_target_memset_async Routine	641
32	Interoperability Routines	643
32.1	omp_get_num_interop_properties Routine	644
32.2	omp_get_interop_int Routine	644
32.3	omp_get_interop_ptr Routine	645
32.4	omp_get_interop_str Routine	646
32.5	omp_get_interop_name Routine	647
32.6	omp_get_interop_type_desc Routine	648
32.7	omp_get_interop_rc_desc Routine	649
33	Memory Management Routines	651
33.1	Memory Space Retrieving Routines	651
33.1.1	omp_get_devices_memspace Routine	652
33.1.2	omp_get_device_memspace Routine	653
33.1.3	omp_get_devices_and_host_memspace Routine	654
33.1.4	omp_get_device_and_host_memspace Routine	654
33.1.5	omp_get_devices_all_memspace Routine	655
33.2	omp_get_memspace_num_resources Routine	656
33.3	omp_get_memspace_pagesize Routine	657
33.4	omp_get_submemspace Routine	658

33.5	OpenMP Memory Partitioning Routines	659
33.5.1	<code>omp_init_mempartitioner</code> Routine	659
33.5.2	<code>omp_destroy_mempartitioner</code> Routine	661
33.5.3	<code>omp_init_mempartition</code> Routine	662
33.5.4	<code>omp_destroy_mempartition</code> Routine	663
33.5.5	<code>omp_mempartition_set_part</code> Routine	664
33.5.6	<code>omp_mempartition_get_user_data</code> Routine	665
33.6	<code>omp_init_allocator</code> Routine	666
33.7	<code>omp_destroy_allocator</code> Routine	667
33.8	Memory Allocator Retrieving Routines	668
33.8.1	<code>omp_get_devices_allocator</code> Routine	669
33.8.2	<code>omp_get_device_allocator</code> Routine	670
33.8.3	<code>omp_get_devices_and_host_allocator</code> Routine	671
33.8.4	<code>omp_get_device_and_host_allocator</code> Routine	671
33.8.5	<code>omp_get_devices_all_allocator</code> Routine	672
33.9	<code>omp_set_default_allocator</code> Routine	673
33.10	<code>omp_get_default_allocator</code> Routine	674
33.11	Memory Allocating Routines	675
33.11.1	<code>omp_alloc</code> Routine	677
33.11.2	<code>omp_aligned_alloc</code> Routine	678
33.11.3	<code>omp_calloc</code> Routine	679
33.11.4	<code>omp_aligned_calloc</code> Routine	680
33.11.5	<code>omp_realloc</code> Routine	681
33.12	<code>omp_free</code> Routine	682
33.13	Groupprivate Memory Routines	683
33.13.1	<code>omp_get_dyn_groupprivate_ptr</code> Routine	684
33.13.2	<code>omp_get_dyn_groupprivate_size</code> Routine	685
33.13.3	<code>omp_get_groupprivate_limit</code> Routine	686
34	Lock Routines	688
34.1	Lock Initializing Routines	689
34.1.1	<code>omp_init_lock</code> Routine	689
34.1.2	<code>omp_init_nest_lock</code> Routine	690
34.1.3	<code>omp_init_lock_with_hint</code> Routine	690

34.1.4	omp_init_nest_lock_with_hint Routine	691
34.2	Lock Destroying Routines	692
34.2.1	omp_destroy_lock Routine	693
34.2.2	omp_destroy_nest_lock Routine	693
34.3	Lock Acquiring Routines	694
34.3.1	omp_set_lock Routine	695
34.3.2	omp_set_nest_lock Routine	696
34.4	Lock Releasing Routines	697
34.4.1	omp_unset_lock Routine	697
34.4.2	omp_unset_nest_lock Routine	698
34.5	Lock Testing Routines	699
34.5.1	omp_test_lock Routine	699
34.5.2	omp_test_nest_lock Routine	700
35	Thread Affinity Routines	702
35.1	omp_get_proc_bind Routine	702
35.2	omp_get_num_places Routine	703
35.3	omp_get_place_num_procs Routine	703
35.4	omp_get_place_proc_ids Routine	704
35.5	omp_get_place_num Routine	705
35.6	omp_get_partition_num_places Routine	705
35.7	omp_get_partition_place_nums Routine	706
35.8	omp_set_affinity_format Routine	707
35.9	omp_get_affinity_format Routine	708
35.10	omp_display_affinity Routine	709
35.11	omp_capture_affinity Routine	710
36	Execution Control Routines	712
36.1	omp_get_cancellation Routine	712
36.2	Resource Relinquishing Routines	713
36.2.1	omp_pause_resource Routine	713
36.2.2	omp_pause_resource_all Routine	714
36.3	Timing Routines	715
36.3.1	omp_get_wtime Routine	715

36.3.2	<code>omp_get_wtick</code> Routine	715
36.4	<code>omp_display_env</code> Routine	716
37	Tool Support Routines	718
37.1	<code>omp_control_tool</code> Routine	718
V	OMPT	720
38	OMPT Overview	721
38.1	OMPT Interfaces Definitions	721
38.2	Activating a First-Party Tool	721
38.2.1	<code>omp_start_tool</code> Procedure	721
38.2.2	Determining Whether to Initialize a First-Party Tool	723
38.2.3	Initializing a First-Party Tool	724
38.2.4	Monitoring Activity on the Host with OMPT	727
38.2.5	Tracing Activity on Target Devices	728
38.3	Finalizing a First-Party Tool	732
39	OMPT Data Types	733
39.1	OMPT Predefined Identifiers	733
39.2	OMPT <code>any_record_omp</code> Type	734
39.3	OMPT <code>buffer</code> Type	735
39.4	OMPT <code>buffer_cursor</code> Type	736
39.5	OMPT <code>callback</code> Type	736
39.6	OMPT <code>callbacks</code> Type	737
39.7	OMPT <code>cancel_flag</code> Type	739
39.8	OMPT <code>data</code> Type	739
39.9	OMPT <code>dependence</code> Type	740
39.10	OMPT <code>dependence_type</code> Type	741
39.11	OMPT <code>device</code> Type	742
39.12	OMPT <code>device_time</code> Type	742
39.13	OMPT <code>dispatch</code> Type	743
39.14	OMPT <code>dispatch_chunk</code> Type	743
39.15	OMPT <code>frame</code> Type	744

39.16	OMPT frame_flag Type	745
39.17	OMPT hwid Type	746
39.18	OMPT id Type	746
39.19	OMPT interface_fn Type	747
39.20	OMPT mutex Type	748
39.21	OMPT native_mon_flag Type	748
39.22	OMPT parallel_flag Type	749
39.23	OMPT record Type	750
39.24	OMPT record_abstract Type	751
39.25	OMPT record_native Type	752
39.26	OMPT record_ompt Type	752
39.27	OMPT scope_endpoint Type	753
39.28	OMPT set_result Type	754
39.29	OMPT severity Type	755
39.30	OMPT start_tool_result Type	756
39.31	OMPT state Type	757
39.32	OMPT subvolume Type	759
39.33	OMPT sync_region Type	760
39.34	OMPT target Type	761
39.35	OMPT target_data_op Type	762
39.36	OMPT target_map_flag Type	764
39.37	OMPT task_flag Type	765
39.38	OMPT task_status Type	766
39.39	OMPT thread Type	767
39.40	OMPT wait_id Type	768
39.41	OMPT work Type	768
40	General Callbacks and Trace Records	770
40.1	Initialization and Finalization Callbacks	771
40.1.1	initialize Callback	771
40.1.2	finalize Callback	772
40.1.3	thread_begin Callback	772
40.1.4	thread_end Callback	773
40.2	error Callback	774

40.3	Parallelism Generation Callback Signatures	775
40.3.1	parallel_begin Callback	775
40.3.2	parallel_end Callback	776
40.3.3	masked Callback	777
40.4	Work Distribution Callback Signatures	778
40.4.1	work Callback	778
40.4.2	dispatch Callback	780
40.5	Tasking Callback Signatures	781
40.5.1	task_create Callback	781
40.5.2	task_schedule Callback	782
40.5.3	implicit_task Callback	783
40.6	cancel Callback	785
40.7	Synchronization Callback Signatures	786
40.7.1	dependences Callback	786
40.7.2	task_dependence Callback	787
40.7.3	OMPT sync_region Type	788
40.7.4	sync_region Callback	789
40.7.5	sync_region_wait Callback	789
40.7.6	reduction Callback	790
40.7.7	OMPT mutex_acquire Type	790
40.7.8	mutex_acquire Callback	792
40.7.9	lock_init Callback	792
40.7.10	OMPT mutex Type	792
40.7.11	lock_destroy Callback	793
40.7.12	mutex_acquired Callback	794
40.7.13	mutex_released Callback	794
40.7.14	nest_lock Callback	795
40.7.15	flush Callback	795
40.8	control_tool Callback	796
41	Device Callbacks and Tracing	798
41.1	device_initialize Callback	798
41.2	device_finalize Callback	799
41.3	device_load Callback	800

41.4	device_unload Callback	801
41.5	buffer_request Callback	801
41.6	buffer_complete Callback	802
41.7	target_data_op_emi Callback	803
41.8	target_emi Callback	806
41.9	target_map_emi Callback	808
41.10	target_submit_emi Callback	810
42	General Entry Points	813
42.1	function_lookup Entry Point	813
42.2	enumerate_states Entry Point	814
42.3	enumerate_mutex_impls Entry Point	815
42.4	set_callback Entry Point	816
42.5	get_callback Entry Point	817
42.6	get_thread_data Entry Point	818
42.7	get_num_procs Entry Point	819
42.8	get_num_places Entry Point	819
42.9	get_place_proc_ids Entry Point	820
42.10	get_place_num Entry Point	821
42.11	get_partition_place_nums Entry Point	821
42.12	get_proc_id Entry Point	822
42.13	get_state Entry Point	822
42.14	get_parallel_info Entry Point	823
42.15	get_task_info Entry Point	825
42.16	get_task_memory Entry Point	827
42.17	get_target_info Entry Point	828
42.18	get_num_devices Entry Point	828
42.19	get_unique_id Entry Point	829
42.20	finalize_tool Entry Point	829
43	Device Tracing Entry Points	831
43.1	get_device_num_procs Entry Point	831
43.2	get_device_time Entry Point	832
43.3	translate_time Entry Point	832

43.4	set_trace_ompt Entry Point	833
43.5	set_trace_native Entry Point	834
43.6	get_buffer_limits Entry Point	835
43.7	start_trace Entry Point	836
43.8	pause_trace Entry Point	837
43.9	flush_trace Entry Point	837
43.10	stop_trace Entry Point	838
43.11	advance_buffer_cursor Entry Point	838
43.12	get_record_type Entry Point	839
43.13	get_record_ompt Entry Point	840
43.14	get_record_native Entry Point	841
43.15	get_record_abstract Entry Point	842

VI OMPD 843

44 OMPD Overview 844

44.1	OMP Interfaces Definitions	845
44.2	Thread and Signal Safety	845
44.3	Activating a Third-Party Tool	845
44.3.1	Enabling Runtime Support for OMPD	845
44.3.2	ompd_dll_locations	845
44.3.3	ompd_dll_locations_valid Breakpoint	846

45 OMPD Data Types 847

45.1	OMP addr Type	847
45.2	OMP address Type	847
45.3	OMP address_space_context Type	848
45.4	OMP callbacks Type	848
45.5	OMP device Type	850
45.6	OMP device_type_sizes Type	851
45.7	OMP frame_info Type	852
45.8	OMP icv_id Type	852
45.9	OMP rc Type	853
45.10	OMP seg Type	854

45.11	OMPD scope Type	855
45.12	OMPD size Type	856
45.13	OMPD team_generator Type	856
45.14	OMPD thread_context Type	857
45.15	OMPD thread_id Type	858
45.16	OMPD wait_id Type	858
45.17	OMPD word Type	859
45.18	OMPD Handle Types	859
45.18.1	OMPD address_space_handle Type	859
45.18.2	OMPD parallel_handle Type	860
45.18.3	OMPD task_handle Type	860
45.18.4	OMPD thread_handle Type	860
46	OMPD Callback Interface	861
46.1	Memory Management of OMPD Library	861
46.1.1	alloc_memory Callback	862
46.1.2	free_memory Callback	862
46.2	Accessing Program or Runtime Memory	863
46.2.1	symbol_addr_lookup Callback	863
46.2.2	OMPD memory_read Type	865
46.2.3	write_memory Callback	867
46.3	Context Management and Navigation	868
46.3.1	get_thread_context_for_thread_id Callback	868
46.3.2	sizeof_type Callback	870
46.4	Device Translating Callbacks	871
46.4.1	OMPD device_host Type	871
46.4.2	device_to_host Callback	872
46.4.3	host_to_device Callback	872
46.5	print_string Callback	873
47	OMPD Routines	874
47.1	OMPD Library Initialization and Finalization	874
47.1.1	ompd_initialize Routine	874
47.1.2	ompd_get_api_version Routine	875

47.1.3	ompd_get_version_string Routine	876
47.1.4	ompd_finalize Routine	877
47.2	Process Initialization and Finalization	877
47.2.1	ompd_process_initialize Routine	877
47.2.2	ompd_device_initialize Routine	878
47.2.3	ompd_get_device_thread_id_kinds Routine	880
47.3	Address Space Information	881
47.3.1	ompd_get_omp_version Routine	881
47.3.2	ompd_get_omp_version_string Routine	881
47.4	Thread Handle Routines	882
47.4.1	ompd_get_thread_in_parallel Routine	882
47.4.2	ompd_get_thread_handle Routine	883
47.4.3	ompd_get_thread_id Routine	884
47.4.4	ompd_get_device_from_thread Routine	885
47.5	Parallel Region Handle Routines	886
47.5.1	ompd_get_curr_parallel_handle Routine	886
47.5.2	ompd_get_enclosing_parallel_handle Routine	887
47.5.3	ompd_get_task_parallel_handle Routine	888
47.6	Task Handle Routines	889
47.6.1	ompd_get_curr_task_handle Routine	889
47.6.2	ompd_get_generating_task_handle Routine	890
47.6.3	ompd_get_scheduling_task_handle Routine	891
47.6.4	ompd_get_task_in_parallel Routine	891
47.6.5	ompd_get_task_function Routine	892
47.6.6	ompd_get_task_frame Routine	893
47.7	Handle Comparing Routines	894
47.7.1	ompd_parallel_handle_compare Routine	894
47.7.2	ompd_task_handle_compare Routine	895
47.7.3	ompd_thread_handle_compare Routine	896
47.8	Handle Releasing Routines	896
47.8.1	ompd_rel_address_space_handle Routine	897
47.8.2	ompd_rel_parallel_handle Routine	897
47.8.3	ompd_rel_task_handle Routine	898

47.8.4	<code>ompd_rel_thread_handle</code> Routine	898
47.9	Querying Thread States	899
47.9.1	<code>ompd_enumerate_states</code> Routine	899
47.9.2	<code>ompd_get_state</code> Routine	900
47.10	Display Control Variables	901
47.10.1	<code>ompd_get_display_control_vars</code> Routine	901
47.10.2	<code>ompd_rel_display_control_vars</code> Routine	902
47.11	Accessing Scope-Specific Information	903
47.11.1	<code>ompd_enumerate_icvs</code> Routine	903
47.11.2	<code>ompd_get_icv_from_scope</code> Routine	904
47.11.3	<code>ompd_get_icv_string_from_scope</code> Routine	905
47.11.4	<code>ompd_get_tool_data</code> Routine	906
48	OMP Breakpoint Symbol Names	908
48.1	<code>ompd_bp_thread_begin</code> Breakpoint	908
48.2	<code>ompd_bp_thread_end</code> Breakpoint	908
48.3	<code>ompd_bp_device_begin</code> Breakpoint	909
48.4	<code>ompd_bp_device_end</code> Breakpoint	909
48.5	<code>ompd_bp_parallel_begin</code> Breakpoint	909
48.6	<code>ompd_bp_parallel_end</code> Breakpoint	910
48.7	<code>ompd_bp_teams_begin</code> Breakpoint	911
48.8	<code>ompd_bp_teams_end</code> Breakpoint	911
48.9	<code>ompd_bp_task_begin</code> Breakpoint	912
48.10	<code>ompd_bp_task_end</code> Breakpoint	912
48.11	<code>ompd_bp_target_begin</code> Breakpoint	912
48.12	<code>ompd_bp_target_end</code> Breakpoint	913
VII	Appendices	914
A	OpenMP Implementation-Defined Behaviors	915
B	Features History	927
B.1	Deprecated Features	927
B.2	Version 6.0 to 6.1 Differences	928

B.3	Version 5.2 to 6.0 Differences	928
B.4	Version 5.1 to 5.2 Differences	935
B.5	Version 5.0 to 5.1 Differences	938
B.6	Version 4.5 to 5.0 Differences	941
B.7	Version 4.0 to 4.5 Differences	944
B.8	Version 3.1 to 4.0 Differences	946
B.9	Version 3.0 to 3.1 Differences	947
B.10	Version 2.5 to 3.0 Differences	948
C	Nesting of Regions	951
D	Conforming Compound Directive Names	953
	Index	957

List of Figures

38.1 First-Party Tool Activation Flow Chart [723](#)

List of Tables

3.1	ICV Scopes and Descriptions	118
3.2	ICV Initial Values	121
3.3	Ways to Modify and to Retrieve ICV Values	124
3.4	ICV Override Relationships	128
4.1	Predefined Place-list Abstract Names	131
4.2	Available Field Types for Formatting OpenMP Thread Affinity Information	141
4.3	Reservation Types for OMP_THREADS_RESERVE	145
5.1	Syntactic Properties for Clauses , Arguments and Modifiers	163
8.1	Implicitly Declared C/C++ Reduction Identifiers	243
8.2	Implicitly Declared Fortran Reduction Identifiers	244
8.3	Implicitly Declared C/C++ Induction Identifiers	245
8.4	Implicitly Declared Fortran Induction Identifiers	245
11.1	Map-Type Decay of Map Type Combinations	297
13.1	Predefined Memory Spaces	315
13.2	Allocator Traits	316
13.3	Predefined Allocators	319
18.1	Affinity-related Symbols used in this Section	408
19.1	work OMPT types for Worksharing-Loop	434
20.1	task_create Callback Flags Evaluation	447
26.1	Routine Argument Properties	555
26.2	Required Values of the interop_property OpenMP Type	562
26.3	Required Values for the interop_rc OpenMP Type	563
26.4	Allowed Key-Values for alloctrail OpenMP Type	567
26.5	Standard Tool Control Commands	586
38.1	OMPT Callback Interface Runtime Entry Point Names and Their Type Signatures	726
38.2	Callbacks for which set_callback Must Return ompt_set_always	728
38.3	OMPT Tracing Interface Runtime Entry Point Names and Their Type Signatures	730

41.1 Association of dev1 and dev2 arguments for target data operations	806
45.1 Mapping of Scope Type and OMPD Handles	856

1

Part I

2

Definitions

1 Overview of the OpenMP API

The collection of compiler [directives](#), library [routines](#), [environment variables](#), and [tool](#) support that this document describes collectively define the specification of the OpenMP Application Program Interface (OpenMP API) for C, C++ and Fortran [base 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](#), [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 [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, [data races](#), or deadlocks. [Compliant implementations](#) also are not required to check for any code sequences that cause a program to be classified as a [non-conforming program](#). 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 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 includes several **directives**. 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 **routine**. Execution is composed of **nested regions** since a given **region** may encounter additional **constructs** and **routines**. References to **regions**, particularly **explicit regions** or **nested regions**, that correspond to a specific type of **construct** or **routine** usually include the name of that **construct** or **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 multiple **devices** for a given execution of an **OpenMP program**. Concurrent execution 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 execute **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**.

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

A **task** that encounters a **target** construct generates a new **target task**; its **region** encloses the **target region**. The **target task** is complete after the **target region** completes execution. When a **target task** executes, an **initial thread** executes the enclosed **target region**. The **initial thread** executes sequentially, as if the **target region** is part of an **initial task region** that an **implicit parallel region** generates. The **initial thread** may execute on the requested **target device**, if it is available. If the **target device** does not exist or the implementation does not support it, all **target**

regions associated with that device execute on the host device. Otherwise, the implementation ensures that the **target region** executes as if it were executed in the **data environment** of the **target device** unless an **if clause** is present and the **if clause** expression evaluates to *false*.

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

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

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

When any **team** encounters a **partitioned worksharing construct**, the work inside the **construct** is divided into work partitions, each of which is executed by one member of the **team**, instead of the work being executed redundantly by each **thread**. An **implicit barrier** occurs at the end of any **region** that corresponds to a **worksharing construct** for which the **nowait** clause is not specified. Redundant execution of code by every **thread** in the **team** resumes after the end of the **worksharing construct**. Regions that correspond to **team-executed constructs**, including all **worksharing regions** and **barrier regions**, are executed by the current **team** such that all **threads** in the **team** execute the **team-executed regions** in the same order.

When a **loop** construct is encountered, the **logical iterations** of the **collapse-affected loops**, which are the **affected loops** as specified by the **collapse** clause, are executed in the context of its **encountering threads**, as determined according to its **binding region**. If the **loop** region binds to a **teams region**, the **region** is encountered by the set of **primary thread** that execute the **teams region**. If the **loop** region binds to a **parallel region**, the **region** is encountered by the **team** that execute the **parallel region**. Otherwise, the **region** is encountered by a single **thread**. If the

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

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

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

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

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

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

If a **task** encounters a **cancel** construct with any other *cancel-directive-name* clause, it activates **cancellation** of the innermost enclosing **region** of the type specified and the **thread** continues execution at the end of that **region**. **Tasks** check if **cancellation** has been activated for their **region** at **cancellation points** and, if so, also resume execution at the end of the canceled **region**.

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

OpenMP specifies circumstances that cause **error termination**. If **compile-time error termination** is specified, the effect is as if the program encounters an **error** directive on which a **severity** clause specifies a *sev-level* argument of **fatal** and an **at** clause specifies an *action-time* argument of **compilation**. If **runtime error termination** is specified, the effect is as if the program encounters an **error** directive on which a **severity** clause specifies a *sev-level* argument of **fatal** and an **at** clause specifies an *action-time* argument of **execution**.

A **construct** that creates a **data environment** creates it at the time that the **construct** is encountered. The description of a **construct** defines whether it creates a **data environment**. Synchronization **constructs** and **routines** are available in the OpenMP API to coordinate **tasks** and their data accesses. In addition, **routines** and **environment variables** are available to control or to query the runtime environment of **OpenMP programs**. The scope of OpenMP synchronization mechanisms may be limited to the **contention group** of the **encountering task**. Except where explicitly specified, any effect of the mechanisms between **contention groups** is **implementation defined**. Section 1.3 details the OpenMP **memory** model, including the effect of these features.

The OpenMP specification makes no guarantee that input or output to the same file is synchronous when executed in parallel. In this case, the programmer is responsible for synchronizing input and output processing with the assistance of **synchronization constructs** or **routines**.

Each **native thread** that enables the execution of a **task** by an **OpenMP thread** executes on a **hardware thread**. A **hardware thread** executes a stream of instructions defined by a given **task region**, so that only one **OpenMP thread** may execute on a **hardware thread** at a time. A set of consecutive **hardware threads** may form a **progress unit**. **Hardware threads** execute distinct streams of instructions unless they are part of the same **progress unit**. **Threads** that execute in the same **progress unit** may execute from a common stream of instructions, with serialized execution of **diverging code paths** that occur due to conditional statements. A program that relies on concurrent execution of such **diverging code paths** for the purposes of synchronization may deadlock.

All concurrency semantics defined by the **base language** with respect to **base language threads** apply to **OpenMP threads**, unless otherwise specified. An **OpenMP thread** *makes progress* when it performs a **flush** operation, performs input or output processing, terminates, or makes progress as defined by the **base language**. **OpenMP threads** will eventually make progress in the absence of dependence cycles, unless otherwise specified by the **base language**. A dependence cycle may be implicitly introduced between **synchronizing threads** where concurrent execution is not guaranteed. **Threads** may therefore not make progress if the program includes **synchronizing threads** that

descend from different [initial teams](#) formed by a [teams construct](#) or if the program includes [synchronizing divergent threads](#) from the same [team](#) that execute on the same [progress unit](#). The generation and execution of [explicit tasks](#) by [threads](#) in the current [team](#) does not prevent any of the [threads](#) from making progress if executing the [explicit tasks](#) as [included tasks](#) would ensure that they make progress.

Each [device](#) is identified by a [device number](#). The [device number](#) for the [host device](#) is the value of the total number of [non-host devices](#), while each [non-host device](#) has a unique [device number](#) that is greater than or equal to zero and less than the [device number](#) for the [host device](#). Additionally, the [predefined identifier](#) [omp_initial_device](#) can be used as an alias for the [host device](#), the [predefined identifier](#) [omp_default_device](#) can be used as an alias for specifying the value of the [default-device-var](#) ICV, and the [predefined identifier](#) [omp_invalid_device](#) can be used to specify an invalid [device number](#). A [conforming device number](#) is either a non-negative integer that is less than or equal to the value returned by [omp_get_num_devices](#) or equal to [omp_initial_device](#), [omp_default_device](#), or [omp_invalid_device](#).

A [signal handler](#) may only execute [directives](#) and [routines](#) that have the [async-signal-safe](#) property.

1.3 Memory Model

1.3.1 Structure of the OpenMP Memory Model

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

A [directive](#) that accepts [data-sharing attribute clauses](#) determines two kinds of access to [variables](#) used in the associated [structured block](#) of the [directive](#): [shared variables](#) and [private variables](#). Each [variable](#) referenced in the [structured block](#) has an [original variable](#), which is the [variable](#) by the same name that exists in the [OpenMP program](#) immediately outside the [construct](#). Each reference to a [shared variable](#) in the [structured block](#) becomes a reference to the [original variable](#). For each [private variable](#) referenced in the [structured block](#), a new version of the [original variable](#) (of the same type and size) is created in [memory](#) for each [task](#) or [SIMD lane](#) that executes code associated with the [directive](#). Creation of the new version does not alter the value of the [original variable](#). However, attempts to access the [original variable](#) from within the [region](#) that corresponds to the [directive](#) result in [unspecified behavior](#); see [Section 7.3.3](#) for additional details. References to a [private variable](#) in the [structured block](#) refer to the [private](#) version of the [original variable](#) for the

current [task](#) or [SIMD lane](#). The relationship between the value of the [original variable](#) and the initial or final value of the [private](#) version depends on the exact [clause](#) that specifies it. Details of this issue, as well as other issues with privatization, are provided in [Chapter 7](#).

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.

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

Two [memory](#) operations are considered unordered if the order in which they must complete, as seen by their affected [threads](#), is not specified by the [memory](#) consistency guarantees listed in [Section 1.3.6](#). If multiple [threads](#) write to the same [memory](#) unit (defined consistently with the above access considerations) then a [data race](#) occurs if the writes are unordered. Similarly, if at least one [thread](#) reads from a [memory](#) unit and at least one [thread](#) writes to that same [memory](#) unit then a [data race](#) occurs if the read and write are unordered. If a [data race](#) occurs then the result of the [OpenMP program](#) is [unspecified behavior](#).

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

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

1.3.2 Device Data Environments Overview

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

implicit **target_data** region of the **device**; whether this happens is **implementation defined**.

When an **original storage block** is mapped to a **device data environment** and a **corresponding storage block** is not present in the **device data environment**, a new **corresponding storage block** (of the same type and size as the **original storage block**) is created in the **device data environment**. Conversely, the **original storage block** becomes the **corresponding storage block** of the new **storage block** in the **device data environment** of the **device** that performs a **mapping operation**.

The **corresponding storage block** in the **device data environment** may share storage with the **original storage block**. Writes to the **corresponding storage block** may alter the value of the **original storage block**. Section 1.3.6 discusses the impact of this possibility on **memory consistency**. When a **task** executes in the context of a **device data environment**, references to the **original storage block** refer to the **corresponding storage block** in the **device data environment**. If an **original storage block** is not currently mapped and a **corresponding storage block** does not exist in the **device data environment** then accesses to the **original storage block** result in **unspecified behavior** unless the **unified_shared_memory** clause is specified on a **requires** directive for the **compilation unit**.

The relationship between the value of the **original storage block** and the initial or final value of the **corresponding storage block** depends on the *map-type*. Details of this issue, as well as other issues with mapping a **variable**, are provided in Section 11.3.

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

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

1.3.3 Memory Management

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

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

1.3.4 The Flush Operation

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

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

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

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

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

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

If a **flush** is a **release flush** or **acquire flush**, it can enforce consistency between the views of **memory** of two synchronizing **threads**. A **release flush** guarantees that any prior operation that writes or reads a **shared variable** will appear to be completed before any operation that writes or reads the same **shared variable** and follows an **acquire flush** with which the **release flush** synchronizes (see Section 1.3.5 for more details on **flush** synchronization). A **release flush** will propagate the values of all **shared variables** in its **temporary view** to **memory** prior to the **thread** performing any subsequent **atomic operation** that may establish a synchronization. An **acquire flush** will discard any value of a **shared variable** in its **temporary view** to which the **thread** has not written since last performing a **release flush**, and it will load any value of a **shared variable** propagated by a **release flush** that **synchronizes with** it (according to the **synchronizes-with relation**) into its **temporary view** so that it may be subsequently read. Therefore, **release flushes** and **acquire flushes** may also be used to guarantee that a value written to a **variable** by one **thread** may be read by a second **thread**. To accomplish this, the programmer must ensure that the second **thread** has not written to the **variable** 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 23 and in Chapter 34, 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.

1.3.5 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 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**.

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

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

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

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

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

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

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

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

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

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

6 1.3.6 OpenMP Memory Consistency

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

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

25 The flush operation can be specified using the flush directive, and is also implied at various
26 locations in an OpenMP program; see Section 23.8.6 for details.

27
28 Note – Since flushes by themselves cannot prevent data races, explicit flushes are only useful in
29 combination with non-sequentially consistent atomic constructs.
30

31 OpenMP programs that:

- 32 • 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 1.3.1](#),

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.

1.4 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.

1.4.1 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 *OpenMP program* 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 38.2.4](#) 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` routine for `tool` control, all other routines in the OMPT interface are intended for use only by `tools`. For that reason, OMPT includes a Fortran binding only for `omp_control_tool`; all other OMPT functionality is supported with C syntax only.

1.4.2 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 44, Chapter 45, Chapter 46, Chapter 47, and Chapter 48 describe OMPD in detail.

1.5 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 `routine` names, and it 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 42. All appendices as well as text designated as a note or comment (see Section 1.7) 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](#).

1.6 Normative References

- ISO/IEC 9899:1990, *Information Technology - Programming Languages - C*.
This OpenMP API specification refers to ISO/IEC 9899:1990 as C90.
- ISO/IEC 9899:1999, *Information Technology - Programming Languages - C*.
This OpenMP API specification refers to ISO/IEC 9899:1999 as C99.
- ISO/IEC 9899:2011, *Information Technology - Programming Languages - C*.
This OpenMP API specification refers to ISO/IEC 9899:2011 as C11.
- ISO/IEC 9899:2018, *Information Technology - Programming Languages - C*.
This OpenMP API specification refers to ISO/IEC 9899:2018 as C18.
- ISO/IEC 9899:2024, *Information Technology - Programming Languages - C*.
This OpenMP API specification refers to ISO/IEC 9899:2024 as C23.
- ISO/IEC 14882:1998, *Information Technology - Programming Languages - C++*.
This OpenMP API specification refers to ISO/IEC 14882:1998 as C++98.
- ISO/IEC 14882:2011, *Information Technology - Programming Languages - C++*.
This OpenMP API specification refers to ISO/IEC 14882:2011 as C++11.
- ISO/IEC 14882:2014, *Information Technology - Programming Languages - C++*.
This OpenMP API specification refers to ISO/IEC 14882:2014 as C++14.
- ISO/IEC 14882:2017, *Information Technology - Programming Languages - C++*.
This OpenMP API specification refers to ISO/IEC 14882:2017 as C++17.
- ISO/IEC 14882:2020, *Information Technology - Programming Languages - C++*.
This OpenMP API specification refers to ISO/IEC 14882:2020 as C++20.
- ISO/IEC 14882:2024, *Information Technology - Programming Languages - C++*.
This OpenMP API specification refers to ISO/IEC 14882:2024 as C++23.
- ISO/IEC 1539:1980, *Information Technology - Programming Languages - Fortran*.
This OpenMP API specification refers to ISO/IEC 1539:1980 as Fortran 77.
- ISO/IEC 1539:1991, *Information Technology - Programming Languages - Fortran*.
This OpenMP API specification refers to ISO/IEC 1539:1991 as Fortran 90.
- ISO/IEC 1539-1:1997, *Information Technology - Programming Languages - Fortran*.
This OpenMP API specification refers to ISO/IEC 1539-1:1997 as Fortran 95.

- ISO/IEC 1539-1:2004, *Information Technology - Programming Languages - Fortran*. This OpenMP API specification refers to ISO/IEC 1539-1:2004 as Fortran 2003.
- ISO/IEC 1539-1:2010, *Information Technology - Programming Languages - Fortran*. This OpenMP API specification refers to ISO/IEC 1539-1:2010 as Fortran 2008.
- ISO/IEC 1539-1:2018, *Information Technology - Programming Languages - Fortran*. This OpenMP API specification refers to ISO/IEC 1539-1:2018 as Fortran 2018.
- ISO/IEC 1539-1:2023, *Information Technology - Programming Languages - Fortran*. This OpenMP API specification refers to ISO/IEC 1539-1:2023 as Fortran 2023.
- Where this OpenMP API specification refers to C, C++ or Fortran, reference is made to the [base language](#) supported by the implementation.

1.7 Organization of this Document

The remainder of this document is structured as normative chapters that define the [directives](#), including their syntax and semantics, the [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++ ▲

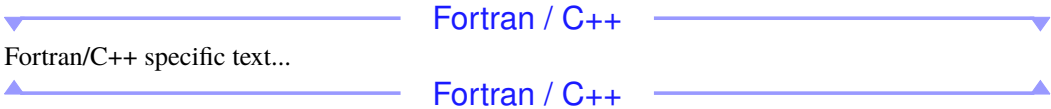
Text that applies only to programs for which the [base language](#) is Fortran is shown as follows:

▼ Fortran ▼

Fortran specific text...

▲ Fortran ▲

1 Text that applies only to programs for which the [base language](#) is Fortran or C++ is shown as
2 follows:

3  Fortran / C++
Fortran/C++ specific text...

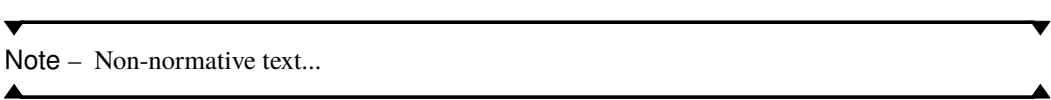
4 Where an entire page consists of [base language](#) specific text, a marker is shown at the top of the
5 page. For Fortran-specific text, the marker is:

 Fortran (cont.)

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

 C/C++ (cont.)

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

9 
Note – Non-normative text...

12 COMMENT: Non-normative text...

2 Glossary

A | B | C | D | E | F | G | H | I | L | M | N | O | P | R | S | T | U | V | W | Z

A

abstract name

A [conceptual abstract name](#) or a [numeric abstract name](#). [131](#), [35](#), [79](#), [131](#), [134](#), [137](#), [916](#), [929](#)

access group

A group of [tasks](#) that has the same value for the grouping criterion defined by an [access group type](#). [19](#), [329](#), [330](#), [683](#), [684](#), [919](#)

access group type

A criterion for grouping [tasks](#) using a specific grouping property. Each group of [tasks](#) classified under that criterion is called an [access group](#). All [tasks](#) in an [access group](#) have the same value for the grouping property of the [access group type](#). The [access group type](#) indicates the scope and access properties of [groupprivate variables](#) and [dynamic groupprivate blocks](#) in [constructs](#), [directives](#) and [routines](#). A [task](#) can belong to at most one [access group](#) for the same [access group type](#). [19](#), [57](#), [329](#), [564](#), [683](#), [684](#), [687](#)

accessible device

The [host device](#) or any [non-host device](#) accessible for execution. [122](#), [142–144](#), [377](#)

accessible storage

A [storage block](#) that may be accessed by a given [thread](#). [305](#), [626](#)

acquire flush

A [flush](#) that has the [acquire flush property](#). [10](#), [11](#), [12](#), [94](#), [104](#), [516](#), [519](#), [521–524](#)

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](#). [19](#), [54](#), [519](#)

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](#). [19](#), [77](#), [103](#), [132](#), [133](#), [136](#), [596](#), [916](#), [923](#), [944](#)

active parallel region

A [parallel region](#) comprised of [implicit tasks](#) that are being executed by a [team](#) to which multiple [threads](#) are assigned. [19](#), [108](#), [118](#), [119](#), [136](#), [275](#), [592](#), [596–600](#), [915](#), [918](#), [948](#), [950](#)

active target region

A [target region](#) that is executed on a [device](#) other than the [device](#) that encountered the [target](#) construct. [127](#)

address range

The addresses of a contiguous set of [storage locations](#). [52](#), [72](#), [101](#), [626](#)

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](#). [20](#), [53](#), [82](#), [97](#), [112](#), [148](#), [149](#), [376](#), [626](#), [724](#), [848](#), [859](#), [864](#), [866](#), [868–871](#), [875](#), [878](#), [879](#), [881](#), [882](#), [884](#), [900](#), [902](#), [903](#)

address space context

A [tool context](#) that refers to an [address space](#) within an [OpenMP process](#). [848](#)

address space handle

A [handle](#) that refers to an [address space](#) within an [OpenMP process](#). [856](#), [878–880](#), [886](#), [897](#)

affected iteration

A [logical iteration](#) of the [affected loops](#) of a [loop-nest-associated directive](#). [62](#), [96](#), [99](#), [401](#)

affected loop

A loop from a [canonical loop nest](#), or a **DO CONCURRENT** loop in Fortran, that is affected by a given [loop-nest-associated directive](#). [208](#), [4](#), [20](#), [64](#), [69](#), [70](#), [111](#), [117](#), [158](#), [208–210](#), [216–218](#), [226](#), [231](#), [232](#), [252](#), [259–261](#), [270](#), [271](#), [388](#), [389](#), [391–393](#), [396](#), [398](#), [399](#), [443](#), [920](#), [943](#)

affected loop nest

The subset of [canonical loop nests](#) of an [associated loop sequence](#) that are selected by the [looprange](#) clause. [213](#), [35](#), [94](#), [210](#), [388](#), [393](#)

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](#). [8](#), [20](#), [41](#), [115](#), [167](#), [224](#), [276](#), [307](#), [465](#), [915](#)

aligned-memory-allocating routine

A [memory-management routine](#) that has the [aligned-memory-allocating-routine property](#).
[675](#), [676](#), [678](#), [680](#)

aligned-memory-allocating-routine property

The [property](#) that a [memory-allocating routine](#) ensures the allocated [memory](#) is aligned with respect to an *alignment* argument. [675](#), [21](#), [678](#), [680](#)

all-constituents property

The [property](#) that a [clause](#) applies to all [leaf constructs](#) that permit it when the [clause](#) appears on a [compound directive](#). [163](#), [163](#), [548](#)

all-contention-group-tasks binding property

The [binding property](#) that the [binding task set](#) is [all tasks](#) in the [contention group](#). [554](#),
[689–691](#), [693](#), [695–700](#)

all-data-environments clause

A [clause](#) that has the [all-data-environments property](#). [75](#), [234](#), [237](#)

all-data-environments property

The [property](#) that a [data-sharing attribute clause](#) affects any [data environment](#) for which it is specified, including [minimal data environments](#). [21](#), [234](#), [236](#), [256](#)

all-device-tasks binding property

The [binding property](#) that the [binding task set](#) is [all tasks](#) on a specified [device](#). [713](#)

all-device-threads binding property

The [binding property](#) that the [binding thread set](#) is [all threads](#) on the [current device](#). The effect of executing a [construct](#) or a [routine](#) with this [property](#) is not related to any specific [region](#) that corresponds to any other [construct](#) or [routine](#). [555](#), [606](#), [615](#), [651–659](#), [661–667](#),
[669–672](#), [703](#), [704](#), [819](#), [820](#)

allocator

A [memory allocator](#). [21](#), [147](#), [316–327](#), [375](#), [483](#), [566](#), [567](#), [575](#), [576](#), [578](#), [660](#), [661](#),
[666–668](#), [674–677](#), [682](#), [919](#), [931](#), [932](#), [937](#)

allocator structured block

A [context-specific structured block](#) that may be associated with an [allocators](#) directive.
[192](#), [326](#)

allocator trait

A [trait](#) of an [allocator](#). [147](#), [316](#), [318](#), [319](#), [322](#), [324](#), [568](#), [570](#), [572](#), [660](#), [666](#), [667](#), [919](#), [931](#),
[932](#), [942](#)

all-privatizing property

The [property](#) that a [clause](#), when it appears on a [combined construct](#) or a [composite construct](#), applies to all [constituent constructs](#) to which it applies for which a [data-sharing attribute clause](#) may create a [private](#) copy of the same [list item](#). [163](#), [323](#), [548](#)

all tasks

All [tasks](#) participating in the [OpenMP program](#) or in a specified limiting context. [19](#), [21](#), [22](#), [29](#), [250](#), [279](#), [317](#), [329](#), [554](#), [714](#)

all-tasks binding property

The [binding property](#) that the [binding task set](#) is [all tasks](#). [714](#), [713](#), [714](#)

all threads

All [OpenMP threads](#) participating in the [OpenMP program](#). A specific usage of the term may be explicitly limited to a limiting context, such as [all threads](#) on a given [device](#) or an [OpenMP thread pool](#). [8](#), [13](#), [21](#), [22](#), [29](#), [232](#), [514](#), [555](#), [651](#), [715](#), [819](#), [820](#)

all-threads binding property

The [binding property](#) that the [binding thread set](#) is [all threads](#). The effect of executing a [construct](#) or a [routine](#) with this [property](#) is not related to any specific [region](#) that corresponds to any other [construct](#) or [routine](#). [555](#)

ancestor thread

For a given [thread](#), its [parent thread](#) or one of the [ancestor threads](#) of its [parent thread](#). [22](#), [598](#), [599](#), [609](#), [934](#), [949](#), [950](#)

antecedent task

A [task](#) that must complete before its [dependent tasks](#) can be executed. [527](#), [43](#), [52](#), [61](#), [88](#), [105](#), [458](#), [524](#), [527](#), [529](#), [788](#)

argument list

A [list](#) that is used as an argument of a [directive](#), [clause](#), or [modifier](#). [161](#), [47](#), [52](#), [65](#), [67](#), [82](#), [85](#), [88](#), [90](#), [111](#), [115](#), [162](#), [163](#), [165](#), [167](#), [215](#), [220](#), [272](#), [273](#), [277](#)

array base

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.

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 [array base](#) is: `(*p0).x0[k1].p1->p2[k2].x1[k3].x2`.

More examples for C/C++:

- The [array base](#) for $x[i]$ and for $x[i:n]$ is x , if x is an array or pointer.
- The [array base](#) 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 [array base](#) 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.

Examples for Fortran:

- The [array base](#) for $x(i)$ and for $x(i:j)$ is x .

[22](#), [23](#), [171](#), [172](#), [235–237](#), [246](#), [290](#), [300](#), [301](#)

array element

A single member of an array as defined by the [base language](#). [23](#), [71](#), [115](#), [236](#), [237](#), [239](#), [246](#), [259](#), [272](#), [273](#), [289](#), [294](#), [300](#), [310](#), [311](#)

array item

An array, an [array section](#), or an [array element](#). [549](#)

array section

A designated subset of the elements of an array that is specified using a subscript notation that can select more than one [array element](#). [22–24](#), [26–28](#), [36](#), [37](#), [39](#), [76](#), [100](#), [103](#), [115](#), [117](#), [143](#), [166](#), [169–172](#), [222](#), [234–237](#), [239](#), [240](#), [242](#), [246](#), [259](#), [272](#), [273](#), [293](#), [300–302](#), [308](#), [310](#), [414](#), [464](#), [529](#), [549](#), [929](#), [938](#), [941–943](#), [945](#), [947](#)

array shaping

A mechanism that reinterprets the region of memory to which an expression that has a type of pointer to T as an n -dimensional array of type T . [97](#), [941](#)

assignable OpenMP type instance

An instance of an [OpenMP type](#) to which an assignment can be performed. [188](#), [188](#)

assigned list item

A [list item](#) to which assignment is performed as the result of a [data-motion clause](#). [311–313](#)

assigned thread

A [thread](#) that has been assigned an [implicit task](#) of a [parallel region](#). [3](#), [4](#), [89](#), [107](#), [109](#), [408](#), [409](#), [433](#), [590](#)

assigning map type

A [map-type](#) for which the [mapping operations](#) may include an assignment operation. [296](#)

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](#). [7](#), [24](#)

associated loop nest

The associated [canonical loop nest](#), or **DO CONCURRENT** loop in Fortran, of a [loop-nest-associated directive](#). [69](#), [70](#), [208](#), [211–213](#), [388](#), [391](#)

associated loop sequence

The associated [canonical loop sequence](#) of a [loop-sequence-associated directive](#). [20](#), [213](#), [388](#), [393](#)

associated memory space

The [associated memory space](#) of a [memory allocator](#) is the [memory space](#) that is specified when the [memory allocator](#) is created. [24](#), [73](#), [316](#), [319](#)

assumed-size array

For C/C++, an [array section](#) for which the *length* is absent and the size of the dimensions is not known. For Fortran, an [assumed-size array](#) in the [base language](#). [24](#), [72](#), [103](#), [117](#), [170](#), [172](#), [203](#), [217](#), [218](#), [223](#), [234](#), [237](#), [293](#), [294](#), [296](#), [300](#), [555](#), [931](#), [949](#)

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. [24](#), [379](#), [382](#), [936](#), [939](#)

assumption scope

The scope for which the invariants specified by an [assumption directive](#) must hold. [379–386](#)

asynchronous device routine

A [routine](#) that has the [asynchronous-device routine property](#). [525](#), [623](#), [624](#), [638](#), [639](#), [642](#)

asynchronous-device routine property

The [property](#) of a [device routine](#) that it performs its operation asynchronously. [24](#), [624](#), [637](#), [638](#), [641](#)

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](#). [24](#), [25](#), [770](#), [803](#), [813](#)

async-signal-safe entry point

An [entry point](#) that has the [async-signal-safe property](#). [813](#)

async-signal-safe property

The [property](#) of a [routine](#) or [entry point](#) that it is [async signal safe](#). [7](#), [24](#), [813](#), [818](#), [819](#), [821–823](#), [825](#), [827–829](#)

atomic captured update

An [atomic update](#) that captures the read value into a separate [storage location](#). [35](#), [112](#), [114](#), [198](#), [511](#), [515](#), [947](#)

atomic conditional update

An [atomic update](#) for which the write to the [storage location](#) is conditionally applied. [35](#), [112](#), [114](#), [196](#), [511](#), [512](#), [515–518](#), [940](#)

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](#). [8](#), [11–13](#), [25](#), [91](#), [94](#), [97](#), [303](#), [319](#), [492](#), [516–518](#), [522](#), [940](#)

atomic read

An [atomic operation](#) that reads from but does not write to a given [storage location](#). [91](#), [195](#), [508](#), [515](#)

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](#). [8](#), [12](#), [319](#), [514](#)

atomic structured block

A [context-specific structured block](#) that may be associated with an [atomic directive](#). [193](#), [31](#), [91](#), [114](#), [117](#), [193](#), [198](#), [514–516](#), [930](#)

atomic update

An [atomic operation](#) that reads from and writes to a given [storage location](#). [25](#), [112](#), [114](#), [195](#), [509](#), [511](#), [515](#), [516](#), [518](#), [947](#)

atomic write

An [atomic operation](#) that writes to but does not read from a given [storage location](#). [117](#), [195](#), [510](#), [515](#)

attached pointer

A pointer [variable](#) or [referring pointer](#) in a [device data environment](#) that, as a result of a [mapping operation](#), points to a given [data entity](#) that also exists in the [device data environment](#). [87](#), [299](#), [304](#), [311](#), [484](#)

attach-ineligible

An attribute of a pointer for which [pointer attachment](#) may not be performed. [298](#), [299](#), [301](#)

automatic storage duration

For C/C++, the lifetime of a [variable](#) or object with automatic storage duration, as defined by the [base language](#). For Fortran, the lifetime of a [variable](#), including implied-do, **FORALL**, and **DO CONCURRENT** indices, that is neither a [variable](#) that has [static storage duration](#) nor a dummy argument without the **VALUE** attribute. For [referencing variables](#), this refers to the lifetime of the [referring pointer](#) unless explicitly specified otherwise. [216](#), [219](#), [221](#)

available device

An [available non-host device](#) or, where explicitly specified, the [host device](#). [142](#), [144](#), [336](#), [656](#), [673](#), [714](#)

available non-host device

A [non-host device](#) that can be used for the current [OpenMP program](#) execution. [26](#), [142](#)

B

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](#). [4](#), [6](#), [26](#), [51](#), [60](#), [287](#), [403](#), [421](#), [423–425](#), [427](#), [428](#), [433](#), [468](#), [495–497](#), [502](#), [516](#), [520–522](#), [541](#), [610](#), [713](#), [728](#), [758](#), [759](#), [789](#), [790](#), [934](#), [951](#)

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 [referenced pointee](#), the address of the first [storage location](#) of its [referenced pointee](#); 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](#). [52](#), [234](#), [237](#), [300](#), [631](#), [685](#)

base array

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](#). 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](#).

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 array](#) is: `(*p0).x0[k1].p1->p2[k2].x1[k3].x2`.

22, 26, 549

base function

A [procedure](#) that is declared and defined in the [base language](#). 42, 55, 94, 116, 339, 346–350, 352–354, 919

base language

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

[Section 1.6](#) lists the current [base languages](#) for the OpenMP API. 2, 3, 6, 8, 12, 13, 15–18, 23, 24, 26–28, 38, 39, 41, 46, 49, 51, 54, 55, 57, 83, 88–90, 95, 96, 101, 102, 104, 151, 155, 157, 159, 160, 165, 167, 168, 170, 171, 173, 188–190, 194, 201, 206, 208, 221, 236–238, 241, 246, 248, 258, 263, 264, 267, 274, 291, 301, 308, 309, 319, 320, 322, 326, 327, 348, 352, 354, 379, 430, 515, 536, 553, 555, 584, 585, 915, 936, 941

base-language code

Code that is part of the [base program](#), exclusive of any [directives](#) or [implementation code](#). 29, 42, 49, 64, 112, 151, 156, 188, 200, 937

base language thread

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, 27

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 storage of the data object referenced by the lvalue expression or [array section](#) is part of the storage of the [implicit array](#) for that lvalue pointer expression. For Fortran, the last data pointer that appears in the designator for a given [variable](#) or [array section](#), where the storage of the data object referenced by the [variable](#) or [array section](#) is part of the storage 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 *pi* have a pointer type declaration and identifiers *xi* have an array type declaration, the [base pointer](#) is: (*p0).x0[k1].p1->p2.

22, 26–28, 38, 76, 217, 237, 259, 294, 298, 299, 301, 305, 345, 456, 457, 482, 483, 548, 549, 928

base program

A program written in a [base language](#). 2, 27, 82

base referencing variable

For C++, a [referencing variable](#) that is used by a given lvalue expression or [array section](#) to refer indirectly to its storage, where the storage of the data object referenced by the lvalue

expression or [array section](#) is part of the storage of the [referenced pointee](#) of the [referencing variable](#). For Fortran, the last [referencing variable](#) that appears in the designator for a given [variable](#) or [array section](#), where the storage of the data object referenced by the [variable](#) or [array section](#) is part of the storage of the [referenced pointee](#) of the [referencing variable](#). [28](#), [217](#), [298](#), [482](#)

base referring pointer

For a given [data entity](#) that has a [base referencing variable](#), the [referring pointer](#) of the [base referencing variable](#). [217](#), [299](#), [928](#)

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` are arrays, all have the [base variable](#) `x`.
- The data objects `p(i)`, `p(:n)`, `p(i)%y(j)` and `p(i)%y(:n)`, where `p` is a pointer and `p(i)%y` is an array, 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](#).

[26](#), [28](#), [276](#), [289](#), [305](#), [456](#), [457](#), [483](#), [548](#), [549](#)

binding implicit task

The [implicit task](#) of the [current team](#) assigned to the [encountering thread](#). [28](#), [29](#), [59](#), [127](#), [407](#), [674](#), [675](#)

binding-implicit-task binding property

The [binding property](#) that the [binding task set](#) is the [binding implicit task](#). [673](#)–[675](#)

binding property

A [property](#) of a [construct](#) or a [routine](#) that determines the [binding region](#), [binding task set](#) and/or [binding thread set](#). [21](#), [22](#), [28](#), [49](#), [50](#), [55](#), [554](#)

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](#). [4](#), [28](#), [29](#), [84](#), [210](#), [431](#), [442–445](#), [495](#), [534](#), [536](#), [537](#), [540](#), [544](#), [555](#), [707](#), [709](#), [910](#), [911](#), [913](#), [924](#), [925](#), [952](#)

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](#). [21](#), [22](#), [28](#), [29](#), [55](#), [124](#), [355](#), [455](#), [474](#), [476](#), [478](#), [481](#), [486](#), [488](#), [498](#), [502](#), [554](#), [623](#), [674](#), [675](#), [714](#), [813](#), [910–913](#)

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](#). [5](#), [21](#), [22](#), [28](#), [29](#), [49](#), [50](#), [84](#), [86](#), [94](#), [110](#), [116](#), [210](#), [229](#), [232](#), [402](#), [412](#), [416](#), [417](#), [421](#), [423–425](#), [427](#), [428](#), [431–433](#), [440](#), [442–446](#), [449](#), [450](#), [455](#), [459](#), [466](#), [493](#), [495](#), [499](#), [502](#), [514–516](#), [519](#), [526](#), [534](#), [536](#), [540](#), [541](#), [544](#), [554](#), [555](#), [651](#), [707](#), [709](#), [813](#), [819](#), [820](#), [924](#), [925](#), [933](#), [934](#)

block-associated directive

A [directive](#) for which its associated [base-language code](#) is a [structured block](#). [156](#), [38](#), [84](#), [154](#), [156–158](#), [191](#), [326](#), [355](#), [385](#), [402](#), [412](#), [421](#), [424–426](#), [428](#), [431](#), [446](#), [455](#), [478](#), [481](#), [493](#), [498](#), [514](#), [536](#)

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. [207](#)

C

callback

A [tool callback](#). [xxviii](#), [14](#), [15](#), [30](#), [33](#), [46](#), [74–76](#), [79–81](#), [83](#), [85](#), [87](#), [88](#), [94](#), [104](#), [113](#), [114](#), [249](#), [294](#), [363](#), [369](#), [404](#), [413](#), [421](#), [425–428](#), [430](#), [432](#), [434](#), [440](#), [447](#), [451](#), [467](#), [469](#), [473](#), [475](#), [477](#), [479](#), [482](#), [483](#), [487](#), [494–498](#), [500](#), [517](#), [520](#), [529](#), [533](#), [535](#), [536](#), [542](#), [610](#), [623](#), [624](#), [628](#), [629](#), [631](#), [632](#), [634–636](#), [638–642](#), [689–695](#), [697–701](#), [719](#), [721](#), [722](#), [724](#), [725](#), [727–732](#), [745](#), [750](#), [755](#), [756](#), [763](#), [770–807](#), [809](#), [811](#), [813](#), [814](#), [817](#), [818](#), [830](#), [831](#),

833–836, 838, 840, 844, 845, 849, 850, 854, 861–871, 873, 875, 877, 880, 882, 899, 902, 903, 905, 925–927, 935, 940, 941

callback dispatch

The processing of 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](#). [30](#), [755](#), [835](#)

callback registration

A process that makes a [tool callback](#) available to an OpenMP implementation to enable [callback dispatch](#). [94](#), [725](#), [727](#)

canceled taskgroup set

A [taskgroup set](#) that has been canceled. [541](#), [541](#)

cancellable construct

A [construct](#) that has the [cancellable property](#). [539](#), [540](#), [544](#)

cancellable property

The [property](#) that a [construct](#) may be subject to [cancellation](#). [539](#), [30](#), [402](#), [426](#), [435](#), [436](#), [498](#)

cancellation

An action that cancels (that is, aborts) a [region](#) and causes the execution of [implicit tasks](#) or [explicit tasks](#) to proceed to the end of the canceled [region](#). [541](#), [5](#), [6](#), [26](#), [30](#), [142](#), [423](#), [495](#), [496](#), [521](#), [524](#), [539–544](#), [712](#), [785](#), [946](#)

cancellation point

A point at which [implicit tasks](#) and [explicit tasks](#) check if [cancellation](#) has been activated. If [cancellation](#) has been activated, they perform the [cancellation](#). [540](#), [5](#), [6](#), [114](#), [119](#), [142](#), [469](#), [495](#), [496](#), [521](#), [524](#), [541–544](#), [767](#)

candidate

A [replacement candidate](#). [341](#), [346](#)

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. [746](#)

canonical loop nest

A loop nest that complies with the rules and restrictions defined in [Section 6.4.1](#). [201](#), [20](#), [24](#), [29](#), [31](#), [55](#), [68–70](#), [78](#), [157](#), [202](#), [206–208](#), [211](#), [213](#), [231](#), [269](#), [386](#), [388](#), [391–393](#), [397](#), [399](#), [400](#), [438](#), [551](#), [933](#), [941](#)

canonical loop sequence

A sequence of [canonical loop nests](#) that complies with the rules and restrictions defined in [Section 6.4.2](#). [207](#), [24](#), [55](#), [69](#), [70](#), [157](#), [202](#), [208](#), [213](#), [388](#), [389](#), [393](#), [396](#), [930](#), [932](#)

capture structured block

An [atomic structured block](#) that may be associated with an [atomic directive](#) that expresses capture semantics. [197](#), [35](#), [114](#), [197](#), [517](#)

C/C++-only property

The [property](#) that an OpenMP feature or a [routine](#) argument is only supported in C/C++. [555](#), [556](#), [708](#), [710](#), [718](#), [734–737](#), [739–757](#), [759–762](#), [764–769](#), [771–778](#), [780–783](#), [785–790](#), [792–796](#), [798–803](#), [806](#), [808](#), [810](#), [813–823](#), [825](#), [827–829](#), [831–842](#), [847](#), [848](#), [850–860](#)

C/C++ pointer property

The [property](#) that a [routine](#) argument has a pointer type in C/C++ but not the **POINTER** attribute in Fortran. [555](#), [575](#), [576](#), [594](#), [659](#), [661–665](#), [684](#), [689–699](#), [701](#), [718](#)

child task

A [task](#) is a [child task](#) of its [generating task region](#). The [region](#) of a [child task](#) is not part of its [generating task region](#), unless the [child task](#) is an [included task](#). [31](#), [43](#), [52](#), [61](#), [98](#), [106](#), [111](#), [499](#), [522](#), [528](#), [531](#), [579](#)

chunk

A contiguous non-empty subset of the [collapsed iterations](#) of a [loop-collapsing construct](#). [96](#), [137](#), [433](#), [434](#), [437–440](#), [442](#), [449](#), [551](#), [595](#), [744](#), [780](#), [925](#)

class type

For C++, the type of any [variable](#) declared with one of the **class**, **struct**, or **union** keywords. [221](#), [222](#), [228](#), [229](#), [231](#), [232](#), [243](#), [248](#), [254](#), [258](#), [276](#), [285–288](#), [302](#), [304](#), [484](#)

clause

A mechanism to specify customized [directive](#) behavior. [xxviii](#), [4–6](#), [8](#), [9](#), [20–22](#), [32](#), [34](#), [36](#), [39–45](#), [47–51](#), [53](#), [55](#), [56](#), [58](#), [59](#), [63](#), [70–73](#), [75](#), [78](#), [79](#), [81–84](#), [88](#), [89](#), [92–97](#), [103](#), [104](#), [106](#), [112](#), [113](#), [119](#), [122](#), [125](#), [127–130](#), [132](#), [135](#), [147](#), [151–153](#), [155](#), [156](#), [161–165](#), [167–169](#), [171–175](#), [178](#), [179](#), [184–187](#), [208](#), [209](#), [211–213](#), [215–229](#), [231–239](#), [242](#), [246–248](#), [250–263](#), [265–276](#), [279](#), [281–291](#), [293–313](#), [315](#), [320–332](#), [336](#), [338](#), [339](#), [341–365](#), [367–384](#), [386–401](#), [403](#), [405–407](#), [411–425](#), [427–429](#), [433](#), [434](#), [437–447](#), [449](#), [450](#), [452–465](#), [470–479](#), [481–484](#), [486–492](#), [494](#), [499–522](#), [524–539](#), [541–543](#), [548–551](#), [555](#), [581](#), [588](#), [590](#), [603](#), [606](#), [610](#), [619](#), [620](#), [624](#), [628](#), [631](#), [667](#), [668](#), [674–677](#), [686](#), [702](#), [740](#), [742](#), [767](#), [774](#), [786](#), [787](#), [809](#), [918–922](#), [928–934](#), [936–949](#), [951](#), [952](#)

clause set

A set of [clauses](#) for which restrictions on their use or other [properties](#) of their use on a given [directive](#) are specified. [163](#), [32](#), [34](#), [51](#), [94](#), [113](#), [164](#), [215](#), [373](#), [380](#), [450](#)

clause group

A [clause set](#) for which restrictions or properties related to their use on all [directives](#) are specified. [161](#), [164](#), [361](#), [373](#), [380](#), [504](#), [508](#), [510](#), [537](#), [539](#), [932](#)

clause-list trait

A [trait](#) that is defined with [properties](#) that match the [clauses](#) that may be specified for a given [directive](#). [335](#), [336](#), [338](#)

closely nested construct

A [construct](#) nested inside another [construct](#) with no other [construct](#) nested between them. [430](#), [432](#), [444](#), [542](#), [544](#)

closely nested region

A [region](#) nested inside another [region](#) with no [parallel region](#) nested between them. [86](#), [257](#), [423](#), [445](#), [542](#), [544](#), [948](#)

code block

A contiguous region of [memory](#) that contains code of an [OpenMP program](#) to be executed on a [device](#). [473](#)

collapse-affected loop

For a [loop-collapsing construct](#), a loop that is affected by the [collapse](#) clause. [4](#), [32](#), [69](#), [106](#), [209](#), [210](#), [220](#), [266](#), [418](#), [433](#), [438](#), [439](#), [442](#), [443](#), [453](#), [454](#), [537](#), [933](#)

collapsed iteration space

The [logical iteration space](#) of the [collapse-affected loops](#) of a [loop-collapsing construct](#). [209](#), [32](#), [266](#), [270](#), [419](#), [434](#), [438](#), [440](#), [442](#)

collapsed iteration

A [logical iteration](#) of the [collapse-affected loops](#) of a [loop-collapsing construct](#). [31–33](#), [35](#), [62](#), [69](#), [96](#), [117](#), [210](#), [220](#), [243](#), [258](#), [260–262](#), [269–271](#), [416](#), [417](#), [420](#), [423](#), [433](#), [434](#), [437–440](#), [442](#), [443](#), [449](#), [523](#), [536](#), [537](#), [551](#), [779](#), [780](#)

collapsed logical iteration

A [collapsed iteration](#). [210](#), [220](#)

collapsed loop

A loop resulting from a [loop-collapsing construct](#) that iterates over the [collapsed iteration space](#). [209](#), [260](#), [918](#)

collective step expression

An expression in terms of a [step expression](#) and a [collector](#) that eliminates recursive calculation in an [induction operation](#). [62](#), [33](#), [243](#)

collector

A binary operator used to eliminate recursion in an [induction operation](#). [62](#), [33](#), [269](#)

collector expression

An [OpenMP stylized expression](#) that evaluates to the value of the [collective step expression](#) of a [collapsed iteration](#). [243](#), [62](#), [178](#), [243](#), [245](#), [267](#), [269](#)

combined construct

A [construct](#) that is a shortcut for specifying one [construct](#) immediately nested inside a [leaf construct](#). [550](#), [22](#), [33](#), [34](#), [944](#), [945](#)

combined directive

A [compound directive](#) that is used to form a [combined construct](#). [33](#), [35](#), [545](#)

combined-directive name

The name of a [combined directive](#). [545](#)

combiner

A binary operator used by a [reduction operation](#). [248](#), [92](#), [188](#), [252](#)

combiner expression

An [OpenMP stylized expression](#) that specifies how a [reduction](#) combines partial results into a single value. [239](#), [92](#), [178](#), [239](#), [240](#), [247](#), [251](#), [263](#), [265](#), [270](#), [927](#)

common-field property

The [property](#) that a field has a name that is used in more than one [OpenMP type](#), or in more than one [OMPD type](#), or in more than one [OMPT type](#). [751](#), [752](#)

common-type-callback property

The [property](#) that a [callback](#) has a type that at least one other [callback](#) has. [789](#), [790](#), [792–794](#), [866](#), [872](#)

compatible context selector

A [context selector](#) that matches the [OpenMP context](#) in which a [directive](#) is encountered. [340](#), [340–342](#), [346](#)

compatible map type

A [map-type](#) that is consistent with the [data-motion attribute](#) of a given [data-motion clause](#). [310](#), [312](#), [313](#)

compatible property

The [property](#) that a [clause](#), an argument, a [modifier](#), or a [clause set](#) does not have the [exclusive property](#). [163](#)

compilation unit

For C/C++, a translation unit. For Fortran, a program unit. [9](#), [44](#), [158](#), [277–280](#), [282](#), [322](#), [323](#), [325](#), [369](#), [372–374](#), [378](#), [385](#), [483](#), [628](#), [667](#), [668](#), [677](#), [918](#)

compile-time error termination

[Error termination](#) that is performed during compilation. [6](#), [373](#), [407](#), [920](#)

complete tile

A [tile](#) that has $\prod_k s_k$ [logical iterations](#), where s_k are the [list items](#) of the [sizes clause](#) on the [construct](#). [399](#), [86](#)

complex modifier

A [modifier](#) that may take at least one argument when it is specified. [162](#), [34](#), [162](#), [165](#), [173](#)

complex property

The [property](#) that a [modifier](#) is a [complex modifier](#). [185](#), [490](#)

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](#). [15](#), [2](#), [5](#), [15–17](#), [34](#), [42](#), [59](#), [113](#), [139](#), [151](#), [438](#), [439](#), [516](#), [553](#), [688](#), [721](#), [814](#), [844](#), [845](#), [922](#)

composite construct

A [construct](#) that is a shortcut for composing a series or nesting of multiple [constructs](#), but that does not have the semantics of a [combined construct](#). [22](#), [269](#), [297](#), [551](#), [931](#), [934](#)

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 one [directive](#) inside the other that would otherwise be [non-conforming](#). If the [composite directive](#) adds semantics not included in its [constituent directives](#), the effects of the [constituent directives](#) may occur either as a nesting of the [directives](#) or as a sequence of the [directives](#). [34](#), [35](#), [99](#), [478](#), [546](#)

composite-directive name

The [directive name](#) of a [composite directive](#). [545](#)

compound construct

A [construct](#) that corresponds to a [compound directive](#). [35](#), [63](#), [81](#), [84](#), [98](#), [179](#), [184](#), [253](#), [335](#), [537](#), [547–551](#), [930](#), [946](#), [952](#), [953](#)

compound directive

A [combined directive](#) or a [composite directive](#). [21](#), [33](#), [35](#), [36](#), [66](#), [163](#), [545](#), [546](#), [548](#)

compound-directive name

The [directive name](#) of a [compound directive](#). [545](#), [47](#), [545](#), [547](#), [934](#), [953](#)

compound target construct

A [compound construct](#) for which [target](#) is a [constituent construct](#). [289](#), [290](#), [549](#)

conceptual abstract name

An [abstract name](#) that refers to an [implementation defined](#) abstraction that is relevant to the execution model described by this specification. [131](#), [19](#), [79](#), [87](#), [131](#)

conditional-update-capture structured block

A [capture structured block](#) that may be associated with an [atomic directive](#) that expresses an [atomic conditional update](#) operation with capture semantics. [197](#), [197](#), [198](#), [517](#), [927](#)

conditional-update structured block

An [update structured block](#) that may be associated with an [atomic directive](#) that expresses an [atomic conditional update](#) operation that is not also an [atomic captured update](#) operation. [196](#), [195–197](#), [517](#)

conforming device number

A [device number](#) that may be used in a [conforming program](#). [7](#), [144](#), [316](#), [338](#), [339](#), [472](#), [567](#), [612](#), [619–623](#), [652](#), [669](#), [687](#), [714](#), [928](#)

conforming program

An [OpenMP program](#) that follows all rules and restrictions of the OpenMP specification. [2](#), [15](#), [34](#), [35](#), [78](#), [81](#), [113](#), [341](#), [388](#), [438](#)

C-only property

The [property](#) that an OpenMP feature is only supported in C. [721](#), [737](#), [848](#), [853](#), [855](#), [857](#), [862](#), [863](#), [865–868](#), [870–878](#), [880–906](#)

consistent schedules

The [loop schedules](#) of two [affected loop nests](#) are [consistent](#) if for each assignment of a [thread](#) to a [collapsed iteration](#) that results from the schedule of one loop nest, the behavior is as if the same [thread](#) is assigned to the corresponding [collapsed iteration](#) of the other loop nest. [210](#), [35](#), [210](#), [399](#), [423](#)

constant property

The [property](#) that an expression, including one that is used as the argument of a [clause](#), a [modifier](#) or a [routine](#), is a compile-time constant. [165](#), [54](#), [95](#), [155](#), [163](#), [166](#), [186–188](#), [209–213](#), [273](#), [315](#), [320](#), [322](#), [324](#), [328](#), [339](#), [361](#), [368](#), [371](#), [374–379](#), [382–384](#), [390](#), [394](#), [395](#), [400](#), [401](#), [419](#), [420](#), [459–461](#), [463](#), [492](#), [504–512](#), [538](#), [539](#), [554](#), [932](#)

constituent construct

For a given [construct](#), a [construct](#) that corresponds to one of the [constituent directives](#) of the [executable directive](#). [22](#), [35](#), [81](#), [98](#), [179](#), [184](#), [253](#), [380](#), [534](#), [536](#), [547–549](#), [934](#)

constituent directive

For a given [directive](#) and its set of [leaf directives](#), a [leaf directive](#) in the set or a [compound directive](#) that is a shortcut for composing two or more members of that set for which the [directive names](#) are consecutively listed. [34](#), [36](#), [163](#), [179](#), [297](#), [478](#), [479](#), [548](#), [551](#), [930](#)

constituent-directive name

The [directive name](#) of a [constituent directive](#). [545](#), [545](#), [551](#), [953](#)

construct

An [executable directive](#), 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](#). [2–7](#), [15](#), [19–22](#), [25](#), [28](#), [30](#), [32–37](#), [40](#), [43–46](#), [48–56](#), [58–61](#), [63–66](#), [70](#), [71](#), [76–79](#), [81](#), [83–86](#), [88](#), [89](#), [92–99](#), [102–110](#), [113](#), [114](#), [116](#), [117](#), [119](#), [120](#), [123](#), [125](#), [127](#), [128](#), [136](#), [137](#), [142](#), [153](#), [156](#), [159](#), [165](#), [167](#), [173](#), [175](#), [179](#), [184–188](#), [197](#), [198](#), [206](#), [209](#), [210](#), [212](#), [215–221](#), [223–226](#), [228](#), [229](#), [231–237](#), [239](#), [247](#), [250](#), [252–254](#), [256–258](#), [260](#), [261](#), [266](#), [269–271](#), [275](#), [276](#), [287](#), [289](#), [290](#), [293](#), [294](#), [297–302](#), [304–307](#), [310](#), [320](#), [321](#), [324–328](#), [330](#), [335](#), [345](#), [349](#), [351](#), [355–359](#), [374–376](#), [380](#), [381](#), [383](#), [390–404](#), [406](#), [412–417](#), [420–433](#), [435–437](#), [439](#), [440](#), [442–447](#), [449–451](#), [453–466](#), [469–479](#), [481–484](#), [486–489](#), [492–496](#), [498–500](#), [502–517](#), [519–526](#), [528](#), [529](#), [531–537](#), [539–545](#), [547–551](#), [581](#), [603–605](#), [621–623](#), [716](#), [722](#), [729](#), [730](#), [744](#), [750](#), [759](#), [767](#), [771](#), [775](#), [779](#), [783–786](#), [796](#), [798](#), [809](#), [856](#), [857](#), [910](#), [911](#), [920–922](#), [930–934](#), [936–949](#), [951–953](#)

construct selector set

A [selector set](#) that may match the [construct trait set](#). [338](#), [335](#), [338–340](#), [346](#), [347](#)

construct trait set

The [trait set](#) that, at a given point in an [OpenMP program](#), consists of all enclosing [constructs](#) up to an enclosing [target construct](#). [335](#), [36](#), [37](#), [335](#), [336](#), [338](#), [340](#), [358](#)

containing array

For C/C++, a non-subscripted array (a [containing array](#)) to which a series of zero or more array subscript operators and, when specified, dot (i.e., '.') operators are applied to yield a given lvalue expression or [array section](#) for which storage is contained by the array. 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 selectors and, when specified, component selectors are applied to yield a given [variable](#) or [array section](#) for which storage is contained by the array.

COMMENT: An array is a [containing array](#) of itself. 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 [containing arrays](#) are: $(*p0).x0[k1].p1 \rightarrow p2[k2].x1$ and $(*p0).x0[k1].p1 \rightarrow p2[k2].x1[k3].x2$.

[26, 28, 36, 37, 169, 294, 302](#)

containing structure

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](#). 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](#).

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

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 [containing structures](#) are: $(*p0).x0[k1].p1$, $(*p0).x0[k1].p1.p2[k2]$ and $(*p0).x0[k1].p1.p2[k2].x1[k3]$

[28, 37, 217, 294, 298, 300, 302, 303](#)

contention group

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](#). [3–6, 21, 29, 66, 83, 96, 102, 119, 120, 133, 137, 145, 279, 317, 329, 377, 405, 411, 473, 493, 514, 554, 564, 591, 604, 605, 621, 622, 683, 688, 922, 931, 939](#)

context-matching construct

A [construct](#) that has the [context-matching property](#). [338](#)

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](#). [37, 355, 402, 412, 417, 435, 436, 481](#)

context selector

The specification of [traits](#) that a [directive variant](#) or [function variant](#) requires in the current [OpenMP context](#) in order for that variant to be selected. [337, 33, 48, 49, 100, 337, 339–342,](#)

346–348, 352–354, 372, 919, 939

context-specific structured block

Structured blocks that conform to specific syntactic forms and restrictions that are required for certain block-associated directives. 191, 21, 25, 55, 192, 193

core

A physically indivisible hardware execution unit on a device onto which one or more hardware threads may be mapped via distinct execution contexts. 65, 78, 100, 131, 751

corresponding list item

For a privatization clause, a new list item that derives from an original list item. For a data-mapping attribute clause, a list item in a device data environment that corresponds to an original list item. 70, 71, 232, 237, 282, 287, 288, 293–295, 299, 301–305, 310–312, 324, 327, 363, 378, 481, 486, 631, 931

corresponding pointer

For a given pointer variable or a given referring pointer, the corresponding variable or handle that exists in a device data environment. 84, 299

corresponding pointer initialization

For a given data entity that has a base pointer or referring pointer, an assignment to the base pointer or referring 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. 299, 482

corresponding storage

For a given storage block, its corresponding storage block. For a given mapped variable, the corresponding storage of its original storage block. 38, 72, 86, 97, 235, 296, 300–305, 311, 483, 625, 765, 922

corresponding storage block

A storage block that contains the storage of one or more variables in a device data environment that corresponds to mapped variables in an original storage block. 8, 9, 38, 71, 293, 303, 304

C pointer

For C/C++, a base language pointer variable. For Fortran, a scalar variable of type C_PTR. 45, 114, 235, 350

current device

The device on which the current task is executing. 8, 10, 21, 24, 46, 59, 74, 105, 148, 336, 455, 471, 555, 598, 600, 603, 651, 668, 675, 676, 707–709, 817, 828

current task

For a given [thread](#), the [task](#) corresponding to the [task region](#) that it is executing. [38, 39, 49, 59, 293, 316, 349, 498, 499, 588, 590, 592–594, 596–598, 600, 609, 612, 613, 618, 702](#)

current task region

The [region](#) that corresponds to the [current task](#). [5, 107, 417, 447, 466, 495, 499, 540, 541, 889](#)

current team

All [threads](#) in the [team](#) executing the innermost enclosing [parallel region](#). [28, 29, 84, 96, 107, 109, 120, 219, 417, 421, 422, 424, 425, 427, 428, 433, 455, 462, 466, 495, 498, 499, 534, 536, 540, 544, 599, 610, 758](#)

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](#). [29, 317](#)

D

data-copying property

The [property](#) that a [clause](#) copies a [list item](#) from one [data environment](#) to other [data environments](#). [285, 286](#)

data entity

For C/C++, a data object that is referenced by a given lvalue expression or [array section](#). For Fortran, a data entity as defined by the [base language](#). [25, 26, 28, 38, 40, 41, 44, 56, 58, 60, 61, 89, 92, 93, 98, 109, 114, 345](#)

data environment

The [variables](#) associated with the execution of a given [region](#). [4, 6, 8, 9, 21, 39–41, 44, 49, 59, 72, 75, 78, 84, 105, 118–120, 124, 127, 128, 215, 234, 256, 287, 293, 310, 446, 449, 457, 465, 474, 476, 481, 486, 623, 827, 936, 947](#)

data-environment attribute

A [data-sharing attribute](#) or a [data-mapping attribute](#). [39, 215, 274, 331](#)

data-environment attribute clause

A [clause](#) that explicitly determines the [data-environment attributes](#) of the [list items](#) in its [list](#) argument. [215, 40, 307, 364, 365, 419, 457, 458, 465](#)

data-environment attribute property

The [property](#) that a [clause](#) is a [data-environment clause](#). [225–227](#), [230](#), [233–236](#), [251](#), [255](#), [257](#), [259](#), [280](#), [281](#), [283](#), [292](#), [326](#), [331](#)

data-environment clause

A [clause](#) that is a [data-environment attribute clause](#) or otherwise affects the [data environment](#). [215](#), [40](#), [215](#), [464](#)

data-mapping attribute

The relationship of a [data entity](#) in a given [device data environment](#) to the version of that entity in the [enclosing data environment](#). [215](#), [39](#), [40](#), [52](#), [60](#), [218](#), [289](#), [307](#), [942](#)

data-mapping attribute clause

A [clause](#) that explicitly determines the [data-mapping attributes](#) of the [list items](#) in its [list](#) argument. [215](#), [8](#), [38](#), [40](#), [52](#), [78](#), [281](#), [289](#), [327](#), [474](#), [476](#), [481](#), [930](#)

data-mapping attribute property

The [property](#) that a [clause](#) is a [data-mapping clause](#). [281](#), [292](#)

data-mapping clause

A [clause](#) that is a [data-mapping attribute clause](#) or otherwise affects the [data environment](#) of the [target device](#). [215](#), [40](#), [72](#), [215](#)

data-mapping construct

A [construct](#) that has the [data-mapping property](#). [49](#), [71](#), [217](#), [293](#), [297](#), [302](#), [304](#), [479](#)

data-mapping property

The [property](#) of a [construct](#) on which a [data-mapping attribute clause](#) may be specified. [40](#), [474](#), [476](#), [478](#), [481](#)

data-motion attribute

The data-movement relationship between a given [device data environment](#) and the version of that [data entity](#) in the [enclosing data environment](#). [33](#), [310](#)

data-motion attribute property

The [property](#) that a [clause](#) is a [data-motion clause](#). [311](#), [313](#)

data-motion clause

A [clause](#) that specifies data movement between a [device](#) set that is specified by the [construct](#) on which it appears. [23](#), [33](#), [40](#), [291](#), [308](#), [310–313](#), [486](#), [938](#)

data race

A condition in which different [threads](#) access the same memory location such that the accesses are unordered and at least one of the accesses is a write. [Data races](#) produce [unspecified behavior](#). [8](#), [2](#), [8](#), [9](#), [13](#), [14](#), [41](#), [226](#), [228](#), [231](#), [250](#), [258](#), [287](#), [304](#), [311](#), [319](#), [421](#), [440](#), [516](#)

data-sharing attribute

For a given [data entity](#) in a [data environment](#), an attribute that determines the scope in which the entity is visible (i.e., its name provides access to its storage) and/or the lifetime of the entity. A [variable](#) that is part of an [aggregate variable](#) cannot have a particular [data-sharing attribute](#) independent of the other components, except for static data members of C++ classes. [215](#), [39](#), [41](#), [44](#), [52](#), [53](#), [56](#), [58](#), [60](#), [61](#), [64–66](#), [88](#), [89](#), [92](#), [98](#), [109](#), [114](#), [215–220](#), [223](#), [224](#), [289](#), [307](#), [474](#), [476](#), [478](#), [481](#), [486](#), [548](#), [918](#), [938](#), [942](#)

data-sharing attribute clause

A [clause](#) that explicitly determines the [data-sharing attributes](#) of the [list items](#) in its [list](#) argument. [215](#), [7](#), [21](#), [22](#), [41](#), [52](#), [84](#), [163](#), [215](#), [218](#), [220–223](#), [225](#), [237](#), [238](#), [324](#), [327](#), [443](#), [446](#), [449](#), [481](#), [483](#), [550](#), [930](#), [945](#)

data-sharing attribute property

The [property](#) that a [clause](#) is a [data-sharing clause](#). [225–227](#), [230](#), [233–236](#), [251](#), [255](#), [257](#), [259](#), [326](#), [465](#)

data-sharing clause

A [clause](#) that is a [data-sharing attribute clause](#). [215](#), [41](#), [215](#), [217](#), [218](#)

declaration-associated directive

A [declarative directive](#) for which its associated [base language](#) code is a [procedure](#) declaration. [157](#), [156–159](#), [351](#), [358](#), [364](#), [365](#), [932](#)

declaration sequence

For C/C++, a sequence of [base language](#) declarations, including definitions, that appear in the same scope. For Fortran, a sequence of zero or more statements in the specification part. The sequence may include other [directives](#) that are associated with the declarations. [353](#), [367](#), [386](#)

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](#). [41](#), [42](#), [51](#), [62](#), [116](#), [155](#), [156](#), [158–160](#), [165](#), [262](#), [266](#), [274](#), [278](#), [307](#), [321](#), [351](#), [353](#), [358](#), [364](#), [366](#), [380](#), [470](#), [929](#)

declare target directive

A [declarative directive](#) that has the [declare-target property](#). [8](#), [71](#), [78](#), [217](#), [239](#), [279](#), [289](#), [302](#), [335](#), [362–364](#), [367](#), [368](#), [373](#), [377](#), [378](#), [482](#), [483](#), [584](#), [920](#), [937](#), [943](#)

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](#). [42](#), [364](#), [366](#)

declare variant directive

A [declarative directive](#) that declares a [function variant](#) for a given [base function](#). [49](#), [335](#), [346](#), [347](#), [353](#), [355](#), [919](#), [939](#), [943](#)

default mapper

The [mapper](#) that is used for a [map clause](#) for which the [mapper modifier](#) is not explicitly specified. [88](#), [291](#)

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](#).

[42](#), [112](#), [134](#), [149](#), [949](#)

delimited directive

A [directive](#) for which the associated [base-language code](#) is explicitly delimited by the use of a required paired [end directive](#). [158](#), [156](#), [158](#), [344](#), [353](#), [366](#), [386](#)

dependence

An ordering relation between two instances of executable code that must be enforced by a [compliant implementation](#). [524](#), [43](#), [48](#), [105](#), [185](#), [455](#), [524–529](#), [532](#), [535](#), [624](#), [740](#), [741](#), [782](#), [787](#), [788](#)

dependence-compatible task

Two [tasks](#) between which a [task dependence](#) may be established. [527](#), [88](#), [105](#), [111](#), [524](#), [528](#), [529](#), [531](#), [579](#)

dependent task

A [task](#) that because of a [task dependence](#) cannot be executed until its [antecedent tasks](#) have completed. [527](#), [22](#), [103](#), [105](#), [468](#), [479](#), [500](#), [522](#), [524](#), [527–529](#), [624](#), [767](#), [788](#)

depend object

An [OpenMP object](#) that supplies user-computed [dependences](#) to [depend](#) clauses. [578](#), [185](#), [455](#), [501](#), [525](#), [526](#), [529](#), [530](#), [624](#), [786](#), [787](#), [944](#)

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](#). [43](#), [160](#), [197](#), [198](#), [263](#), [553](#), [623](#), [735](#), [738](#), [763](#), [804](#), [807](#), [809](#), [811](#), [915](#), [927](#), [928](#), [935–939](#), [941](#), [944](#)

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](#). [37](#), [39](#), [43](#), [450](#), [468](#), [522](#), [523](#), [541](#)

detachable task

An [explicit task](#) that only completes after an associated [event](#) variable that represents an *allow-completion event* is fulfilled and execution of the associated [structured block](#) has completed. [465](#), [446](#), [457](#), [522–524](#), [558](#), [610](#), [943](#)

device

An implementation-defined logical execution engine.

COMMENT: A [device](#) could have one or more [processors](#).

[3](#), [4](#), [7–9](#), [20–22](#), [24](#), [29](#), [32](#), [38](#), [40](#), [42–46](#), [49](#), [54](#), [58](#), [61](#), [73](#), [77](#), [78](#), [81](#), [86](#), [100](#), [103–106](#), [112](#), [118–120](#), [127](#), [130](#), [131](#), [142–144](#), [148](#), [186](#), [235](#), [281](#), [282](#), [289](#), [293](#), [310](#), [311](#), [317–319](#), [329](#), [335](#), [336](#), [338](#), [340](#), [349](#), [362](#), [363](#), [376](#), [377](#), [456](#), [470](#), [473](#), [475](#), [477](#), [481–484](#), [486](#), [514](#), [556](#), [584](#), [585](#), [592](#), [610](#), [612](#), [614–617](#), [619–623](#), [625–628](#), [631–634](#), [640](#), [651](#), [653–656](#), [667](#), [668](#), [670–673](#), [675](#), [676](#), [687](#), [707](#), [713](#), [714](#), [717](#), [728–731](#), [733](#), [735](#), [736](#), [742](#), [747](#), [751](#), [770](#), [798–805](#), [811](#), [813](#), [814](#), [820](#), [828](#), [829](#), [831–838](#), [840–842](#), [847](#), [851](#), [854](#), [861](#), [864](#), [871](#), [875](#), [879–882](#), [886](#), [909](#), [913](#), [915](#), [920](#), [922](#), [925](#), [929](#), [931](#), [932](#), [934](#), [935](#), [938–940](#), [942](#), [944–946](#)

device address

An address of an object that may be referenced on a [target device](#). [8](#), [8](#), [45](#), [64](#), [114](#), [233–236](#), [345](#), [349](#), [376](#), [377](#), [628](#), [915](#), [938](#), [942](#)

device-affecting construct

A [construct](#) that has the [device-affecting property](#). [483](#), [620](#), [622](#), [951](#)

device-affecting property

The [property](#) that a [device construct](#) can modify the state of the [device data environment](#) of a specified [target device](#). [44](#), [474](#), [476](#), [478](#), [481](#), [486](#)

device-associated property

The [property](#) of a [clause](#) that a [device](#) must be associated with the [construct](#) on which it appears. [233–236](#)

device construct

A [construct](#) that has the [device property](#). [2](#), [44](#), [45](#), [58](#), [64](#), [104](#), [105](#), [114](#), [144](#), [294](#), [372](#), [373](#), [471](#), [761](#), [786](#), [807](#), [811](#), [941](#), [946](#)

device data environment

The initial [data environment](#) associated with a [device](#). [8](#), [8](#), [9](#), [25](#), [38](#), [40](#), [44](#), [58](#), [70–74](#), [86](#), [89](#), [114](#), [127](#), [215](#), [234](#), [236](#), [237](#), [256](#), [282](#), [289](#), [293](#), [294](#), [296](#), [299–305](#), [310](#), [311](#), [349](#), [363](#), [378](#), [474](#), [476](#), [481](#), [484](#), [486](#), [619](#), [621](#), [622](#), [625](#), [628](#), [629](#), [631](#), [633](#), [640](#), [805](#), [915](#), [930](#), [934](#)

device global requirement clause

A [requirement clause](#) that has the [device global requirement property](#). [372](#)

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](#). [44](#), [373](#), [375–379](#)

device-information property

The [property](#) of a [routine](#) that it provides or modifies information about a specified [device](#) that supports use of the [device](#) in an [OpenMP program](#). [612](#), [44](#), [612–621](#)

device-information routine

A [routine](#) that has the [device-information property](#). [612](#), [612](#)

device-local attribute

For a given [device](#), a [data-sharing attribute](#) of a [data entity](#) that it has [static storage duration](#) and is visible only to [tasks](#) that execute on that [device](#). [281](#), [44](#), [216](#), [219](#)

device-local variable

A [variable](#) that has the [device-local attribute](#) with respect to a given [device](#). [281](#), [8](#), [294](#), [362](#), [363](#), [378](#), [915](#)

device-memory-information routine

A [routine](#) that has the [device-memory-information routine property](#). [624](#), [623](#)

device-memory-information routine property

The [property](#) of a [device memory routine](#) that it enables operations on [memory](#) that is associated with the specified [devices](#) but does not itself directly operate on that [memory](#). [624](#), [45](#), [624–626](#)

device memory routine

A [device routine](#) that has the [device memory routine property](#). [623](#), [45](#), [104](#), [105](#), [584](#), [623](#), [624](#), [687](#), [805](#), [918](#), [946](#)

device memory routine property

The [property](#) that a [device routine](#) operates on or otherwise enables operations on [memory](#) that is associated with the specified [devices](#). [623](#), [45](#), [624–627](#), [629–631](#), [634](#), [635](#), [637](#), [638](#), [640](#), [641](#)

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](#). [7](#), [7](#), [35](#), [118](#), [119](#), [122](#), [123](#), [130](#), [142–144](#), [319](#), [471](#), [481](#), [562](#), [613](#), [614](#), [616](#), [618–622](#), [631](#), [633](#), [714](#), [717](#), [799–801](#), [805](#), [807](#), [828](#), [934](#)

device pointer

An [implementation defined handle](#) that refers to a [device address](#) and is represented by a [C pointer](#). [8](#), [64](#), [114](#), [233](#), [234](#), [345](#), [349](#), [376](#), [624](#), [627](#), [628](#), [631–634](#), [675](#), [915](#), [940](#)

device procedure

A [procedure](#) that can be executed on a [target device](#), as part of a [target region](#). [105](#), [283](#), [362](#), [363](#), [372](#), [373](#), [377](#), [378](#)

device property

The [property](#) of a [construct](#) that it accepts the [device](#) clause. [44](#), [364](#), [366](#), [474](#), [476](#), [478](#), [481](#), [486](#), [488](#)

device region

A [region](#) that corresponds to a [device construct](#). [740](#), [747](#), [770](#), [805](#), [807](#), [809](#), [811](#)

device routine

An [OpenMP API routine](#) that may require access to one or more specified [devices](#). [24](#), [45](#), [144](#)

device selector set

A [selector set](#) that may match the [device trait set](#). [338](#), [338–340](#)

device-specific environment variable

An alternative [OpenMP environment variable](#) that controls the behavior of the program only with respect to a particular [device](#) or set of [devices](#). [122](#), [123](#), [130](#), [131](#), [142](#), [938](#)

device-tracing callback

A [callback](#) that has the [device-tracing property](#). [798](#)

device-tracing entry point

An [entry point](#) that has the [device-tracing property](#). [798](#), [799](#)

device-tracing property

The [property](#) that an [entry point](#) or [callback](#) is part of the [OMPT](#) tracing interface and, so, is used to control the collection of [trace records](#) on a [device](#). [798](#), [46](#), [798–803](#), [806](#), [808](#), [810](#)

device trait set

The [trait set](#) that consists of [traits](#) that define the characteristics of the [device](#) that the compiler determines will be the [current device](#) during program execution at a given point in the [OpenMP program](#). [336](#), [45](#), [335–337](#), [354](#)

device-translating callback

A [callback](#) that has the [device-translating property](#). [870](#), [872](#)

device-translating property

The [property](#) that a [callback](#) translates data between the formats used for the [device](#) on which the [third-party tool](#) and [OMPD library](#) run and the [device](#) on which the [OpenMP program](#) runs. [870](#), [46](#), [872](#)

directive

A [base language](#) mechanism to specify [OpenMP program](#) behavior. [2](#), [3](#), [6–9](#), [13](#), [15](#), [17](#), [19](#), [21](#), [22](#), [24](#), [25](#), [27](#), [29](#), [31–37](#), [41](#), [42](#), [46–51](#), [55](#), [58](#), [59](#), [62](#), [64](#), [66](#), [70](#), [71](#), [75](#), [81](#), [82](#), [91](#), [93](#), [97](#), [99](#), [103](#), [106](#), [110](#), [112](#), [114](#), [115](#), [117](#), [119](#), [130](#), [147](#), [151–165](#), [167–169](#), [171](#), [175](#), [179](#), [187](#), [188](#), [192](#), [193](#), [195–197](#), [203](#), [206–211](#), [213](#), [215](#), [216](#), [218–220](#), [223](#), [225](#), [231](#), [238](#), [246–248](#), [253](#), [256](#), [260–264](#), [266](#), [267](#), [269–274](#), [276–280](#), [282](#), [283](#), [289](#), [291](#), [293](#), [299](#), [303](#), [307–309](#), [315](#), [317](#), [318](#), [321–327](#), [332](#), [335](#), [336](#), [338](#), [339](#), [341–345](#), [352–355](#), [357](#), [359](#), [360](#), [363–374](#), [376](#), [379](#), [380](#), [385–390](#), [397](#), [400](#), [403](#), [406](#), [407](#), [413](#), [418](#), [420](#), [428](#), [430](#), [443](#), [446](#), [449](#), [451](#), [454](#), [465](#), [471](#), [472](#), [474–479](#), [481](#), [483](#), [486](#), [489](#), [490](#), [494](#), [501–503](#), [508](#), [516](#), [520](#), [521](#), [523](#), [525](#), [539](#), [543](#), [544](#), [546](#), [555](#), [581](#), [584](#), [585](#), [628](#), [667](#), [668](#), [674](#), [675](#), [677](#), [688](#), [771](#), [774](#), [917](#), [918](#), [920](#), [921](#), [927–934](#), [936–939](#), [942](#), [943](#), [945](#), [947](#), [949](#), [951](#)

directive name

The name of a [directive](#) or a corresponding [construct](#). [34–36](#), [47](#), [66](#), [99](#), [153](#), [166](#), [178](#), [179](#), [184](#), [185](#), [187](#), [211](#), [212](#), [223](#), [225–227](#), [230](#), [233–236](#), [251](#), [255–257](#), [260](#), [265](#), [268](#), [269](#),

272, 273, 281, 283, 285, 287, 292, 306, 312, 313, 320, 321, 324, 327, 329, 331, 332, 342,
343, 347, 348, 350, 356–358, 361, 362, 368, 370, 371, 374–384, 389, 391, 394, 397, 401,
406, 410, 411, 415, 416, 419, 420, 422, 437, 441, 444, 452, 453, 459–465, 470–472, 490,
492, 501, 503–509, 511–513, 526, 527, 531, 532, 538, 539, 545, 547

directive-name list

An [argument list](#) that consists of [directive-name list items](#). [165](#)

directive-name list item

A [list item](#) that is a [directive name](#). [165](#), [47](#)

directive-name separator

Characters used to separate the [directive names](#) of [leaf constructs](#) in a [compound-directive name](#). A [directive-name separator](#) is either [white space](#) or, in Fortran, a plus sign (i.e., '+'); a given instance of a [compound-directive name](#) must use the same character for all [directive-name separators](#). [545](#), [47](#), [545–547](#)

directive specification

The [directive specifier](#) and list of [clauses](#) that specify a given [directive](#). [153](#), [47](#), [153](#), [166](#)

directive-specification list

An [argument list](#) that consists of [directive-specification list items](#). [165](#)

directive-specification list item

A [list item](#) that is a [directive specification](#). [165](#), [47](#), [167](#)

directive specifier

The [directive name](#) and, where permitted, the [directive](#) arguments that are specified for a given [directive](#). [153](#), [47](#)

directive variant

A [directive specification](#) that can be used in a [metadirective](#). [341](#), [37](#), [94](#), [341–344](#), [943](#)

divergent threads

Two [threads](#) are [divergent](#) if one executes a [diverging code path](#) and the other does not due to a conditional statement. [7](#), [47](#), [379](#)

diverging code path

For a given pair of [threads](#), the [region](#) of a [structured block sequence](#) that is executed by only one of the [threads](#). [6](#), [47](#)

doacross-affected loop

For a [worksharing-loop construct](#) in which a stand-alone [ordered directive](#) is closely nested, a loop that is affected by its [ordered clause](#). [48](#), [212](#), [388](#), [534](#), [535](#), [932](#)

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. [524](#), [48](#), [100](#), [532](#), [535](#), [740](#)

doacross iteration

A [logical iteration](#) of a [doacross loop nest](#). [48](#), [100](#), [523](#), [524](#), [532](#), [535](#)

doacross iteration space

The [logical iteration space](#) of a [doacross loop nest](#). [48](#), [532](#)

doacross logical iteration

A [doacross iteration](#). [532](#)

doacross loop nest

The [doacross-affected loops](#) of a [worksharing-loop construct](#) in which a stand-alone [ordered construct](#) is closely nested. [48](#), [532](#), [534](#), [535](#), [945](#), [946](#)

dynamic context selector

Any [context selector](#) that is not a [static context selector](#). [354](#)

dynamic groupprivate block

An unnamed, [groupprivate storage block](#) with [dynamic storage duration](#) that is instantiated from a [groupprivate-instantiating construct](#). [19](#), [48](#), [56](#), [57](#), [103](#), [216](#), [329](#), [330](#), [683](#), [685–687](#), [919](#)

dynamic-groupprivate-information routine

A [groupprivate-information routine](#) that has the [dynamic-groupprivate-information-routine property](#). [330](#), [683](#), [684](#)

dynamic-groupprivate-information-routine property

The [property](#) of a [groupprivate-information routine](#) that it provides information about [dynamic groupprivate blocks](#). [48](#), [683–685](#)

dynamic replacement candidate

A [replacement candidate](#) that may be selected at runtime to replace a given [metadirective](#). [341](#), [341](#), [342](#), [346](#)

dynamic storage duration

For C/C++, the lifetime of an object with dynamic storage duration, as defined by the [base language](#). For Fortran, the lifetime of a data object that is dynamically allocated with the **ALLOCATE** statement or some other language mechanism. [48](#), [216](#), [219](#)

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. [337](#), [115](#), [335](#), [337](#), [338](#)

E

effective context selector

The resulting [context selector](#) that must be satisfied for a given [function variant](#) to be selected, as determined by the [match clauses](#) of all **begin declare_variant** directives that delimit a [base-language code region](#) that encloses the [declare variant directive](#). [353](#), [353](#), [354](#)

effective [map](#) clause set

The set of all [map clauses](#) that apply to a [data-mapping construct](#), including any implicit [map clauses](#) and [map clauses](#) applied by [mappers](#). [302](#), [302](#), [303](#)

enclosing context

For C/C++, the innermost scope enclosing a [directive](#). For Fortran, the innermost scoping unit enclosing a [directive](#). [49](#), [75](#), [84](#), [98](#), [218](#), [219](#), [252](#), [253](#), [258](#), [259](#), [263](#), [267](#), [288](#), [341](#), [357](#), [358](#), [427](#), [430](#), [432](#), [440](#), [441](#), [943](#)

enclosing data environment

For a given [directive](#), the [data environment](#) of its [enclosing context](#). [40](#), [53](#), [58](#), [64](#), [65](#), [96](#), [114](#), [456](#), [457](#)

encountering device

For a given [construct](#), the [device](#) on which the [encountering task](#) of the [construct](#) executes. [234](#), [304](#), [310](#), [312](#), [313](#), [483](#), [922](#)

encountering task

For a given [region](#), the [current task](#) of the [encountering thread](#). [6](#), [49](#), [106](#), [310](#), [351](#), [357](#), [369](#), [403](#), [412](#), [413](#), [434](#), [447](#), [451](#), [456](#), [463](#), [465](#), [482](#), [488](#), [496](#), [497](#), [500](#), [502](#), [540–542](#), [554](#), [599](#), [607](#), [608](#), [694](#), [697–701](#), [730](#), [770](#), [785](#), [786](#), [807](#), [826](#), [827](#), [910](#), [911](#)

encountering-task binding property

The [binding property](#) that the [binding thread set](#) is the [encountering task](#). [554](#)

encountering thread

For a given [region](#), the [thread](#) that encounters the corresponding [construct](#), [structured block sequence](#), or [routine](#). [4](#), [5](#), [28](#), [29](#), [49](#), [50](#), [61](#), [93](#), [252](#), [402](#), [407–409](#), [412](#), [421](#), [422](#), [443](#), [445–447](#), [481](#), [489](#), [519](#), [526](#), [554](#), [598](#), [599](#), [609](#), [614](#), [705](#), [707](#), [709–711](#), [713](#), [719](#), [750](#), [796](#), [813](#), [819](#), [821–823](#), [826](#), [828](#), [934](#)

encountering-thread binding property

The [binding property](#) that the [binding thread set](#) is the [encountering thread](#). [554](#)

end-clause property

The [property](#) that a [clause](#) may appear on an [end directive](#). [153](#), [286](#), [501](#)

end directive

For a given [directive](#), a paired [directive](#) that lexically delimits the code associated with that [directive](#). [153](#), [36](#), [42](#), [50](#), [153](#), [156–159](#), [164](#), [192](#), [193](#), [197](#), [344](#), [353](#), [354](#), [367](#), [494](#), [936](#), [937](#)

ending address

The address of the last [storage location](#) of a [list item](#) or, for a [mapped variable](#), of its [original list item](#). [52](#), [72](#), [300](#)

entry point

A [runtime entry point](#). [24](#), [25](#), [46](#), [81](#), [725](#), [727](#), [729–731](#), [736](#), [745](#), [747](#), [755](#), [771](#), [798](#), [799](#), [802](#), [813–842](#), [925](#), [926](#), [935](#)

enumeration

A type or any variable of a type that consists of a specified set of named integer values. For C/C++, an [enumeration](#) type is specified with the **enum** specifier. For Fortran, an [enumeration](#) type is specified by either (1) a named integer constant that is used as the integer kind of a set of named integer constants that have unique values or (2) a C-interoperable enumeration definition. [50](#), [556](#), [559–561](#), [564](#), [565](#), [568](#), [570](#), [574](#), [577–580](#), [582](#), [584](#), [585](#), [587](#), [737](#), [739](#), [741](#), [743](#), [745](#), [748–750](#), [752–755](#), [757](#), [760–762](#), [764–768](#), [816](#), [853](#), [855](#), [856](#), [903](#)

environment variable

Unless specifically stated otherwise, an [OpenMP environment variable](#). [2](#), [6](#), [121](#), [122](#), [130–140](#), [142–150](#), [716](#), [717](#), [901](#), [916](#), [917](#), [928](#), [929](#), [938](#), [940](#), [941](#), [944](#), [946–948](#)

error termination

A **fatal** action preformed in response to an error. [6](#), [34](#), [95](#), [407](#), [932](#)

event

A point of interest in the execution of a [thread](#) or a [task](#). [10](#), [11](#), [14](#), [15](#), [30](#), [43](#), [93](#), [105](#), [110](#), [111](#), [249](#), [294](#), [363](#), [369](#), [403](#), [404](#), [412](#), [413](#), [421](#), [424–428](#), [430](#), [432](#), [434](#), [440](#), [446](#), [447](#), [450](#), [451](#), [457](#), [465–467](#), [469](#), [473](#), [475–477](#), [479](#), [482](#), [483](#), [486](#), [487](#), [494–498](#), [500](#), [516](#), [517](#), [520](#), [522–524](#), [529](#), [533](#), [535](#), [536](#), [541](#), [542](#), [558](#), [606](#), [609](#), [610](#), [623](#), [624](#), [628](#), [629](#), [631–636](#), [638–642](#), [689–701](#), [719](#), [721](#), [724](#), [725](#), [727–729](#), [735](#), [751](#), [753](#), [755](#), [767](#), [770](#), [772](#), [783–785](#), [787](#), [789–791](#), [793](#), [797](#), [798](#), [802–804](#), [807](#), [809](#), [811](#), [813](#), [817](#), [818](#), [824](#), [833](#), [834](#), [836](#), [840–842](#), [844](#), [908](#), [910](#), [911](#), [913](#), [925](#), [934](#)

exception-aborting directive

A [directive](#) that has the [exception-aborting property](#). [383](#), [917](#)

exception-aborting property

For C++, the [property](#) of a [directive](#) that whether an exception that occurs in its associated [region](#) is caught or results in a [runtime error termination](#) is [implementation defined](#). [51](#), [152](#), [481](#)

exclusive property

The [property](#) that a [clause](#), an argument, or a [modifier](#) may not be specified when, (respectively), a different [clause](#), argument or [modifier](#) is specified. When applied to a [clause set](#), the [property](#) applies only to [clauses](#) within that set. [164](#), [34](#), [163–165](#), [260](#), [269](#), [324](#), [360](#), [400](#), [424](#), [446](#), [449](#), [504](#), [508](#), [539](#)

exclusive scan computation

A [scan computation](#) for which the value read does not include the updates performed in the same [logical iteration](#). [272](#), [273](#), [942](#)

executable directive

A [directive](#) in an executable context that results in [implementation code](#) or prescribes the manner in which any associated user code must execute. [3](#), [36](#), [62](#), [66](#), [69](#), [70](#), [100](#), [103](#), [115](#), [152](#), [155](#), [156](#), [158](#), [191](#), [203](#), [326](#), [342](#), [355](#), [369](#), [370](#), [392](#), [393](#), [395](#), [397](#), [398](#), [400](#), [402](#), [412](#), [417](#), [421](#), [424–426](#), [428](#), [431](#), [435](#), [436](#), [439](#), [442](#), [446](#), [449](#), [455](#), [466](#), [474](#), [476](#), [478](#), [481](#), [486](#), [488](#), [493](#), [495](#), [498](#), [499](#), [514](#), [518](#), [525](#), [534](#), [536](#), [540](#), [544](#)

explicit barrier

A [barrier](#) that is specified by a [barrier](#) construct. [495](#)

explicitly associated directive

A [declarative directive](#) for which its associated [base language](#) declarations are explicitly specified in a [variable list](#) or [extended list](#) argument. [156](#), [156](#), [159](#), [274](#), [278](#), [321](#), [364](#)

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](#). [289](#)

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](#). [218](#), [215](#), [218](#), [224](#)

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. [3](#), [3](#), [101](#), [152](#), [153](#), [432](#), [466](#), [713](#), [830](#)

explicit task

A [task](#) that is not an [implicit task](#). [5](#), [5](#), [7](#), [26](#), [30](#), [43](#), [52](#), [54](#), [85](#), [96](#), [106](#), [119](#), [252](#), [253](#), [403](#), [407](#), [446](#), [447](#), [449–451](#), [467](#), [495](#), [522](#), [523](#), [544](#), [606](#), [713](#), [744](#), [782](#), [826](#), [893](#), [943](#), [945](#), [949](#)

explicit task region

A [region](#) that corresponds to an [explicit task](#). [8](#), [93](#), [225](#), [447](#), [547](#), [607](#), [936](#)

exporting task

A [task](#) that permits one of its [child tasks](#) to be an [antecedent task](#) of a [task](#) for which it is a [preceding dependence-compatible task](#). [531](#), [111](#), [447](#), [457](#), [528](#), [531](#), [579](#)

extended address range

For a given [original list item](#), the [address range](#) that starts from the minimum of its [starting address](#) and its [base address](#) and ends with maximum of its [ending address](#) and its [base address](#). [300](#), [72](#), [300](#)

extended list

An [argument list](#) that consists of [extended list items](#). [165](#), [51](#)

extended list item

A [variable list item](#) or the name of a [procedure](#). [165](#), [52](#), [167](#), [168](#)

extension trait

A [trait](#) that is [implementation defined](#). [337](#), [335](#)

F

finalized taskgraph record

A [taskgraph record](#) in which all information required for a [replay execution](#) has been saved. [456](#), [73](#), [456](#)

final task

A [task](#) that generates [included final tasks](#) when it encounters [task-generating constructs](#) on which the [final clause](#) may be specified. [462](#), [53](#), [119](#), [447](#), [456](#), [457](#), [459](#), [462](#), [463](#), [465](#), [608](#), [948](#)

first-party tool

A [tool](#) that executes in the [address space](#) of the program that it is monitoring. [721](#), [14](#), [30](#), [80](#), [148](#), [719](#), [721](#), [723](#), [935](#), [944](#)

firstprivate attribute

For a given [construct](#), a [data-sharing attribute](#) of a [variable](#) that implies the [private attribute](#), and additionally the [variable](#) is initialized with the value of the [variable](#) that has the same name in the [enclosing data environment](#) of the [construct](#). [228](#), [53](#), [217–219](#), [290](#), [307](#), [456](#), [482](#), [936](#), [945](#)

firstprivate variable

A [private variable](#) that has the [firstprivate attribute](#) with respect to a given [construct](#). [450](#), [457](#), [922](#)

flat-memory-copying property

The [property](#) that a [memory-copying routine](#) copies a unidimensional, contiguous [storage block](#). [632](#), [53](#), [634](#), [637](#)

flat-memory-copying routine

A [routine](#) that has the [flat-memory-copying property](#). [632](#), [633](#), [635](#), [638](#)

flatten-affected loop

A [transformation-affected loop](#) of a [flatten](#) construct. [392](#)

flattened iteration space

The [logical iteration space](#) of a [flattened loop](#). [392](#), [392](#)

flattened loop

The [generated loop](#) of a [flatten](#) construct. [392](#), [53](#), [392](#), [920](#)

flush

An operation that a [thread](#) performs to enforce consistency between its view of [memory](#) and the view of [memory](#) of any other [threads](#). [6](#), [10–14](#), [19](#), [53](#), [60](#), [94](#), [101](#), [110](#), [423](#), [492](#), [514](#), [519–522](#), [941](#), [948](#)

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](#). [11](#), [941](#)

flush-set

The set of [variables](#) upon which a [strong flush](#) operates. [10](#), [10](#)

foreign execution context

A context that is instantiated from a [foreign runtime environment](#) in order to facilitate execution on a given [device](#). [54](#), [186](#), [488](#), [489](#), [562](#), [939](#)

foreign runtime environment

A runtime environment that exists outside the OpenMP runtime with which the OpenMP implementation may interoperate. [54](#), [64](#), [88](#), [488](#), [491](#), [559](#), [562](#)

foreign runtime identifier

A [base language string literal](#) or a [constant](#) expression of integer [OpenMP type](#) that represents a [foreign runtime environment](#). [188](#), [489](#), [491](#), [922](#), [934](#)

foreign task

An instance of executable code that is executed in a [foreign execution context](#). [186](#), [457](#), [489](#), [922](#)

Fortran-only property

The [property](#) that an OpenMP feature is only supported in Fortran. [554](#)

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. [30](#), [54](#), [744–746](#), [770](#), [826](#), [852](#), [893](#), [894](#)

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. [54](#), [102](#), [109](#), [119](#), [136](#), [145](#), [146](#), [407](#), [408](#), [468](#), [608](#), [609](#), [759](#), [921](#), [929](#), [934](#)

free property

The [property](#) that a [modifier](#) can appear in any position in a *modifier-specification-list*. [163](#)

function

A [routine](#) or [procedure](#) that returns a type that can be the right-hand side of a [base language](#) assignment operation. [159](#), [160](#), [166](#), [322](#), [349](#), [354](#), [589–591](#), [593](#), [595](#), [597–604](#), [606–608](#), [613–619](#), [621](#), [624–627](#), [630](#), [631](#), [634](#), [635](#), [637](#), [638](#), [640](#), [641](#), [644–649](#), [652–658](#),

664–666, 669–672, 674, 677–681, 684–686, 699, 700, 702, 703, 705, 708, 710, 712–715,
718, 721, 771, 796, 813–823, 825, 827–829, 831–834, 836–842, 862, 863, 865–868,
870–878, 880–906

function dispatch

A [base function](#) call for which [variant substitution](#) may be controlled. 192

function-dispatch structured block

A [context-specific structured block](#) that may be associated with a [dispatch](#) directive. 192,
192, 193, 335, 348, 350, 355

function variant

A definition of a [procedure](#) that may be used as an alternative to the [base language](#) definition.
37, 42, 49, 94, 116, 335, 346–353, 355, 357, 488, 939, 943

G

generally-composable property

The [property](#) of a [loop-transforming construct](#) that it may use [directives](#) other than
[loop-transforming directives](#) in its [apply](#) clauses. 390, 392, 395, 400

generated loop

A loop that is generated by a [loop-transforming construct](#) and is one of the resulting loops
that replace the [construct](#). 388, 53, 56, 61, 79, 110, 202, 208, 210, 388–390, 393, 396, 397,
399, 400, 451, 932

generated loop nest

A [canonical loop nest](#) that is generated by a [loop-transforming construct](#). 388, 389

generated loop sequence

A [canonical loop sequence](#) that is generated by a [loop-transforming construct](#). 388

generated task

The [task](#) that is generated as a result of the [generating task](#) encountering a [task-generating
construct](#). 5, 127, 218, 446, 447, 449, 450, 454, 459, 460, 462, 464, 488, 489, 499, 500, 502,
528, 529, 531, 782, 786, 787

generating task

For a given [region](#), the [task](#) for which execution by a [thread](#) generated the [region](#). 29, 55, 56,
127, 355, 447, 474, 476, 478, 481, 486, 488, 523, 524, 623, 890

generating-task binding property

The [binding property](#) that the [binding task set](#) is the [generating task](#). 623, 627, 629–631, 634,
635, 637, 638, 640, 641

generating task region

For a given [region](#), the [region](#) that corresponds to its [generating task](#). [31](#), [61](#), [112](#), [890](#)

global

A program aspect such as a scope that covers the whole [OpenMP program](#). [59](#), [118–120](#), [122](#), [130](#), [322](#), [945](#)

grid loop

The [generated loops](#) of a [tile](#) or [stripe construct](#) that iterate over cells of a grid superimposed over the [logical iteration space](#), with spacing determined by the [sizes clause](#). [79](#), [397–399](#), [920](#), [933](#)

groupprivate absolute limit

The maximum amount of [memory](#) available in a [groupprivate memory space](#), inclusive of [groupprivate variables](#) and any [groupprivate](#) copy of an instantiated [dynamic groupprivate block](#). [329](#), [329](#), [687](#)

groupprivate attribute

For a given group of [tasks](#), a [data-sharing attribute](#) of a [data entity](#) that it is visible and/or accessible only to those [tasks](#). [278](#), [48](#), [56](#), [57](#), [216](#), [219](#), [279](#), [280](#), [329](#), [330](#)

groupprivate effective limit

The amount of remaining [memory](#) available in a [groupprivate memory space](#) upon encountering a [groupprivate-instantiating construct](#), exclusive of any [groupprivate variables](#) that will be used in the corresponding [region](#) of that [construct](#). [329](#), [329](#)

groupprivate-information routine

A [memory-management routine](#) that has the [groupprivate-information-routine property](#). [48](#), [683](#), [928](#)

groupprivate-information-routine property

The [property](#) of a [memory-management routine](#) that it provides information about [groupprivate variables](#) and [dynamic groupprivate blocks](#). [56](#), [683–686](#)

groupprivate-instantiating construct

A [construct](#) that has the [groupprivate-instantiating property](#). [48](#), [56](#), [216](#), [329](#), [330](#), [685–687](#)

groupprivate-instantiating property

The [property](#) that a [construct](#) can instantiate a [dynamic groupprivate block](#) via a [dyn_groupprivate clause](#). [56](#), [412](#), [481](#)

groupprivate memory space

For a given set of [tasks](#) that are grouped according to an [access group type](#), a [memory space](#) that contains visible [groupprivate variables](#) and, potentially, a [groupprivate](#) copy of an instantiated [dynamic groupprivate block](#). [329](#), [56](#), [103](#), [329](#), [687](#), [919](#)

groupprivate variable

A [variable](#) that has the [groupprivate attribute](#) with respect to a given group of [tasks](#). [278](#), [19](#), [56](#), [57](#), [279](#), [280](#), [294](#), [362](#), [364](#), [365](#), [367](#), [433](#), [482](#), [687](#)

H

handle

An opaque reference that uniquely identifies an abstraction. [20](#), [38](#), [45](#), [57](#), [76](#), [77](#), [81](#), [85](#), [93](#), [97](#), [106](#), [116](#), [186](#), [295](#), [316](#), [317](#), [567](#), [568](#), [651](#), [657](#), [658](#), [666–668](#), [674–676](#), [735](#), [736](#), [746](#), [823](#), [848](#), [854](#), [856](#), [857](#), [859–862](#), [869](#), [878](#), [879](#), [881–884](#), [887–894](#), [896](#), [900](#), [904](#), [905](#), [915](#), [926](#)

handle-comparing property

The [property](#) that a [routine](#) compares two [handle](#) arguments. [894](#), [57](#), [894–896](#)

handle-comparing routine

A [routine](#) that has the [handle-comparing property](#). [894](#), [894](#), [926](#)

handle property

The [property](#) that a type is used to represent [handles](#). [859](#), [57](#), [848](#), [857](#), [859](#), [860](#)

handle-releasing property

The [property](#) that a [routine](#) releases a [handle](#). [896](#), [57](#), [897](#), [898](#)

handle-releasing routine

A [routine](#) that has the [handle-releasing property](#). [896](#), [896](#)

handle type

An [OpenMP type](#), [OMPD type](#), or [OMPT type](#) that has the [handle property](#). [859](#)

happens before

For an event *A* to happen before an event *B*, *A* must precede *B* in [happens-before order](#). [13](#), [13](#)

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](#). [13](#), [57](#), [318](#), [319](#), [377](#), [489](#), [941](#)

hard pause

An instance of a [resource-relinquishing routine](#) that specifies that the OpenMP state is not required to persist. [584](#), [584](#), [585](#)

hardware thread

An indivisible hardware execution unit on which only one [OpenMP thread](#) can execute at a time. [6](#), [6](#), [38](#), [90](#), [131](#), [134](#), [554](#), [616](#), [751](#)

has-device-addr attribute

For a given [device construct](#), a [data-sharing attribute](#) of a [data entity](#) that refers to an object in a [device data environment](#) that is the same object to which the [data entity](#) of the same name in the [enclosing data environment](#) of the [construct](#) refers. [235](#)

host address

An address of an object that may be referenced on the [host device](#). [58](#), [377](#), [940](#)

host device

The [device](#) on which the [OpenMP program](#) begins execution. [3–5](#), [7](#), [9](#), [19](#), [26](#), [58](#), [61](#), [78](#), [103](#), [123](#), [130](#), [131](#), [139](#), [141](#), [144](#), [311](#), [318](#), [336](#), [376](#), [470](#), [474–477](#), [482](#), [484](#), [487](#), [584](#), [602](#), [605](#), [614](#), [618–622](#), [625](#), [627](#), [631](#), [652](#), [654](#), [655](#), [669](#), [671](#), [672](#), [687](#), [714](#), [717](#), [721](#), [725](#), [727](#), [729](#), [730](#), [747](#), [819](#), [820](#), [831–833](#), [842](#), [856](#), [878–880](#), [886](#), [928](#), [934](#), [943](#)

host pointer

A pointer that refers to a [host address](#). [376](#), [377](#), [625](#), [627](#), [631–633](#), [940](#)

I

ICV

An [internal control variable](#). [118](#), [7](#), [58](#), [59](#), [63](#), [82](#), [87](#), [118](#), [121–125](#), [127–130](#), [132–134](#), [136–140](#), [142–150](#), [275](#), [321](#), [338](#), [356](#), [375](#), [406–409](#), [412](#), [415](#), [423](#), [434](#), [438](#), [446](#), [449](#), [457](#), [463](#), [471](#), [473](#), [474](#), [476](#), [481](#), [486](#), [521](#), [524](#), [540](#), [541](#), [557](#), [583](#), [588](#), [590–598](#), [600–608](#), [612–615](#), [619–622](#), [674](#), [675](#), [702–704](#), [706–712](#), [716](#), [717](#), [723](#), [724](#), [820](#), [822](#), [845](#), [853](#), [857](#), [883](#), [892](#), [903–905](#), [921](#), [922](#), [924](#), [925](#), [928](#), [934](#), [936](#), [938](#), [940](#), [941](#), [947–949](#)

ICV-defaulted clause

A [clause](#) that has the [ICV-defaulted property](#). [457](#)

ICV-defaulted property

The [property](#) of a [clause](#) that if it is not explicitly specified on a [directive](#) then the behavior is as if it were specified with an argument that is the value of an [ICV](#). [58](#), [321](#), [471](#)

ICV modifying property

The [property](#) of a [routine](#) or [clause](#) that its effect includes modifying the value of an [ICV](#). [472](#), [588](#), [592](#), [593](#), [596](#), [601](#), [604](#), [612](#), [620](#), [621](#), [707](#)

ICV retrieving property

The [property](#) of a [routine](#) that its effect includes returning the value of an [ICV](#). [590](#), [591](#), [593](#), [594](#), [597](#), [600–604](#), [606–608](#), [613](#), [615](#), [619](#), [621](#), [702–706](#), [708](#), [712](#)

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](#). [118](#), [59](#), [118](#), [122](#), [124](#), [127](#), [130](#), [457](#), [474](#), [476](#), [481](#), [486](#)

idle thread

An [unassigned thread](#) that is not currently executing any [task](#). [467](#), [759](#)

immediately nested construct

A [construct](#) is an [immediately nested construct](#) of another [construct](#) if it is immediately nested within the other [construct](#) with no intervening statements or [directives](#). [59](#), [104](#), [325](#), [414](#), [934](#)

imperfectly nested loop

A nested loop that is not a [perfectly nested loop](#). [943](#)

implementation code

Implicit code that is introduced by the OpenMP implementation. [27](#), [41](#), [51](#), [93](#), [744](#)

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](#). [6](#), [8](#), [9](#), [16](#), [35](#), [43](#), [45](#), [51](#), [52](#), [77](#), [90](#), [93](#), [103](#), [113](#), [121](#), [122](#), [128](#), [131–134](#), [136–140](#), [142](#), [144](#), [145](#), [148–152](#), [161](#), [209](#), [220](#), [233](#), [235](#), [275](#), [282](#), [315–319](#), [329](#), [332](#), [336](#), [337](#), [339](#), [341](#), [342](#), [346](#), [347](#), [352](#), [359](#), [362](#), [369](#), [371](#), [372](#), [392](#), [398–401](#), [403](#), [405](#), [407–410](#), [412](#), [415](#), [417](#), [420](#), [424](#), [427](#), [434](#), [438](#), [440](#), [450](#), [457](#), [473](#), [483](#), [489](#), [491](#), [516](#), [553–555](#), [559](#), [561](#), [566](#), [578](#), [582](#), [594–596](#), [617](#), [631](#), [633](#), [634](#), [644](#), [648](#), [688](#), [704](#), [707](#), [709–711](#), [717](#), [719](#), [721](#), [725](#), [727](#), [729](#), [744](#), [751](#), [755](#), [758](#), [790](#), [805](#), [815](#), [820–823](#), [845](#), [873](#), [894](#), [899](#), [903](#), [915–926](#), [936](#), [941](#), [947](#)

implementation selector set

A [selector set](#) that may match the [implementation trait set](#). [338](#), [338](#), [340](#)

implementation trait set

The **trait set** that consists of **traits** that describe the functionality supported by the OpenMP implementation at a given point in the **OpenMP program**. 336, 59, 335, 337, 354

implicit array

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

COMMENT: For C/C++, the implicit array for pointer *p* with type *T* (*)[10] consists of all accessible elements *p*[*i*][*j*], for all *i* and *j*=0,1,...,9.

26, 27, 294

implicit barrier

A **barrier** that is specified as part of the semantics of a **construct** other than the **barrier construct**. 4–6, 403, 425, 427, 429, 431, 440, 468, 496, 497, 502, 541, 758

implicit flush

A **flush** that is specified as part of the semantics of a **construct** or **routine** other than the **flush construct**. 12, 104, 522, 944

implicitly determined data-mapping attribute

A **data-mapping attribute** that applies to a **data entity** for which no **data-mapping attribute** is otherwise determined. 289, 289, 305–307, 764

implicitly determined data-sharing attribute

A **data-sharing attribute** that applies to a **data entity** for which no **data-sharing attribute** is otherwise determined. 218, 98, 215, 218, 219, 223, 224, 289, 290, 306, 307, 945

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**. 3–5, 37, 60, 63, 97, 135, 279, 407, 413, 444, 466, 467, 602, 605, 620, 622, 713, 856, 951

implicit task

A **task** generated by an **implicit parallel region** or generated when a **parallel construct** is encountered during execution. 3, 4, 8, 20, 23, 28, 30, 37, 39, 52, 61, 63, 83, 85, 89, 102, 107, 108, 118–120, 127, 128, 219, 228, 252, 285, 287, 288, 402–404, 407–409, 423, 424, 426–434, 440, 521, 523, 544, 706, 744, 770, 784, 822, 826, 856, 891–893

implicit task region

A [region](#) that corresponds to an [implicit task](#). [3](#), [128](#), [784](#)

importing task

A [task](#) that permits a [preceding dependence-compatible task](#) to be an [antecedent task](#) of one of its [child tasks](#). [531](#), [111](#), [447](#), [457](#), [528](#), [531](#), [579](#)

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](#). [60](#), [598](#), [599](#)

inactive target region

A [target region](#) that is executed on the same [device](#) that encountered the [target construct](#). [127](#)

included task

A [task](#) for which execution is sequentially included in the [generating task region](#). That is, an [included task](#) is an [undelayed task](#) and executed by the [encountering thread](#). [7](#), [31](#), [53](#), [61](#), [93](#), [446](#), [459](#), [462](#), [474](#), [476](#), [479](#), [481](#), [486](#), [488](#), [499](#), [502](#), [623](#)

inclusive scan computation

A [scan computation](#) for which the value read includes the updates performed in the same [logical iteration](#). [272](#), [272](#), [942](#)

index-set splitting

The splitting of the [logical iteration space](#) into partitions that each are executed by a [generated loop](#). [396](#), [933](#)

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++), a dummy procedure (Fortran), or a procedure pointer (Fortran) that refers to the host version of the [procedure](#). [368](#)

induction

A use of an [induction operation](#). [62](#), [238](#)

induction attribute

For a given [loop-nest-associated construct](#), a [data-sharing attribute](#) of a [data entity](#) that implies the [private attribute](#) and for which the value is updated according to an [induction operation](#). [257](#), [66](#)

induction expression

A [collector expression](#) or an [inductor expression](#). [239](#), [239](#)

induction identifier

An [OpenMP identifier](#) that specifies an [inductor OpenMP operation](#) to use in an [induction](#). [238](#), [238](#), [239](#), [245–248](#), [258](#), [266](#), [267](#)

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 in 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](#). [33](#), [61](#), [62](#), [66](#), [101](#), [114](#), [178](#), [238](#), [242](#), [258](#), [269](#), [930](#)

induction variable

A [variable](#) for which an [induction operation](#) determines its values. [62](#), [242](#), [266](#)

inductor

A binary operator used by an [induction operation](#). [62](#), [242](#)

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](#). [242](#), [62](#), [178](#), [242](#), [243](#), [245](#), [247](#), [258](#), [267](#), [268](#)

informational directive

A [directive](#) that is neither [declarative](#) nor [executable](#), but otherwise conveys user code properties to the compiler. [369](#), [115](#), [155](#), [372](#), [380](#), [385](#), [386](#)

initialization phase

The portion of an [affected iteration](#) that includes all statements that initialize [private variables](#) prior to the [input phase](#) and [scan phase](#) of a [scan computation](#). [269](#), [269](#), [271](#), [273](#), [931](#)

initializer

An [OpenMP operation](#) that uses an [initializer expression](#). [248](#), [62](#), [92](#), [243](#), [244](#), [248](#), [252](#)

initializer expression

An [OpenMP stylized expression](#) that determines the [initializer](#) for the [private](#) copies of [list items](#) in a [reduction clause](#). [240](#), [62](#), [92](#), [178](#), [241–243](#), [247](#), [251](#), [263](#), [265](#), [270](#), [363](#)

initial task

An **implicit task** associated with an **implicit parallel region**. 4, 5, 29, 63, 97, 127, 128, 252, 407, 412, 413, 432, 440, 466, 467, 473, 482, 523, 703, 730, 744, 784, 811, 820, 826, 913

initial task region

A **region** that corresponds to an **initial task**. 3, 118, 119, 521, 523, 592, 598, 600

initial team

The **team** that comprises an **initial thread** executing an **implicit parallel region**. 4, 7, 107, 119, 412, 440, 442, 601, 857

initial thread

The **thread** that executes an **implicit parallel region**. 3, 4, 63, 86, 89, 109, 136, 138, 275, 412, 413, 431, 440, 444, 466, 467, 521, 523, 768, 916, 918

innermost-leaf property

The **property** that a **clause** applies to the innermost **leaf construct** that permits it when it appears on a **compound construct**. 163, 185, 226, 233, 259, 272, 273, 286, 465, 508–512, 526, 537, 538, 548

input map type

The **map type** specified in a **map clause** specified on a **construct** to which **map-type decay** is applied to determine an **output map type**. 297, 72, 85, 112, 297

input phase

The portion of a **logical iteration** that contains all computations that update a **list item** for which a **scan computation** is performed. 269, 62, 114, 269, 270, 272, 273

input place partition

The **place partition** that is used to determine the *place-partition-var* and *place-assignment-var* ICVs and the **place** assignments of the **implicit tasks** of a **parallel region**. 407, 407–409, 411

intent(in) property

The **property** that a **routine** argument is an **intent (in)** dummy argument in Fortran. In C/C++, the memory pointed to by the argument is not written by the runtime but must be readable. 555, 616, 617, 625, 626, 630, 632, 634, 636–638, 644–649, 652–659, 661–667, 669–673, 707, 709, 710, 716, 722, 751, 759, 774–778, 781, 785, 786, 788, 791, 792, 795, 796, 798, 800, 803, 806, 808, 813, 863, 865, 867, 869, 871, 873, 874, 883

intent(out) property

The **property** that a **routine** argument is an **intent (out)** dummy argument in Fortran. In C/C++, the memory pointed to by the argument is not read by the runtime but must be

writeable. [555](#), [644–646](#), [659](#), [662](#), [664](#), [684](#), [708](#), [710](#), [815](#), [816](#), [876](#), [882](#), [899](#), [901](#), [903](#), [905](#)

internal control variable

A conceptual [variable](#) that specifies runtime behavior of a set of [threads](#) or [tasks](#) in an [OpenMP program](#). [118](#), [58](#), [915](#)

interoperability object

An [OpenMP object](#) of [interop OpenMP type](#), which is an [opaque type](#). These objects represent information that supports interaction with [foreign runtime environments](#). [559](#), [64](#), [185](#), [345](#), [351](#), [357](#), [488–491](#), [559](#), [563](#), [643](#), [650](#), [923](#), [935](#), [939](#)

interoperability property

A [property](#) associated with an [interoperability object](#). [488](#), [64](#), [561](#), [643–646](#), [648](#), [649](#)

interoperability-property-retrieving property

The [property](#) that a [routine](#) retrieves an [interoperability property](#) from an [interoperability object](#). [643](#), [64](#), [644–646](#)

interoperability-property-retrieving routine

A [routine](#) that has the [interoperability-property-retrieving property](#). [643](#), [643](#), [645–647](#)

interoperability routine

A [routine](#) that has the [interoperability-routine property](#). [643](#), [488](#), [561](#), [563](#), [643](#), [650](#)

interoperability-routine property

The [property](#) that a [routine](#) provides a mechanism to inspect the [properties](#) associated with an [interoperability object](#). [643](#), [64](#), [644–649](#)

intervening code

For two consecutive [affected loops](#) of a [loop-nest-associated construct](#), user code that appears inside the [loop body](#) of the outer [affected loop](#) but outside the [loop body](#) of the inner [affected loop](#). [203](#), [86](#), [203](#), [209](#), [210](#), [454](#)

intrinsic identifier

Unless otherwise specified, an [OpenMP identifier](#) that represents a special [variable](#) and that is defined for use in [directives](#) but not in [base-language code](#). [188](#), [178](#), [239](#), [241–243](#), [396](#)

is-device-ptr attribute

For a given [device construct](#), a [data-sharing attribute](#) of a [variable](#) that implies the [private attribute](#), and additionally the [variable](#) is initialized with a [device address](#) that corresponds to the [device pointer variable](#) of the same name in the [enclosing data environment](#) of the [construct](#). [233](#)

ISO C binding property

The [property](#) of a [routine](#) that its Fortran version has the **BIND (C)** attribute. [65](#), [575](#), [576](#), [623–627](#), [629–631](#), [634](#), [635](#), [637](#), [638](#), [640](#), [641](#), [657](#), [662](#), [664](#), [665](#), [677–682](#)

ISO C property

The [property](#) that a [routine](#) argument has the **BIND (C)** attribute in Fortran. If any argument of a [routine](#) has the [ISO C property](#) then the [routine](#) has the [ISO C binding property](#). [555](#), [65](#), [575](#), [625–627](#), [629](#), [630](#), [632](#), [634](#), [636–638](#), [640](#), [641](#), [662](#), [664](#), [677–682](#), [684](#), [687](#), [796](#), [800](#), [803](#)

iteration count

The number of times that the [loop body](#) of a given loop is executed. [208](#), [208–210](#), [266](#), [392](#), [397](#), [401](#), [918](#), [920](#)

iterator

A programming mechanism to specify a set of values. [173](#), [174](#), [201](#), [209](#), [294](#), [418](#), [938](#), [949](#)

iterator specifier

A tuple that specifies an *iterator-identifier* and its associated [iterator value set](#). [172](#), [65](#), [166](#), [173](#)

iterator-specifier list

An [argument list](#) that consists of [iterator-specifier list items](#). [165](#)

iterator-specifier list item

A [list item](#) that is an [iterator specifier](#). [165](#), [65](#)

iterator value set

The set of values that correspond to a given instance of an *iterator modifier*. [173](#), [65](#), [173–175](#)

L

last-level cache

The last cache in a [memory](#) hierarchy that is used by a set of [cores](#). [131](#)

lastprivate attribute

For a given [construct](#), a [data-sharing attribute](#) of a [variable](#) that implies the [private attribute](#), and additionally, the final value of the [variable](#) may be assigned to the [variable](#) that has the same name in the [enclosing data environment](#) of the [construct](#). [230](#), [65](#), [216](#)

lastprivate variable

A [private variable](#) that has the [lastprivate attribute](#) with respect to a given [construct](#). [942](#)

leaf construct

For a given [construct](#), a [construct](#) that corresponds to one of the [leaf directives](#) of the [executable directive](#). [21](#), [33](#), [47](#), [63](#), [84](#), [179](#), [335](#), [536](#), [537](#), [548–551](#), [952](#)

leaf directive

For a given [directive](#), the [directive](#) itself if it is not a [compound directive](#), or a [directive](#) from which the [compound directive](#) is composed that is not itself a [compound directive](#). [36](#), [66](#), [99](#), [546](#)

leaf-directive name

The [directive name](#) of a [leaf directive](#). [545](#), [545–547](#), [953](#)

league

The set of [teams](#) formed by a [teams construct](#), each of which is associated with a different [contention group](#). [4](#), [107](#), [119](#), [252](#), [412](#), [413](#), [440](#), [442](#), [601](#), [750](#), [784](#)

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 $k \in \{1, \dots, m - 1\}$. [398](#), [399](#)

linear attribute

For a given [loop-nest-associated construct](#), a [data-sharing attribute](#) of a [variable](#) that is equivalent to an [induction attribute](#) for which the [induction operation](#) is a linear recurrence, where the binary operator \oplus is $+$ and the [step expression](#) s is a loop-invariant integer expression. [260](#), [66](#)

linear variable

A [private variable](#) that has the [linear attribute](#) with respect to a given [construct](#). [260](#)

list

A comma-separated set. [22](#), [39–41](#), [66](#), [87](#), [161](#), [162](#), [165](#), [167](#), [363](#), [367](#), [405](#), [464](#), [724](#), [916](#)

list item

A member of a [list](#). [22](#), [23](#), [34](#), [38–41](#), [47](#), [50](#), [52](#), [62](#), [63](#), [65](#), [67](#), [70–73](#), [75](#), [78](#), [82](#), [84](#), [85](#), [88–90](#), [100](#), [103](#), [111](#), [115](#), [144](#), [161](#), [163](#), [165](#), [167–169](#), [172–174](#), [178](#), [215–217](#), [219–223](#), [225](#), [226](#), [228](#), [229](#), [231–237](#), [239](#), [240](#), [242–244](#), [246–250](#), [252–262](#), [269–273](#), [276](#), [277](#), [279–284](#), [286–289](#), [293–305](#), [308–311](#), [322–324](#), [326](#), [332](#), [333](#), [345](#), [349](#), [350](#), [355](#), [357](#), [362–367](#), [380](#), [381](#), [389–391](#), [396](#), [397](#), [399](#), [419](#), [441](#), [443](#), [450](#), [456–458](#), [464](#), [465](#), [474](#), [476](#), [479](#), [481–484](#), [486](#), [519](#), [520](#), [528–530](#), [541](#), [542](#), [548–551](#), [904](#), [918](#), [928](#), [929](#), [931](#), [932](#), [936](#), [937](#), [942](#), [949](#)

local static variable

A [variable](#) with [static storage duration](#) that for C/C++ has block scope and for Fortran is declared in the specification part of a [procedure](#) or **BLOCK** construct. [316](#), [320](#)

locator list

An [argument list](#) that consists of [locator list items](#). [165](#), [163](#), [310](#), [457](#)

locator list item

A [list item](#) that refers to [storage locations](#) in [memory](#) and is one of the items specifically identified in [Section 5.2.1](#). [166](#), [67](#), [166–168](#), [185](#), [455–457](#), [525–528](#), [530](#)

lock

An OpenMP [variable](#) that is used in [lock routines](#) to enforce mutual exclusion. [67](#), [68](#), [75–77](#), [82](#), [99](#), [112](#), [113](#), [469](#), [516](#), [521](#), [524](#), [578](#), [581](#), [688–692](#), [694–701](#), [759](#), [768](#), [795](#), [816](#), [823](#), [924](#), [946](#)

lock-acquiring property

The [property](#) that a [routine](#) may acquire a [lock](#) by putting it into the [locked state](#). [694](#), [67](#), [688](#), [695](#), [696](#)

lock-acquiring routine

A [routine](#) that has the [lock-acquiring property](#). [694](#), [469](#), [688](#), [694](#), [695](#), [699](#), [791–794](#)

lock-destroying property

The [property](#) that a [routine](#) destroys a [lock](#) by putting it into the [uninitialized state](#). [692](#), [67](#), [693](#)

lock-destroying routine

A [routine](#) that has the [lock-destroying property](#). [692](#), [692–694](#), [793](#), [794](#)

locked state

The [lock state](#) that indicates the [lock](#) has been set by some [task](#). [688](#), [67](#), [68](#), [697](#)

lock-initializing property

The [property](#) that a [routine](#) initializes a [lock](#) by putting it into the [unlocked state](#). [689](#), [67](#), [689–691](#)

lock-initializing routine

A [routine](#) that has the [lock-initializing property](#). [689](#), [689–692](#), [791](#), [792](#)

lock property

The [property](#) that a [routine](#) operates on [locks](#). [688](#), [68](#)

lock-releasing property

The [property](#) that a [routine](#) may unset a [lock](#) by returning it to the [unlocked state](#). [697](#), [68](#), [688](#), [697](#), [698](#)

lock-releasing routine

A [routine](#) that has the [lock-releasing property](#). [697](#), [469](#), [688](#), [697](#), [793](#), [794](#)

lock routine

A [routine](#) that has the [lock property](#). [688](#), [67](#), [554](#), [688](#), [924](#)

lock state

The state of a [lock](#) that determines if it can be set. [688](#), [67](#), [112](#), [113](#), [688](#), [697–699](#)

lock-testing property

The [property](#) that a [routine](#) that may set a [lock](#) by putting it into the [locked state](#) does not suspend execution of the [task](#) that executes the [routine](#) if it cannot set the [lock](#). [699](#), [68](#), [699](#), [700](#)

lock-testing routine

A [routine](#) that has the [lock-testing property](#). [699](#), [699](#), [792–794](#)

logical iteration

An instance of the executed [loop body](#) of a [canonical loop nest](#), or a **DO CONCURRENT** loop in Fortran, denoted by a number in the [logical iteration space](#) of the loops that indicates an order in which the [logical iteration](#) would be executed relative to the other [logical iterations](#) in a sequential execution. [4](#), [20](#), [32](#), [34](#), [48](#), [51](#), [61](#), [63](#), [68](#), [69](#), [94](#), [96](#), [101](#), [110](#), [114](#), [178](#), [209](#), [210](#), [252](#), [386](#), [388](#), [392](#), [393](#), [395–399](#), [401](#), [419](#), [449–453](#), [744](#), [780](#), [920](#), [921](#), [938](#), [939](#), [943](#), [945](#), [949](#)

logical iteration space

For a [canonical loop nest](#), or a **DO CONCURRENT** loop in Fortran, the sequence $0, \dots, N - 1$ where N is the number of distinct [logical iterations](#). [209](#), [392](#), [32](#), [48](#), [53](#), [56](#), [61](#), [68](#), [110](#), [178](#), [209](#), [391](#), [396–399](#)

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. [66](#), [68](#), [90](#), [210](#), [398](#), [399](#), [938](#)

logical iteration vector space

The set of [logical iteration vectors](#) that each correspond to a [logical iteration](#) of a [canonical loop nest](#). [210](#), [398](#), [399](#)

loop body

A [structured block](#) that encompasses the executable statements that are iteratively executed by a loop statement. [202](#), [64](#), [65](#), [68](#), [396](#), [454](#)

loop-collapsing construct

A [loop-nest-associated construct](#) for which some number of outer loops of the [associated loop nest](#) may be [collapse-affected loops](#). [31](#), [32](#), [210](#), [220](#), [260](#), [416](#)

loop-iteration variable

For a loop of a [canonical loop nest](#), *var* as defined in [Section 6.4.1](#). A C++ range-based **for**-statement has no [loop-iteration variable](#). [69](#), [175](#), [201](#), [205–210](#), [216–218](#), [231](#), [260](#), [261](#), [388](#), [443](#), [454](#), [533](#), [549](#), [551](#), [949](#)

loop-iteration vector

An n -tuple (i_1, \dots, i_n) that identifies a [logical iteration](#) of the [affected loops](#) of a [loop-nest-associated directive](#), where n is the number of [affected loops](#) and i_k is the value of the [loop-iteration variable](#) of the k^{th} [affected loop](#), from outermost to innermost. [69](#), [208](#), [209](#), [532](#), [533](#)

loop-iteration vector space

The set of [loop-iteration vectors](#) that each corresponds to a [logical iteration](#) of the [affected loops](#) of a [loop-nest-associated directive](#). [209](#), [208](#), [209](#)

loop-nest-associated construct

A [loop-nest-associated directive](#) and its [associated loop nest](#). [61](#), [64](#), [66](#), [69](#), [94](#), [96](#), [99](#), [117](#), [158](#), [210](#), [259](#), [261](#), [389](#), [390](#), [398](#), [399](#), [423](#), [532](#), [551](#)

loop-nest-associated directive

An [executable directive](#) for which the associated user code must be a [canonical loop nest](#). [157](#), [20](#), [24](#), [69](#), [156](#), [157](#), [203](#), [208](#), [216](#), [217](#), [258](#), [261](#), [388](#), [389](#), [392](#), [393](#), [395](#), [397](#), [398](#), [400](#), [417](#), [435](#), [436](#), [439](#), [442](#), [449](#), [537](#)

loop nest depth

For a [canonical loop nest](#), the maximal number of loops, including the outermost loop, that can be affected by a [loop-nest-associated directive](#). [68](#), [208](#), [211](#), [213](#), [391](#)

loop schedule

The manner in which the [collapsed iterations](#) of [affected loops](#) are to be distributed among a set of [threads](#) that cooperatively execute the [affected loops](#). [210](#), [35](#), [94](#), [96](#), [210](#), [416](#), [423](#), [433](#), [440](#), [443](#), [937](#)

loop-sequence-associated construct

A [loop-sequence-associated directive](#) and its associated [canonical loop sequence](#). [70](#), [213](#)

loop-sequence-associated directive

An [executable directive](#) for which the associated user code must be a [canonical loop sequence](#). [157](#), [24](#), [69](#), [156](#), [388](#), [392](#)

loop sequence length

For a [canonical loop sequence](#), the number of consecutive [canonical loop nests](#) regardless of their nesting into blocks. [208](#), [213](#)

loop-sequence-transforming construct

A [loop-sequence-associated construct](#) with the [loop-transforming property](#). [388](#)

loop-transforming construct

A [loop-transforming directive](#) and its [associated loop nest](#) or associated [canonical loop sequence](#). [388](#), [55](#), [78](#), [111](#), [202](#), [208](#), [210](#), [386–391](#), [396](#), [451](#), [932](#), [933](#), [936](#), [939](#)

loop-transforming directive

A [directive](#) with the [loop-transforming property](#). [55](#), [70](#), [111](#), [388](#), [390](#), [391](#), [397](#)

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 [affected loops](#). [70](#), [386](#), [392](#), [393](#), [395](#), [397](#), [398](#), [400](#)

loosely structured block

For Fortran, 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. [102](#), [157](#)

M

map-entering clause

A [map clause](#) that, if it appears on a [map-entering construct](#), specifies that the reference counts of [corresponding list items](#) are increased and, as a result, those [list items](#) may enter the [device data environment](#). [296](#), [70](#), [304](#), [378](#), [475](#)

map-entering construct

A [construct](#) that has the [map-entering property](#). [70](#), [296](#), [299](#), [300](#), [302](#), [303](#), [547](#), [584](#), [928](#)

map-entering map type

A [map-type](#) that specifies the [clause](#) on which it is specified is a [map-entering clause](#). [296](#), [296](#), [297](#)

map-entering property

A [property](#) of a [construct](#) that it may include [mapping operations](#) that allocate storage on the [target device](#) and that result in assignment to the [corresponding list item](#) from the [original list item](#). [70](#), [297](#), [474](#), [478](#), [481](#)

map-exiting clause

A [map clause](#) that, if it appears on a [map-exiting construct](#), specifies that the reference counts of [corresponding list items](#) are decreased and, as a result, those [list items](#) may exit the [device data environment](#). [296](#), [71](#), [477](#)

map-exiting construct

A [construct](#) that has the [map-exiting property](#). [71](#), [296](#), [303](#), [547](#)

map-exiting map type

A [map-type](#) that specifies the [clause](#) on which it is specified is a [map-exiting clause](#). [296](#), [296](#), [297](#)

map-exiting property

A [property](#) of a [construct](#) that it may include [mapping operations](#) that release storage on the [target device](#) and that result in assignment from the [corresponding list item](#) to the [original list item](#). [71](#), [297](#), [476](#), [478](#), [481](#)

mappable storage block

A [storage block](#), derived from the [list items](#) of [map clauses](#) specified on a [data-mapping construct](#), for which a [corresponding storage block](#) in a [device data environment](#) is created, removed, or otherwise referenced by the [construct](#). [299](#), [302](#), [303](#), [305](#), [311](#)

mappable type

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

For C, the type must be a complete type.

For C++, the type must be a complete type; in addition, for class types:

- All member functions accessed in any [target region](#) must appear in a [declare target directive](#).

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

- All type-bound procedures accessed in any [target region](#) must appear in a [declare_target directive](#).

COMMENT: Pointer types are [mappable types](#) but the memory block to which the pointer refers is not mapped.

71, 283, 284, 294, 311

mapped address range

For a given [original list item](#), the [address range](#) that starts from its [starting address](#) and ends with its [ending address](#). [300](#), [72](#), [300](#)

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. [38](#), [50](#), [71](#), [72](#), [84](#), [100](#), [484](#), [584](#)

mapper

An operation that defines how [variables](#) of given type are to be mapped or updated with respect to a [device data environment](#). [42](#), [49](#), [114](#), [188](#), [289](#), [291](#), [297](#), [301](#), [302](#), [308–314](#)

mapper identifier

An [OpenMP identifier](#) that specifies the name of a [user-defined mapper](#). [291](#), [291](#), [310](#)

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](#). [9](#), [23](#), [25](#), [71](#), [72](#), [97](#), [294](#), [296](#), [304](#), [378](#), [585](#), [759](#), [764](#), [765](#), [931](#), [932](#)

map type

A categorization of a [data-mapping clause](#) that determines whether the [mapping operations](#) that result from that [clause](#) include assignments between the [original storage](#) and [corresponding storage](#) of its [list items](#). [63](#), [85](#), [112](#), [293](#), [303](#)

map-type decay

A process applied to [input map type](#), according to an [underlying map type](#), that results in an [output map type](#). [297](#), [63](#), [85](#), [293](#), [297](#), [301](#), [479](#)

map-type-modifying property

The [property](#) that a [modifier](#) that combines with a [map-type](#) to determine details of a [mapping operation](#). [292](#), [293](#), [301](#)

matchable candidate

A [mapped variable](#) for which [corresponding storage](#) was created in a [device data environment](#). [300](#), [72](#), [300](#)

matched candidate

A [matchable candidate](#) that, due to a matching [mapped address range](#) or [extended address range](#), may determine the lower bound and length to use for a given [assumed-size array](#) that is

a [list item](#) in a [map](#) clause. [300](#), [103](#), [234](#), [293](#), [300](#), [936](#)

matching taskgraph record

A [finalized taskgraph record](#) that has a matching value for the scalar expression that identifies a [taskgraph](#) region. [456](#), [94](#), [455–459](#)

memory

A storage resource for storing and retrieving [variables](#) that are accessible by [threads](#). [7](#), [6–11](#), [13](#), [19](#), [21](#), [32](#), [45](#), [53](#), [56](#), [65](#), [67](#), [73–75](#), [78](#), [91](#), [94](#), [101](#), [103](#), [104](#), [108](#), [110](#), [117](#), [119](#), [147](#), [168](#), [231](#), [281](#), [315–319](#), [329](#), [376](#), [377](#), [504–507](#), [514](#), [519](#), [529](#), [564](#), [575](#), [582](#), [623](#), [628](#), [629](#), [633](#), [640](#), [651](#), [660](#), [664](#), [668](#), [675–677](#), [682](#), [745](#), [800](#), [804](#), [805](#), [827](#), [849](#), [854](#), [861–866](#), [868](#), [875](#), [882](#), [901](#), [903](#), [905](#), [915](#), [931](#), [932](#), [934](#), [935](#), [940–942](#), [946](#), [948](#)

memory-allocating routine

A [memory-management routine](#) that has the [memory-allocating-routine](#) property. [675](#), [21](#), [74](#), [91](#), [117](#), [675–677](#), [682](#)

memory-allocating-routine property

The [property](#) that a [memory-management routine](#) allocates [memory](#). [675](#), [73](#), [677–681](#)

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](#). [9](#), [9](#), [21](#), [24](#), [73](#), [74](#), [119](#), [295](#), [316–325](#), [329](#), [375](#), [483](#), [566](#), [668](#), [674–676](#), [682](#), [919](#), [931](#), [935](#), [942](#)

memory-allocator-retrieving property

The [property](#) that a [memory-management routine](#) retrieves a [memory allocator](#) handle. [668](#), [73](#), [669–672](#)

memory-allocator-retrieving routine

A [memory-management routine](#) that has the [memory-allocator-retrieving](#) property. [668](#), [668–673](#)

memory-copying property

The [property](#) that a [routine](#) copies [memory](#) from the [device data environment](#) of one [device](#) to the [device data environment](#) of another [device](#). [632](#), [73](#), [634](#), [635](#), [637](#), [638](#)

memory-copying routine

A [routine](#) that has the [memory-copying](#) property. [632](#), [53](#), [91](#), [468](#), [633](#), [634](#)

memory-management routine

A [routine](#) that has the [memory-management-routine](#) property. [651](#), [21](#), [56](#), [73–75](#), [651](#), [657](#), [658](#)

memory-management-routine property

The [property](#) that a [routine](#) manages [memory](#) on the [current device](#). [651](#), [73](#), [652–659](#), [661–667](#), [669–674](#), [677–682](#), [684–686](#)

memory part

A [storage block](#) that resides on a single storage resource within a [memory space](#). [74](#)

memory partition

A definition of how a [memory allocator](#) divides the allocated memory into [memory parts](#) and the storage resources on which it allocates those [memory parts](#). [74](#), [318](#), [573](#), [575–577](#), [660](#), [662–666](#)

memory partitioner

An [OpenMP object](#) that represents mechanisms to create and to destroy [memory partitions](#). [74](#), [317](#), [318](#), [568](#), [574](#), [576](#), [659–661](#), [663](#), [665](#), [666](#)

memory-partitioning property

The [property](#) that a [memory-management routine](#) creates or destroys or otherwise affects [memory partitions](#) or [memory partitioners](#). [659](#), [74](#), [659](#), [661–665](#)

memory-partitioning routine

A [memory-management routine](#) that has the [memory-partitioning property](#). [659](#)

memory-reading callback

A [callback](#) that has the [memory-reading property](#). [865](#), [866](#), [867](#)

memory-reading property

The [property](#) that a [callback](#) reads [memory](#) from an [OpenMP program](#). [865](#), [74](#), [866](#)

memory-reallocating routine

A [memory-management routine](#) that has the [memory-reallocating-routine property](#). [676](#), [676](#), [677](#), [681](#)

memory-reallocating-routine property

The [property](#) that a [memory-allocating routine](#) deallocates [memory](#) in addition to allocating it. [74](#), [681](#)

memory-setting property

The [property](#) that a [routine](#) fills [memory](#) in a [device data environment](#) with a specified value. [639](#), [74](#), [640](#), [641](#)

memory-setting routine

A [routine](#) that has the [memory-setting property](#). [639](#), [468](#), [640–642](#)

memory space

A representation of storage resources from which [memory](#) can be allocated or deallocated. More than one [memory space](#) may exist. [651](#), [9](#), [24](#), [57](#), [74](#), [75](#), [105](#), [147](#), [295](#), [315](#), [318](#), [328](#), [329](#), [575](#), [576](#), [651](#), [652](#), [657–659](#), [664](#), [666](#), [667](#), [669](#), [919](#), [935](#), [942](#)

memory-space-retrieving property

The [property](#) that a [memory-management routine](#) retrieves a [memory space](#) handle. [651](#), [75](#), [652–655](#)

memory-space-retrieving routine

A [memory-management routine](#) that has the [memory-space-retrieving property](#). [651](#), [651–656](#)

mergeable task

A [task](#) that may be a [merged task](#) if it is an [underrferred task](#). [460](#), [105](#), [448](#), [460](#), [488](#), [499](#)

merged task

A [task](#) with a [minimal data environment](#). [75](#), [448](#), [460](#), [469](#), [479](#), [744](#), [807](#), [912](#)

metadirective

A [directive](#) that conditionally resolves to another [directive](#). [341](#), [47](#), [48](#), [94](#), [156](#), [341–344](#), [380](#), [919](#), [936](#), [937](#), [939](#), [943](#)

minimal data environment

A [data environment](#) of a [task](#) that, inclusive of ICVs, is the same as that of its [enclosing context](#), with the exception of [list items](#) in [all-data-environments clauses](#) that are specified on the [task-generating construct](#) that generated the [task](#). [21](#), [75](#), [234](#), [237](#)

modifier

A mechanism to customize [clause](#) behavior for its specified arguments. [xxviii](#), [22](#), [34](#), [36](#), [42](#), [51](#), [54](#), [65](#), [72](#), [78](#), [82](#), [83](#), [88](#), [89](#), [93](#), [94](#), [99](#), [112](#), [113](#), [129](#), [162–165](#), [167](#), [173](#), [175](#), [179](#), [186](#), [224](#), [231](#), [232](#), [237](#), [248](#), [251](#), [261](#), [271](#), [291](#), [293](#), [294](#), [296](#), [297](#), [300–302](#), [309–311](#), [327–329](#), [332](#), [349](#), [350](#), [359](#), [360](#), [365](#), [433](#), [438](#), [440](#), [455–458](#), [479](#), [488](#), [490](#), [491](#), [525](#), [534](#), [548](#), [549](#), [765](#), [919](#), [921](#), [922](#), [928](#), [930–934](#), [936–944](#)

mutex-acquiring callback

A [callback](#) that has the [mutex-acquiring property](#). [791](#)

mutex-acquiring property

The [property](#) of a [callback](#) that it is dispatched when attempting to acquire mutually-exclusive access for a [mutual-exclusion construct](#) or when initializing or attempting to acquire a [lock](#). [791](#), [75](#), [792](#)

mutex-execution callback

A [callback](#) that has the [mutex-execution property](#). [793](#)

mutex-execution property

The [property](#) of a [callback](#) that it is dispatched when mutually-exclusive access is acquired or released for a [mutual-exclusion construct](#) or when a [lock](#) is acquired, released, or destroyed. [793](#), [76](#), [793](#), [794](#)

mutual-exclusion construct

A [construct](#) that has the [mutual-exclusion property](#). [75](#), [76](#), [791–794](#)

mutual-exclusion property

The [property](#) that a [construct](#) provides mutual-exclusion semantics. [76](#), [493](#), [514](#), [534](#), [536](#)

mutually exclusive tasks

[Tasks](#) that may be executed in any order, but not at the same time. [468](#), [528](#)

N

named-handle property

The [property](#) that a [handle](#) is an integer kind in Fortran that is distinguished by the name of the [handle](#). [558](#), [573](#), [578–580](#)

named parameter list item

A [parameter list item](#) that is the name of a parameter of a [procedure](#). [166](#), [166](#), [331](#), [332](#)

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->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: p_0 , $(*p_0).x_0[k_1].p_1$, and $(*p_0).x_0[k_1].p_1->p_2$.

[76](#), [169](#)

name-list trait

A [trait](#) that is defined with [properties](#) that match the names that identify particular instances of the [trait](#) that are effective at a given point in an [OpenMP program](#). [335](#), [336](#), [338](#), [340](#)

native thread

An execution entity upon which an [OpenMP thread](#) may be implemented. [3](#), [5](#), [6](#), [77](#), [82](#), [83](#), [90](#), [109](#), [120](#), [138](#), [139](#), [403](#), [413](#), [416](#), [744](#), [758](#), [759](#), [768](#), [771](#), [773](#), [774](#), [803](#), [813](#), [845](#), [857](#), [858](#), [864](#), [884–886](#), [896](#), [908](#)

native thread context

A [tool context](#) that refers to a [native thread](#). [850](#), [864](#), [865](#), [868](#), [869](#)

native thread handle

A [handle](#) that refers to a [native thread](#). [856](#), [883–886](#), [896](#), [899](#)

native thread identifier

An identifier for a [native thread](#) defined by a [native thread](#) implementation. [141](#), [850](#), [858](#), [869](#), [880](#), [884](#), [885](#)

native trace format

A format for [implementation defined trace records](#) that may be [device-specific](#). [77](#), [729](#), [730](#), [840](#), [842](#)

native trace record

A [trace record](#) in a [native trace format](#). [730](#), [751](#), [752](#), [840](#), [842](#)

nestable lock

A [lock](#) that can be acquired (i.e., set) multiple times by the same [task](#) before being released (i.e., unset). [688](#), [77](#), [524](#), [580](#), [688](#), [689](#), [696](#), [759](#), [795](#), [823](#)

nestable lock property

The [property](#) that a [routine](#) operates on [nestable locks](#). [688](#), [77](#), [690](#), [691](#), [693](#), [696](#), [698](#), [700](#)

nestable lock routine

A [routine](#) that has the [nestable lock property](#). [688](#), [580](#)

nested construct

A [construct](#) (lexically) enclosed by another [construct](#). [215](#)

nested parallelism

A condition in which more than one level of parallelism is [active](#) at a point in the execution of an [OpenMP program](#). [4](#), [941](#)

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](#). [3](#), [37](#), [77](#), [86](#), [386](#), [423](#)

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. [220](#), [38](#), [89](#), [90](#), [114](#), [220](#), [221](#), [226](#), [228](#), [229](#), [231](#), [233](#), [234](#), [258](#), [260](#), [270](#), [293](#), [299](#), [303](#), [304](#), [949](#)

NUMA domain

A [device](#) partition in which the closest [memory](#) to all [cores](#) is the same [memory](#) and is at a similar distance from the [cores](#). [131](#)

non-negative property

The [property](#) that an expression, including one that is used as the argument of a [clause](#), a [modifier](#) or a [routine](#), has a value that is greater than or equal to zero. [165](#), [122](#), [133](#), [134](#), [137](#), [143](#), [145–147](#), [163](#), [166](#), [209](#), [316](#), [328](#), [339](#), [396](#), [402](#), [412](#), [463](#), [561](#), [596](#), [602](#), [620](#), [658](#), [704](#), [719](#), [797](#), [820](#), [822](#), [923](#), [924](#)

non-conforming program

An [OpenMP program](#) that is not a [conforming program](#). [2](#), [34](#), [42](#), [113](#), [219](#), [276](#), [449](#), [469](#), [525](#), [688](#), [932](#)

non-constant property

The [property](#) that an expression, including one that is used as the argument of a [clause](#), a [modifier](#) or a [routine](#), is not a compile-time constant. [338](#)

non-host declare target directive

A [declare target directive](#) that does not specify a [device_type](#) clause with **host**. [362](#)

non-host device

A [device](#) that is not the [host device](#). [7](#), [19](#), [26](#), [103](#), [120](#), [122](#), [123](#), [130](#), [142](#), [143](#), [346](#), [376](#), [379](#), [403](#), [445](#), [470](#), [484](#), [614](#), [687](#), [714](#), [717](#), [879](#), [880](#), [886](#), [920](#), [928](#)

non-null pointer

A pointer that is not [NULL](#). [643](#), [685](#), [722](#), [724](#), [729](#), [771](#), [772](#), [842](#)

non-null value

A value that is not [NULL](#). [676](#), [756](#), [825](#), [826](#), [846](#), [864](#), [865](#), [868](#), [900](#)

non-property trait

A [trait](#) that is specified without additional [properties](#). [335](#), [336](#), [340](#)

nonrectangular-compatible property

The [property](#) that the transformation defined by a [loop-transforming construct](#) is compatible with [non-rectangular loops](#) and therefore will not yield a non-conforming [canonical loop nest](#) due to their presence. [388](#), [389](#), [393](#)

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. 78, 205, 207, 210, 212, 259, 261, 389, 392–394, 398, 399, 439, 442, 453, 920, 941

non-sequentially consistent atomic construct

An **atomic construct** for which the **seq_cst** clause is not specified 13

NULL

A null pointer. For C/C++, the value **NULL** or the value **nullptr**. For Fortran, the disassociated pointer for variables that have the **POINTER** attribute or the value **C_NULL_PTR** for variables of type **C_PTR**. 78, 148, 349, 610, 617, 625–630, 632, 633, 640, 648, 649, 676, 682, 683, 685, 708, 710, 711, 719, 722, 724, 729, 770, 783, 784, 789, 790, 797, 799, 800, 805, 807, 814, 817, 818, 823–827, 846, 864, 865, 868, 873, 901, 925, 926

numeric abstract name

An **abstract name** that refers to a quantity associated with a **conceptual abstract name**. 131, 19, 87, 131–133, 137, 928

O

offsetting loop

The outer **generated loops** of a **stripe construct** that determine the offsets within the grid cells used for each execution of the **grid loops**. 397, 397, 398, 920

OMPD

An interface that helps a **third-party tool** inspect the OpenMP state of a program that has begun execution. 844, 2, 14, 15, 79, 110, 119, 150, 189, 190, 844–847, 851, 852, 856, 858, 861, 864, 869, 874–878, 884, 908

OMPD callback

A **callback** that has the **OMPD** property. 189, 190, 851, 854, 856, 859, 861, 864, 866, 868, 869

OMPD library

A dynamically loadable library that implements the **OMPD** interface. 844, 15, 46, 844–851, 854, 857–859, 861–871, 873–880, 882, 896, 899, 901, 903, 905

OMPD property

The **property** that a **callback**, **routine** or type is included in **OMPD** and its namespace, which implies it has the **ompd_** prefix. 79, 80, 847, 848, 850–860, 862, 863, 865–868, 870–878, 880–906

OMPD routine

A [routine](#) that has the [OMPD property](#). [854](#), [856](#), [859](#), [874](#), [875](#), [877–879](#), [884](#), [885](#), [887–891](#), [904–906](#)

OMPD type

A type that has the [OMPD property](#). [189](#), [33](#), [57](#), [83](#), [85](#), [189](#), [190](#), [847–849](#), [851–865](#), [868–871](#), [873](#)

OMPT

An interface that helps a [first-party tool](#) monitor the execution of an [OpenMP program](#). [721](#), [2](#), [14](#), [15](#), [46](#), [80](#), [100](#), [148](#), [149](#), [190](#), [496](#), [585](#), [714](#), [721–725](#), [727–730](#), [747](#), [751](#), [752](#), [758](#), [770–772](#), [798](#), [813](#), [814](#), [830](#), [831](#), [840](#), [841](#), [906](#), [925](#), [935](#)

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. [719](#), [724](#), [725](#), [732](#)

OMPT callback

A [callback](#) that has the [OMPT property](#). [190](#), [727](#), [736](#), [739](#), [770](#), [814](#), [830](#)

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](#). [719](#), [723–725](#), [771](#)

OMPT interface state

A state that indicates the permitted interactions between a [first-party tool](#) and the OpenMP implementation. [80](#), [719](#), [723–725](#), [732](#), [771](#)

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](#). [723](#), [724](#)

OMPT property

The [property](#) that a [callback](#), [runtime entry point](#) or type is included in [OMPT](#) and its namespace, which implies it has the `ompt_` prefix. [80](#), [81](#), [721](#), [722](#), [734–737](#), [739–757](#), [759–762](#), [764–769](#), [771–778](#), [780–796](#), [798–803](#), [806](#), [808](#), [810](#), [813–825](#), [827–829](#), [831–842](#)

OMPT-tool finalizer

An implementation of the [finalize](#) callback. [732](#), [466](#), [722](#), [772](#)

OMPT-tool initializer

An implementation of the [initialize](#) callback. [721](#), [466](#), [722](#), [724](#), [725](#), [727](#), [771](#)

OMPT type

A type that has the [OMPT property](#). [xxviii](#), [189](#), [33](#), [57](#), [83](#), [85](#), [190](#), [434](#), [721](#), [722](#), [724](#), [727](#), [729](#), [730](#), [732](#), [733](#), [735](#), [736](#), [738–740](#), [742–756](#), [758](#), [760–764](#), [766–769](#), [771–774](#), [776–780](#), [782–802](#), [804](#), [805](#), [807](#), [809](#), [811](#), [814–824](#), [826–842](#), [852](#), [858](#), [893](#), [899](#), [906](#), [925](#), [927](#), [935](#), [937](#), [940](#)

once-for-all-constituents property

The [property](#) that a [clause](#) applies once for all [constituent constructs](#) to which it applies when it appears on a [compound construct](#). [163](#), [210](#), [211](#), [548](#)

opaque property

The [property](#) that an [OpenMP type](#) is opaque, which implies that objects of that type may only be accessed, modified and destroyed through OpenMP [directives](#), [routines](#), [callbacks](#) and [entry points](#). Further, an object of an [opaque type](#) can be copied without affecting, or copying, its underlying state. Destruction of an [OpenMP object](#), which by definition has an [opaque type](#), destroys the state to which all copies of the object refer. All [handles](#) have [opaque types](#). [81](#), [558](#), [559](#), [573](#), [578–580](#), [644–649](#), [735](#), [736](#), [742](#), [798](#), [802](#), [839–841](#), [869](#), [878–882](#), [886](#), [887](#), [889](#), [892](#), [894–901](#), [903–906](#)

opaque type

A type that has the [opaque property](#). [64](#), [81](#), [82](#), [558](#), [559](#), [573](#), [574](#), [578](#), [580](#)

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/>. [81](#), [143](#), [336](#), [489](#), [559](#), [561](#)

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. [45](#), [82](#), [95](#), [118](#), [130](#), [239](#), [376](#), [377](#), [384](#), [553](#), [606](#), [651](#), [712](#), [718](#), [923](#)

OpenMP architecture

The architecture on which a [region](#) executes. [82](#), [724](#)

OpenMP context

The execution context of an [OpenMP program](#) as represented by a set of [traits](#), including active [constructs](#), execution [devices](#), OpenMP functionality supported by the implementation

and any available dynamic values. [335](#), [33](#), [37](#), [100](#), [188](#), [335](#), [337](#), [338](#), [340–342](#), [346–348](#), [352](#), [354](#), [358](#), [372](#), [561](#), [919](#), [939](#)

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](#). [130](#), [46](#), [50](#), [118](#), [123](#), [130](#), [901](#), [947](#)

OpenMP identifier

An identifier that has a specialized purpose for use in [OpenMP programs](#), as defined by this specification. [188](#), [62](#), [64](#), [72](#), [88](#), [92](#), [95](#), [162](#), [168](#), [188](#), [190](#), [239](#)

OpenMP lock variable

A [lock](#). [688](#)

OpenMP object

Any object of an [opaque type](#) that allows programmers to save, to manipulate and to use state related to the OpenMP API. [43](#), [64](#), [73](#), [74](#), [81](#), [525](#), [799](#), [802](#), [831](#), [839](#), [841](#), [842](#)

OpenMP operation

When used as a [list item](#), a special expression that returns an object of a specified [OpenMP types](#). Otherwise, an operation that is applied to a [list item](#) according to the semantics of a [directive](#), [clause](#), or [modifier](#). [168](#), [62](#), [82](#), [92](#), [166](#), [168](#), [188](#), [350](#), [425](#), [519](#)

OpenMP operation list

An [argument list](#) that consists of [OpenMP operation list items](#). [165](#), [168](#)

OpenMP operation list item

A [list item](#) that is an [OpenMP operation](#). [165](#), [82](#)

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. [20](#), [82](#), [847](#), [848](#), [857](#), [864](#), [874](#), [875](#), [878](#), [879](#)

OpenMP program

A program that consists of a [base program](#) that is annotated with [OpenMP directives](#) or that calls [OpenMP API routines](#). [3](#), [5–9](#), [13](#), [14](#), [19](#), [22](#), [26](#), [32](#), [35](#), [36](#), [44–46](#), [49](#), [56](#), [58–60](#), [64](#), [74](#), [76–78](#), [80–82](#), [93](#), [95](#), [111](#), [113](#), [118](#), [120](#), [130](#), [141](#), [151](#), [152](#), [168](#), [188](#), [219](#), [223](#), [250](#), [261](#), [276](#), [282](#), [308](#), [309](#), [315](#), [316](#), [335–337](#), [342](#), [377](#), [386](#), [413](#), [423](#), [463](#), [483](#), [484](#), [492](#), [493](#), [517](#), [518](#), [520](#), [525](#), [602](#), [605](#), [612](#), [620](#), [622](#), [633](#), [688](#), [702](#), [712](#), [714](#), [715](#), [718](#), [719](#),

721, 723, 724, 727, 745, 746, 770, 797, 817, 824, 829, 830, 836, 844–846, 849, 854, 857,
863, 865, 868, 871, 872, 908, 915, 948, 951

OpenMP property

The [property](#) that a [routine](#), [callback](#) or type is in the OpenMP namespace, which implies it has the `omp_` prefix. 83, 556–562, 564–566, 568, 570, 573–580, 582–585, 587, 593, 594, 644–649, 652–659, 661–667, 669–673, 677–682, 684, 685, 687, 689–699, 701, 718

OpenMP stylized expression

A [base language](#) expression that is subject to restrictions that enable its use within an OpenMP implementation. 33, 62, 162, 190, 239

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](#). 5–7, 22, 58, 77, 82, 86, 108, 109, 135, 137, 139, 588, 803, 880, 883–887, 889, 892, 900, 908, 921

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, 5, 22, 95, 96, 109, 462, 468

OpenMP type

A type that has the [OpenMP property](#) or a type that is an [OMPD type](#) or an [OMPT type](#). 188, 23, 33, 54, 57, 64, 81, 82, 85, 162, 165, 166, 168, 186–190, 209, 351, 395, 489, 530, 539, 553, 554, 556, 558, 559, 561, 563, 564, 566, 568, 570, 572–575, 577–587, 643, 797, 923, 938, 940

optional property

The [property](#) that a [clause](#), a [modifier](#) or an argument is optional and thus may be omitted. If any argument of a [routine](#) has the [optional property](#) then the [routine](#) has the [overloaded property](#). 83, 161–163, 166, 211, 273, 342, 343, 351, 358, 361, 364, 368, 374–379, 382–384, 389, 400, 401, 411, 420, 437, 441, 459–461, 493, 501, 503–512, 518, 531, 538, 555, 637, 638, 641, 644–646, 683–685, 687

order-concurrent-nestable construct

A [construct](#) that has the [order-concurrent-nestable](#) property. 416, 951

order-concurrent-nestable property

The [property](#) that a [construct](#) or [routine](#) generates a [region](#) that may be a [strictly nested region](#) of a [region](#) that was generated by a [construct](#) on which an [order clause](#) with an *ordering*

argument of **concurrent** is specified. 83, 84, 392, 393, 395, 397, 398, 400, 402, 417, 442, 514

order-concurrent-nestable routine

A **routine** that has the **order-concurrent-nestable** property. 416, 951

original list item

The instance of a **list item** in the **data environment** of the **enclosing context**. 38, 50, 52, 71, 72, 84, 100, 220–222, 225, 228, 229, 231–235, 237, 241, 246, 247, 250, 252–256, 258–260, 270, 282, 286, 293, 295, 303–305, 310, 311, 313, 363, 378, 437, 440, 442, 464, 486, 931, 949

original list-item updating clause

A **clause** that has the **original list-item updating** property 542

original list-item updating property

The **property** that a **clause** includes an effect of updating the value of the **original list item** when the **region** for which it is specified is completed. 84, 230, 251, 255, 257

original pointer

An **original list item** that corresponds to a **corresponding pointer**. 304

original storage

The storage of a given **mapped variable**. 8, 72, 97, 294, 304, 305, 765

original storage block

A **storage block** that contains the storage of one or more **mapped variables** in a **data environment**. 8, 9, 38, 303

original variable

For a **variable** that is referenced in the **structured block** that is associated with a **block-associated directive** that accepts **data-sharing attribute clauses**, the **variable** by the same name that exists immediately outside the **construct**. 7, 7, 8

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**. 535

outermost-leaf property

The **property** that a **clause** applies to the outermost **leaf construct** that permits it when it appears on a **compound construct**. 163, 235, 285, 501, 503, 548

output map type

The [map type](#) that results when [map-type decay](#) is applied to an [input map type](#). [297](#), [63](#), [72](#), [112](#), [293](#), [297](#), [301](#)

overlapping type name

An [OpenMP type](#) for which its name has the [overlapping type-name property](#). [780](#)

overlapping type-name property

The [property](#) that an [OpenMP type](#) name is used for both an ordinary [OpenMP type](#) (possibly an [OMPD type](#) or an [OMPT type](#)) and for a [callback](#) in the same name space; which type is intended should be apparent from the context in this document. [85](#), [743](#), [748](#), [760](#), [768](#), [778](#), [780](#), [788](#), [791](#), [792](#)

overloaded property

The [property](#) that a [routine](#) has an overloaded C++ interface. [83](#), [85](#), [676–686](#)

overloaded routine

A [routine](#) that has the [overloaded property](#). [676](#), [682](#), [683](#)

P

parallel handle

A [handle](#) that refers to a [parallel region](#). [860](#), [856](#), [887](#), [888](#), [895](#), [897](#)

parallelism-generating construct

A [construct](#) that has the [parallelism-generating property](#). [232](#), [384](#), [388](#), [546](#)

parallelism-generating property

The [property](#) that a [construct](#) enables parallel execution by generating one or more [teams](#), [explicit tasks](#), or [SIMD instructions](#). [85](#), [402](#), [412](#), [417](#), [446](#), [449](#), [474](#), [476](#), [478](#), [481](#), [486](#)

parallel region

A [region](#) that has a set of associated [implicit tasks](#) and an associated [team](#) of [threads](#) that execute those [tasks](#). [4](#), [5](#), [20](#), [23](#), [32](#), [39](#), [54](#), [61](#), [85](#), [87](#), [102](#), [106](#), [107](#), [117](#), [119](#), [128](#), [135](#), [139](#), [287](#), [288](#), [407](#), [421](#), [423–425](#), [427](#), [428](#), [433](#), [443](#), [444](#), [446](#), [449](#), [495–498](#), [522](#), [547](#), [556](#), [588–590](#), [740](#), [747](#), [770](#), [784](#), [789](#), [790](#), [824–826](#), [856](#), [859](#), [860](#), [883](#), [887](#), [888](#), [892](#), [895](#), [925](#), [926](#), [947](#), [949](#)

parameter list

An [argument list](#) that consists of [parameter list items](#). [165](#)

parameter list item

A [list item](#) that identifies one or more parameters of a [procedure](#). [166](#), [76](#), [85](#), [165](#), [167](#), [178](#)

parent device

For a given **target region**, the **device** on which the corresponding **target construct** was encountered. [256](#), [376](#), [471](#), [481](#)

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**. [4](#), [22](#), [86](#)

partial tile

A **tile** that is not a **complete tile**. [399](#), [399](#)

partitioned construct

A **construct** that has the **partitioned property**. [423](#), [86](#), [546](#)

partitioned property

The **property** of a **construct** that it 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**. [86](#), [424](#), [426](#), [428](#), [431](#), [435](#), [436](#), [439](#), [442](#)

partitioned worksharing construct

A **construct** that is both a **partitioned construct** and a **worksharing construct**. [4](#), [86](#)

partitioned worksharing region

A **region** that corresponds to a **partitioned worksharing construct**. [951](#)

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. [203](#), [59](#), [271](#), [392–394](#), [397–399](#), [534](#), [949](#)

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. [377](#), [8](#), [915](#)

place

An unordered set of **processors** on a **device**. [133](#), [4](#), [63](#), [87](#), [108](#), [119](#), [120](#), [131](#), [134–136](#), [407–411](#), [703–706](#), [820–822](#), [916](#), [921](#), [929](#)

place-assignment group

A logical group of [places](#) and positions from the *place-assignment-var* ICV that is used to define a set of assignments of [threads](#) to [places](#) according to a given [thread affinity](#) policy. [408](#), [408](#), [409](#)

place-count abstract name

A [numeric abstract name](#) that refers to a quantity associated with a [place-list abstract name](#). [131](#)

place list

The ordered [list](#) that describes all OpenMP [places](#) available to the execution environment. [87](#), [134](#), [412](#), [703](#), [820](#), [916](#), [929](#)

place-list abstract name

A [conceptual abstract name](#) that refers to a set of hardware abstractions of a given category that may be used to specify each [place](#) in a [place list](#). [131](#), [87](#), [131](#), [134](#)

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](#). [408](#), [408](#), [705](#), [706](#), [821](#), [822](#)

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](#). [63](#), [108](#), [120](#), [409](#), [410](#)

pointer association query

A query to the association status of a pointer via comparison to zero in C/C++ or by calling the **ASSOCIATED** intrinsic with one argument in Fortran. [483](#)

pointer attachment

The process of making a pointer variable an [attached pointer](#). [26](#), [293](#), [299](#), [305](#), [928](#)

pointer property

The [property](#) that a [routine](#) or [callback](#) either returns a pointer type in C/C++ and is an assumed-size array in Fortran or has an argument that has such a type. [555](#), [616](#), [617](#), [636–638](#), [641](#), [646–649](#), [652](#), [654](#), [658](#), [666](#), [669](#), [671](#), [704](#), [706–710](#), [722](#), [740](#), [751](#), [759](#), [771–778](#), [780–782](#), [784–788](#), [791](#), [792](#), [795](#), [796](#), [798](#), [800–803](#), [806](#), [808](#), [810](#), [813](#), [815](#), [816](#), [818](#), [820](#), [821](#), [823–825](#), [827](#), [828](#), [831–842](#), [863](#), [865](#), [867](#), [869–871](#), [873–875](#), [878–883](#), [885–906](#)

pointer-to-pointer property

The [property](#) that a [routine](#) or [callback](#) either returns a pointer-to-pointer type in C/C++ or has an argument that has such a type. [555](#), [801](#), [808](#), [815](#), [816](#), [824](#), [825](#), [827](#), [828](#), [862](#), [869](#), [876](#), [878–880](#), [882](#), [883](#), [885–892](#), [899](#), [903](#), [905](#)

positive property

The [property](#) that an expression, including one that is used as the argument of a [clause](#), a [modifier](#) or a [routine](#), has a value that is greater than zero. [165](#), [132–134](#), [137](#), [138](#), [163](#), [166](#), [210–213](#), [316](#), [317](#), [320](#), [324](#), [332](#), [390](#), [391](#), [394](#), [395](#), [401](#), [406](#), [411](#), [415](#), [419](#), [420](#), [437](#), [441](#), [452](#), [453](#), [472](#), [567](#), [588](#), [603](#), [604](#), [622](#), [625](#), [636](#), [638](#), [652](#), [654](#), [667](#), [669](#), [671](#), [759](#), [833](#), [916](#), [917](#), [920–924](#)

post-modified property

The [property](#) of a [clause](#) that its [modifiers](#) must appear after its arguments. [161](#), [163](#), [165](#), [223](#), [259](#), [305](#), [331](#)

preceding dependence-compatible task

For a given [task](#), a [dependence-compatible task](#) that may be its [antecedent task](#). [528](#), [52](#), [61](#), [528](#)

predecessor task

For a given [task](#), an [antecedent task](#) of that [task](#), or any [predecessor task](#) of any of its [antecedent tasks](#). [527](#), [88](#), [475](#), [477](#), [482](#), [486](#), [499](#), [528](#)

predefined default mapper

The [default mapper](#) that is used if no [default mapper](#) that is a [user-defined mapper](#) is visible for the type of a given [list item](#). [291](#), [236](#), [291](#), [295](#), [299](#), [301](#), [310](#), [311](#)

predefined identifier

Unless otherwise specified, an [OpenMP identifier](#) that is defined for use in arbitrary [base language](#) expressions. [188](#), [7](#), [188](#), [553](#), [554](#), [612](#), [717](#), [733](#), [876](#), [923](#), [928](#)

predetermined data-sharing attribute

A [data-sharing attribute](#) that applies regardless of the [clauses](#) that are specified on a given [construct](#), unless explicitly specified otherwise. [216](#), [215–218](#), [223](#), [224](#), [289](#), [307](#), [482](#), [548](#), [949](#)

preference specification

The specification of a set of preferences for interoperating with a [foreign runtime environment](#). [490](#), [89](#), [166](#), [491](#), [922](#)

preference specification list

An [argument list](#) that consists of [preference specification list items](#). [165](#)

1 **preference specification list item**

2 A [list item](#) that is a [preference specification](#). [165](#), [88](#), [490](#)

3 **pre-modified property**

4 The [property](#) of a [clause](#) that its [modifiers](#) must appear before its arguments. [161](#), [165](#)

5 **preprocessed code**

6 For C/C++, a sequence of preprocessing tokens that result from the first six phases of
7 translation, as defined by the [base language](#). For Fortran, a sequence of source lines. [354](#), [939](#)

8 **present storage**

9 A [storage block](#) that exists in a given [device data environment](#). [294](#), [301–305](#)

10 **primary thread**

11 An [assigned thread](#) that has [thread number](#) 0. A [primary thread](#) may be an [initial thread](#) or
12 the [thread](#) that encounters a [parallel construct](#), forms a [team](#), generates a set of [implicit](#)
13 [tasks](#), and then executes one of those [tasks](#) as [thread number](#) 0. [4](#), [4](#), [5](#), [29](#), [54](#), [61](#), [86](#), [89](#),
14 [108](#), [109](#), [274](#), [285](#), [286](#), [402](#), [403](#), [408](#), [410](#), [422](#), [424](#), [523](#), [590](#), [824](#), [949](#)

15 **private attribute**

16 For a given [construct](#), a [data-sharing attribute](#) of a [data entity](#) that its lifetime is limited to that
17 of the corresponding [region](#) and it is visible only to a single [task](#) generated by the [construct](#) or
18 to a single [SIMD lane](#) used by the [construct](#). [220](#), [7](#), [8](#), [22](#), [53](#), [61](#), [62](#), [64](#), [65](#), [89](#), [90](#), [92](#), [114](#),
19 [163](#), [215–217](#), [219](#), [221](#), [229](#), [232](#), [234](#), [237](#), [240](#), [241](#), [246](#), [252](#), [253](#), [256](#), [270](#), [271](#), [288](#),
20 [307](#), [324](#), [388](#), [424](#), [541](#), [548](#), [949](#)

21 **private-only variable**

22 A [variable](#) that has a [private attribute](#) and no other [data-sharing attribute](#) with respect to a
23 given [construct](#). [226](#), [457](#)

24 **private variable**

25 A [variable](#) that has a [private attribute](#) with respect to a given [construct](#). [7](#), [7](#), [8](#), [53](#), [62](#), [65](#), [66](#),
26 [89](#), [92](#), [221](#), [222](#), [269](#), [271](#), [285](#), [287](#), [429](#), [430](#), [432](#), [437](#), [441](#), [442](#), [930](#)

27 **privatization clause**

28 The [clause](#) that may result in [private variables](#) that are [new list items](#). [215](#), [38](#), [78](#), [90](#), [223](#),
29 [234](#)

30 **privatization property**

31 The [property](#) that a [clause](#) privatizes [list items](#). [226](#), [227](#), [230](#), [233](#), [234](#), [251](#), [255](#), [257](#), [259](#),
32 [465](#)

privatized list item

A [list item](#) that appears in the [argument list](#) of a [privatization clause](#), resulting in one or more [private new list items](#). [220](#), [220–223](#), [225](#), [226](#), [233](#), [252](#)

procedure

A function (for C/C++ and Fortran) or subroutine (for Fortran). [15](#), [27](#), [30](#), [36](#), [41](#), [42](#), [45](#), [52](#), [54](#), [55](#), [61](#), [67](#), [76](#), [85](#), [93](#), [97](#), [102](#), [109](#), [110](#), [116](#), [149](#), [153](#), [158](#), [165–167](#), [178](#), [189](#), [193](#), [219](#), [220](#), [225](#), [226](#), [229](#), [239](#), [260–263](#), [267](#), [290](#), [293](#), [298](#), [299](#), [309](#), [311](#), [332](#), [335](#), [336](#), [339](#), [346](#), [347](#), [352](#), [353](#), [359–368](#), [420](#), [432](#), [433](#), [469](#), [470](#), [482](#), [484](#), [553](#), [575–577](#), [660](#), [663](#), [665](#), [666](#), [721](#), [722](#), [724](#), [732](#), [744](#), [746](#), [756](#), [826](#), [834](#), [849](#), [854](#), [856](#), [859](#), [864](#), [865](#), [869](#), [920](#), [939](#), [943](#)

procedure property

The [property](#) that a [routine](#) argument has a function pointer type in C/C++ and a procedure type in Fortran. [555](#), [659](#), [836](#)

processor

An [implementation defined](#) hardware unit on which one or more [threads](#) can execute. [43](#), [86](#), [120](#), [134](#), [139](#), [615](#), [704](#), [819–822](#), [831](#), [915](#), [916](#), [947](#)

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 $k \in \{1, \dots, n\}$. [399](#)

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 *sequenced-before* relation between operations performed by the same [thread](#).

[12](#), [13](#), [90](#), [100](#)

progress group

A group of consecutive [threads](#) in a [team](#) that may execute on the same [progress unit](#). [411](#), [411](#)

progress unit

An [implementation defined](#) set of consecutive [hardware threads](#) on which [native threads](#) may execute a common stream of instructions. [6](#), [6](#), [7](#), [90](#), [411](#), [554](#), [616](#)

property

A characteristic of an OpenMP feature. [xxviii](#), [21](#), [22](#), [24](#), [25](#), [28](#), [30–37](#), [39–42](#), [44–46](#), [48](#), [50](#), [51](#), [53–59](#), [63–65](#), [67](#), [68](#), [70–81](#), [83–99](#), [104](#), [106–110](#), [112–117](#), [162–164](#), [172](#), [178](#), [184–187](#), [210–213](#), [223](#), [225–227](#), [230](#), [233–237](#), [250](#), [251](#), [255–257](#), [259](#), [260](#), [262](#), [265](#), [266](#), [268](#), [269](#), [272–274](#), [278](#), [281](#), [283](#), [285](#), [287](#), [291](#), [292](#), [296](#), [297](#), [299](#), [306–308](#), [311–313](#), [320](#), [321](#), [323](#), [324](#), [326–329](#), [331](#), [332](#), [336–338](#), [340](#), [342–344](#), [347](#), [348](#), [350](#), [351](#), [353](#), [355–358](#), [361](#), [362](#), [364](#), [366](#), [368–372](#), [374–386](#), [389](#), [391–398](#), [400–402](#), [406](#), [410–412](#), [415–422](#), [424–426](#), [428](#), [431](#), [435–437](#), [439](#), [441](#), [442](#), [444](#), [446](#), [449](#), [452–455](#), [458–466](#), [470–472](#), [474](#), [476](#), [478](#), [481](#), [486](#), [488–490](#), [492](#), [493](#), [495](#), [498](#), [499](#), [501–514](#), [518](#), [525–527](#), [531](#), [532](#), [534](#), [536](#), [538–540](#), [544](#), [548](#), [555–566](#), [568](#), [570](#), [573–580](#), [582–585](#), [587–604](#), [606–610](#), [612–622](#), [624–627](#), [629–632](#), [634–638](#), [640](#), [641](#), [643–649](#), [652–659](#), [661–667](#), [669–674](#), [677–682](#), [684–687](#), [689–710](#), [712–716](#), [718](#), [721](#), [722](#), [734–737](#), [739–757](#), [759–762](#), [764–769](#), [771–778](#), [780–796](#), [798–803](#), [806](#), [808](#), [810](#), [813–825](#), [827–829](#), [831–842](#), [847](#), [848](#), [850–860](#), [862](#), [863](#), [865–906](#), [923](#), [931](#)

pure property

The [property](#) that a [directive](#) has no observable side effects or state, yielding the same result every time it is encountered. [153](#), [262](#), [266](#), [269](#), [274](#), [278](#), [307](#), [321](#), [342](#), [344](#), [351](#), [358](#), [364](#), [369](#), [385](#), [386](#), [392](#), [393](#), [395](#), [397](#), [398](#), [400](#), [417](#), [929](#), [936](#)

R

raw-memory-allocating routine

A [memory-allocating routine](#) that has the [raw-memory-allocating-routine](#) property. [675](#), [675](#), [676](#), [678](#)

raw-memory-allocating-routine property

The [property](#) that a [memory-allocating routine](#) returns a pointer to uninitialized [memory](#). [675](#), [91](#), [677](#), [678](#)

read-modify-write operation

An [atomic operation](#) that reads and then writes to a given [storage location](#). [11](#)

read structured block

An [atomic structured block](#) that may be associated with an [atomic directive](#) that expresses an [atomic read](#) operation. [194](#), [195](#), [197](#), [517](#)

rectangular-memory-copying property

The [property](#) of a [memory-copying routine](#) that the [memory](#) that it copies forms a rectangular subvolume. [632](#), [91](#), [635](#), [638](#)

rectangular-memory-copying routine

A [routine](#) with the [rectangular-memory-copying](#) property. [632](#), [633](#), [636](#), [639](#), [760](#), [805](#), [924](#)

reduction

A use of a [reduction operation](#). [33](#), [92](#), [106](#), [188](#), [238](#), [239](#), [241](#), [243](#), [244](#), [248–253](#), [255](#), [450](#), [930](#), [936](#), [939](#), [942](#), [945](#), [947](#)

reduction attribute

For a given [construct](#), a [data-sharing attribute](#) of a [data entity](#) that implies the [private attribute](#) and for which a partial result is computed in the context of a [reduction](#) computation. [248](#), [92](#)

reduction clause

A [reduction-scoping clause](#) or a [reduction-participating clause](#). [238](#), [62](#), [220](#), [223](#), [238–240](#), [246–251](#), [253](#), [255](#), [256](#), [263](#)

reduction expression

A [combiner expression](#) or an [initializer expression](#). [239](#), [239](#)

reduction identifier

An [OpenMP identifier](#) that specifies a [combiner OpenMP operation](#) to use in a [reduction](#). [238](#), [188](#), [238](#), [239](#), [243](#), [244](#), [246–248](#), [251](#), [263](#), [264](#), [450](#), [931](#)

reduction operation

An operation that applies a [combiner](#) and an associated [initializer](#) to a set of values. [33](#), [92](#), [96](#), [115](#), [178](#), [238](#)

reduction-participating clause

A [clause](#) that defines the participants in a [reduction](#). [238](#), [92](#), [238](#), [250](#), [252](#), [256](#)

reduction-participating property

The [property](#) that a [clause](#) is a [reduction-participating clause](#). [251](#), [255](#)

reduction-scoping clause

A [clause](#) that defines the [region](#) in which a [reduction](#) is computed. [238](#), [92](#), [238](#), [250](#), [252](#), [255](#), [256](#), [450](#)

reduction-scoping property

The [property](#) that a [clause](#) is a [reduction-scoping clause](#). [251](#), [255](#)

reduction variable

A [private variable](#) that has the [reduction attribute](#) with respect to a given [construct](#). [248](#), [248](#)

referenced pointee

For a given [referencing variable](#), the referenced data object to which the [referring pointer](#) points. [26](#), [28](#), [93](#), [235](#), [237](#), [298](#), [303](#), [310](#)

referencing variable

For C++, a [data entity](#) that is a reference. For Fortran, a [data entity](#) that is an allocatable [variable](#) or a data pointer. [26–28](#), [92](#), [93](#), [115](#), [215](#), [217](#), [235](#), [237](#), [282](#), [293](#), [298](#), [303](#), [310](#)

referring pointer

If a given [referencing variable](#) is a Fortran data pointer, the pointer object that is pointer associated with the [referenced pointee](#); otherwise, an associated [implementation defined handle](#) through which the [referenced pointee](#) is made accessible. [25](#), [26](#), [28](#), [38](#), [92](#), [215](#), [217](#), [237](#), [282](#), [298](#), [303](#), [482](#)

region

All code encountered during a specific instance of the execution of a given [construct](#), [structured block sequence](#) or [routine](#). A [region](#) includes any code in called [procedures](#) 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 a [routine](#). During the execution of an [OpenMP program](#), a [construct](#) may give rise to many [regions](#). [3–9](#), [12](#), [13](#), [19–22](#), [26](#), [29–32](#), [38](#), [39](#), [43](#), [45](#), [47–52](#), [55](#), [56](#), [60](#), [61](#), [63](#), [71](#), [73](#), [77](#), [81](#), [83–86](#), [89](#), [92–95](#), [97–99](#), [101](#), [104–110](#), [112](#), [116–120](#), [125](#), [127](#), [131–133](#), [136](#), [139](#), [152](#), [159](#), [198](#), [199](#), [203](#), [210](#), [215](#), [219–222](#), [229](#), [232](#), [235](#), [238](#), [239](#), [247](#), [250](#), [252–257](#), [274–276](#), [285](#), [286](#), [294](#), [299](#), [300](#), [302–304](#), [311](#), [317–319](#), [322](#), [325](#), [327](#), [329](#), [345](#), [355–357](#), [362](#), [375](#), [376](#), [383](#), [386](#), [402](#), [403](#), [406](#), [407](#), [409](#), [412–414](#), [416–418](#), [421](#), [423–433](#), [440](#), [443–447](#), [449](#), [450](#), [453](#), [455–459](#), [465–471](#), [474](#), [476](#), [478](#), [479](#), [481–484](#), [486–488](#), [492–500](#), [514–517](#), [519–524](#), [526](#), [533](#), [534](#), [536](#), [537](#), [539–545](#), [554](#), [555](#), [564](#), [585](#), [588–590](#), [592](#), [596](#), [598](#), [600–603](#), [605](#), [608](#), [610](#), [612–614](#), [616–623](#), [640](#), [651](#), [667](#), [668](#), [674](#), [675](#), [677](#), [685–687](#), [689–696](#), [698–702](#), [707](#), [709–711](#), [713](#), [714](#), [717–719](#), [727](#), [728](#), [730](#), [740](#), [744](#), [750](#), [758](#), [759](#), [761](#), [768](#), [770](#), [776–779](#), [784](#), [789–794](#), [805](#), [807](#), [811](#), [816](#), [823](#), [824](#), [828](#), [879](#), [888](#), [908–911](#), [913](#), [915](#), [917](#), [918](#), [920](#), [922](#), [924](#), [925](#), [930](#), [932–934](#), [936](#), [939](#), [942](#), [945](#), [946](#), [948](#), [949](#), [951](#), [952](#)

region endpoint

An [event](#) that indicates the beginning or end of a [region](#) that may be of interest to a [tool](#). [727](#), [754](#)

region-invariant property

The [property](#) that an expression, including one that is used as the argument of a [clause](#), a [modifier](#) or a [routine](#), has a value that is invariant for the associated [region](#). [165](#), [163](#), [257](#), [260](#), [332](#), [402](#), [412](#), [437](#), [441](#)

1 registered callback

2 A [callback](#) for which [callback registration](#) has been performed. [14](#), [30](#), [80](#), [725](#), [727](#), [925](#)

3 release flush

4 A [flush](#) that has the [release flush property](#). [10](#), [11](#), [12](#), [94](#), [104](#), [516](#), [519](#), [521–524](#)

5 release flush property

6 A [flush](#) with the [release flush property](#) orders [memory](#) operations that precede the [flush](#) before
7 [memory](#) operations performed by a different [thread](#) with which it synchronizes. [53](#), [94](#), [519](#)

8 release sequence

9 A set of modifying [atomic operations](#) that are associated with a [release flush](#) that may
10 establish a [synchronizes-with relation](#) between the [release flush](#) and an [acquire flush](#). [11](#), [12](#),
11 [522](#)

12 repeatable property

13 The [property](#) that a [clause](#) or [modifier](#) may appear more than once in a given context with
14 which it is associated. [163](#), [185](#)

15 replacement candidate

16 A [directive variant](#) or [function variant](#) that may be selected to replace a [metadirective](#) or [base](#)
17 [function](#). [341](#), [30](#), [48](#), [341](#), [342](#), [346](#), [348](#), [352](#), [919](#)

18 replayable construct

19 A [task-generating construct](#) that an implementation must record into a [taskgraph record](#), if
20 one is recorded. [455](#), [94](#), [96](#), [106](#), [237](#), [455–458](#), [461](#)

21 replay execution

22 An execution of a given [taskgraph region](#) that entails executing [replayable constructs](#) that
23 are saved in a [matching taskgraph record](#). [456](#), [52](#), [96](#), [106](#), [237](#), [455–457](#), [922](#), [930](#)

24 reproducible schedule

25 A [loop schedule](#) for the [affected loop nest](#) of a given [loop-nest-associated construct](#) that does
26 not change between different executions of the [construct](#) that have the same [binding thread](#)
27 [set](#) and have the same number of [logical iterations](#). [423](#), [210](#), [416](#), [433](#), [440](#), [443](#), [937](#)

28 required property

29 The [property](#) that a [clause](#), a [modifier](#), an argument, or at least one member of a [clause set](#) is
30 required and, thus, may not be omitted. [164](#), [161](#), [163](#), [164](#), [186](#), [250](#), [251](#), [255–257](#), [265](#),
31 [268](#), [269](#), [342](#), [347](#), [348](#), [373](#), [380](#), [391](#), [396](#), [478](#), [486](#), [488](#), [525](#), [532](#), [539](#), [555](#)

32 reservation type

33 A [thread-reservation type](#). [145](#)

reserved locator

An [OpenMP identifier](#) that represents system storage that is not necessarily bound to any [base language](#) storage item. [168](#), [166–168](#), [527](#), [529](#), [938](#)

reserved thread

A [thread](#) in an [OpenMP thread pool](#) that must have a particular [thread-reservation type](#) when executing a [task](#). [145](#)

resource-relinquishing property

The [property](#) that a [routine](#) relinquishes some (or all) resources that the [OpenMP program](#) is currently using. [712](#), [95](#), [713](#), [714](#)

resource-relinquishing routine

A [routine](#) that has the [resource-relinquishing property](#). [712](#), [58](#), [100](#), [584](#), [585](#), [712](#), [713](#)

reverse-offload region

A [region](#) that is associated with a [target construct](#) that specifies a [device clause](#) with the [ancestor device-modifier](#). [362](#), [943](#)

routine

Unless specifically stated otherwise, an [OpenMP API routine](#). [xxviii](#), [2](#), [3](#), [6](#), [7](#), [14](#), [15](#), [17](#), [19](#), [21](#), [22](#), [24](#), [25](#), [28](#), [31](#), [36](#), [44](#), [45](#), [50](#), [53](#), [54](#), [57](#), [59](#), [60](#), [63–65](#), [67](#), [68](#), [73](#), [74](#), [77–81](#), [83–85](#), [87](#), [88](#), [90](#), [91](#), [93](#), [95](#), [99](#), [108](#), [113](#), [115](#), [118](#), [123–125](#), [132](#), [142](#), [150](#), [275](#), [317](#), [329](#), [330](#), [416](#), [483](#), [553–555](#), [557](#), [575](#), [577](#), [581](#), [583](#), [584](#), [586–610](#), [612–633](#), [635](#), [636](#), [638–647](#), [649–676](#), [678–683](#), [685–718](#), [722](#), [725](#), [770](#), [771](#), [780](#), [786](#), [795](#), [814](#), [826](#), [845](#), [854](#), [861](#), [874–896](#), [899–906](#), [923–925](#), [928](#), [933–936](#), [939–944](#), [946](#), [948–951](#)

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. [725](#), [24](#), [50](#), [80](#), [95](#), [721](#), [729](#), [771](#), [813](#)

runtime error termination

An [error termination](#) that is performed during execution. [6](#), [51](#), [152](#), [293](#), [304](#), [310](#), [329](#), [407](#), [470](#), [471](#), [620](#), [622](#), [623](#), [714](#), [917](#)

S

[safesync-compatible expression](#)

An expression that is [omp_curr_progress_width](#), a [constant](#) expression, or an expression for which all operands are [safesync-compatible expressions](#). [95](#), [411](#)

saved data environment

For a given [replayable construct](#) that is recorded in a [taskgraph record](#), an associated [enclosing data environment](#) that is also saved in the record for possible use in a [replay execution](#) of the [construct](#). [456](#), [106](#), [237](#), [455](#), [457](#)

scalar variable

For C/C++, a scalar-variable, as defined by the [base language](#). For Fortran, a scalar variable with enum, enumeration, assumed, or intrinsic type, excluding character type, as defined by the [base language](#). [190](#), [194](#), [200](#), [205](#), [217](#), [220](#), [224](#), [232](#), [290](#), [306](#), [805](#), [918](#), [945](#)

scan computation

A computation performed in the [logical iterations](#) of a loop nest that yields a set of values that are a running total, as defined by a [reduction operation](#), over an input set of values. [270](#), [51](#), [61–63](#), [96](#), [114](#), [238](#), [252](#), [254](#), [269](#), [270](#)

scan phase

The portion of an [affected iteration](#) that includes all statements that read the result of a [scan computation](#). [269](#), [62](#), [269–273](#)

schedulable task

A member of the [schedulable task set](#) of a [thread](#). [469](#)

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. [96](#), [467](#), [468](#)

schedule specification

The specification of a [loop schedule](#) for a given [loop-nest-associated construct](#), which includes but is not limited to the [schedule type](#) and [chunk](#) size. [423](#), [96](#), [210](#), [423](#)

schedule-specification clause

A [clause](#) that has the [schedule-specification property](#). [423](#)

schedule-specification property

The [property](#) of a [clause](#) that it defines, in part or in full, the [schedule specification](#) of a given [loop-nest-associated construct](#). [96](#), [416](#), [437](#), [441](#)

schedule type

The part of a [schedule specification](#) that identifies the method by which the [collapsed iterations](#) are distributed to [threads](#). [96](#), [120](#), [128](#), [137](#), [434](#), [438](#), [557](#), [594](#), [595](#), [923](#)

scope handle

A [handle](#) that refers to an OpenMP scope. [855](#), [904–906](#)

segment

A portion of an [address space](#) associated with a set of address ranges. [20](#), [855](#)

selector set

Unless specifically stated otherwise, a [trait selector set](#). [36](#), [45](#), [59](#), [104](#), [115](#), [339](#)

self map

A [mapping operation](#) for which the [corresponding storage](#) is the same as its [original storage](#). [304](#), [86](#), [303](#), [304](#), [378](#), [932](#)

semantic requirement set

A logical set of semantic [properties](#) maintained by a [task](#) that is updated by [directives](#) in the scope of the [task region](#). [345](#), [349](#), [351](#), [355](#), [357](#), [502](#)

separated construct

A [construct](#) for which its associated [structured block](#) is split into multiple [structured block sequences](#) by a [separating directive](#). [158](#), [97](#), [158](#), [159](#), [269–271](#)

separating directive

A [directive](#) that splits a [structured block](#) that is associated with a [construct](#), the [separated construct](#), into multiple [structured block sequences](#). [158](#), [97](#), [156](#), [158](#), [159](#), [269–271](#), [428](#)

sequentially consistent atomic operation

An [atomic operation](#) that is specified by an [atomic construct](#) for which the [seq_cst clause](#) is specified. [13](#), [947](#)

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.

[97](#), [275](#), [707](#), [709](#)

shape-operator

For C/C++, an [array shaping](#) operator that reinterprets a pointer expression as an array with one or more specified dimensions. [169](#), [168](#), [169](#), [310](#), [464](#), [529](#), [941](#)

shared attribute

For a given [construct](#), a [data-sharing attribute](#) of a [data entity](#) that its lifetime is not limited to that of the corresponding [region](#) and, if the [data entity](#) is a [variable](#), it is visible to all [tasks](#) generated by the [construct](#) in addition to being visible in the [enclosing context](#) of the [construct](#) if declared outside the [construct](#). [225](#), [8](#), [98](#), [215–220](#), [225](#), [252](#), [253](#), [258](#), [259](#), [447](#), [450](#), [474](#), [476](#), [481](#), [486](#), [918](#)

shared variable

A [variable](#) that has the [shared attribute](#) with respect to a given [construct](#). [7](#), [7](#), [10–12](#), [14](#), [508–511](#)

sharing task

A [task](#) for which the [implicitly determined data-sharing attribute](#) is [shared](#) unless explicitly specified otherwise. [218](#), [98](#), [478](#)

sharing-task property

The [property](#) that a [task-generating construct](#) generates [sharing tasks](#). [478](#)

sibling task

Two [tasks](#) are each a [sibling task](#) of the other if they are [child tasks](#) of the same [task region](#). [98](#), [528](#)

signal

A software interrupt delivered to a [thread](#). [24](#), [98](#), [845](#)

signal handler

A function called asynchronously when a [signal](#) is delivered to a [thread](#). [7](#), [24](#), [745](#), [813](#), [845](#)

SIMD

Single Instruction, Multiple Data, a lock-step parallelization paradigm. [260](#), [261](#), [331](#), [335](#), [359](#), [360](#), [420](#), [919](#), [920](#), [947](#)

SIMD chunk

A set of iterations executed concurrently, each by a [SIMD lane](#), by a single [thread](#) by means of [SIMD instructions](#). [417](#), [99](#), [359](#), [417](#), [419](#), [945](#)

SIMD construct

A [simd construct](#) or a [compound construct](#) for which the [simd construct](#) is a [constituent construct](#). [438](#)

SIMD instruction

A single machine instruction that can operate on multiple data elements. [3](#), [85](#), [98](#), [99](#), [332](#), [417](#), [928](#)

SIMDizable construct

A [construct](#) that has the [SIMDizable property](#). 417, 951

SIMDizable property

The [property](#) that a [construct](#) may be encountered during execution of a [simd region](#). 99, 392, 393, 395, 397, 398, 400, 417, 442, 443, 514, 536

SIMD lane

A software or hardware mechanism capable of processing one data element from a [SIMD instruction](#). 5, 7, 8, 89, 98, 220, 221, 226, 250, 252, 258, 260–262, 417, 420, 921

SIMD loop

A loop that includes at least one [SIMD chunk](#). 331, 359

SIMD-partitionable construct

A [construct](#) that has the [SIMD-partitionable property](#). 546

SIMD-partitionable property

The [property](#) of a [loop-nest-associated construct](#) that it partitions the set of [affected iterations](#) such that each partition can be divided into [SIMD chunks](#). 99, 435, 436, 439, 449

simple-composite directive

A [directive](#) that has the [simple-composite property](#). 99, 546

simple-composite-directive name

A [directive name](#) of a [simple-composite directive](#). 545, 547

simple-composite property

The [property](#) of a [directive](#) that it is a [composite directive](#) but the [directive name](#) is not formed from the [directive names](#) of its [leaf directives](#). 99, 478

simple lock

A [lock](#) that cannot be set if it is already owned by the [task](#) trying to set it. 688, 99, 580, 688, 695

simple lock property

The [property](#) that a [routine](#) operates on [simple locks](#). 688, 99, 689, 690, 693, 695, 697, 699

simple lock routine

A [routine](#) that has the [simple lock property](#). 688, 580

simple modifier

A [modifier](#) that can never take an argument when it is specified. 162, 162, 163, 165

1 **simply contiguous array section**

2 An [array section](#) that can be determined to have contiguous storage at compile time. In
3 Fortran, this determination may result from the specification of the **CONTIGUOUS** attribute
4 on the declaration of the array. [220](#), [918](#)

5 **simply happens before**

6 For an event *A* to simply happen before an event *B*, *A* must precede *B* in [simply](#)
7 [happens-before order](#). [12](#), [12](#), [13](#)

8 **simply happens-before order**

9 An ordering relation that is consistent with [program order](#) and the [synchronizes-with relation](#).
10 [12](#), [57](#), [100](#)

11 **sink iteration**

12 A [doacross iteration](#) for which executable code, because of a [doacross dependence](#), cannot
13 execute until executable code from the [source iteration](#) has completed. [532](#), [48](#)

14 **socket**

15 The physical location to which a single chip of one or more [cores](#) of a [device](#) is attached. [131](#)

16 **soft pause**

17 An instance of a [resource-relinquishing routine](#) that specifies that the OpenMP state is
18 required to persist. [584](#), [585](#)

19 **source iteration**

20 A [doacross iteration](#) for which executable code must complete execution before executable
21 code from another [doacross iteration](#) can execute due to a [doacross dependence](#). [532](#), [48](#), [100](#)

22 **stand-alone directive**

23 A [unassociated directive](#) that is also an [executable directive](#). [156](#), [159](#)

24 **standard trace format**

25 A format for [OMPT trace records](#). [729](#), [735](#), [753](#), [840](#), [925](#)

26 **starting address**

27 The address of the first [storage location](#) of a [list item](#) or, for a [mapped variable](#) of its [original](#)
28 [list item](#). [52](#), [72](#), [300](#)

29 **static context selector**

30 The [context selector](#) for which [traits](#) in the [OpenMP context](#) can be fully determined at
31 compile time. [48](#), [341–343](#), [346](#)

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. [8](#), [26](#), [44](#), [67](#), [109](#), [216](#), [217](#), [219](#), [224](#), [241](#), [274](#), [277](#), [279](#), [283](#), [284](#), [289](#), [301](#), [302](#), [312](#), [316](#), [320](#), [322](#), [362](#), [363](#), [377](#), [378](#), [456](#), [457](#), [482](#), [915](#)

step expression

A loop-invariant expression used by an [induction operation](#). [33](#), [62](#), [66](#), [175](#), [242](#), [243](#), [247](#), [266](#)

storage block

The physical storage that corresponds to an [address range](#) in [memory](#). [9](#), [19](#), [38](#), [48](#), [53](#), [71](#), [74](#), [84](#), [89](#), [101](#), [115](#), [293](#), [305](#), [330](#), [483](#), [922](#)

storage location

A [storage block](#) in [memory](#). [7–9](#), [20](#), [25](#), [26](#), [50](#), [67](#), [91](#), [100](#), [193](#), [198–200](#), [234](#), [235](#), [256](#), [259](#), [260](#), [300](#), [319](#), [377](#), [419](#), [456](#), [514–520](#), [528](#), [529](#), [628](#), [740](#), [918](#), [922](#)

strictly nested region

A [region](#) nested inside another [region](#) with no other [explicit region](#) nested between them. [83](#), [108](#), [413](#), [414](#), [416](#), [440](#), [445](#), [602](#), [605](#), [620](#), [622](#), [933](#), [951](#)

strictly structured block

A single Fortran **BLOCK** construct, with a single entry at the top and a single exit at the bottom. [102](#), [157](#), [430](#)

string literal

For C/C++, a string literal. For Fortran, a character literal constant. [54](#), [143](#), [489](#), [491](#)

striping

The reordering of [logical iterations](#) of a loop that follows a grid while skipping [logical iterations](#) in-between. [397](#), [933](#)

strong flush

A [flush](#) that has the [strong flush property](#). [10](#), [10](#), [11](#), [13](#), [54](#), [516](#), [519](#)

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](#). [53](#), [101](#), [519](#)

structure

A [structure](#) is a [variable](#) that contains one or more [variables](#) that may have different types. This includes [variables](#) that have a **struct** type in C/C++, [variables](#) that have a **class** type in C++, and [variables](#) that have a derived type and are not arrays in Fortran. [37](#), [102](#), [220](#), [236](#), [289](#), [291](#), [293](#), [295](#), [301–303](#), [311–313](#), [483](#), [566](#), [722](#), [724](#), [732](#), [740](#), [743](#), [744](#), [751–753](#), [756](#), [759](#), [770–772](#), [780](#), [786](#), [787](#), [826](#), [840](#), [847](#), [848](#), [851](#), [852](#), [859](#), [918](#), [945](#)

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](#). [191](#), [3](#), [7](#), [29](#), [36](#), [38](#), [43](#), [69](#), [84](#), [97](#), [112](#), [113](#), [136](#), [157–159](#), [191](#), [192](#), [203](#), [207](#), [234–237](#), [285–288](#), [359](#), [388](#), [401–403](#), [413](#), [421](#), [424–427](#), [429–432](#), [434](#), [440](#), [446](#), [447](#), [455](#), [459](#), [467](#), [478](#), [479](#), [494](#), [499](#), [522](#), [524](#), [536](#), [610](#), [730](#), [750](#), [767](#), [770](#), [780](#), [793](#), [794](#), [912](#), [921](#)

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. [29](#), [47](#), [50](#), [93](#), [97](#), [104](#), [158](#), [191](#), [203](#), [207](#), [231](#), [232](#), [269–273](#), [427](#), [428](#), [921](#)

structured parallelism

Parallel execution through the [implicit tasks](#) of (possibly nested) [parallel regions](#) by the set of [structured threads](#) in a [contention group](#). [145](#), [146](#)

structured thread

A [thread](#) that is assigned to a [team](#) and is not a [free-agent thread](#). [96](#), [102](#), [109](#), [120](#), [145](#), [405](#), [929](#)

structure element

A single data member (for C/C++) or data component (for Fortran) of a [structure](#), as defined by the base language, for which storage is contained in the storage of that [structure](#). [71](#), [115](#), [217](#), [289](#), [291](#), [293](#), [326](#), [483](#), [942](#)

subroutine

A [procedure](#) for which a call cannot be used as the right-hand side of a [base language](#) assignment operation. [575](#), [576](#), [588](#), [592–594](#), [596](#), [601](#), [604](#), [609](#), [612](#), [620](#), [621](#), [629](#), [659](#), [661–663](#), [667](#), [673](#), [682](#), [689–691](#), [693](#), [695–698](#), [704](#), [706](#), [707](#), [709](#), [716](#), [736](#), [747](#), [772–778](#), [780–783](#), [785–790](#), [792–795](#), [798–803](#), [806](#), [808](#), [810](#), [829](#), [835](#)

subsidiary directive

A [directive](#) that is not an [executable directive](#) and that appears only as part of a [construct](#). [156](#), [159](#), [269](#), [271](#), [427](#), [428](#), [449](#), [454](#), [455](#), [933](#)

substitute array section

An [array section](#) that is substituted for an [assumed-size array list item](#) in a [map](#) clause based on a [matched candidate](#). [300](#), [300](#)

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](#). [5](#), [5](#), [468](#), [469](#)

successful groupprivate fallback situation

The situation in which the *fallback-mode* of the *fallback-modifier* in the [dyn_groupprivate](#) clause is specified as [default_mem](#) or [null](#), the [dynamic groupprivate block](#) cannot be instantiated using [memory](#) from the [groupprivate memory space](#), and the fallback behavior is executed successfully. [685](#)

successor task

For a given [task](#), a [dependent task](#) of that [task](#), or any [successor task](#) of a [dependent task](#) of that [task](#). [527](#), [103](#)

supported active levels

An [implementation defined](#) maximum number of [active levels](#) of parallelism. [596](#), [915](#)

supported device

The [host device](#) or any [non-host device](#) supported by the implementation, including any [device](#)-related requirements specified by the [requires](#) directive. [122](#), [142–144](#), [470](#)

synchronization construct

A [construct](#) that orders the completion of code executed by different [threads](#). [492](#), [2](#), [6](#), [542](#), [786](#)

synchronization hint

An indicator of the expected dynamic behavior or suggested implementation of a synchronization mechanism. [581](#), [492](#), [581](#), [582](#), [688](#), [924](#), [944](#)

synchronizes with

For an event *A* to synchronize with an event *B*, a [synchronizes-with relation](#) must exist from *A* to *B*. [12](#), [11](#), [12](#), [19](#), [522–524](#)

synchronizes-with relation

An asymmetric relation that relates a [release flush](#) to an [acquire flush](#), or, for C/C++, any pair of events *A* and *B* such that *A* “synchronizes with” *B* according to the [base language](#), and establishes [memory](#) consistency between their respective executing [threads](#). [10](#), [94](#), [100](#), [103](#)

synchronizing-region callback

A [callback](#) that has the [synchronizing-region](#) property. [789](#), [790](#)

synchronizing-region property

The [property](#) that a [callback](#) indicates the beginning or end of a synchronization-related [region](#). [789](#), [104](#), [789](#), [790](#)

synchronizing threads

Two [threads](#) are [synchronizing](#) if the completion of a [structured block sequence](#) by one of the [threads](#) requires that it first observes a modification by the other [thread](#), including the modification to an internal synchronization [variable](#) that an implementation performs for [implicit flush](#) synchronization as described in [Section 1.3.5](#). [6](#), [7](#), [104](#), [379](#), [411](#)

T

target-consistent clause

A [clause](#) for which all expressions that are specified on it are [target-consistent](#) expressions. [414](#)

target-consistent expression

An expression that has the [target-consistent](#) property. [104](#), [414](#)

target-consistent property

The [property](#) of an expression that its evaluation results in the same value when used on an [immediately nested construct](#) of a [target construct](#) as when specified on that [target construct](#). [104](#), [184](#), [323](#), [328](#), [415](#), [472](#)

target device

A [device](#) with respect to which the current [device](#) performs an operation, as specified by a [device construct](#) or [device memory routine](#). [471](#), [3](#), [4](#), [14](#), [40](#), [43–45](#), [71](#), [105](#), [118](#), [119](#), [235](#), [237](#), [256](#), [294](#), [296](#), [303–305](#), [310](#), [312](#), [313](#), [324](#), [336](#), [378](#), [470](#), [471](#), [473](#), [474](#), [476](#), [482](#), [487](#), [612](#), [613](#), [623](#), [624](#), [628](#), [629](#), [631](#), [632](#), [721](#), [725](#), [728–730](#), [746](#), [747](#), [798](#), [799](#), [805](#), [807](#), [811](#), [828](#), [831–833](#), [835](#), [842](#), [925](#), [931](#), [942](#)

target_device selector set

A [selector set](#) that may match the [target device trait set](#). [338](#), [338–340](#), [939](#)

target device trait set

The [trait set](#) that consists of [traits](#) that define the characteristics of a [device](#) that the implementation supports. [336](#), [104](#), [335–338](#), [340](#), [929](#)

target memory space

A [memory space](#) that is associated with at least one [device](#) that is not the current [device](#) when it is created. [651](#), [318](#), [667](#), [668](#)

target task

A [mergeable 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](#). [3](#), [256](#), [294](#), [474–477](#), [481–483](#), [486](#), [487](#), [521](#), [523](#), [623](#), [624](#), [634](#), [640](#), [744](#), [782](#), [786](#), [805](#), [807](#), [811](#), [826](#)

target variant

A version of a [device procedure](#) that can only be executed as part of a [target region](#). [335](#)

task

A specific instance of executable code and its [data environment](#) that the OpenMP implementation can schedule for execution by a [team](#). [3–9](#), [19](#), [22](#), [29](#), [31](#), [39](#), [42–44](#), [51–53](#), [55–57](#), [59–62](#), [64](#), [67](#), [68](#), [75–77](#), [83](#), [85](#), [88](#), [89](#), [93](#), [95–99](#), [103](#), [105–107](#), [109–114](#), [118–120](#), [127](#), [136](#), [137](#), [186](#), [220](#), [221](#), [225](#), [226](#), [228](#), [249](#), [250](#), [252](#), [255–258](#), [274](#), [279](#), [294](#), [300](#), [302](#), [304](#), [305](#), [316](#), [317](#), [329](#), [330](#), [345](#), [402](#), [404](#), [405](#), [408](#), [411](#), [413](#), [421](#), [424](#), [426–428](#), [430](#), [432](#), [434](#), [440](#), [446–450](#), [452–456](#), [458–460](#), [462–465](#), [467–469](#), [473](#), [478](#), [479](#), [488](#), [489](#), [493–496](#), [498–500](#), [502](#), [514](#), [516](#), [517](#), [522–524](#), [527–529](#), [533](#), [535](#), [536](#), [541](#), [542](#), [544](#), [551](#), [554](#), [564](#), [579](#), [591](#), [595](#), [604–606](#), [610](#), [621](#), [622](#), [683](#), [684](#), [688–697](#), [744](#), [745](#), [747](#), [758](#), [766–768](#), [770](#), [781–786](#), [788](#), [813](#), [826](#), [827](#), [856](#), [859](#), [860](#), [889–891](#), [893–895](#), [912](#), [919](#), [921](#), [922](#), [933](#), [934](#), [939](#), [947–950](#)

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 *allow-completion* [event](#) that is created for the [task](#) has been fulfilled. [107](#), [446](#)

task dependence

A [dependence](#) between two [dependence-compatible tasks](#): the [dependent task](#) and an [antecedent task](#). The [task dependence](#) is fulfilled when the [antecedent task](#) has completed. [524](#), [42](#), [43](#), [105](#), [106](#), [111](#), [468](#), [525](#), [528](#), [529](#), [531](#), [579](#), [606](#), [624](#), [740–742](#), [934](#), [939](#), [947](#)

task-generating construct

A [construct](#) that has the [task-generating property](#). [5](#), [53](#), [55](#), [75](#), [93](#), [94](#), [98](#), [106](#), [127](#), [136](#), [216](#), [218](#), [219](#), [447](#), [455](#), [457](#), [462](#), [478](#), [528](#), [529](#), [547](#), [930](#), [933](#), [942](#), [951](#)

task-generating property

The [property](#) that a [construct](#) generates one or more [explicit tasks](#) that are [child tasks](#) of the [encountering task](#). [105](#), [446](#), [449](#), [474](#), [476](#), [478](#), [481](#), [486](#)

taskgraph-altering clause

A [clause](#) that has the [taskgraph-altering property](#). [456](#), [457](#)

taskgraph-altering property

The [property](#) of a [clause](#) that if it appears on a [replayable construct](#), it affects the resulting number of [tasks](#) or the resulting [task dependences](#) in a [replay execution](#) of a [taskgraph record](#). [106](#), [452](#), [453](#), [527](#)

taskgraph record

For a given [taskgraph construct](#) that is encountered on a given [device](#), a data structure that contains a sequence of recorded [replayable constructs](#), with their respective [saved data environments](#), that are encountered while executing the corresponding [taskgraph region](#). [455](#), [52](#), [94](#), [96](#), [106](#), [455–458](#), [922](#)

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](#). [30](#), [106](#), [498](#), [541](#)

task handle

A [handle](#) that refers to a [task region](#). [856](#), [889–892](#), [895](#), [898](#)

task-inherited clause

A [clause](#) that has the [task-inherited property](#). [454](#)

task-inherited property

The [property](#) of a [clause](#) that if it appears on a [task_iteration directive](#), it will be inherited by the [tasks](#) that are generated by a [task-generating construct](#). [106](#), [464](#), [527](#)

taskloop-affected loop

A [collapse-affected loop](#) of a [taskloop construct](#). [175](#), [451](#), [454](#)

task priority

A hint for the [task](#) execution order of [tasks](#) generated by a [construct](#). [463](#), [146](#), [463](#), [944](#), [946](#)

task reduction

A [reduction](#) that is performed over a set of [tasks](#) that may include [explicit tasks](#). [255](#), [252](#), [256](#), [942](#)

task region

A [region](#) consisting of all code encountered during the execution of a [task](#). [4–6](#), [8](#), [39](#), [43](#), [86](#), [98](#), [103](#), [110](#), [114](#), [228](#), [275](#), [402](#), [403](#), [412](#), [468](#), [469](#), [474](#), [476](#), [478](#), [486](#), [521](#), [523](#), [541](#), [608](#), [697](#), [740](#), [744](#), [747](#), [782](#), [826](#), [888](#), [893](#), [912](#)

task scheduling point

A point during the execution of the [current task region](#) at which the [task](#) can be suspended to be resumed later; or the point of [task completion](#), after which the executing [thread](#) may switch to a different [task](#). [467](#), [5](#), [103](#), [249](#), [274](#), [403](#), [447](#), [466–469](#), [495](#), [496](#), [498](#), [499](#), [515](#), [521](#), [633](#), [640](#), [767](#), [783](#), [947](#)

task synchronization construct

A [taskwait](#), a [taskgroup](#), or a [barrier](#) construct. [5](#), [446](#), [468](#)

team

A set of one or more [assigned threads](#) assigned to execute the set of [implicit tasks](#) of a [parallel region](#). [4](#), [3](#), [4](#), [7](#), [20](#), [26](#), [39](#), [61](#), [63](#), [66](#), [83](#), [85](#), [89](#), [90](#), [102](#), [105](#), [108](#), [109](#), [112](#), [114](#), [117](#), [119](#), [128](#), [136](#), [137](#), [252](#), [254](#), [259](#), [261](#), [275](#), [285–287](#), [379](#), [402](#), [403](#), [407–413](#), [415](#), [421](#), [423–425](#), [427–429](#), [433](#), [434](#), [437–440](#), [442](#), [444](#), [473](#), [493](#), [495](#), [496](#), [515](#), [522](#), [523](#), [536](#), [543](#), [589](#), [590](#), [601](#), [603](#), [619](#), [620](#), [750](#), [758](#), [776](#), [784](#), [811](#), [824](#), [825](#), [857](#), [883](#), [887](#), [888](#), [892](#), [917](#), [921](#), [922](#), [938](#), [940](#), [948–951](#)

team-executed construct

A [construct](#) that has the [team-executed property](#). [4](#)

team-executed property

The [property](#) that a [construct](#) gives rise to a [team-executed region](#). [107](#), [424–426](#), [428](#), [435](#), [436](#), [442](#), [495](#)

team-executed region

A [region](#) that is executed by all or none of the [threads](#) in the [current team](#). [4](#), [107](#), [951](#)

team-generating construct

A [construct](#) that has the [team-generating property](#). [951](#)

team-generating property

The [property](#) that a [construct](#) generates a [parallel region](#). [107](#), [402](#), [412](#), [481](#)

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](#). [107](#), [120](#), [442](#), [603](#), [784](#)

teams-nestable construct

A **construct** that has the **teams-nestable** property. 414, 951

teams-nestable property

The **property** that a **construct** or **routine** generates a **region** that may be a **strictly nested region** of a **teams** region. 108, 392, 393, 395, 397, 398, 400, 402, 439, 442, 601, 602

teams-nestable routine

A **routine** that has the **teams-nestable** property. 414, 951

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**. 4, 116, 136

temporary view

The state of **memory** that is accessible to a particular **thread**. 7, 7, 10, 11, 519

third-party tool

A **tool** that executes as a separate process from the process that it is monitoring and potentially controlling. 844, 15, 46, 79, 119, 844–846, 848, 849, 851, 854, 857–859, 861, 863, 864, 869, 871, 874, 875, 880, 908, 944

thread

Unless specifically stated otherwise, an **OpenMP thread**. 3–8, 10–15, 19, 20, 22, 23, 25, 26, 29, 35, 39, 41, 47, 50, 51, 53–55, 63, 64, 69, 73, 83, 85–87, 89, 90, 94–96, 98, 101–105, 107–110, 112–114, 116–120, 122, 132, 133, 137–139, 141, 145, 146, 152, 210, 228, 229, 249, 250, 252, 254, 259, 261, 274–276, 285–288, 294, 316–319, 363, 369, 370, 377, 383, 402–413, 421, 423–434, 437–440, 443, 445–447, 449–451, 455, 459, 462, 467–469, 473, 475, 477, 482, 483, 487, 492–498, 500, 502, 514–517, 519–524, 529, 533, 535–537, 540–544, 554, 581, 582, 588–593, 599, 604, 605, 610, 621, 622, 628, 629, 631–634, 640, 689–702, 705, 716, 719, 721, 725, 730, 740, 750, 758, 759, 768, 773, 774, 776, 780, 784, 791, 795, 807, 813, 819, 821, 823–827, 830, 841, 842, 849, 858–861, 864, 865, 868, 869, 874, 883, 884, 887–891, 893, 900, 908, 916–918, 920–922, 932, 933, 939, 944, 947–951

thread affinity

A binding of **threads** to **places** within the current **place partition**. 407, 87, 118, 119, 135, 136, 139–141, 275, 407–410, 702, 710, 711, 916, 921, 941, 946

thread-exclusive construct

A **construct** that has the **thread-exclusive** property. 951

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. [108](#), [493](#), [536](#)

thread-limiting construct

A [construct](#) that has the [thread-limiting property](#). [152](#)

thread-limiting property

For C++, the [property](#) that a [construct](#) limits the [threads](#) that can catch an exception thrown in the corresponding [region](#) to the [thread](#) that threw the exception. [109](#), [402](#), [412](#), [421](#), [424–426](#), [446](#), [481](#), [493](#), [536](#)

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 [thread number](#) is the value `omp_unassigned_thread`. [402](#), [89](#), [109](#), [120](#), [275](#), [402](#), [408](#), [411](#), [422](#), [438](#), [590](#), [598](#), [784](#), [826](#), [883](#), [949](#)

thread-pool-worker thread

A [thread](#) in an [OpenMP thread pool](#) that is not the [initial thread](#). [768](#)

threadprivate attribute

For a given [OpenMP thread](#), a [data-sharing attribute](#) of a [data entity](#) that it has [static storage duration](#), or thread storage duration for C/C++, and is visible only to [tasks](#) that are executed by the [thread](#). [274](#), [109](#), [216](#), [219](#), [275](#), [276](#), [278](#), [286](#), [288](#), [949](#)

threadprivate memory

The set of [threadprivate variables](#) associated with each [thread](#). [7](#), [276](#), [469](#), [918](#)

threadprivate variable

A [variable](#) that has the [threadprivate attribute](#) with respect to a given [OpenMP thread](#). [274](#), [109](#), [274–278](#), [285](#), [416](#), [433](#), [483](#)

thread-reservation type

A categorization of a [thread](#) as either a [structured thread](#) or a [free-agent thread](#). [145](#), [94](#), [95](#)

thread-safe procedure

A [procedure](#) that performs the intended function even when executed concurrently (by multiple [native threads](#)). [15](#)

thread-selecting construct

A [construct](#) that has the [thread-selecting property](#). [546](#), [547](#)

thread-selecting property

The [property](#) that a [construct](#) selects a subset of [threads](#) that can execute the corresponding [region](#) from the [binding thread set](#) of the [region](#). [109](#), [421](#), [424](#)

thread-set

The set of [threads](#) for which a [flush](#) may enforce [memory](#) consistency. [10](#), [10](#), [12](#), [13](#), [514](#), [519](#), [521](#)

thread state

The state associated with a [thread](#), which may be represented by 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. [5](#), [14](#), [721](#), [725](#), [758](#), [815](#), [823](#), [899](#), [900](#), [925](#)

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](#). [5](#), [4](#), [402](#), [459](#), [468](#)

tile

For a [tile](#) directive, the [logical iteration space](#) of the [tile loops](#). [34](#), [86](#), [110](#), [391](#), [399](#), [401](#), [928](#)

tile loop

The inner [generated loops](#) of a [tile](#) construct that iterate over the [logical iterations](#) that correspond to a [tile](#). [398](#), [110](#), [398](#), [399](#), [401](#), [920](#), [933](#)

tool

Code that can observe and/or modify the execution of an application. [2](#), [14](#), [15](#), [17](#), [53](#), [93](#), [95](#), [108](#), [110](#), [120](#), [148–150](#), [473](#), [479](#), [585–587](#), [635](#), [636](#), [638](#), [639](#), [641](#), [642](#), [713](#), [718](#), [719](#), [721–725](#), [727–731](#), [740](#), [745](#), [747](#), [751](#), [756](#), [758](#), [770–774](#), [776–780](#), [782–805](#), [807](#), [809](#), [811](#), [813–824](#), [826–842](#), [861–871](#), [873–884](#), [887](#), [888](#), [894](#), [896–903](#), [905](#), [906](#), [925](#)

tool callback

A [procedure](#) that a [tool](#) provides to an OpenMP implementation to invoke when an associated [event](#) occurs. [14](#), [29](#), [30](#), [496](#), [533](#), [551](#), [729](#), [771](#), [836](#), [925](#)

tool context

An opaque reference provided by a [tool](#) to an [OMPD](#) library. A [tool context](#) uniquely identifies an abstraction. [20](#), [77](#), [110](#), [862](#), [868](#)

tool defined

Behavior that must be documented by the [tool](#) implementation, and is allowed to vary among different compliant [tools](#). [586](#), [719](#), [797](#)

trace record

A data structure in which to store information associated with an occurrence of an [event](#). [46](#), [77](#), [100](#), [114](#), [189](#), [190](#), [729–731](#), [735](#), [751](#), [753](#), [770](#), [787](#), [799–805](#), [807](#), [809](#), [811](#), [831](#), [833](#), [834](#), [836](#), [838–842](#), [925](#), [927](#)

trait

An aspect of an OpenMP implementation or the execution of an [OpenMP program](#). [9](#), [21](#), [32](#), [37](#), [46](#), [49](#), [52](#), [60](#), [76](#), [78](#), [81](#), [100](#), [105](#), [111](#), [115](#), [142–144](#), [147](#), [315–320](#), [325](#), [327](#), [335–340](#), [354](#), [372](#), [567](#), [568](#), [575](#), [660](#), [666](#), [667](#), [675](#), [676](#), [919](#), [929](#), [931](#), [932](#), [939](#), [943](#)

trait selector

A member of a [trait selector set](#). [337](#), [335](#), [338–341](#), [343](#), [347](#), [354](#)

trait selector set

A set of [traits](#) that are specified to match the [trait set](#) at a given point in an [OpenMP program](#). [337](#), [97](#), [111](#), [339](#)

trait set

A grouping of related [traits](#). [335](#), [36](#), [46](#), [49](#), [60](#), [105](#), [111](#), [335](#), [338](#), [340](#)

transformation-affected loop

For a [loop-transforming construct](#), an [affected loop](#) that is replaced according to the semantics of the constituent [loop-transforming directive](#). [53](#), [210](#), [386](#), [388](#), [392–401](#)

translation unit

For C/C++, a translation unit in the base language. For Fortran, a source text file, after processing all **INCLUDE** lines, that contains one or more program units. [282](#), [918](#)

transparent task

A [task](#) for which [child tasks](#) are visible to external [dependence-compatible tasks](#) for the purposes of establishing [task dependences](#). Unless otherwise specified, a [transparent task](#) is both an [importing task](#) and an [exporting task](#). [531](#), [111](#), [457](#)

type-name list

An [argument list](#) that consists of [type-name list items](#). [165](#), [173](#), [263](#), [264](#), [308](#)

type-name list item

A [list item](#) that is the name of a type. [166](#), [111](#), [166–168](#), [266](#), [267](#)

U

ultimate property

The [property](#) that a [clause](#) or an argument must be the lexically last [clause](#) or argument to appear on the [directive](#). For a [modifier](#), the [property](#) that it must be the lexically last [modifier](#) to appear on a pre-modified [clause](#) or that it must be the lexically first [modifier](#) to appear on a post-modified [clause](#). [164](#), [163–165](#), [213](#), [250](#), [251](#), [255–257](#), [332](#), [343](#), [415](#), [437](#), [441](#)

unassigned thread

A [thread](#) that is not currently assigned to any [team](#). [3](#), [3](#), [4](#), [54](#), [59](#), [96](#), [109](#), [462](#), [468](#), [590](#), [759](#)

unassociated directive

A [directive](#) that is not directly associated with any [base-language code](#). [156](#), [100](#), [156–158](#), [262](#), [266](#), [307](#), [344](#), [369](#), [372](#), [385](#), [386](#), [454](#), [466](#), [474](#), [476](#), [486](#), [488](#), [495](#), [499](#), [518](#), [525](#), [534](#), [540](#), [544](#)

unconditional-update structured block

An [update structured block](#) that may be associated with an [atomic directive](#) that expresses an [atomic update](#) operation that is neither an [atomic captured update](#) or [atomic conditional update](#) operation. [195](#), [195](#), [196](#)

underrferred 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 [underrferred task](#) is completed. [447](#), [61](#), [75](#), [112](#), [447](#), [450](#), [457](#), [460](#), [524](#)

undefined

For [variables](#), the property of not being [defined](#); that is, the [variable](#) does not have a valid value. [9](#), [150](#), [542](#), [768](#), [818](#), [821](#), [822](#), [824](#), [826–828](#)

underlying map type

The [map type](#) that determines which [output map type](#) results from an [input map type](#). [297](#), [72](#), [297](#)

unified address space

An [address space](#) that is used by all [devices](#). [376](#)

uninitialized state

The [lock state](#) that indicates the [lock](#) must be initialized before it can be set. [67](#), [660](#), [663](#), [689](#), [692](#), [695](#), [699](#)

union

A **union** is a type that defines one or more fields that overlap in memory, so only one of the fields can be used at any given time. For C/C++, implemented using union types. For Fortran, implemented using derived types. [113](#), [734](#), [735](#), [739](#)

unique property

The **property** that a **clause**, a **modifier**, or an argument may appear at most once in a given context with which it is associated. For a **clause set**, each member of the **clause set** may appear at most once in the given context. [164](#), [163](#), [164](#), [172](#), [178](#), [184–187](#), [210–213](#), [223](#), [225–227](#), [230](#), [233–236](#), [251](#), [255–257](#), [260](#), [265](#), [268](#), [269](#), [272](#), [273](#), [281](#), [283](#), [285](#), [287](#), [291](#), [292](#), [297](#), [299](#), [305](#), [306](#), [312](#), [313](#), [320](#), [321](#), [324](#), [327](#), [329](#), [331](#), [332](#), [342](#), [343](#), [347](#), [348](#), [350](#), [356–358](#), [360–362](#), [368](#), [370](#), [371](#), [373–384](#), [389](#), [391](#), [394](#), [396](#), [397](#), [400](#), [401](#), [406](#), [410](#), [411](#), [415](#), [416](#), [419](#), [420](#), [422](#), [437](#), [441](#), [444](#), [452](#), [453](#), [458–465](#), [470–472](#), [490](#), [492](#), [501](#), [503–513](#), [525–527](#), [531](#), [532](#), [537–539](#)

unit of work

In **constructs** that use **units of work**, one or more 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**. [113](#), [428](#), [429](#), [431](#), [432](#), [779](#)

unlocked state

The **lock state** that indicates the **lock** can be set by any **task**. [688](#), [67](#), [68](#), [688](#), [689](#), [692](#), [695–699](#)

unsigned property

The **property** that a **routine** or **callback** either returns an unsigned type in C/C++ or has an argument that has such a type. [722](#), [775](#), [784](#), [791](#), [808](#), [810](#), [833](#)

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**. **Unspecified behavior** may result from:

- Issues that this specification documents as having **unspecified behavior**;
- A **non-conforming program**; or
- A **conforming program** exhibiting an **implementation defined** behavior.

[7–9](#), [34](#), [41](#), [59](#), [113](#), [130](#), [153](#), [223](#), [228](#), [235](#), [241](#), [246](#), [276](#), [280](#), [309](#), [317](#), [325](#), [376](#), [379](#), [463](#), [464](#), [482](#), [484](#), [497](#), [530](#), [542](#), [581](#), [612–614](#), [616–623](#), [628](#), [631](#), [632](#), [643](#), [650](#), [667](#), [668](#), [677](#), [682](#), [685](#), [687](#), [688](#), [706](#), [707](#), [709–711](#), [717](#), [830](#)

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](#). [5](#), [105](#), [276](#), [447](#), [459](#), [468](#), [523](#), [947](#)

untraced-argument property

The [property](#) of an argument of a [callback](#) that it is omitted from the corresponding [trace record](#) of the [callback](#). [773](#), [775](#), [781](#), [796](#), [803](#), [806](#), [810](#)

update-capture structured block

A [capture structured block](#) that may be associated with an [atomic directive](#) that expresses an [atomic captured update](#) operation that is not also an [atomic conditional update](#) operation. [197](#), [197](#), [198](#), [517](#)

update structured block

An [atomic structured block](#) that may be associated with an [atomic directive](#) that expresses an [atomic update](#) operation. [195](#), [35](#), [112](#), [197](#), [517](#)

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](#). [270](#), [114](#), [270](#)

use-device-addr attribute

For a given [device construct](#), a [data-sharing attribute](#) of a [data entity](#) that refers to an object in a [device data environment](#) that corresponds to the [data entity](#) of the same name in the [enclosing data environment](#) of the [construct](#) if such an object exists, and otherwise refers to the entity in the [enclosing data environment](#). [236](#)

use-device-ptr attribute

For a given [device construct](#), a [data-sharing attribute](#) of a [C pointer variable](#) that implies the [private attribute](#), and additionally the [variable](#) is initialized to be a [device pointer](#) that refers to the [device address](#) that corresponds to the value of a [C pointer](#) of the same name in the [enclosing data environment](#) of the [construct](#). [234](#)

user-defined cancellation point

A [cancellation point](#) that is specified by a [cancellation point construct](#). [544](#), [544](#)

user-defined induction

An [induction operation](#) that is defined by a [declare_induction directive](#). [266](#), [178](#), [267–269](#), [930](#)

user-defined mapper

A [mapper](#) that is defined by a [declare_mapper directive](#). [308](#), [72](#), [88](#), [188](#), [301](#), [308](#), [310](#), [936](#)

user-defined reduction

A [reduction operation](#) that is defined by a [declare_reduction](#) directive. [263](#), [178](#), [263](#), [265](#), [543](#), [947](#)

user selector set

A [selector set](#) that may match [traits](#) in the [dynamic trait set](#). [338](#), [338–340](#)

utility directive

A [directive](#) that facilitates interactions with the compiler and/or supports code readability. A [utility directive](#) is an [informational directive](#) except when specified to be an [executable directive](#). [369](#), [115](#), [156](#), [369](#), [370](#), [386](#)

V

value property

The [property](#) that a [routine](#) parameter does not have a pointer type in C/C++ and has the **VALUE** attribute in Fortran. [555](#), [575](#), [625–627](#), [629](#), [630](#), [632](#), [634](#), [636–638](#), [640](#), [641](#), [677–682](#), [684](#), [685](#), [687](#), [759](#), [796](#), [800](#), [803](#)

variable

A [referencing variable](#) or 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](#).

[7–13](#), [15](#), [16](#), [20](#), [25–28](#), [31](#), [37–39](#), [41](#), [42](#), [44](#), [53](#), [54](#), [57](#), [62](#), [64–67](#), [72](#), [73](#), [76](#), [84](#), [89](#), [93](#), [98](#), [101](#), [102](#), [104](#), [109](#), [112](#), [114](#), [115](#), [118](#), [156](#), [158](#), [159](#), [166–168](#), [173](#), [178](#), [186–190](#), [192](#), [194](#), [204](#), [206](#), [210](#), [215–219](#), [221–225](#), [228](#), [229](#), [231](#), [232](#), [236](#), [238–243](#), [247](#), [254](#), [258](#), [259](#), [262](#), [263](#), [267](#), [274–280](#), [282](#), [283](#), [285–291](#), [293–295](#), [297](#), [301](#), [302](#), [306–309](#), [315](#), [316](#), [319–324](#), [326](#), [327](#), [339](#), [342](#), [346](#), [348](#), [357](#), [358](#), [362](#), [363](#), [365–367](#), [377](#), [378](#), [388](#), [406](#), [412](#), [422](#), [424](#), [433](#), [437](#), [442](#), [443](#), [447](#), [450](#), [456](#), [457](#), [462](#), [465](#), [470](#), [474](#), [476](#), [479](#), [481–484](#), [486](#), [519](#), [520](#), [531](#), [533](#), [548–550](#), [688](#), [729](#), [768](#), [805](#), [815](#), [816](#), [818](#), [823–828](#), [845](#), [846](#), [858](#), [862](#), [864](#), [880](#), [915](#), [918](#), [919](#), [931](#), [937](#), [938](#), [942](#), [945](#), [949](#)

variable list

An [argument list](#) that consists of [variable list items](#). [165](#), [51](#), [248](#), [324](#)

variable list item

For C/C++, a [list item](#) that is a [variable](#) or an [array section](#); for Fortran, a [list item](#) that is a named item specifically identified in [Section 5.2.1](#). [166](#), [52](#), [115](#), [165–168](#), [293](#), [456](#), [457](#)

variant-generating directive

A [declarative directive](#) that has the [variant-generating property](#). 342

variant-generating property

The [property](#) that a [declarative directive](#) generates a variant of a [procedure](#). 116, 358, 364, 366

variant substitution

The replacement of a call to a [base function](#) by a call to a [function variant](#). 55, 346, 355, 939

W

wait identifier

A unique [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. 768, 768, 823

white space

A non-empty sequence of space and/or horizontal tab characters. 47, 130, 138, 140, 153, 154, 159–162, 176, 177, 546, 927

work distribution

The manner in which execution of a [region](#) that corresponds to a [work-distribution construct](#) is assigned to [threads](#). 211

work-distribution construct

A [construct](#) that has the [work-distribution property](#). 423, 2, 86, 116, 117, 228, 229, 232, 253, 423, 424, 443, 778

work-distribution property

The [property](#) that a [construct](#) is cooperatively executed by [threads](#) in the [binding thread set](#) of the corresponding [region](#). 116, 424–426, 428, 431, 435, 436, 439, 442

work-distribution region

A [region](#) that corresponds to a [work-distribution construct](#). 229, 232, 423, 424

worker thread

Unless specifically stated otherwise, a [team-worker thread](#). 109, 403

worksharing construct

A [construct](#) that has the [worksharing property](#). 423, 4, 86, 117, 229, 252–254, 259, 261, 427, 433, 443, 497, 543, 547, 548, 758

worksharing-loop construct

A [construct](#) that has the [worksharing-loop property](#). [433](#), [48](#), [117](#), [137](#), [253](#), [258](#), [434–438](#), [534–537](#), [541](#), [543](#), [546](#), [779](#), [921](#), [931](#), [937](#), [943](#), [945](#), [949](#), [952](#)

worksharing-loop property

The [property](#) of a [worksharing construct](#) that it is a [loop-nest-associated construct](#) that distributes the [collapsed iterations](#) of the [affected loops](#) among the [threads](#) in the [team](#). [117](#), [435](#), [436](#), [549](#)

worksharing-loop region

A [region](#) that corresponds to a [worksharing-loop construct](#). [433](#), [120](#), [128](#), [433](#), [434](#), [534](#), [536](#), [952](#)

worksharing property

The [property](#) of a [construct](#) that it is a [work-distribution construct](#) that is executed by the [team](#) of the innermost enclosing [parallel region](#) and includes, by default, an implicit barrier. [116](#), [424–426](#), [428](#), [435](#), [436](#), [442](#)

worksharing region

A [region](#) that corresponds to a [worksharing construct](#). [423](#), [4](#), [229](#), [252](#), [423](#), [496](#), [522](#), [779](#), [939](#), [951](#)

write structured block

An [atomic structured block](#) that may be associated with an [atomic directive](#) that expresses an [atomic write](#) operation. [195](#), [195](#), [197](#), [517](#)

Z

zeroed-memory-allocating routine

A [memory-allocating routine](#) that has the [zeroed-memory-allocating-routine property](#). [675](#), [675](#), [679](#), [680](#)

zeroed-memory-allocating-routine property

The [property](#) that a [memory-allocating routine](#) returns a pointer to [memory](#) that has been set to zero. [675](#), [117](#), [679](#), [680](#)

zero-length array section

An [array section](#) that does not include any elements of the array. [171](#), [246](#), [293](#), [529](#)

zero-offset assumed-size array

An [assumed-size array](#) for which the lower bound is zero. [234](#), [290](#), [301](#)

3 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 a distinct copy of the [ICV](#) exists for each [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 3.2](#).

3.1 ICV Descriptions

Section [3.1](#) shows the [ICV scope](#) and description of each [ICV](#).

TABLE 3.1: ICV Scopes and Descriptions

ICV	Scope	Description
active-levels-var	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
affinity-format-var	device	Controls the thread affinity format when displaying thread affinity
available-devices-var	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 third-party 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	Boolean that is true if a given task is an explicit task , otherwise false
<i>final-task-var</i>	data environment	Boolean that is true if a given task is a final task , otherwise false
<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>free-agent-var</i>	data environment	Boolean that is true if a free-agent thread is currently executing a given task , otherwise false
<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>nteam-var</i>	device	Controls the number of teams requested for encountered teams regions

ICV	Scope	Description
<i>nthreads-var</i>	data environment	Controls the number of threads requested for encountered parallel regions
<i>num-devices-var</i>	global	Number of available non-host devices
<i>num-procs-var</i>	device	The number of processors available on the device
<i>place-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 type
<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

3.2 ICV Initialization

Section 3.2 shows the ICVs, associated environment variables, and initial values.

TABLE 3.2: ICV Initial Values

ICV	Environment Variable	Initial Value
<i>active-levels-var</i>	(none)	0 (zero)
<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	false
<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)	0 (zero)
<i>display-affinity-var</i>	OMP_DISPLAY_AFFINITY	false
<i>dyn-var</i>	OMP_DYNAMIC	implementation defined
<i>explicit-task-var</i>	(none)	false
<i>final-task-var</i>	(none)	false
<i>free-agent-thread-limit-var</i>	OMP_THREAD_LIMIT, OMP_THREADS_RESERVE	See below
<i>free-agent-var</i>	(none)	false
<i>league-size-var</i>	(none)	1 (one)
<i>levels-var</i>	(none)	0 (zero)
<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	0 (zero)
<i>nteam-var</i>	OMP_NUM_TEAMS	0 (zero)
<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

ICV	Environment Variable	Initial Value
<i>structured-thread-limit-var</i>	OMP_THREAD_LIMIT , OMP_THREADS_RESERVE	<i>See below</i>
<i>target-offload-var</i>	OMP_TARGET_OFFLOAD	<i>default</i>
<i>team-generator-var</i>	(none)	<i>0 (zero)</i>
<i>team-num-var</i>	(none)	<i>0 (zero)</i>
<i>team-size-var</i>	(none)	<i>1 (one)</i>
<i>teams-thread-limit-var</i>	OMP_TEAMS_THREAD_LIMIT	<i>0 (zero)</i>
<i>thread-limit-var</i>	OMP_THREAD_LIMIT	<i>implementation defined</i>
<i>thread-num-var</i>	(none)	<i>0 (zero)</i>
<i>tool-libraries-var</i>	OMP_TOOL_LIBRARIES	<i>empty string</i>
<i>tool-var</i>	OMP_TOOL	<i>enabled</i>
<i>tool-verbose-init-var</i>	OMP_TOOL_VERBOSE_INIT	<i>disabled</i>
<i>wait-policy-var</i>	OMP_WAIT_POLICY	<i>implementation defined</i>

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 21.2); the semantic and precedence is described in Chapter 4.

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 the value returned by **omp_get_initial_device**.

- The value of the *nthreads-var* ICV is a list.

- The value of the *bind-var* ICV is a list.

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

Cross References

- `OMP_AFFINITY_FORMAT`, see [Section 4.3.5](#)
- `OMP_ALLOCATOR`, see [Section 4.4.1](#)
- `OMP_AVAILABLE_DEVICES`, see [Section 4.3.7](#)
- `OMP_CANCELLATION`, see [Section 4.3.6](#)
- `OMP_DEBUG`, see [Section 4.6.1](#)
- `OMP_DEFAULT_DEVICE`, see [Section 4.3.8](#)
- `OMP_DISPLAY_AFFINITY`, see [Section 4.3.4](#)
- `OMP_DYNAMIC`, see [Section 4.1.2](#)
- `OMP_MAX_ACTIVE_LEVELS`, see [Section 4.1.5](#)
- `OMP_MAX_TASK_PRIORITY`, see [Section 4.3.11](#)
- `OMP_NUM_TEAMS`, see [Section 4.2.1](#)
- `OMP_NUM_THREADS`, see [Section 4.1.3](#)
- `OMP_PLACES`, see [Section 4.1.6](#)
- `OMP_PROC_BIND`, see [Section 4.1.7](#)
- `OMP_SCHEDULE`, see [Section 4.3.1](#)
- `OMP_STACKSIZE`, see [Section 4.3.2](#)
- `OMP_TARGET_OFFLOAD`, see [Section 4.3.9](#)
- `OMP_TEAMS_THREAD_LIMIT`, see [Section 4.2.2](#)
- `OMP_THREAD_LIMIT`, see [Section 4.1.4](#)
- `OMP_TOOL`, see [Section 4.5.1](#)
- `OMP_TOOL_LIBRARIES`, see [Section 4.5.2](#)
- `OMP_WAIT_POLICY`, see [Section 4.3.3](#)

3.3 Modifying and Retrieving ICV Values

Section 3.3 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 3.3: Ways to Modify and to Retrieve ICV Values

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

ICV	Ways to Modify Value	Ways to Retrieve Value
<i>place-partition-var</i>	(none)	<code>omp_get_partition_num_places</code> , <code>omp_get_partition_place_nums</code> , <code>omp_get_place_num_procs</code> , <code>omp_get_place_proc_ids</code>
<i>run-sched-var</i>	<code>omp_set_schedule</code>	<code>omp_get_schedule</code>
<i>stacksize-var</i>	(none)	(none)
<i>structured-thread-limit-var</i>	(none)	(none)
<i>target-offload-var</i>	(none)	(none)
<i>team-generator-var</i>	(none)	(none)
<i>team-num-var</i>	(none)	<code>omp_get_team_num</code>
<i>team-size-var</i>	(none)	<code>omp_get_num_threads</code>
<i>teams-thread-limit-var</i>	<code>omp_set_device_teams_thread_limit</code>	<code>omp_get_device_teams_thread_limit</code> <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

- `omp_get_active_level` Routine, see [Section 27.17](#)
- `omp_get_affinity_format` Routine, see [Section 35.9](#)
- `omp_get_cancellation` Routine, see [Section 36.1](#)

- `omp_get_default_allocator` Routine, see [Section 33.10](#)
- `omp_get_default_device` Routine, see [Section 30.2](#)
- `omp_get_device_num` Routine, see [Section 30.4](#)
- `omp_get_device_num_teams` Routine, see [Section 30.11](#)
- `omp_get_device_teams_thread_limit` Routine, see [Section 30.13](#)
- `omp_get_dynamic` Routine, see [Section 27.8](#)
- `omp_get_level` Routine, see [Section 27.14](#)
- `omp_get_max_active_levels` Routine, see [Section 27.13](#)
- `omp_get_max_task_priority` Routine, see [Section 29.1.1](#)
- `omp_get_max_teams` Routine, see [Section 28.4](#)
- `omp_get_max_threads` Routine, see [Section 27.4](#)
- `omp_get_num_devices` Routine, see [Section 30.3](#)
- `omp_get_num_procs` Routine, see [Section 30.5](#)
- `omp_get_num_teams` Routine, see [Section 28.1](#)
- `omp_get_num_threads` Routine, see [Section 27.2](#)
- `omp_get_partition_num_places` Routine, see [Section 35.6](#)
- `omp_get_partition_place_nums` Routine, see [Section 35.7](#)
- `omp_get_place_num_procs` Routine, see [Section 35.3](#)
- `omp_get_place_proc_ids` Routine, see [Section 35.4](#)
- `omp_get_proc_bind` Routine, see [Section 35.1](#)
- `omp_get_schedule` Routine, see [Section 27.10](#)
- `omp_get_supported_active_levels` Routine, see [Section 27.11](#)
- `omp_get_team_num` Routine, see [Section 28.3](#)
- `omp_get_teams_thread_limit` Routine, see [Section 28.5](#)
- `omp_get_thread_limit` Routine, see [Section 27.5](#)
- `omp_get_thread_num` Routine, see [Section 27.3](#)
- `omp_in_explicit_task` Routine, see [Section 29.1.2](#)
- `omp_in_final` Routine, see [Section 29.1.3](#)
- `omp_set_affinity_format` Routine, see [Section 35.8](#)

- `omp_set_default_allocator` Routine, see [Section 33.9](#)
- `omp_set_default_device` Routine, see [Section 30.1](#)
- `omp_set_device_num_teams` Routine, see [Section 30.12](#)
- `omp_set_device_teams_thread_limit` Routine, see [Section 30.14](#)
- `omp_set_dynamic` Routine, see [Section 27.7](#)
- `omp_set_max_active_levels` Routine, see [Section 27.12](#)
- `omp_set_num_teams` Routine, see [Section 28.2](#)
- `omp_set_num_threads` Routine, see [Section 27.1](#)
- `omp_set_schedule` Routine, see [Section 27.9](#)
- `omp_set_teams_thread_limit` Routine, see [Section 28.6](#)
- `thread_limit` Clause, see [Section 21.3](#)

3.4 How the Per-Data Environment ICVs Work

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

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

When a [task-generating construct](#) is encountered, each [generated task](#) inherits the value of [nthreads-var](#) from the [nthreads-var](#) value of the [generating task](#). If a [parallel construct](#) is encountered on which a [num_threads](#) clause is specified with a [nthreads](#) list of more than one list item, the value of [nthreads-var](#) for the generated [implicit tasks](#) is the list obtained by deletion of the first item of the [nthreads](#) list. Otherwise, when a [parallel construct](#) is encountered, if the [nthreads-var](#) list of the [generating task](#) contains a single element, the generated [implicit tasks](#) inherit that list as the value of [nthreads-var](#); if the [nthreads-var](#) list of the [generating task](#) contains multiple elements, the generated [implicit tasks](#) inherit the value of [nthreads-var](#) as the list obtained by deletion of the first element from the [nthreads-var](#) value of the [generating task](#). The [bind-var ICV](#) is handled in the same way as the [nthreads-var ICV](#), except that an override list cannot be specified through the [proc_bind](#) clause of an encountered [parallel construct](#).

When a [target construct](#) corresponds to an [active target region](#), the resulting [initial task](#) uses the values of the [data environment](#) scoped [ICVs](#) from the [device data environment ICV](#) values of the [device](#) that will execute the [region](#), unless otherwise specified.

When a [target construct](#) corresponds to an [inactive target region](#), the resulting [initial task](#) uses the values of the [ICVs](#) with [data environment ICV scope](#) from the [data environment](#) of the [task](#) that encountered the [target construct](#), unless otherwise specified.

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

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

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

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

If a **target** construct, **teams** construct, or **parallel** construct is encountered, the *team-generator-var* ICV for the **data environments** of the generated **implicit tasks** is instead set to the value of the appropriate **team** generator type as specified in [Section 45.13](#).

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

Cross References

- OMPD **team_generator** Type, see [Section 45.13](#)

3.5 ICV Override Relationships

Section [3.5](#) shows the override relationships among **construct clauses** and **ICVs**. The table only lists **ICVs** that can be overridden by a **clause**.

TABLE 3.4: ICV Override Relationships

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

1 If a **schedule** clause specifies a **modifier** then that **modifier** overrides any **modifier** that is
2 specified in the *run-sched-var* ICV.

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

5 Cross References

- 6 • **allocate** Clause, see [Section 13.6](#)
- 7 • **allocator** Clause, see [Section 13.4](#)
- 8 • **num_teams** Clause, see [Section 18.2.1](#)
- 9 • **num_threads** Clause, see [Section 18.1.2](#)
- 10 • **proc_bind** Clause, see [Section 18.1.4](#)
- 11 • **schedule** Clause, see [Section 19.6.3](#)
- 12 • **thread_limit** Clause, see [Section 21.3](#)

4 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 3](#)). 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](#). Unless otherwise specified, setting an [OpenMP environment variable](#) to an empty string is [unspecified](#). The assigned values for most [environment variables](#) are strings or integers. In particular, boolean values are specified as the string [true](#) or [false](#). 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](#). These examples demonstrate how to set the [OpenMP environment variables](#) 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 in [Section 3.2](#), [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 its setting 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 through the use of the [_DEV](#) suffix or the absence of a 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 ICV scope](#).

- **Device-specific environment variables** must not specify the **host device**.

4.1 Parallel Region Environment Variables

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

4.1.1 Abstract Name Values

This section defines **abstract names** that must be understood by the execution and runtime environment for the **environment variables** that explicitly allow them. The entities defined by the **abstract names** are **implementation defined**. There are two kinds of **abstract names**: **conceptual abstract names** and **numeric abstract names**.

Conceptual abstract names include **place-list abstract names** that are the strings defined in Table 4.1. If an **environment variable** is set to a value that includes a **place-list abstract name**, the behavior is as if the **place-list abstract name** were replaced with the list of **places** associated with that **abstract name** on each **device** where the **environment variable** is applied.

TABLE 4.1: Predefined Place-list Abstract Names

Abstract Name	Meaning
threads	A set where each place corresponds to a single hardware thread of the device .
cores	A set where each place corresponds to a single core of the device .
ll_caches	A set where each place corresponds to the set of cores for a single last-level cache of the device .
numa_domains	A set where each place corresponds to the set of cores for a single NUMA domain of the device .
sockets	A set where each place corresponds to the set of cores for a single socket of the device .

For each **place-list abstract name** specified in Table 4.1, a corresponding **place-count abstract name** prefixed with **n_** also exists for which the associated value is the number of **places** in the list of **places** specified by the **place-list abstract name**, as described above.

If an **environment variable** is set to a value that includes a **numeric abstract name**, the behavior is as if the **numeric abstract name** were replaced with the value associated with that **numeric abstract name**.

4.1.2 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:

```
export OMP_DYNAMIC=true
```

Cross References

- *dyn-var ICV*, see [Table 3.1](#)
- **omp_get_dynamic** Routine, see [Section 27.8](#)
- **omp_set_dynamic** Routine, see [Section 27.7](#)
- **parallel** Construct, see [Section 18.1](#)

4.1.3 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 3](#) 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** routine and dynamic adjustment of **threads**, and [Section 18.1.1](#) for a complete algorithm that describes how the number of **threads** for a **parallel region** is determined.

The value of this **environment variable** must be a list of **positive** integer values and/or **numeric abstract names**. The values of the list set the number of **threads** to use for **parallel regions** at the corresponding nested levels.

The behavior of the program is **implementation defined** if any value of the list specified in the **OMP_NUM_THREADS** environment variable leads to a number of **threads** that is greater than an implementation can support or if any value is not a **positive** integer.

The **OMP_NUM_THREADS** environment variable sets the *max-active-levels-var ICV* to the number of **active levels** of parallelism that the implementation supports if the **OMP_NUM_THREADS** environment variable is set to a comma-separated list of more than one value. The value of the

max-active-levels-var ICV may be overridden by setting `OMP_MAX_ACTIVE_LEVELS`. See [Section 4.1.5](#) for details.

Example:

```
export OMP_NUM_THREADS=4, 3, 2
export OMP_NUM_THREADS=n_cores, 2
```

Cross References

- `OMP_MAX_ACTIVE_LEVELS`, see [Section 4.1.5](#)
- *nthreads-var* ICV, see [Table 3.1](#)
- `num_threads` Clause, see [Section 18.1.2](#)
- `omp_set_num_threads` Routine, see [Section 27.1](#)
- `parallel` Construct, see [Section 18.1](#)

4.1.4 OMP_THREAD_LIMIT

The `OMP_THREAD_LIMIT` environment variable sets the number of threads to use for a contention group by setting the *thread-limit-var* ICV. The value of this environment variable must be a positive integer or a numeric abstract name. The behavior of the program is implementation defined if the requested value of `OMP_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

- *thread-limit-var* ICV, see [Table 3.1](#)

4.1.5 OMP_MAX_ACTIVE_LEVELS

The `OMP_MAX_ACTIVE_LEVELS` environment variable controls the maximum number of nested active `parallel` regions by setting the initial value of the *max-active-levels-var* ICV. The value of this environment variable must be a non-negative integer. The behavior of the program is implementation defined if the requested value of `OMP_MAX_ACTIVE_LEVELS` is greater than the maximum number of active levels an implementation can support, or if the value is not a non-negative integer.

Cross References

- *max-active-levels-var* ICV, see [Table 3.1](#)

4.1.6 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 a **place-list 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 $\langle \text{lower-bound} \rangle : \langle \text{length} \rangle : \langle \text{stride} \rangle$ notation to represent the following list of numbers: “ $\langle \text{lower-bound} \rangle$, $\langle \text{lower-bound} \rangle + \langle \text{stride} \rangle$, ..., $\langle \text{lower-bound} \rangle + (\langle \text{length} \rangle - 1) * \langle \text{stride} \rangle$.” When $\langle \text{stride} \rangle$ 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 **place-list abstract names** listed in Table 4.1 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* ($\langle \text{len} \rangle$) or *abstract_name* ($\langle \text{len} \rangle : \langle \text{stride} \rangle$) where *abstract_name* is a **place-list abstract name** listed in Table 4.1, *len* denotes the length of the **place list** and *stride* denotes the increment between consecutive **places** in the **place list**. When requesting fewer **places** than available on the system, the determination of which resources of type *abstract_name* are to be included in the **place list** is **implementation defined**. When requesting more resources than available, the length of the **place list** is **implementation defined**.

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 a **place-list abstract name**.

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

$$\begin{aligned}
\langle \text{list} \rangle &\models \langle \text{p-list} \rangle \mid \langle \text{aname} \rangle \\
\langle \text{p-list} \rangle &\models \langle \text{p-interval} \rangle \mid \langle \text{p-list} \rangle, \langle \text{p-interval} \rangle \\
\langle \text{p-interval} \rangle &\models \langle \text{place} \rangle : \langle \text{len} \rangle : \langle \text{stride} \rangle \mid \langle \text{place} \rangle : \langle \text{len} \rangle \mid \langle \text{place} \rangle \mid !\langle \text{place} \rangle \\
\langle \text{place} \rangle &\models \{ \langle \text{res-list} \rangle \} \mid \langle \text{res} \rangle \\
\langle \text{res-list} \rangle &\models \langle \text{res-interval} \rangle \mid \langle \text{res-list} \rangle, \langle \text{res-interval} \rangle \\
\langle \text{res-interval} \rangle &\models \langle \text{res} \rangle : \langle \text{len} \rangle : \langle \text{stride} \rangle \mid \langle \text{res} \rangle : \langle \text{len} \rangle \mid \langle \text{res} \rangle \mid !\langle \text{res} \rangle \\
\langle \text{aname} \rangle &\models \langle \text{word} \rangle (\langle \text{len} \rangle : \langle \text{stride} \rangle) \mid \langle \text{word} \rangle (\langle \text{len} \rangle) \mid \langle \text{word} \rangle \\
\langle \text{word} \rangle &\models \text{sockets} \mid \text{cores} \mid \text{ll_caches} \mid \text{numa_domains} \\
&\quad \mid \text{threads} \mid \text{<implementation-defined abstract name>} \\
\langle \text{res} \rangle &\models \textit{non-negative integer} \\
\langle \text{len} \rangle &\models \textit{positive integer} \\
\langle \text{stride} \rangle &\models \textit{integer}
\end{aligned}$$

Examples:

```

export OMP_PLACES=threads
export OMP_PLACES="threads (4) "
export OMP_PLACES="threads (8:2) "
export OMP_PLACES
    ="{0,1,2,3},{4,5,6,7},{8,9,10,11},{12,13,14,15}"
export OMP_PLACES="{0:4},{4:4},{8:4},{12:4}"
export OMP_PLACES="{0:4}:4:4"

```

where each of the last three definitions corresponds to the same four **places** including the smallest units of execution exposed by the execution environment numbered, in turn, 0 to 3, 4 to 7, 8 to 11, and 12 to 15.

Cross References

- *place-partition-var* ICV, see [Table 3.1](#)

4.1.7 OMP_PROC_BIND

The **OMP_PROC_BIND** environment variable sets the initial value of the *bind-var* ICV. The value of this environment variable is either **true**, **false**, or a comma separated list of **primary**, **close**, or **spread**. The values of the list set the **thread affinity** policy to be used for **parallel regions** at the corresponding nested level. The first value also sets the **thread affinity** policy to be used for **implicit parallel regions**.

If the environment variable is set to **false**, the execution environment may move **OpenMP threads** between OpenMP **places**, **thread affinity** is disabled, and **proc_bind** clauses on **parallel**

constructs are ignored.

Otherwise, the execution environment should not move *team-worker threads* between *places*, *thread affinity* is enabled, and the *initial thread* is bound to the first *place* in the *place-partition-var* ICV prior to the first *active parallel region*, or immediately after encountering the first *task-generating construct*. An *initial thread* that is created by a *teams construct* is bound to the first *place* in its *place-partition-var* ICV before it begins execution of the associated *structured block*. A *free-agent thread* that executes a *task* bound to a *team* is assigned a *place* according to the rules described in Section 18.1.3.

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

The **OMP_PROC_BIND** *environment variable* sets the *max-active-levels-var* ICV to the number of *active levels* of parallelism that the implementation supports if the **OMP_PROC_BIND** *environment variable* is set to a comma-separated list of more than one element. The value of the *max-active-levels-var* ICV may be overridden by setting **OMP_MAX_ACTIVE_LEVELS**. See Section 4.1.5 for details.

Examples:

```
export OMP_PROC_BIND=false
export OMP_PROC_BIND="spread, spread, close"
```

Cross References

- **OMP_MAX_ACTIVE_LEVELS**, see Section 4.1.5
- Controlling OpenMP Thread Affinity, see Section 18.1.3
- *bind-var* ICV, see Table 3.1
- *max-active-levels-var* ICV, see Table 3.1
- *place-partition-var* ICV, see Table 3.1
- **omp_get_proc_bind** Routine, see Section 35.1
- **parallel** Construct, see Section 18.1
- **proc_bind** Clause, see Section 18.1.4
- **teams** Construct, see Section 18.2

4.2 Teams Environment Variables

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

4.2.1 OMP_NUM_TEAMS

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

Cross References

- *ntteams-var* ICV, see Table 3.1
- **teams** Construct, see Section 18.2

4.2.2 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 or a numeric abstract name. 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 neither a positive integer nor one of the allowed abstract names.

Cross References

- *teams-thread-limit-var* ICV, see Table 3.1
- **teams** Construct, see Section 18.2

4.3 Program Execution Environment Variables

This section defines environment variables that affect program execution.

4.3.1 OMP_SCHEDULE

The **OMP_SCHEDULE** environment variable controls the schedule type and chunk size of all worksharing-loop constructs that have the schedule type **runtime**, by setting the value of the *run-sched-var* ICV. The value of this environment variable takes the form [*modifier*:]*kind*[, *chunk*], where:

- *modifier* is one of **monotonic** or **nonmonotonic**;
- *kind* specifies the schedule type and is one of **static**, **dynamic**, **guided**, or **auto**;
- *chunk* is an optional positive integer that specifies the chunk size.

If the *modifier* is not present, the *modifier* is set to **monotonic** if *kind* is **static**; for any other *kind* it is set to **nonmonotonic**.

If *chunk* is present, **white space** may be on either side of the “,”.

The behavior of the program is **implementation defined** if the value of **OMP_SCHEDULE** does not conform to the above format.

Examples:

```
export OMP_SCHEDULE="guided, 4"  
export OMP_SCHEDULE="dynamic"  
export OMP_SCHEDULE="nonmonotonic:dynamic, 4"
```

Cross References

- *run-sched-var* ICV, see [Table 3.1](#)
- **schedule** Clause, see [Section 19.6.3](#)

4.3.2 OMP_STACKSIZE

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

- *size* is a **positive** integer that specifies the size of the stack for **threads**.
- *unit* is **B**, **K**, **M**, or **G** and specifies whether the given size is in Bytes, Kilobytes (1024 Bytes), Megabytes (1024 Kilobytes), or Gigabytes (1024 Megabytes), respectively. If *unit* is present, **white space** may occur between *size* and it, whereas if *unit* is not present then **K** is assumed.

The behavior of the program is **implementation defined** if **OMP_STACKSIZE** does not conform to the above format, or if the implementation cannot provide a stack with the requested size.

Examples:

```
export OMP_STACKSIZE=2000500B  
export OMP_STACKSIZE="3000 k "  
export OMP_STACKSIZE=10M  
export OMP_STACKSIZE=" 10 M "  
export OMP_STACKSIZE="20 m "  
export OMP_STACKSIZE=" 1G"  
export OMP_STACKSIZE=20000
```

Cross References

- *stacksize-var* ICV, see [Table 3.1](#)

4.3.3 OMP_WAIT_POLICY

The **OMP_WAIT_POLICY** environment variable provides a hint to an OpenMP implementation about the desired behavior of waiting **native threads** by setting the *wait-policy-var* ICV. A **compliant implementation** may or may not abide by the setting of the environment variable. The value of this environment variable must be one of the following:

active | **passive**

The **active** value specifies that waiting **native threads** should mostly be active, consuming processor cycles, while waiting. A **compliant implementation** may, for example, make waiting **native threads** spin. The **passive** value specifies that waiting **native threads** should mostly be passive, not consuming processor cycles, while waiting. For example, a **compliant implementation** may make waiting **native threads** yield the processor to other **native threads** or go to sleep. The details of the **active** and **passive** behaviors are **implementation defined**. The behavior of the program is **implementation defined** if the value of **OMP_WAIT_POLICY** is neither **active** nor **passive**.

Examples:

```
export OMP_WAIT_POLICY=ACTIVE
export OMP_WAIT_POLICY=active
export OMP_WAIT_POLICY=PASSIVE
export OMP_WAIT_POLICY=passive
```

Cross References

- *wait-policy-var* ICV, see [Table 3.1](#)

4.3.4 OMP_DISPLAY_AFFINITY

The **OMP_DISPLAY_AFFINITY** environment variable sets the *display-affinity-var* ICV so that the runtime displays formatted affinity information for the **host device**. Affinity information is printed for all **OpenMP threads** in each **parallel region** upon first entering it. Also, if the information accessible by the format specifiers listed in [Table 4.2](#) changes for any **thread** in the **parallel region** then **thread affinity** information for all **threads** in that **region** is again displayed. If the **thread affinity** for each respective **parallel region** at each nesting level has already been displayed and the **thread affinity** has not changed, then the information is not displayed again. **Thread affinity** information for **threads** in the same **parallel region** may be displayed in any order. The value of the **OMP_DISPLAY_AFFINITY** environment variable may be set to one of these values:

true | **false**

The **true** value instructs the runtime to display the **thread affinity** information, and uses the format setting defined in the *affinity-format-var* ICV. The runtime does not display the **thread affinity** information when the value of the **OMP_DISPLAY_AFFINITY** environment variable is **false** or

undefined. For all values of the [environment variable](#) other than **true** or **false**, the display action is [implementation defined](#).

Example:

```
export OMP_DISPLAY_AFFINITY=TRUE
```

For this example, an OpenMP implementation displays [thread affinity](#) information during program execution, in a format given by the [affinity-format-var](#) ICV. The following is a sample output:

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

Cross References

- **OMP_AFFINITY_FORMAT**, see [Section 4.3.5](#)
- Controlling OpenMP Thread Affinity, see [Section 18.1.3](#)
- [affinity-format-var](#) ICV, see [Table 3.1](#)
- [display-affinity-var](#) ICV, see [Table 3.1](#)

4.3.5 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 4.2: Available Field Types for Formatting OpenMP Thread Affinity Information

Short Name	Long Name	Meaning
t	team_num	The value returned by <code>omp_get_team_num</code>
T	num_teams	The value returned by <code>omp_get_num_teams</code>
L	nesting_level	The value returned by <code>omp_get_level</code>
n	thread_num	The value returned by <code>omp_get_thread_num</code>
N	num_threads	The value returned by <code>omp_get_num_threads</code>
a	ancestor_tnum	The value returned by <code>omp_get_ancestor_thread_num</code> with an argument of one less than the value returned by <code>omp_get_level</code>
H	host	The name for the host device on which the OpenMP program is running
P	process_id	The process identifier used by the implementation
i	native_thread_id	The native thread identifier used by the implementation
A	thread_affinity	The list of numerical identifiers, in the format of a comma-separated list of integers or integer ranges, that represent processors on which a thread may execute, subject to OpenMP thread affinity control and/or other external affinity mechanisms

Implementations may define additional field types. If an implementation does not have information for a field type or an unknown field type is part of a field specifier, "undefined" is printed for this field when displaying [thread affinity](#) information.

Example:

```
export OMP_AFFINITY_FORMAT=\
"Thread Affinity: %0.3L %.8n %.15{thread_affinity} %.12H"
```

The above example causes an OpenMP implementation to display [thread affinity](#) information in the following form:

Thread Affinity: 001	0	0-1, 16-17	nid003
Thread Affinity: 001	1	2-3, 18-19	nid003

Cross References

- Controlling OpenMP Thread Affinity, see [Section 18.1.3](#)
- affinity-format-var* ICV, see [Table 3.1](#)

- `omp_get_ancestor_thread_num` Routine, see [Section 27.15](#)
- `omp_get_level` Routine, see [Section 27.14](#)
- `omp_get_num_teams` Routine, see [Section 28.1](#)
- `omp_get_num_threads` Routine, see [Section 27.2](#)
- `omp_get_thread_num` Routine, see [Section 27.3](#)

4.3.6 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 `cancel` constructs 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` Construct, see [Section 24.2](#)
- *cancel-var* ICV, see [Table 3.1](#)

4.3.7 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 devices* and the `omp_get_num_devices` routine returns the number of *devices* stored in the ICV.

The value of this environment variable set to `none` or an empty string implies no *non-host devices* are available; otherwise, it must be a comma-separated list. Each item is either a *trait* specification as specified in the following or `*`. A `*` expands to all *non-host accessible devices* that are *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*), **vendor**(*vendor-name*), and **uid**(*uid-string*), where the names are as specified in [Section 15.1](#) 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. Parentheses can be used for grouping, but are optional except that **&&** and **| |** may not appear in the same grouping level. The unary **!** operator inverts the meaning of the immediately following **trait** or parenthesized group.

Each **trait** specification or ***** yields a (possibly zero-sized) array of **non-host devices** with the lowest array element, if it exists, having index zero. The C/C++ syntax **[index]** can be used to select an element and the [array section](#) syntax for C/C++ as specified in [Section 5.2.5](#) can be used to specify a subset of elements. Any array element specified by the subscript that is outside the bounds of the array resulting from the **trait** specification or ***** is silently excluded.

Example:

Four GPUs are **accessible** and **supported**, with unique identifiers represented as **<uid-gpu0>**, ..., **<uid-gpu3>**.

```
export OMP_AVAILABLE_DEVICES="kind(gpu) "  
export OMP_AVAILABLE_DEVICES="uid(<uid-gpu0>), kind(gpu) "  
export OMP_AVAILABLE_DEVICES="uid(<uid-gpu1>), kind(gpu) [:2] "
```

where the above **OMP_AVAILABLE_DEVICES** assignments select:

- All GPUs;
- All GPUs with **device** **<uid-gpu0>** assigned **device number** 0; and
- **Device** **<uid-gpu1>**, which is assigned **device number** 0, and two other GPUs.

Cross References

- Device Directives and Clauses, see [Chapter 21](#)
- *available-devices-var* ICV, see [Table 3.1](#)

4.3.8 OMP_DEFAULT_DEVICE

The **OMP_DEFAULT_DEVICE** environment variable sets the initial value of the *default-device-var* ICV. The value of this environment variable must be a comma-separated list, each item being either a **non-negative** integer value that denotes the **device number**, a **trait** specification with an optional subscript selector, or one of the following case-insensitive **string literals**: **initial** to specify the

host device, **invalid** to specify the device number **omp_invalid_device**, or **default** to set the ICV as if this environment variable was not specified (see Section 1.2).

The **trait** specification is as described for **OMP_AVAILABLE_DEVICES** (see Section 4.3.7), except that in addition the **trait device_num(device number)** may be specified and **host** is permitted as *kind-name*. The device numbers yielded by the **trait** specification are sorted in ascending order by device number and form a set; the array-element syntax as described for **OMP_AVAILABLE_DEVICES** can be used to select an element from this set. If an item is an empty set, non-existing element, or does not evaluate to an available device, the next item is evaluated; otherwise, the *default-device-var* ICV is set to the first value of the set. However, **initial**, **invalid**, and **default** always match. If none of the list items match, the *default-device-var* ICV is set to **omp_invalid_device**.

Example:

Four GPUs are accessible and supported, with unique identifiers represented as <uid-gpu0>, ..., <uid-gpu3>. The default device is set to device <uid-gpu0>.

```
export OMP_DEFAULT_DEVICE="uid(<uid-gpu0>)"
```

Cross References

- Device Directives and Clauses, see Chapter 21
- *default-device-var* ICV, see Table 3.1

4.3.9 OMP_TARGET_OFFLOAD

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

mandatory | **disabled** | **default**

The **mandatory** value specifies that the effect of any device construct or device routine that uses a device that is not an available device or a supported device, 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 1.2.

Example:

```
export OMP_TARGET_OFFLOAD=mandatory
```

Cross References

- Device Directives and Clauses, see Chapter 21
- Device Memory Routines, see Chapter 31
- *target-offload-var* ICV, see Table 3.1

4.3.10 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.

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 4.3. 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 4.3: Reservation Types for `OMP_THREADS_RESERVE`

Reservation Type	Meaning	Default Value
<code>structured</code>	Threads reserved for structured threads	1
<code>free_agent</code>	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 4.1.

Algorithm 4.1 Initial `structured-thread-limit-var` and `free-agent-thread-limit-var` ICVs Values

```

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

```

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

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

Examples:

```
export OMP_THREADS_RESERVE=4
export OMP_THREADS_RESERVE="structured(4) "
export OMP_THREADS_RESERVE="structured"
export OMP_THREADS_RESERVE="structured(2), free_agent(2) "
```

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

Cross References

- `free-agent-thread-limit-var` ICV, see [Table 3.1](#)
- `structured-thread-limit-var` ICV, see [Table 3.1](#)
- `parallel` Construct, see [Section 18.1](#)
- `threadset` Clause, see [Section 20.8](#)

4.3.11 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:

```
export OMP_MAX_TASK_PRIORITY=20
```

Cross References

- `max-task-priority-var` ICV, see [Table 3.1](#)

4.4 Memory Allocation Environment Variables

This section defines [environment variables](#) that affect [memory](#) allocations.

4.4.1 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](#).

$$\begin{aligned}\langle \text{allocator} \rangle & \models \langle \text{predef-allocator} \rangle \mid \langle \text{predef-mem-space} \rangle \mid \langle \text{predef-mem-space} \rangle : \langle \text{traits} \rangle \\ \langle \text{traits} \rangle & \models \langle \text{trait} \rangle = \langle \text{value} \rangle \mid \langle \text{trait} \rangle = \langle \text{value} \rangle, \langle \text{traits} \rangle \\ \langle \text{predef-allocator} \rangle & \models \text{one of the predefined } \textit{allocators} \text{ from Table 13.3} \\ \langle \text{predef-mem-space} \rangle & \models \text{one of the predefined } \textit{memory spaces} \text{ from Table 13.1} \\ \langle \text{trait} \rangle & \models \text{one of the } \textit{allocator trait names} \text{ from Table 13.2} \\ \langle \text{value} \rangle & \models \text{one of the allowed values from Table 13.2} \mid \text{non-negative integer} \\ & \mid \langle \text{predef-allocator} \rangle\end{aligned}$$

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:

```
export OMP_ALLOCATOR=omp_high_bw_mem_alloc
export OMP_ALLOCATOR="omp_large_cap_mem_space:alignment=16,\
pinned=true"
export 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 13.2](#)
- *def-allocator-var* ICV, see [Table 3.1](#)

4.5 OMPT Environment Variables

This section defines [environment variables](#) that affect operation of the [OMPT tool](#) interface.

4.5.1 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:

```
export OMP_TOOL=enabled
```

Cross References

- OMPT Overview, see [Chapter 38](#)
- [tool-var](#) ICV, see [Table 3.1](#)

4.5.2 OMP_TOOL_LIBRARIES

The [OMP_TOOL_LIBRARIES environment variable](#) sets the [tool-libraries-var ICV](#) to a list of [tool](#) libraries that are considered for use on a [device](#) on which an OpenMP implementation is being initialized. The value of this [environment variable](#) must be a list of names of dynamically-loadable libraries, separated by an implementation specific, platform typical separator. Whether the value of this [environment variable](#) is case sensitive is [implementation defined](#).

If the [tool-var ICV](#) is not [enabled](#), the value of [tool-libraries-var](#) is ignored. Otherwise, if [ompt_start_tool](#) is not visible in the [address space](#) on a [device](#) where OpenMP is being initialized or if [ompt_start_tool](#) returns [NULL](#), an OpenMP implementation will consider libraries in the [tool-libraries-var](#) list in a left-to-right order. The OpenMP implementation will search the list for a library that meets two criteria: it can be dynamically loaded on the [current device](#) and it defines the symbol [ompt_start_tool](#). If an OpenMP implementation finds a suitable library, no further libraries in the list will be considered.

Example:

```
export OMP_TOOL_LIBRARIES=libtoolXY64.so:/usr/local/lib/  
libtoolXY32.so
```

Cross References

- OMPT Overview, see [Chapter 38](#)
- *tool-libraries-var* ICV, see [Table 3.1](#)
- `ompt_start_tool` Procedure, see [Section 38.2.1](#)

4.5.3 OMP_TOOL_VERBOSE_INIT

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

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

If `OMP_TOOL_VERBOSE_INIT` is set to any value other than case insensitive `disabled`, `stdout`, or `stderr`, the value is interpreted as a filename and the OpenMP runtime will try to log to a file with prefix *filename*. If the value is interpreted as a filename, whether it is case sensitive is *implementation defined*. If opening the logfile fails, the output will be redirected to `stderr`. If `OMP_TOOL_VERBOSE_INIT` is not *defined*, the default value for *tool-verbose-init-var* is `disabled`. Support for logging to `stdout` or `stderr` is *implementation defined*. Unless *tool-verbose-init-var* is `disabled`, the OpenMP runtime will log the steps of the *tool* activation process defined in [Section 38.2.2](#) to a file with a name that is constructed using the provided filename prefix. The format and detail of the log is *implementation defined*. At a minimum, the log will contain one of the following:

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

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

Example:

```
export OMP_TOOL_VERBOSE_INIT=disabled
export OMP_TOOL_VERBOSE_INIT=STDERR
export OMP_TOOL_VERBOSE_INIT=ompt_load.log
```

Cross References

- OMPT Overview, see [Chapter 38](#)
- *tool-verbose-init-var* ICV, see [Table 3.1](#)

4.6 OMPD Environment Variables

This section defines [environment variables](#) that affect operation of the [OMPD tool](#) interface.

4.6.1 OMP_DEBUG

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

[enabled](#) | [disabled](#)

If [OMP_DEBUG](#) is set to any value other than [enabled](#) or [disabled](#) then the behavior is [implementation defined](#).

Example:

```
export OMP_DEBUG=enabled
```

Cross References

- Enabling Runtime Support for OMPD, see [Section 44.3.1](#)
- OMPD Overview, see [Chapter 44](#)
- *debug-var* ICV, see [Table 3.1](#)

4.7 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 36.4](#)). The value of the [OMP_DISPLAY_ENV](#) environment variable may be set to one of these values:

[true](#) | [false](#) | [verbose](#)

If the [environment variable](#) is set to [true](#), the effect is as if the [omp_display_env](#) routine is called with the *verbose* argument set to *false* at the beginning of the program. If the [environment variable](#) is set to [verbose](#), the effect is as if the [omp_display_env](#) routine is called with the *verbose* argument set to *true* at the beginning of the program. If the [environment variable](#) is [undefined](#) or set to [false](#), the runtime does not display any information. For all values of the [environment variable](#) other than [true](#), [false](#), and [verbose](#), the displayed information is unspecified.

Example:

```
export OMP_DISPLAY_ENV=true
```

For the output of the above example, see [Section 36.4](#).

Cross References

- [omp_display_env](#) Routine, see [Section 36.4](#)

5 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

Restrictions on **OpenMP programs** include:

- 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**.

C / C++

- The use of **omp** as the first preprocessing token of a pragma **directive** must be for OpenMP **directives** that are defined in this specification; OpenMP reserves these uses for OpenMP **directives**.
- The use of **omp** as the attribute namespace of an attribute specifier, or as the optional namespace qualifier within a **sequence** attribute, must be for OpenMP **directives** that are defined in this specification; OpenMP reserves these uses for such **directives**.
- The use of **omp_x** as the first preprocessing token of a pragma **directive** must be for **implementation defined** extensions to the OpenMP **directives**; OpenMP reserves these uses for such extensions.
- The use of **omp_x** as the attribute namespace of an attribute specifier, or as the optional namespace qualifier within a **sequence** attribute, must be for **implementation defined** extensions to the OpenMP **directives**; OpenMP reserves these uses for such extensions.

C / C++

Fortran

- In free form source files, the **!\$omp** sentinel must be used for OpenMP **directives** that are defined in this specification; OpenMP reserves these uses for such **directives**.
- In fixed form source files, sentinels that end with **omp** must be used for OpenMP **directives** that are defined in this specification; OpenMP reserves these uses for such **directives**.
- In free form source files, the **!\$omp_{px}** sentinel must be used for **implementation defined** extensions to the OpenMP **directives**; OpenMP reserves these uses for such extensions.
- In fixed form source files, sentinels that end with **om_{px}** must be used for **implementation defined** extensions to the OpenMP **directives**; OpenMP reserves these uses for such extensions.

Fortran

- A **clause** name must be the name of a **clause** that is defined in this specification except for those that begin with **omp_{px}_**, which may be used for **implementation defined** extensions and which OpenMP reserves for such extensions.
- OpenMP reserves names that begin with the **omp_**, **ompt_** and **ompd_** prefixes for names defined in this specification so **OpenMP programs** must not declare names that begin with them.
- OpenMP reserves names that begin with the **omp_{px}_** prefix for **implementation defined** extensions so **OpenMP programs** must not declare names that begin with it.

C++

- **OpenMP programs** must not declare a namespace with the **omp**, **omp_{px}**, **ompt** or **ompd** names, as these are reserved for the OpenMP implementation.

C++

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++

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

5.1 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**. Thus, the *directive-specification* is:

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

where the *directive-specifier* is:

```
directive-name
```

or for argument-modified directives:

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

where *directive-name* is the **directive name** of the **directive**.

Some **directives** specify a paired **end directive**. If the *directive-name* of such a **directive** starts with **begin**, the **end directive** has the same **directive name** except **begin** is replaced with **end**. If the *directive-name* does not start with **begin**, unless otherwise specified the **directive name** of the **end directive** is **end directive-name**.

Some **directives** have underscores in their *directive-name*. Some of those **directives** are explicitly specified alternatively to allow the underscores in their *directive-name* to be replaced with **white space**. In addition, if a *directive-name* starts with either **begin** or **end** then it is separated from the rest of the *directive-name* by **white space**.

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")
```

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.

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. The **directive** attribute must not be specified inside a **sequence** attribute unless it specifies a **block-associated directive**.

Note – This example shows the expected transformation:

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

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

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

C / C++

C++

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

C++

Fortran

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

sentinel directive-specification

All **directives** must begin with a **directive sentinel**. The format of a sentinel differs between fixed form and free form source files, as described in [Section 5.1.1](#) and [Section 5.1.2](#). In order to simplify the presentation, free form is used for the syntax of **directives** for Fortran throughout this document, except as noted.

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

Fortran

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

- **declarative directive**;
- **executable directive**;
- **informational directive**;

- `metadirective`;
- `subsidiary directive`; or
- `utility directive`.

Base-language code can be associated with `directives`. A `directive` may be categorized by its base-language code association as one of the following:

- `block-associated directive`;
- `declaration-associated directive`;
- `delimited directive`;
- `explicitly associated directive`;
- `loop-nest-associated directive`;
- `loop-sequence-associated directive`;
- `separating directive`; or
- `unassociated directive`.

A `directive` and its associated `base-language code` (if any) constitute a syntactic formation that follows the syntax given below unless otherwise specified. The *end-directive* in a specified formation refers to the paired `end directive` for the `directive`. A `construct` is a formation for an `executable directive`. An `end directive` is considered a `subsidiary directive` of a `construct` if it is the `end directive` of that `construct`.

`Unassociated directives` are not directly associated with any `base-language code`. The resulting formation therefore has the following syntax:

`directive`

`Unassociated directives` that are `declarative directives` declare identifiers for use in other `directives`. `Unassociated directives` that are `executable directives` are `stand-alone directives`.

`Explicitly associated directives` are `declarative directives` that take a `variable` or extended list as a `directive` or `clause` argument that indicates the declarations with which the `directive` is associated. As a result, `explicitly associated directives` have the same syntax as the formation for `unassociated directives`.

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

`directive`
`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 directive** can be a **DO CONCURRENT** loop where it is explicitly allowed.

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

Fortran

A **declaration-associated directive** is directly associated with a **base language** declaration.

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 an **unassociated directive** as the associated declaration is determined directly from the specification part in which the **directive** appears.

Fortran

Fortran / C++

If a **directive** appears in the specification part of a module then the behavior is as if that **directive**, with the **variables**, types and **procedures** that have **PRIVATE** accessibility omitted, appears in the specification part of any **compilation unit** that references the module unless otherwise specified.

Fortran / C++

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 the associated **structured block** of a **block-associated directive** (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 **affected 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** is an **unassociated directive**, 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 an **executable directive**. For a **block-associated directive**, the resulting formation consists of the attribute specifier and a **structured block** to which the specifier applies. If the **directives** are **separating directives** or **delimited directives** then the resulting formation is as specified above for those associations except that the attribute specifier for each **directive**, including the **end directive**, applies to a null statement.

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

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

C++

or

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

C++

An **explicitly associated directive** may alternatively be expressed with an attribute specifier that also uses the **decl** attribute, applies to a **variable** and/or function declaration, and omits the **variable** list or extended list argument. The effect is as if the omitted list argument is the list of declared **variables** and/or functions to which the attribute specifier applies.

Formations that result from **directives** that are specified as attribute specifiers and are **declaration-associated directives** or use the **decl** attribute are specified as follows. If the **directives** are **declaration-associated directives** 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:

C / C++

- A *directive-name* must not include **white space** except where explicitly allowed.

C / C++

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

Fortran

- 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.
- The **directive** and **subsidiary directives** of a **construct** must all use the attribute syntax or must all use 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**.



5.1.1 Free Source Form Directives

The following sentinels are recognized in free form source files:

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

The sentinel can appear in any column as long as it is preceded only by **white space**. It must appear as a single word with no intervening **white space**. Fortran free form line length and **white space** rules apply to the **directive** line. The syntax that allows **white space** to be optional has been **deprecated**. 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.



5.1.2 Fixed Source Form Directives

The following sentinels are recognized in fixed form source files:

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

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. The syntax that allows **white space** to be optional has been **deprecated**. Initial **directive** lines must have a space or a zero in column 6, and continuation **directive** lines must have a character other than a space or a zero in column 6.

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



5.2 Clause Format

This section defines the format and categories of OpenMP [clauses](#). [Clauses](#) are specified as part of a *directive-specification*. [Clauses](#) have the [optional property](#) and, thus, may be omitted from a *directive-specification* unless otherwise specified, in which case they have the [required property](#). 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 group](#). 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 `ompx` prefix for use on any OpenMP [directive](#), and the format and semantics of any such [clause](#) is [implementation defined](#).

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
```

Unless otherwise specified, modified arguments have the [pre-modified property](#), in which case the format is:

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

Some modified arguments are explicitly specified to have the [post-modified property](#), in which case the format is:

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

For many [clauses](#), *clause-argument-list* is an OpenMP [argument list](#), which is a comma-separated [list](#) of a specific kind of [list items](#) (see [Section 5.2.1](#)), in which case the format of *clause-argument-list* is:

```
argument-name
```

For all other [clauses](#), *clause-argument-list* is a comma-separated [list](#) of arguments so the format is:

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

In most of these cases, the [list](#) only has a single item so the format of *clause-argument-list* is again:

```
argument-name
```

In all cases, [white space](#) in *clause-argument-list* is [optional](#).

A *modifier-specification-list* is a comma-separated [list](#) of [clause](#) argument [modifiers](#) for which the format is:

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

[Clause](#) argument [modifiers](#) may be [simple modifiers](#) or [complex modifier](#). Many [clause](#) argument [modifiers](#) are [simple modifiers](#), for which the format of *modifier-specification* is:

```
modifier-name
```

The format of a [complex modifier](#) is:

```
modifier-name[ (modifier-parameter-specification) ]
```

where *modifier-parameter-specification* is a comma-separated [list](#) of arguments as defined above for *clause-argument-list*. The position of each *modifier-argument-name* in the [list](#) is significant. The *modifier-parameter-specification* and parentheses are required unless every *modifier-argument-name* is [optional](#) and omitted, in which case the format of the [complex modifier](#) is identical to that of a [simple modifier](#):

```
modifier-name
```

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

- *keyword*: An OpenMP keyword;
- *OpenMP identifier*: An [OpenMP identifier](#);
- *OpenMP argument list*: An OpenMP [argument list](#);
- *expression*: An expression of some [OpenMP type](#); or
- *OpenMP stylized expression*: An [OpenMP stylized expression](#).

A particular lexical instantiation of an argument specifies a parameter of the [clause](#), while a lexical instantiation of a [modifier](#) and its parameters affects how or when the argument is applied.

The order of arguments must match the order in the *clause-specification* or *modifier-specification*. The order of [modifiers](#) in a *clause-argument-specification* is not significant unless otherwise specified.

General syntactic [properties](#) govern the use of [clauses](#), [clause](#) and [directive](#) arguments, and [modifiers](#) in a [directive](#). These [properties](#) are summarized in [Table 5.1](#), along with the respective default [properties](#) for [clauses](#), arguments and [modifiers](#).

TABLE 5.1: 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 on a clause with the post- modified property)	free	free	free	free

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

Arguments of **directives**, **clauses** and **modifiers** are never **repeatable**. Instead, **argument lists** are used whenever the corresponding semantics may be specified for multiple **list items** that serve as the arguments of the **directives**, **clauses** or **modifiers**.

Some **clause properties** determine the **constituent directives** to which they apply when specified on **compound directives**. A **clause** with the **all-constituents property** applies to all **constituent directives** of any **compound directive** on which it is specified. Unless otherwise specified, a **clause** has the **all-constituents property**. That is, the **all-constituents property** is a default **clause property**. A **clause** with the **once-for-all-constituents property** applies to the **directive** once, before any of the **constituent directives** are applied. A **clause** with the **innermost-leaf property** applies to the innermost **constituent directive** to which it may be applied. A **clause** with the **outermost-leaf property** applies to the outermost **constituent directive** to which it may be applied. A **clause** with the **all-privatizing property** applies to all **constituent directives** that permit the **clause** and to which a **data-sharing attribute clause** that may create a **private** copy of the same **list item** is applied.

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

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

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

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

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

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

Restrictions

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

- A **clause** with the **required property** for a **directive** must appear on the **directive**.
- A **clause** with the **unique property** for a **directive** may appear at most once on the **directive**.
- A **clause** with the **exclusive property** for a **directive** must not appear if a **clause** with a different *clause-name* also appears on the **directive**.
- An **ultimate clause**, that is one that has the **ultimate property** for a **directive**, must be the lexically last **clause** to appear on the **directive**.
- If a **clause set** has the **required property**, at least one **clause** in the set must be present on the **directive** for which the **clause set** is specified.
- If a **clause** is a member of a **clause set** that has the **unique property** for a **directive** then the **clause** has the **unique property** for that **directive** regardless of whether it has the **unique property** when it is not part of such a **clause set**.
- If one **clause** of a **clause set** with the **exclusive property** appears on a **directive**, no other **clauses** with a different *clause-name* in that **clause set** may appear on the **directive**.
- An argument with the **required property** must appear in the *clause-specification*, unless otherwise specified.
- An argument with the **unique property** may appear at most once in a *clause-argument-specification*.
- An argument with the **exclusive property** must not appear if an argument with a different *argument-name* appears in the *clause-argument-specification*.
- A **modifier** with the **required property** must appear in the *clause-argument-specification*.
- A **modifier** with the **unique property** may appear at most once in a *clause-argument-specification*.

- A **modifier** with the **exclusive property** must not appear if a **modifier** with a different *modifier-name* also appears in the *clause-argument-specification*.
- If a **clause** has the **pre-modified property**, a **modifier** with the **ultimate property** must be the last **modifier** in any *clause-argument-specification* in which any **modifier** appears.
- If a **clause** has the **post-modified property**, a **modifier** with the **ultimate property** must be the first **modifier** in any *clause-argument-specification* in which any **modifier** appears.
- A **modifier** that is an expression must neither lexically match the name of a **simple modifier** defined for the **clause** that is an OpenMP keyword nor *modifier-name parenthesized-tokens*, where *modifier-name* is the *modifier-name* of a **complex modifier** defined for the **clause** and *parenthesized-tokens* is a token sequence that starts with **(** and ends with **)**.
- An argument or parameter with the **constant property** must be a compile-time constant.
- An argument or parameter with the **positive property** must be greater than zero.
- An argument or parameter with the **non-negative property** must be greater than or equal to zero.
- An argument or parameter with the **region-invariant property** must have the same value throughout any given execution of the **construct** or, for **declarative directives**, execution of the **procedure** with which the declaration is associated.

Cross References

- Directive Format, see [Section 5.1](#)
- OpenMP Argument Lists, see [Section 5.2.1](#)
- OpenMP Stylized Expressions, see [Section 6.2](#)
- OpenMP Types and Identifiers, see [Section 6.1](#)

5.2.1 OpenMP Argument Lists

The OpenMP API defines several kinds of **lists**, each of which can be used as syntactic instances of **directive**, **clause** and **modifier** arguments. These comma-separated **argument lists** allow the corresponding semantics to apply to multiple **list items**. In any **argument list** the separation of **list items** has precedence for commas over any **base language** semantics for commas. Thus, application of **base language** semantics for commas to any expression in an **argument list** may require the use of parentheses.

A **list** of any **OpenMP type** consists of a comma-separated collection of one or more expressions of that **OpenMP type**. A **parameter list** consists of a comma-separated collection of one or more **parameter list items**. A **variable list** consists of a comma-separated collection of one or more **variable list items**. An **extended list** consists of a comma-separated collection of one or more **extended list items**, each of which is a **variable list item** or the name of a **procedure**. A **locator list**

consists of a comma-separated collection of one or more [locator list items](#). A [type-name list](#) consists of a comma-separated collection of one or more [type-name list items](#). A [directive-name list](#) consists of a comma-separated collection of one or more [directive-name list items](#), each of which is a [directive name](#). A [directive-specification list](#) consists of a comma-separated collection of one or more [directive-specification list items](#), each of which is a [directive specification](#). A [preference specification list](#) consists of a comma-separated collection of one or more [preference specification list items](#), each of which is a [preference specification](#) as defined in [Section 22.1.3](#). An [OpenMP operation list](#) consists of a comma-separated collection of one or more [OpenMP operation list items](#), each of which is a [OpenMP operation](#) defined in [Section 5.2.3](#). An [iterator-specifier list](#) consists of a comma-separated collection of one or more [iterator-specifier list items](#), each of which is an [iterator specifier](#) defined in [Section 5.2.6](#).

A [parameter list item](#) can be one of the following:

- A [named parameter list item](#);
- The position of a parameter in a parameter specification specified by a [positive](#) integer, where 1 represents the first parameter; or
- A parameter range specified by $lb : ub$ where both lb and ub must be an expression of integer [OpenMP type](#) with the [constant property](#) and the [positive property](#).

In both lb and ub , an expression using [omp_num_args](#), that enables identification of parameters relative to the last argument of the call, can be used with the form:

$$\text{omp_num_args} [\pm \text{logical_offset}]$$

where logical_offset is an expression of integer [OpenMP type](#) with the [constant property](#) and the [non-negative property](#). The lb and ub expressions are both [optional](#). If lb is not specified the first element of the range will be 1. If ub is not specified the last element of the range will be [omp_num_args](#). The effect of a specified range of $lb..ub$ is as if the parameters $lb^{th}, (lb + 1)^{th}, \dots, ub^{th}$ had been specified individually.

C / C++

A [named parameter list item](#) is the name of a [function](#) parameter. A [variable list item](#) is a [variable](#) or an [array section](#). A [locator list item](#) is a [reserved locator](#), an [array section](#), or any lvalue expression including [variables](#). A [type-name list item](#) is a type name.

C / C++

Fortran

A [named parameter list item](#) is a dummy argument of a subroutine or function. A [variable list item](#) is one of the following:

- a [variable](#) that is not coindexed and that is not a substring;
- an [array section](#) that is not coindexed and that does not contain an element that is a substring;
- a named constant;
- a [procedure](#) pointer;

- an associate name that may appear in a **variable** definition context; or
- a common block name (enclosed in slashes).

A **locator list item** is a **variable list item**, a function reference with data pointer result, or a **reserved locator**. A **type-name list item** is a type specifier in the base language.

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

Fortran

Restrictions

The restrictions to **argument lists** are as follows:

- All **list items** must be visible, according to the scoping rules of the **base language**.
- Unless otherwise specified, OpenMP **list items** other than **parameter 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**.
- Unless otherwise specified, any given **parameter list item** can only be specified once across all **clauses** of the same type in a given **directive**.
- The *directive-specifier* and the **clauses** in a **directive-specification list 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

- A named constant or a **procedure** pointer can appear as a **list item** only in **clauses** where it is explicitly allowed.
- Unless otherwise specified, a **variable** that is part of an **aggregate variable** must not be a **variable list item** or an **extended list item**.

- Unless otherwise specified, an assumed-type [variable](#) must not be a [variable list item](#), an [extended list item](#), or a [locator list item](#).
- A [type-name list item](#) must not specify an abstract type or be either **CLASS (*)** or **TYPE (*)**.
- Since common block names cannot be accessed by use association or host association, a common block name specified in a [clause](#) must be declared to be a common block in the same scoping unit in which the [clause](#) appears.

Fortran

5.2.2 Reserved Locators

On some [directives](#), some [clauses](#) accept the use of [reserved locators](#) as special [OpenMP identifiers](#) that represent system storage not necessarily bound to any [base language](#) storage item. The [reserved locators](#) are:

```
omp_all_memory
```

The [reserved locator](#) **omp_all_memory** is an [OpenMP identifier](#) that denotes a [list item](#) treated as having storage that corresponds to the storage of all other objects in [memory](#).

Restrictions

Restrictions to the [reserved locators](#) are as follows:

- [Reserved locators](#) may only appear in [clauses](#) and [directives](#) where they are explicitly allowed and may not otherwise be referenced in an [OpenMP program](#).

5.2.3 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 return an object of the *<generic_name>* [OpenMP type](#) (see [Section 6.1](#)). In general, the format of an [OpenMP operation](#) is the following:

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

C / C++

5.2.4 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++

5.2.5 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 : ]
```

```

1      [ : length ]
2      [ : : stride ]
3      [ : : ]
4      [ : ]

```

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](#):

```

23      a[0:6]
24      a[0:6:1]
25      a[1:10]
26      a[1:]
27      a[:10:2]
28      b[10][:][:]
29      b[10][:][:0]
30      c[42][0:6][:]
31      c[42][0:6:2][:]
32      c[1:10][42][0:6]

```

```

1      s.c[:100]
2      p->y[:10]
3      this->a[:N]
4      (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 [array bases](#).

The following are examples that are non-conforming [array sections](#):

```

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

```

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

C / C++

Fortran

Fortran has built-in support for [array sections](#) although some restrictions apply to their use in OpenMP [directives](#), as enumerated at the end of this section.

Fortran

Restrictions

Restrictions to [array sections](#) are as follows:

- An [array section](#) can appear only in [clauses](#) for which it is explicitly allowed.
- A *stride* expression may not be specified unless otherwise stated.

C / C++

- 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 **array base** 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.

C / C++

C++

- If the type of the **array base** 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.

C++

Fortran

- If a stride expression is specified, it must be positive.
- The upper bound for the last dimension of a dummy **assumed-size 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.

Fortran

5.2.6 iterator Modifier


Modifiers

Name	Modifies	Type	Properties
<i>iterator</i>	<i>locator-list</i>	Complex, name: iterator Arguments: iterator-specifier list of iterator specifier list item type (<i>default</i>)	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 *iterator value set*. An *iterator-identifier* expands to those values in the *clause* argument for which it is specified. Each *list item* of the *iterator* argument is an *iterator specifier* with this format:



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.



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

An *iterator* only exists in the context of the *clause* argument that its *iterator modifier* modifies. An *iterator* also hides all accessible symbols with the same name in the context of that *clause* argument.

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

C / C++

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

- $i_0 = (\text{iterator-type}) \text{ begin};$
- $i_j = (\text{iterator-type}) (i_{j-1} + \text{step}),$ where $j \geq 1;$ and
- if $\text{step} > 0,$
 - $i_0 < (\text{iterator-type}) \text{ end};$
 - $i_{N-1} < (\text{iterator-type}) \text{ end};$ and
 - $(\text{iterator-type}) (i_{N-1} + \text{step}) \geq (\text{iterator-type}) \text{ end};$
- if $\text{step} < 0,$
 - $i_0 > (\text{iterator-type}) \text{ end};$
 - $i_{N-1} > (\text{iterator-type}) \text{ end};$ and
 - $(\text{iterator-type}) (i_{N-1} + \text{step}) \leq (\text{iterator-type}) \text{ end}.$

C / C++

Fortran

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

- $i_1 = \text{begin};$
- $i_j = i_{j-1} + \text{step},$ where $j \geq 2;$ and
- if $\text{step} > 0,$
 - $i_1 \leq \text{end};$
 - $i_N \leq \text{end};$ and
 - $i_N + \text{step} > \text{end};$
- if $\text{step} < 0,$
 - $i_1 \geq \text{end};$
 - $i_N \geq \text{end};$ and
 - $i_N + \text{step} < \text{end}.$

Fortran

The **iterator value set** will be empty if no possible value complies with the conditions above.

If an *iterator-identifier* appears in a **list item** expression of the modified argument, the effect is as if the **list item** is instantiated within the **clause** for each member of the **iterator value set**, substituting each occurrence of *iterator-identifier* in the **list item** expression with the member of the **iterator value set**. If the **iterator value set** is empty then the effect is as if the **list item** was not specified.

1 **Restrictions**

2 Restrictions to *iterator modifiers* are as follows:

- 3
 - The *iterator-type* must not declare a new type.
 - For each value *i* in an *iterator value set*, the mathematical result of *i* + *step* must be representable in *iterator-type*.

4 C / C++

- 5
 - The *iterator-type* must be an integral or pointer type.
 - The *iterator-type* must not be **const** qualified.

6 C / C++

7 Fortran

- 8
 - The *iterator-type* must be an integer type.

9 Fortran

- 10
 - If the *step expression* of a *range-specification* equals zero, the behavior is unspecified.
 - Each *iterator-identifier* can only be defined once in the *modifier-parameter-specification*.
 - An *iterator-identifier* must not appear in the *range-specification*.
 - If an *iterator modifier* appears in a *clause* that is specified on a **task_iteration** directive then the *loop-iteration variables* of **taskloop-affected loops** of the associated **taskloop construct** must not appear in the *range-specification*.

15 **Cross References**

- 16
 - **affinity** Clause, see [Section 20.10](#)
 - **depend** Clause, see [Section 23.9.5](#)
 - **from** Clause, see [Section 12.2](#)
 - **map** Clause, see [Section 11.3](#)
 - **to** Clause, see [Section 12.1](#)

21 **5.3 Conditional Compilation**

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

23 Fortran

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

25 Fortran

Restrictions

Restrictions to conditional compilation are as follows:

- A **#define** or a **#undef** preprocessing directive in user code must not define or undefine the **_OPENMP** macro name.

Fortran

5.3.1 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):

```
!23456789
!$ iam = omp_get_thread_num() +      &
!$&    index

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

Fortran

5.3.2 Fixed Source Form Conditional Compilation Sentinels

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

!\$ | *\$ | c\$

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

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

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

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

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

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

5.4 Intrinsic Identifiers

Intrinsic Identifiers

Name	Properties
<code>omp_fill</code>	<i>default</i>
<code>omp_num_args</code>	<i>default</i>
<code>omp_in</code>	<i>default</i>
<code>omp_out</code>	<i>default</i>
<code>omp_priv</code>	<i>default</i>
<code>omp_orig</code>	<i>default</i>
<code>omp_var</code>	<i>default</i>
<code>omp_idx</code>	<i>default</i>
<code>omp_step</code>	<i>default</i>

The intrinsic identifier `omp_fill` is a context-specific value that can only be used as a *list item* of the *counts* clause. It represents the number of *logical iterations* of a *logical iteration space* that remain after removing those specified by the other *list items*.

The intrinsic identifier `omp_num_args` can only be used in *parameter list items* and is a context-specific value that evaluates to the number of parameters of the associated declaration plus any variadic arguments that were passed, if any, at a given *procedure* call site.

The intrinsic identifiers `omp_in`, `omp_out`, `omp_priv` and `omp_orig` are special *variables* that are used to specify a *reduction operation* for a *user-defined reduction*. The `omp_in` and `omp_out` intrinsic identifiers can only be used in *combiner expressions*, and the `omp_priv` and `omp_orig` intrinsic identifiers can only be used in *initializer expressions*.

The intrinsic identifiers `omp_var`, `omp_idx` and `omp_step` are special *variables* that are used to specify an *induction operation* for a *user-defined induction*. The `omp_var` intrinsic identifier can only be used in *inductor expressions*, the `omp_idx` intrinsic identifier can only be used in *collector expressions*, and the `omp_step` intrinsic identifier can only be used in *inductor expressions* and *collector expressions*.

5.5 *directive-name-modifier* Modifier

Modifiers

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

Clauses

`absent`, `acq_rel`, `acquire`, `adjust_args`, `affinity`, `align`, `aligned`, `allocate`, `allocator`, `append_args`, `apply`, `at`, `atomic_default_mem_order`, `bind`, `capture`, `collapse`, `collector`, `combiner`, `compare`, `contains`, `copyin`, `copyprivate`, `default`, `defaultmap`, `depend`, `destroy`, `detach`, `device`,

[device_safesync](#), [device_type](#), [dist_schedule](#), [doacross](#),
[dynamic_allocators](#), [enter](#), [exclusive](#), [fail](#), [filter](#), [final](#), [firstprivate](#),
[from](#), [full](#), [grainsize](#), [graph_id](#), [graph_reset](#), [has_device_addr](#), [hint](#), [holds](#),
[if](#), [in_reduction](#), [inbranch](#), [inclusive](#), [indirect](#), [induction](#), [inductor](#), [init](#),
[init_complete](#), [initializer](#), [interop](#), [is_device_ptr](#), [lastprivate](#), [linear](#),
[link](#), [local](#), [map](#), [match](#), [memscope](#), [mergeable](#), [message](#), [no_openmp](#),
[no_openmp_constructs](#), [no_openmp_routines](#), [no_parallelism](#), [nocontext](#),
[nogroup](#), [nontemporal](#), [notinbranch](#), [novariants](#), [nowait](#), [num_tasks](#),
[num_teams](#), [num_threads](#), [order](#), [ordered](#), [otherwise](#), [partial](#), [permutation](#),
[priority](#), [private](#), [proc_bind](#), [read](#), [reduction](#), [relaxed](#), [release](#),
[replayable](#), [reverse_offload](#), [safelen](#), [safesync](#), [schedule](#), [self_maps](#),
[seq_cst](#), [severity](#), [shared](#), [simd](#), [simdlen](#), [sizes](#), [task_reduction](#),
[thread_limit](#), [threads](#), [threadset](#), [to](#), [transparent](#), [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 [clause](#). The *directive-name-modifier* specifies *directive-name*, which is the [directive name](#) of a [directive](#), [construct](#) or [constituent construct](#) to which the [clause](#) applies. If the [directive name](#) is that of a [compound construct](#), then the [leaf constructs](#) to which the [clause](#) applies are determined as specified in [Section 25.2](#). 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* are as follows:

- The *directive-name-modifier* must specify the [directive name](#) of either the [directive](#) on which the [clause](#) appears or a [constituent directive](#) of that [directive](#).

Cross References

- **absent** Clause, see [Section 16.6.1.1](#)
- **acq_rel** Clause, see [Section 23.8.1.1](#)
- **acquire** Clause, see [Section 23.8.1.2](#)
- **adjust_args** Clause, see [Section 15.6.2](#)
- **affinity** Clause, see [Section 20.10](#)
- **align** Clause, see [Section 13.3](#)
- **aligned** Clause, see [Section 14.2](#)
- **allocate** Clause, see [Section 13.6](#)

- **allocator** Clause, see [Section 13.4](#)
- **append_args** Clause, see [Section 15.6.3](#)
- **apply** Clause, see [Section 17.1](#)
- **at** Clause, see [Section 16.2](#)
- **atomic_default_mem_order** Clause, see [Section 16.5.1.1](#)
- **bind** Clause, see [Section 19.8.1](#)
- **capture** Clause, see [Section 23.8.3.1](#)
- **full** Clause, see [Section 17.10.1](#)
- **partial** Clause, see [Section 17.10.2](#)
- **collapse** Clause, see [Section 6.4.5](#)
- **collector** Clause, see [Section 8.16.2](#)
- **combiner** Clause, see [Section 8.15.1](#)
- **compare** Clause, see [Section 23.8.3.2](#)
- **contains** Clause, see [Section 16.6.1.2](#)
- **copyin** Clause, see [Section 10.1](#)
- **copyprivate** Clause, see [Section 10.2](#)
- **default** Clause, see [Section 7.3.1](#)
- **defaultmap** Clause, see [Section 11.4](#)
- **depend** Clause, see [Section 23.9.5](#)
- **destroy** Clause, see [Section 5.8](#)
- **detach** Clause, see [Section 20.11](#)
- **device** Clause, see [Section 21.2](#)
- **device_safesync** Clause, see [Section 16.5.1.7](#)
- **device_type** Clause, see [Section 21.1](#)
- **dist_schedule** Clause, see [Section 19.7.1](#)
- **doacross** Clause, see [Section 23.9.7](#)
- **dynamic_allocators** Clause, see [Section 16.5.1.2](#)
- **enter** Clause, see [Section 9.4](#)
- **exclusive** Clause, see [Section 8.17.2](#)

- 1 • **fail** Clause, see [Section 23.8.3.3](#)
- 2 • **filter** Clause, see [Section 18.5.1](#)
- 3 • **final** Clause, see [Section 20.7](#)
- 4 • **firstprivate** Clause, see [Section 7.3.4](#)
- 5 • **from** Clause, see [Section 12.2](#)
- 6 • **grainsize** Clause, see [Section 20.2.1](#)
- 7 • **graph_id** Clause, see [Section 20.3.1](#)
- 8 • **graph_reset** Clause, see [Section 20.3.2](#)
- 9 • **has_device_addr** Clause, see [Section 7.3.8](#)
- 10 • **hint** Clause, see [Section 23.1](#)
- 11 • **holds** Clause, see [Section 16.6.1.3](#)
- 12 • **if** Clause, see [Section 5.6](#)
- 13 • **in_reduction** Clause, see [Section 8.12](#)
- 14 • **inbranch** Clause, see [Section 15.8.1.1](#)
- 15 • **inclusive** Clause, see [Section 8.17.1](#)
- 16 • **indirect** Clause, see [Section 15.9.3](#)
- 17 • **induction** Clause, see [Section 8.13](#)
- 18 • **inductor** Clause, see [Section 8.16.1](#)
- 19 • **init** Clause, see [Section 5.7](#)
- 20 • **init_complete** Clause, see [Section 8.17.3](#)
- 21 • **initializer** Clause, see [Section 8.15.2](#)
- 22 • **interop** Clause, see [Section 15.7.1](#)
- 23 • **is_device_ptr** Clause, see [Section 7.3.6](#)
- 24 • **lastprivate** Clause, see [Section 7.3.5](#)
- 25 • **linear** Clause, see [Section 8.14](#)
- 26 • **link** Clause, see [Section 9.5](#)
- 27 • **local** Clause, see [Section 9.3](#)
- 28 • **map** Clause, see [Section 11.3](#)
- 29 • **match** Clause, see [Section 15.6.1](#)

- **memscope** Clause, see [Section 23.8.4](#)
- **mergeable** Clause, see [Section 20.5](#)
- **message** Clause, see [Section 16.3](#)
- **no_omp** Clause, see [Section 16.6.1.4](#)
- **no_omp_constructs** Clause, see [Section 16.6.1.5](#)
- **no_omp_routines** Clause, see [Section 16.6.1.6](#)
- **no_parallelism** Clause, see [Section 16.6.1.7](#)
- **nocontext** Clause, see [Section 15.7.3](#)
- **nogroup** Clause, see [Section 23.7](#)
- **nontemporal** Clause, see [Section 18.4.1](#)
- **notinbranch** Clause, see [Section 15.8.1.2](#)
- **novariants** Clause, see [Section 15.7.2](#)
- **nowait** Clause, see [Section 23.6](#)
- **num_tasks** Clause, see [Section 20.2.2](#)
- **num_teams** Clause, see [Section 18.2.1](#)
- **num_threads** Clause, see [Section 18.1.2](#)
- **order** Clause, see [Section 18.3](#)
- **ordered** Clause, see [Section 6.4.6](#)
- **otherwise** Clause, see [Section 15.4.2](#)
- **permutation** Clause, see [Section 17.5.1](#)
- **priority** Clause, see [Section 20.9](#)
- **private** Clause, see [Section 7.3.3](#)
- **proc_bind** Clause, see [Section 18.1.4](#)
- **read** Clause, see [Section 23.8.2.1](#)
- **reduction** Clause, see [Section 8.10](#)
- **relaxed** Clause, see [Section 23.8.1.3](#)
- **release** Clause, see [Section 23.8.1.4](#)
- **replayable** Clause, see [Section 20.6](#)
- **reverse_offload** Clause, see [Section 16.5.1.3](#)

- 1 • **safelen** Clause, see [Section 18.4.2](#)
- 2 • **safesync** Clause, see [Section 18.1.5](#)
- 3 • **schedule** Clause, see [Section 19.6.3](#)
- 4 • **self_maps** Clause, see [Section 16.5.1.6](#)
- 5 • **seq_cst** Clause, see [Section 23.8.1.5](#)
- 6 • **severity** Clause, see [Section 16.4](#)
- 7 • **shared** Clause, see [Section 7.3.2](#)
- 8 • **simd** Clause, see [Section 23.10.3.2](#)
- 9 • **simdlen** Clause, see [Section 18.4.3](#)
- 10 • **sizes** Clause, see [Section 17.2](#)
- 11 • **task_reduction** Clause, see [Section 8.11](#)
- 12 • **thread_limit** Clause, see [Section 21.3](#)
- 13 • **threads** Clause, see [Section 23.10.3.1](#)
- 14 • **threadset** Clause, see [Section 20.8](#)
- 15 • **to** Clause, see [Section 12.1](#)
- 16 • **transparent** Clause, see [Section 23.9.6](#)
- 17 • **unified_address** Clause, see [Section 16.5.1.4](#)
- 18 • **unified_shared_memory** Clause, see [Section 16.5.1.5](#)
- 19 • **uniform** Clause, see [Section 14.1](#)
- 20 • **untied** Clause, see [Section 20.4](#)
- 21 • **update** Clause, see [Section 23.8.2.2](#)
- 22 • **update** Clause, see [Section 23.9.4](#)
- 23 • **use** Clause, see [Section 22.1.2](#)
- 24 • **use_device_addr** Clause, see [Section 7.3.9](#)
- 25 • **use_device_ptr** Clause, see [Section 7.3.7](#)
- 26 • **uses_allocators** Clause, see [Section 13.8](#)
- 27 • **weak** Clause, see [Section 23.8.3.4](#)
- 28 • **when** Clause, see [Section 15.4.1](#)
- 29 • **write** Clause, see [Section 23.8.2.3](#)

5.6 if Clause

Name: if	Properties: target-consistent
-----------------	---

Arguments

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

Modifiers

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

Directives

[cancel](#), [parallel](#), [simd](#), [target](#), [target_data](#), [target_enter_data](#), [target_exit_data](#), [target_update](#), [task](#), [task_iteration](#), [taskgraph](#), [taskloop](#), [teams](#)

Semantics

The effect of the [if clause](#) depends on the [construct](#) to which it is applied. If the [construct](#) is not a [compound 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](#) Construct, see [Section 24.2](#)
- [parallel](#) Construct, see [Section 18.1](#)
- [simd](#) Construct, see [Section 18.4](#)
- [target](#) Construct, see [Section 21.8](#)
- [target_data](#) Construct, see [Section 21.7](#)
- [target_enter_data](#) Construct, see [Section 21.5](#)
- [target_exit_data](#) Construct, see [Section 21.6](#)
- [target_update](#) Construct, see [Section 21.9](#)
- [task](#) Construct, see [Section 20.1](#)
- [task_iteration](#) Directive, see [Section 20.2.3](#)

- **taskgraph** Construct, see [Section 20.3](#)
- **taskloop** Construct, see [Section 20.2](#)
- **teams** Construct, see [Section 18.2](#)

5.7 init Clause

Name: <code>init</code>	Properties: innermost-leaf
--------------------------------	---

Arguments

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

Modifiers

Name	Modifies	Type	Properties
prefer-type	<i>init-var</i>	Complex, name: prefer_type Arguments: prefer-type-specification list of preference specification list item type (default)	complex , unique
<i>depinfo-modifier</i>	<i>init-var</i>	Complex, Keyword: in , inout , inoutset , mutexinoutset , out Arguments: locator-list-item locator list item (default)	complex , unique
<i>interop-type</i>	<i>init-var</i>	Keyword: target , targetsync	repeatable
directive-name-modifier	<i>all arguments</i>	Keyword: <i>directive-name</i> (a directive name)	unique

Directives

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

Semantics

When the [init clause](#) appears on a [depobj construct](#), it specifies that *init-var* is a [depend object](#) for which the state is set to *initialized*. The effect is that *init-var* is set to represent a [dependence type](#) and [locator list item](#) as specified by the name and argument of the [depinfo-modifier](#).

When the [init clause](#) appears on an [interop construct](#), it specifies that *init-var* is an [interoperability object](#) that 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 `omp_interop_none`.

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`.

Restrictions

- *init-var* must not be `constant`.
- If the `init` clause appears on a `depobj` construct, *init-var* must refer to a `variable` of `depend` OpenMP type that is *uninitialized*.
- If the `init` clause appears on a `depobj` construct then the *depinfo-modifier* has the `required` property and otherwise it must not be present.
- If the `init` clause appears on an `interop` construct, *init-var* must refer to a `variable` of `interop` OpenMP type.
- If the `init` clause appears on an `interop` construct, the *interop-type* modifier has the `required` property and each *interop-type* keyword has the `unique` property. Otherwise, the *interop-type* modifier must not be present.
- The *prefer-type* modifier must not be present unless the `init` clause appears on an `interop` construct.

Cross References

- `depobj` Construct, see [Section 23.9.3](#)
- `interop` Construct, see [Section 22.1](#)

5.8 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: <i>directive-name</i> (a <i>directive name</i>)	<i>unique</i>

Directives

depobj, **interop**

Additional information

When the **destroy** clause appears on a **depobj** directive that specifies *depend-object* as a **directive** argument, the *destroy-var* argument may be omitted. If omitted, the effect is as if *destroy-var* refers to the *depend-object* argument.

Semantics

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

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

Restrictions

- *destroy-var* must not be **constant**.
- If the **destroy** clause appears on a **depobj** construct, *destroy-var* must refer to a **variable** of **depend** OpenMP type that is *initialized*.
- If the **destroy** clause appears on an **interop** construct, *destroy-var* must refer to a **variable** of **interop** OpenMP type that is *initialized*.

Cross References

- **depobj** Construct, see [Section 23.9.3](#)
- **interop** Construct, see [Section 22.1](#)

6 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.

6.1 OpenMP Types and Identifiers

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

[Intrinsic identifiers](#) denote special [variables](#) that are part of the implementation and can be used only in [directives](#), generally restricted to certain contexts. [Predefined identifiers](#) can be used in [base-language code](#). Many [predefined identifiers](#) have the [constant property](#), as is indicated where they are defined in this specification. The implementation implicitly declares these [OpenMP identifiers](#) and evaluates them when they are 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 [OpenMP types](#) support the definition of many important OpenMP concepts independently of the [base language](#) in which they are used.

[Assignable OpenMP type instances](#) are 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 ▼

22 An [assignable OpenMP type instance](#) is a [variable](#) or a function reference with data pointer result of
23 that [OpenMP type](#).

▲ Fortran ▲

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

C / C++

Any expression of logical **OpenMP type** is a scalar expression. This document uses *true* as a generic term for a non-zero integer value and *false* as a generic term for an integer value of zero.

C / C++

Fortran

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

Fortran

The integer **OpenMP type** supports integer **variables** and expressions in any **base language**.

C / C++

Any expression of integer **OpenMP type** is an integer expression.

C / C++

Fortran

Any expression of integer **OpenMP type** is a scalar integer expression.

Fortran

The string **OpenMP type** supports character string **variables** and expressions in any **base language**.

C / C++

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

C / C++

Fortran

Any expression of string **OpenMP type** is a character string of default kind.

Fortran

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

Each **OpenMP type** other than those specifically defined in this section has a generic name, *<generic_name>*, by which it is referred throughout this document and that is used to construct the **base language** construct that corresponds to that **OpenMP type**. Some **OpenMP types** are **OMPD types** or **OMPT types**; all of these **OpenMP types** have generic names.

C / C++

Unless otherwise specified, an **OMPD trace record** has a *<generic_name>* **OMPD type**, which corresponds to the type **ompd_record_<generic_name>_t** and an **OMPD callback** has a *<generic_name>* **OMPD type** signature, which corresponds to the type **ompd_callback_<generic_name>_fn_t**. Unless otherwise specified, all other *<generic_name>* **OMPD types** correspond to the type **ompd_<generic_name>_t**.

Unless otherwise specified, an **OMPT trace record** has a `<generic_name> OMPT type`, which corresponds to the type `ompt_record_<generic_name>_t` and an **OMPT callback** has a `<generic_name> OMPT type` signature, which corresponds to the type `ompt_callback_<generic_name>_t`. Unless otherwise specified, all other `<generic_name> OMPT types` correspond to the type `ompt_<generic_name>_t`.

Otherwise, unless otherwise specified, a **variable** of `<generic_name> OpenMP type` is a **variable** of type `omp_<generic_name>_t`.



Unless otherwise specified, the type of an **OMPD trace record** is not defined and the type signature of an **OMPD callback** is not defined. Unless otherwise specified, a **variable** of a `<generic_name> OMPD type` is an integer **scalar variable** of kind `ompd_<generic_name>_kind`.

Unless otherwise specified, the type of an **OMPT trace record** is not defined and the type signature of an **OMPT callback** is not defined. Unless otherwise specified, a **variable** of a `<generic_name> OMPT type` is an integer **scalar variable** of kind `ompt_<generic_name>_kind`.

Otherwise, unless otherwise specified, a **variable** of `<generic_name> OpenMP type` is an integer **scalar variable** of kind `omp_<generic_name>_kind`.



Cross References

- OpenMP Foreign Runtime Identifiers, see [Section 22.1.1](#)
- OpenMP Reduction and Induction Identifiers, see [Section 8.1](#)
- Mapper Identifiers and **mapper** Modifiers, see [Section 11.2](#)

6.2 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. **OpenMP stylized expressions** often use **OpenMP identifiers** that the implementation binds to well-defined internal state.

Cross References

- OpenMP Collector Expressions, see [Section 8.2.4](#)
- OpenMP Combiner Expressions, see [Section 8.2.1](#)
- OpenMP Inductor Expressions, see [Section 8.2.3](#)
- OpenMP Initializer Expressions, see [Section 8.2.2](#)

6.3 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 **setjmp**.
- **longjmp** 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

6.3.1 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** Construct, see [Section 13.7](#)

6.3.2 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** Construct, see [Section 15.7](#)

6.3.3 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 +, *, -, /, &, ^, |, <<, or >>.

- `==` 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 or *variables* that are non-polymorphic scalar pointers and any length type parameter must be constant.
- *e* (expected) and *d* (desired) are either scalar expressions or *scalar variables*.
- *expr* is a scalar expression or *scalar variable*.
- *r* (result) is a scalar logical *variable*.
- *expr-list* is a comma-separated, non-empty list of scalar expressions and *scalar variables*.
- *intrinsic-procedure-name* is one of `MAX`, `MIN`, `IAND`, `IOR`, `IEOR`, `PREVIOUS`, or `NEXT`.
- *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 structured block** can be specified for **atomic directives** that enforce **atomic read** semantics but not capture semantics.

▼ C / C++ ▼

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

```
v = x;
```

▲ C / C++ ▲

▼ Fortran ▼

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

```
v = x
```

```
v => x
```

▲ Fortran ▲

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

▼ C / C++ ▼

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

```
x = expr;
```

▲ C / C++ ▲

▼ Fortran ▼

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

```
x = expr
```

```
x => expr
```

▲ Fortran ▲

An **update structured block** can be specified for **atomic directives** that enforce **atomic update** semantics but not capture semantics. It is further categorized as **unconditional-update structured blocks** or **conditional-update structured blocks**.

▼ C / C++ ▼

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

```
x++;
```

```
x--;
```

```
++x;
```

```
--x;
```

```
x binop= expr;
```

```
x = x binop expr;
```

```
x = expr binop x;
```

▲ C / C++ ▲

Fortran

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

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

Fortran

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

C / C++

A **conditional-update 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;
{ r = x == e; if(r) x = d; }
```

C / C++

Fortran

A **conditional-update 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
x = ( x equalop e ? d : x )
if (x ordop expr) x = expr
if (x ordop expr) then; x = expr; end if
x = ( x ordop expr ? expr : x )
if (expr ordop x) x = expr
if (expr ordop x) then; x = expr; end if
x = ( expr ordop x ? expr : x )
if (associated(x)) x => expr
if (associated(x)) then; x => expr; end if
if (associated(x, e)) x => expr
if (associated(x, e)) then; x => expr; end if
```

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

```
r = cond; if (r) x assign d  
r = cond; if (r) then; x assign d; end if
```

where *assign* is either = or => and *cond* denotes one of the following conditions:

```
x equalop e  
ASSOCIATED (x)  
ASSOCIATED (x, e)
```

For an **atomic** construct with a **read structured block**, **write structured block**, or **update structured block** that consists of only one executable statement (which may include a block of further statements), the paired **end directive** is optional.

Fortran

A **capture structured block** can be specified for **atomic** directives that enforce capture semantics. It is further categorized as **update-capture structured block** or **conditional-update-capture structured block**.

C / C++

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

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

The following form for a **conditional-update-capture structured block** has been **deprecated**:

```
{ r = x == e; if (r) x = d; }
```

C / C++

Fortran

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

```
statement  
capture-statement
```

or

```
capture-statement
statement
```

where *capture-statement* has either of the following forms:

```
v = x
v => x
```

If *statement* is *write-statement* or *update-statement* as specified above then it is an **update-capture structured block** and may be used in **atomic constructs** that enforce **atomic captured update semantics**. If *statement* is *conditional-update-statement* as specified above then it is a **conditional-update-capture structured block**. In addition, a **conditional-update-capture structured block** can have one of the following forms:

```
if (cond) then
  x assign d
else
  v assign x
end if
```

or

```
r = cond
if (r) then
  x assign d
else
  v assign x
endif
```

where *assign* and *cond* are defined as indicated above.

The following form for a **conditional-update-capture structured block** has been **deprecated**:

```
r = cond
if (r) x assign d
```

Fortran

Restrictions

Restrictions to OpenMP **atomic structured blocks** 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*, *r*, *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*.
- The expression *x == e* must be numerically equivalent to *x == (e)*. This requirement is satisfied if the operators in *e* have precedence equal to or greater than *==*, or by using parentheses around *e* or subexpressions of *e*.

C / C++
Fortran

- *x* must not have the **ALLOCATABLE** attribute.
- During the execution of an **atomic region**, multiple syntactic occurrences of *x* must designate the same **storage location**.
- 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*, *r*, *expr*, and *expr-list* (as applicable) may access the same **storage location** as any other symbol in the list.

- In forms that capture the original value of x in v , v may not access the same [storage location](#) as e .
- If *intrinsic-procedure-name* refers to **IAND**, **IOR**, **IEOR**, **PREVIOUS**, or **NEXT** then exactly one expression must appear in *expr-list*.
- The expression $x \text{ operator } \textit{expr}$ must be, depending on its type, either mathematically or logically equivalent to $x \text{ operator } (\textit{expr})$. This requirement is satisfied if the operators in *expr* have precedence greater than *operator*, or by using parentheses around *expr* or subexpressions of *expr*.
- The expression $\textit{expr operator } x$ must be, depending on its type, either mathematically or logically equivalent to $(\textit{expr}) \text{ operator } x$. This requirement is satisfied if the operators in *expr* have precedence equal to or greater than *operator*, or by using parentheses around *expr* or subexpressions of *expr*.
- The expression $x \text{ equalop } e$ must be, depending on its type, either mathematically or logically equivalent to $x \text{ equalop } (e)$. This requirement is satisfied if the operators in e have precedence equal to or greater than *equalop*, or by using parentheses around e or subexpressions of e .
- *intrinsic-procedure-name* must refer to the intrinsic procedure name and not to other program entities.
- *operator* must refer to the intrinsic operator and not to a user-defined operator.
- Assignments must be either all intrinsic assignments or all pointer assignments.
- If the **ASSOCIATED** intrinsic function is referenced in a condition, all assignments must be pointer assignments. If pointer assignments are used, only the **ASSOCIATED** intrinsic function may be referenced in a condition.
- Unless x is a [scalar variable](#) or a function references with scalar data pointer result of non-character intrinsic type, intrinsic assignments, *equalop*, and *ordop* must not be used.
- Arguments to an **ASSOCIATED** intrinsic function must not have zero-sized storage sequences.

Fortran

Cross References

- **atomic** Construct, see [Section 23.8.5](#)

6.4 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.

6.4.1 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 [loop-iteration variable](#).

C++

or

Fortran

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

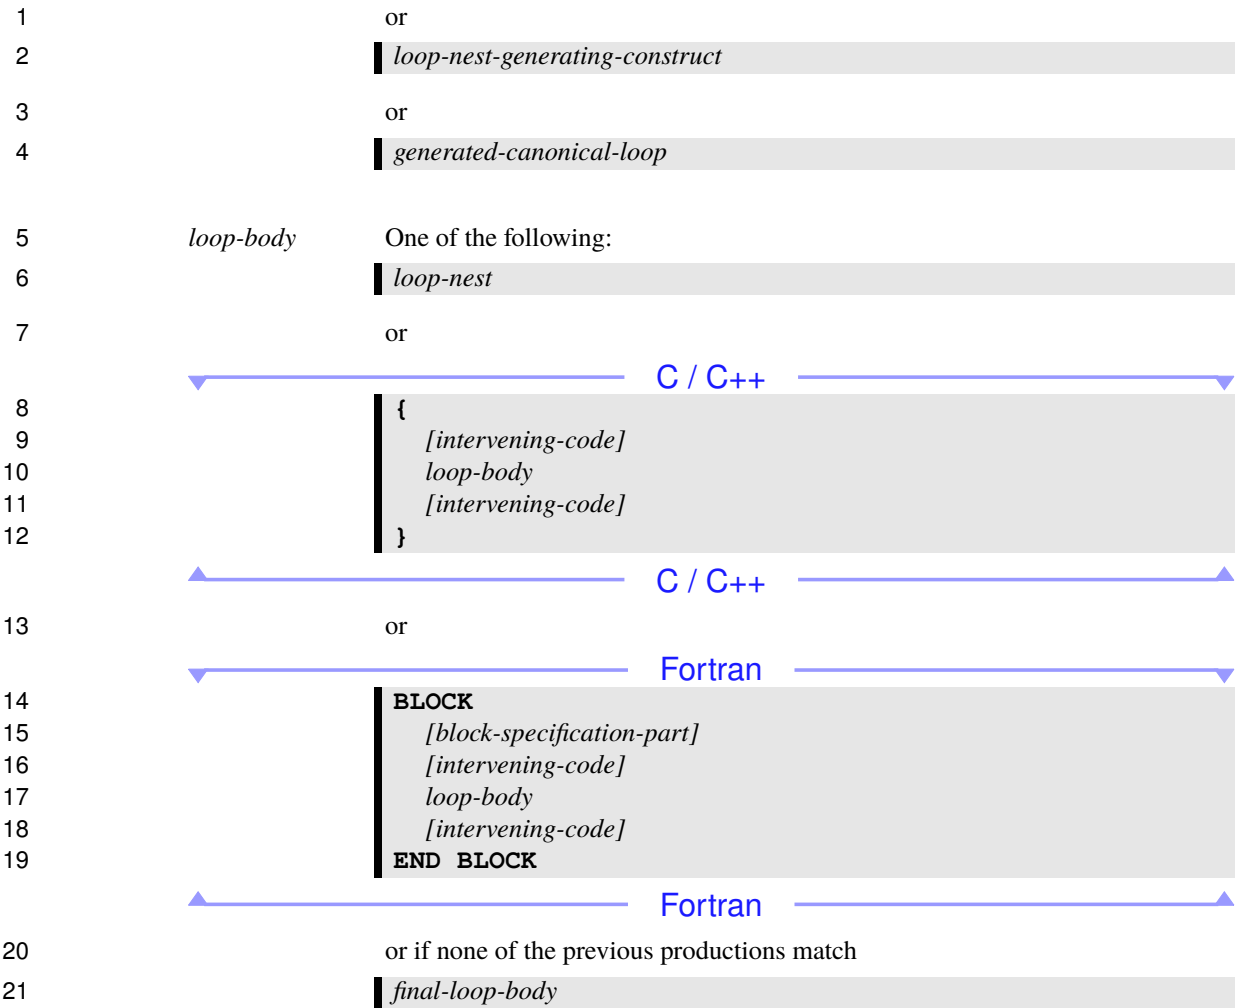
If the *loop-nest* is a *nonblock-do-construct*, it is treated as a *block-do-construct* for each **DO** construct.

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

or

```
BLOCK
    loop-nest
END BLOCK
```

Fortran



22 *loop-nest-generating-construct*

23 A [loop-transforming construct](#) that generates a [canonical loop nest](#), which may

24 be a [canonical loop sequence](#) that contains exactly one [canonical loop nest](#).

25 *generated-canonical-loop*

26 A [generated loop](#) from a [loop-transforming construct](#) that has [canonical loop nest](#)

27 form and for which the [loop body](#) matches *loop-body*.

1	<i>intervening-code</i>	
2		C / C++
3		A non-empty sequence of structured blocks or declarations, referred to as
4		intervening code . It must not contain iteration statements, continue
		statements or break statements that apply to the enclosing loop.
		C / C++
		Fortran
5		A non-empty structured block sequence , referred to as intervening code . It must
6		not contain:
7		• loops;
8		• CYCLE statements;
9		• EXIT statements;
10		• array expressions;
11		• array references with a vector subscript;
12		• assignment statements where the target is an array object;
13		• references to elemental procedures with an array actual argument;
14		• references to procedures where the actual argument is an array that is not
15		simply contiguous and the corresponding dummy argument has the
16		CONTIGUOUS attribute or is an explicit-shape array or assumed-size array .
		Fortran
17		Additionally, intervening code must not contain executable directives or calls to
18		the OpenMP runtime API in its corresponding region . If intervening code is
19		present, then a loop at the same depth within the loop nest is not a perfectly
20		nested loop .
21	<i>final-loop-body</i>	A structured block that terminates the scope of loops in the loop nest. If the loop
22		nest is associated with a loop-nest-associated directive , loops in this structured
23		block cannot be associated with that directive .
		C / C++
24	<i>init-expr</i>	One of the following:
25		<i>var = lb</i>
26		<i>integer-type var = lb</i>
27		<i>pointer-type var = lb</i>
		C
		C

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

Fortran

A scalar variable of integer type.

Fortran

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

lb, ub

One of the following:

Expressions of a type compatible with the type of *var* that are loop invariant with respect to the outermost loop.

or

One of the following:

var-outer

var-outer + *a2*

a2 + *var-outer*

var-outer - *a2*

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

or

If *var* is of an integer type, one of the following:

a2 - *var-outer*

a1 * *var-outer*

a1 * *var-outer* + *a2*

a2 + *a1* * *var-outer*

a1 * *var-outer* - *a2*

a2 - *a1* * *var-outer*

var-outer * *a1*

var-outer * *a1* + *a2*

a2 + *var-outer* * *a1*

var-outer * *a1* - *a2*

a2 - *var-outer* * *a1*

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

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

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

1 *a1, a2, incr* Integer expressions that are loop invariant with respect to the outermost loop of
 2 the loop nest.
 3 If the loop is associated with a **directive**, the expressions are evaluated before the
 4 **construct** formed from that **directive**.

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

▼ C++ ▼

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

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

▲ C++ ▲

11 Restrictions

12 Restrictions to **canonical loop nests** are as follows:

▼ C / C++ ▼

- 13 • If *test-expr* is of the form *var relational-op b* and *relational-op* is < or <= then *incr-expr* must
 14 cause *var* to increase on each iteration of the loop. If *test-expr* is of the form *var*
 15 *relational-op b* and *relational-op* is > or >= then *incr-expr* must cause *var* to decrease on
 16 each iteration of the loop. Increase and decrease are using the order induced by *relational-op*.
- 17 • If *test-expr* is of the form *ub relational-op var* and *relational-op* is < or <= then *incr-expr*
 18 must cause *var* to decrease on each iteration of the loop. If *test-expr* is of the form *ub*
 19 *relational-op var* and *relational-op* is > or >= then *incr-expr* must cause *var* to increase on
 20 each iteration of the loop. Increase and decrease are using the order induced by *relational-op*.
- 21 • If *relational-op* is != then *incr-expr* must cause *var* to always increase by 1 or always
 22 decrease by 1 and the increment must be a constant expression.
- 23 • *final-loop-body* must not contain any **break** statement that would cause the termination of
 24 the innermost loop.

▲ C / C++ ▲

▼ Fortran ▼

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

▲ Fortran ▲

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

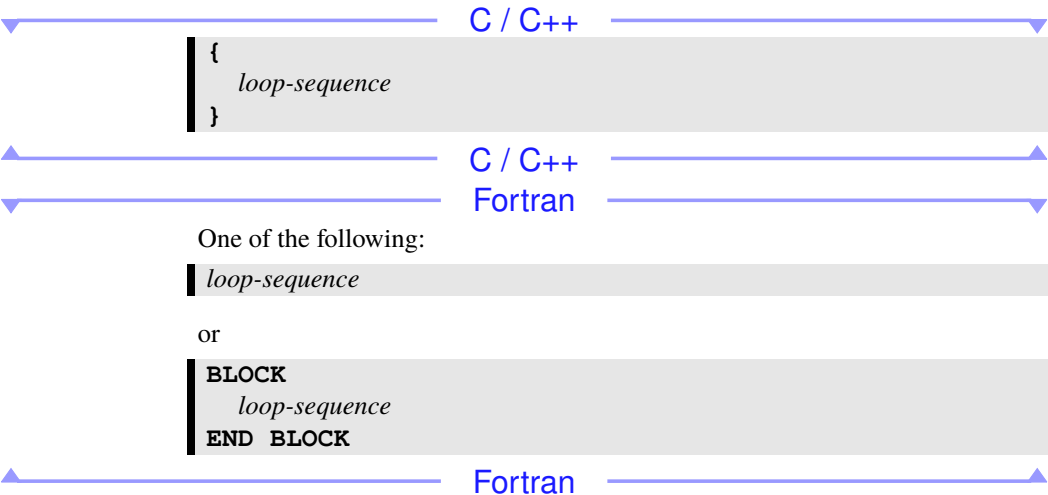
Cross References

- Canonical Loop Sequence Form, see [Section 6.4.2](#)
- Loop-Transforming Constructs, see [Chapter 17](#)
- **threadprivate** Directive, see [Section 9.1](#)

6.4.2 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 6.4.1](#)). The loops must be [bounds-independent loops](#) with respect to *canonical-loop-sequence*.

loop-sequence-generating-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

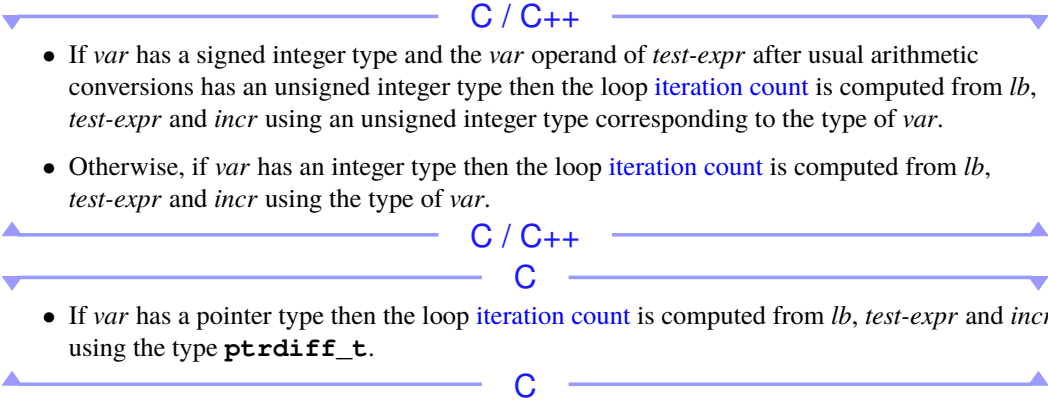
- **looprange** Clause, see [Section 6.4.8](#)
- Canonical Loop Nest Form, see [Section 6.4.1](#)
- Loop-Transforming Constructs, see [Chapter 17](#)

6.4.3 OpenMP Loop-Iteration Spaces and Vectors

A [loop-nest-associated directive](#) affects some number of the outermost loops of an [associated loop nest](#), called the [affected loops](#), in accordance with its specified [clauses](#). These [affected loops](#) and their [loop-iteration variables](#) form an OpenMP [loop-iteration vector space](#). OpenMP [loop-iteration vectors](#) allow other [directives](#) to refer to points in that [loop-iteration vector space](#).

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

The [canonical loop nest](#) form allows the [iteration count](#) of all [affected loops](#) to be computed before executing the outermost loop. For any [affected loop](#), the [iteration count](#) is computed as follows:



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 affected with a **construct** be *n*, where all of the **affected 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 **loop-iteration variables** of the **affected 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 **constant, non-negative** expression of integer **OpenMP type** 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 **constant, non-negative** expression of integer **OpenMP type**.

The iterations of some number of **affected 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 **affected loop**. The number of times that any **intervening code** between any two **collapse-affected 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.

At the beginning of each **collapsed iteration** in a **loop-collapsing construct**, the **loop-iteration variable** or the **variable** declared by *range-decl* of each **collapse-affected loop** has the value that it would have if the **collapse-affected loops** were not associated with any **directive**.

6.4.4 Consistent Loop Schedules

A **loop schedule** for a given **loop-nest-associated construct** assigns a **thread** in the **binding thread set** of that **construct** to a **logical iteration vector** of the **affected loop nest**. If the **loop schedules** of two **loop-nest-associated constructs** are **consistent schedules**, the behavior is as if they produce the same mapping of **logical iteration vectors** to **threads**. In particular, if two **loop-nest-associated construct** have **consistent schedules** and they have the same **binding thread set**, the implementation will guarantee that memory effects of a **logical iteration** in the first loop nest have completed before the execution of the corresponding **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 **schedule specification**;
- The **constructs** have **reproducible schedules**;
- The **affected loops** have identical **logical iteration vector spaces**;
- The two sets of **affected loops** either consist of only rectangular loops or both contain a **non-rectangular loop**; and
- The **loop schedules** of **transformation-affected loops** among any **affected loops** that are **generated loops** of a **loop-transforming construct** are all themselves **consistent**.

6.4.5 collapse Clause

Name: collapse		Properties: once-for-all-constituents, unique
Arguments		
Name	Type	Properties
<i>n</i>	expression of integer type	constant, positive

Modifiers

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

Directives

distribute, **do**, **for**, **loop**, **simd**, **taskloop**

Semantics

The **collapse** clause affects one or more loops of a **canonical loop nest** 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 **loop nest depth**.

Cross References

- **distribute** Construct, see [Section 19.7](#)
- **do** Construct, see [Section 19.6.2](#)
- **for** Construct, see [Section 19.6.1](#)
- **loop** Construct, see [Section 19.8](#)
- **simd** Construct, see [Section 18.4](#)
- **taskloop** Construct, see [Section 20.2](#)

6.4.6 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: <i>directive-name</i> (a <i>directive name</i>)	unique

Directives

do, for

Semantics

The **ordered** clause is used to specify the **doacross-affected loops** for the purpose of identifying cross-iteration dependencies. The argument *n* specifies the number of **doacross-affected loops** 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 **doacross-affected loops** may be **non-rectangular loops**.
- *n* must not evaluate to a value greater than the depth of the **associated loop nest**.
- 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.

Cross References

- **collapse** Clause, see [Section 6.4.5](#)
- **do** Construct, see [Section 19.6.2](#)
- **for** Construct, see [Section 19.6.1](#)

6.4.7 depth Clause

Name: depth	Properties: unique
---------------------------	----------------------------------

Arguments

Name	Type	Properties
<i>depth-expr</i>	expression of integer type	positive, constant

Modifiers

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

Directives

flatten, fuse

Semantics

The **depth** clause specifies the depth to which the **construct** on which the **clause** appears affects any **associated loop nests**.

Restrictions

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

- The *depth-expr* must evaluate to at most the **loop nest depth** of any **associated loop nest**.

Cross References

- **flatten** Construct, see [Section 17.3](#)
- **fuse** Construct, see [Section 17.4](#)

6.4.8 looprange Clause

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

Arguments

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

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 associated **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** Construct, see [Section 17.4](#)
- Canonical Loop Sequence Form, see [Section 6.4.2](#)

1

Part II

2

Data Control

7 Data-Sharing Control

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

[Data-sharing 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 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](#).

7.1 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 7.1.1](#) describes the [data-sharing attribute](#) rules for [variables](#) referenced in a [construct](#).
- [Section 7.1.2](#) describes the [data-sharing attribute](#) rules for [variables](#) referenced in a [region](#), but outside any [construct](#).

For any [variable](#) that is a [referencing variable](#) (including formal arguments passed by reference for C++), the [data-sharing attribute](#) rules apply only to its [referring pointer](#) unless otherwise specified.

7.1.1 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 a [nested construct](#) causes an implicit reference to the [variable](#) in the enclosing [construct](#). Specifying a [variable](#) in a [map clause](#) of an enclosed [construct](#) may cause an implicit

reference to the **variable** in the enclosing **construct**. Such implicit references are also subject to the **data-sharing attribute** rules outlined in this section.

Fortran

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

Fortran

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

- **Variables** with **automatic storage duration** that are declared in a scope inside the **construct** are **private**.
- **Variables** and common blocks (in Fortran) that appear as arguments in **threadprivate directives** or **variables** with the **__Thread_local** (in C) or **thread_local** (in C/C++) storage-class specifier are **threadprivate**.
- **Variables** and common blocks (in Fortran) that appear as arguments in **groupprivate directives** are **groupprivate**.
- **Variables** and common blocks (in Fortran) that appear as **list items** in **local clauses** on **declare_target** directives are **device-local**.
- **Variables** with **static storage duration** that are declared in a scope inside the **construct** are **shared**.
- A **dynamic groupprivate block** instantiated by a **groupprivate-instantiating construct** is **groupprivate**.
- Objects with **dynamic storage duration** are **shared**.
- The **loop-iteration variable** in any **affected loop** of a **loop** or **simd** construct is **lastprivate**.
- The **loop-iteration variable** in any **affected loop** of a **loop-nest-associated directive** is otherwise **private**.

C++

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

C++

Fortran

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

Fortran

C / C++

- Variables with **static storage duration** that are declared in a scope inside the **construct** are **shared**.
- If a **list item** in a **has_device_addr** clause or in a **map** clause on the **target** construct has a **base pointer**, and the **base pointer** is a **scalar variable** that is not a **list item** in a **map** clause on the **construct**, the **base pointer** is **firstprivate**.
- If a **list item** in a **reduction** or **in_reduction** clause on the **construct** has a **base pointer** then the **base pointer** is **private**.
- Static data members are **shared**.
- If a **list item** in a **shared** clause on the **construct** is a **referencing variable** then the **referring pointer** of the **list item** is **firstprivate**.
- If a **list item** in a **map** clause on the **target** construct has a **base referencing variable** that does not have a **containing structure**, the **base referring pointer** is **firstprivate**.
- The **__func__** variable and similar function-local predefined variables are **shared**.

C / C++

Fortran

- **Assumed-size arrays** and named constants are **shared** in **constructs** that are not **data-mapping constructs**.
- A named constant is **firstprivate** in **target** constructs.
- An associate name that may appear in a **variable** definition context is **shared** if its association occurs outside of the **construct** and otherwise it has the same **data-sharing attribute** as the selector with which it is associated.
- If a **list item** in a **map** clause on the **target** construct has a **base referencing variable** that is not the **list item** itself, the **base referring pointer** is **firstprivate** unless that **referencing variable** is a **structure element**, a **list item** in an **enter** clause on a **declare target directive**, or a **list item** in a **map** clause on the **construct** where the semantics of the **clause** apply to its **referring pointer**.

Fortran

- If a **list item** in a **has_device_addr** clause on the **target** construct has a **base referencing variable**, the **base referring pointer** is **firstprivate**.

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

- The **loop-iteration variable** in any affected loop of a **loop-nest-associated directive** may be listed in a **private** or **lastprivate** clause.

- If a **simd** construct has just one **affected loop** then its **loop-iteration variable** may be listed in a **linear** clause with a *linear-step* that is the increment of the **affected 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

- A **loop-iteration variable** of a loop that is not associated with any **directive** may be listed in a **data-sharing attribute clause** on the surrounding **teams**, **parallel** or **task-generating construct**, and on enclosed **constructs**, subject to other restrictions.
- An **assumed-size array** may be listed in a **shared** clause.
- A named constant 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 7.3](#).

Variables with **explicitly determined data-sharing attributes** are those that are referenced in a given **construct** and are listed in a **data-sharing 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 7.3.1](#)).
- In a **parallel** construct, if no **default** clause is present, these **variables** are **shared**.
- If no **default** clause is present on **constructs** that are not **task-generating constructs**, these **variables** reference the **variables** with the same names that exist in the **enclosing context**. If no **default** clause is present on a **task-generating construct** and the **generated task** is a **sharing task**, these **variables** are **shared**.
- In a **target** construct, **variables** that are not mapped after applying **data-mapping attribute** rules (see [Section 11.1](#)) 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 is **shared** if the **variable** is determined to be **shared** by all **implicit tasks** bound to the **current team** in the **enclosing context**.
- 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 7.3.4**.

7.1.2 Variables Referenced in a Region but not in a Construct

The **data-sharing attribute** of a **variable** or object that is referenced in a **region**, but not in the corresponding **construct**, is determined by the first matching rule from the following list.

- **Variables** with **automatic storage duration** that are declared in called **procedures** in the **region** are **private**.
- **Variables** and common blocks (in Fortran) that appear as arguments in **threadprivate directives** or **variables** with the **_Thread_local** (in C) or **thread_local** (in C/C++) storage-class specifier are **threadprivate**.
- **Variables** and common blocks (in Fortran) that appear as arguments in **groupprivate directives** are **groupprivate**.
- **Variables** and common blocks (in Fortran) that appear as **list items** in **local** clauses on **declare_target** directives are **device-local**.
- **Variables** with **static storage duration** are **shared**.
- Objects with **dynamic storage duration** are **shared**.

Fortran

- **Variables** that are accessed by host or use association are **shared**.
- A dummy argument of a called **procedure** 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 **procedure** 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**.

Fortran

7.2 List Item Privatization

Some **data-sharing attribute clauses**, including **reduction clauses**, specify that **list items** that appear in their **argument list** 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**, and the **new list items** have the **private attribute**.

If the **construct** is a **loop-collapsing construct** then, within the same **collapsed logical iteration** of the **collapse-affected 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**.

Unless otherwise specified, 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++

If the type of a **list item** is a reference to a type T then the type will be considered to be T for all purposes of the **clause**.

The order in which any default constructors for different **private variables** of **class type** are called is unspecified. The order in which any destructors for different **private variables** of **class type** are called is unspecified.

C++

Fortran

If any statement of the **construct** references a **list item**, a **new list item** of the same type and type parameters is allocated. This allocation occurs once for each **task** generated by the **construct** and once for each **SIMD lane** used by the **construct**. If the type of the **list item** has default initialization, the **new list item** has default initialization. Otherwise, the initial value of the **new list item** is undefined. The initial status of a **private** pointer is undefined.

For a **list item** or the subobject of a **list item** with the **ALLOCATABLE** attribute:

- If the allocation status is unallocated, the **new list item** or the subobject of the **new list item** will have an initial allocation status of unallocated;
- If the allocation status is allocated, the **new list item** or the subobject of the **new list item** will have an initial allocation status of allocated; and
- If the **new list item** or the subobject of the **new list item** is an array, its bounds will be the same as those of the **original list item** or the subobject of the **original list item**.

A **privatized list item** may be storage-associated with other **variables** when the **data-sharing attribute clause** is encountered. Storage association may exist because of **base language** constructs such as **EQUIVALENCE** or **COMMON**. If A is a **variable** that is **privatized** by a **construct** and B is a **variable** that is storage-associated with A then:

- The contents, allocation, and association status of B are undefined on entry to the **region**;
- Any definition of A , or of its allocation or association status, causes the contents, allocation, and association status of B to become undefined; and

- Any definition of *B*, or of its allocation or association status, causes the contents, allocation, and association status of *A* to become undefined.

A **privatized list item** may be a selector of an **ASSOCIATE**, **SELECT RANK** or **SELECT TYPE** construct. If the construct association is established prior to a **parallel region**, the association between the associate name and the **original list item** will be retained in the **region**.

The dynamic type of a **privatized list item** of a polymorphic type is the declared type.

Finalization of a **list item** of a finalizable type or subobjects of a **list item** of a finalizable type occurs at the end of the **region**. The order in which any final subroutines for different **variables** of a finalizable type are called is unspecified.

Fortran

If a **list item** appears in both **firstprivate** and **lastprivate** clauses, the update required for the **lastprivate** clause occurs after all initializations for the **firstprivate** clause.

Restrictions

The following restrictions apply to any **list item** that is **privatized** unless otherwise specified for a given **data-sharing attribute** clause:

- If a **list item** is an array or **array section**, it must specify contiguous storage.

C++

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

C++

C / C++

- A **variable** that is **privatized** must not have a **const**-qualified type unless it is of **class type** with a **mutable** member. This restriction does not apply to the **firstprivate** clause.
- A **variable** that is **privatized** must not have an incomplete type or be a reference to an incomplete type.

C / C++

Fortran

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

- Assumed-size arrays must not be privatized.
- An optional dummy argument that is not present must not appear as a list item in a privatization clause or be privatized as a result of an implicitly determined data-sharing attribute or predetermined data-sharing attribute.

Fortran

7.3 Data-Sharing Attribute Clauses

Several constructs accept clauses that allow a user to control the data-sharing attributes of variables referenced in the construct. 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 Chapter 8.

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

7.3.1 default Clause

Name: default	Properties: unique, post-modified
---------------	-----------------------------------

Arguments

Name	Type	Properties
data-sharing-attribute	Keyword: firstprivate, none, private, shared	default

Modifiers

Name	Modifies	Type	Properties
variable-category	implicit-behavior	Keyword: aggregate, all, allocatable, pointer, scalar	default
directive-name-modifier	all arguments	Keyword: directive-name (a directive name)	unique

Directives

parallel, **target**, **target_data**, **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 7.1.1.

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 scalar variables.

▲ C / C++ ▲

▼ Fortran ▼

The **scalar** *variable-category* specifies non-pointer and non-allocatable scalar variables. 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 *data-sharing-attribute* is not **none**, the *data-sharing attributes* of the selected variables will be *data-sharing-attribute*. If *data-sharing-attribute* is **none**, the *data-sharing attribute* is not implicitly determined. If *data-sharing-attribute* is **shared** then the clause has no effect on a **target** construct; otherwise, its effect on a **target** construct is equivalent to specifying the **defaultmap** clause with the same *data-sharing-attribute* and *variable-category*. If both the **default** and **defaultmap** clauses are specified on a **target** construct, and their *variable-category* modifiers specify intersecting categories, the **defaultmap** clause has precedence over the **default** clause for variables of those categories.

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

- `defaultmap` Clause, see [Section 11.4](#)
- `parallel` Construct, see [Section 18.1](#)
- `target` Construct, see [Section 21.8](#)
- `target_data` Construct, see [Section 21.7](#)
- `task` Construct, see [Section 20.1](#)
- `taskloop` Construct, see [Section 20.2](#)
- `teams` Construct, see [Section 18.2](#)

7.3.2 shared Clause

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

Arguments

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

Modifiers

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

Directives

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

Semantics

The [shared](#) clause declares one or more [list items](#) to have a [shared attribute](#) in [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 [list items](#) may include assumed-type [variables](#) and [procedure](#) pointers.

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** Construct, see [Section 18.1](#)
- **target_data** Construct, see [Section 21.7](#)
- **task** Construct, see [Section 20.1](#)
- **taskloop** Construct, see [Section 20.2](#)
- **teams** Construct, see [Section 18.2](#)

7.3.3 private Clause

Name: private	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: <i>directive-name</i> (a <i>directive name</i>)	<i>unique</i>

Directives

distribute, **do**, **for**, **loop**, **parallel**, **scope**, **sections**, **simd**, **single**, **target**, **target_data**, **task**, **taskloop**, **teams**

Semantics

The **private** clause specifies that its **list items** are to be **privatized list item** according to [Section 7.2](#). 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 **affected loops** and an **order** clause that specifies **concurrent** is also present. Each **new list item** is a **private-only variable**, unless otherwise specified.

Fortran

The **list items** may include **procedure** pointers.

Fortran

Restrictions

Restrictions to the **private** clause are as specified in [Section 7.2](#).

Cross References

- **distribute** Construct, see [Section 19.7](#)
- **do** Construct, see [Section 19.6.2](#)
- **for** Construct, see [Section 19.6.1](#)
- List Item Privatization, see [Section 7.2](#)
- **loop** Construct, see [Section 19.8](#)
- **parallel** Construct, see [Section 18.1](#)
- **scope** Construct, see [Section 19.2](#)
- **sections** Construct, see [Section 19.3](#)
- **simd** Construct, see [Section 18.4](#)
- **single** Construct, see [Section 19.1](#)
- **target** Construct, see [Section 21.8](#)
- **target_data** Construct, see [Section 21.7](#)
- **task** Construct, see [Section 20.1](#)
- **taskloop** Construct, see [Section 20.2](#)
- **teams** Construct, see [Section 18.2](#)

7.3.4 firstprivate Clause

Name: firstprivate	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>saved</i>	<i>list</i>	Keyword: saved	<i>default</i>
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: <i>directive-name</i> (a directive name)	<i>unique</i>

Directives


`distribute`, `do`, `for`, `parallel`, `scope`, `sections`, `single`, `target`, `target_data`,
`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 7.3.3, except as noted. In addition, the new list item has the **firstprivate** attribute and 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.



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.



If the **firstprivate** clause is on a **target construct** and a variable is of polymorphic type, the behavior is unspecified.

If an **original list item** does not have the **POINTER** attribute:

- If the **firstprivate** clause is not on a **target construct** then the initialization of the **new list items** occurs as if by an assignment statement; and
- If the **firstprivate** clause is on a **target construct** then the initialization of the **new list items** occurs as if either by an intrinsic assignment statement or by one or more defined assignment statements.

If an **original list item** that does not have the **POINTER** attribute has an allocation status of unallocated, the **new list items** will have the same status.

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

The **list items** may include named constants and **procedure** pointers.

Fortran

Restrictions

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

- A **list item** that is **private** within a **parallel region** must not appear in a **firstprivate** clause on a **worksharing construct** if any of the **worksharing regions** that arise from the **worksharing construct** ever bind to any of the **parallel regions** that arise from the **parallel construct**.
- A **list item** that is **private** within a **teams region** must not appear in a **firstprivate** clause on a **distribute construct** if any of the **distribute regions** that arise from the **distribute construct** ever bind to any of the **teams regions** that arise from the **teams construct**.
- A **list item** that appears in a **reduction clause** on a **parallel construct** must not appear in a **firstprivate** clause on a **task** or **taskloop construct** if any of the **task regions** that arise from the **task** or **taskloop construct** ever bind to any of the **parallel regions** that arise from the **parallel construct**.
- A **list item** that appears in a **reduction clause** on a **worksharing construct** must not appear in a **firstprivate** clause on a **task construct** encountered during execution of any of the **worksharing regions** that arise from the **worksharing construct**.

C++

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

C++

Cross References

- **distribute** Construct, see [Section 19.7](#)
- **do** Construct, see [Section 19.6.2](#)
- **for** Construct, see [Section 19.6.1](#)
- **parallel** Construct, see [Section 18.1](#)
- **private** Clause, see [Section 7.3.3](#)
- **scope** Construct, see [Section 19.2](#)
- **sections** Construct, see [Section 19.3](#)
- **single** Construct, see [Section 19.1](#)
- **target** Construct, see [Section 21.8](#)
- **target_data** Construct, see [Section 21.7](#)
- **task** Construct, see [Section 20.1](#)
- **taskloop** Construct, see [Section 20.2](#)
- **teams** Construct, see [Section 18.2](#)

7.3.5 lastprivate Clause

Name: lastprivate	Properties: data-environment attribute, data-sharing attribute, original list-item updating, 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	<i>default</i>
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: <i>directive-name</i> (a 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 7.3.3. In addition, each **new list item** has the **lastprivate** attribute. Further, when a **lastprivate** clause without the **conditional** modifier appears on a **directive** and the **list item** is not a **loop-iteration variable** of any **affected loop**, the value of each **new list item** from the sequentially last iteration of the **affected loops**, or the lexically last **structured block sequence** associated with a **sections** construct, is assigned to the **original list item**. Alternatively, when the **conditional** modifier appears on the **clause** or the **list item** is a **loop-iteration variable** of one of the **affected loops**, if execution of the **canonical loop nest**, when it is not associated with a **directive**, would assign a value to the **list item** then the **original list item** is assigned that value.

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 an assignment statement.

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 **affected 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 **affected loop** has an unspecified value after the **construct** if it would not be assigned a value during execution of the **canonical loop nest** when the loop nest is not associated with a **directive**. 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**. 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** so if an assignment to the **original list item** occurs then other synchronization must ensure that the assignment completes and the **original list item** is flushed to **memory**. In either case, to avoid **data**

1 **races**, concurrent reads or updates of the **original list item** must be synchronized with any update of
2 the **original list item** that occurs as a result of the **lastprivate** clause.

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 by an assignment that does not lexically
5 assign to the **list item** then the value assigned to 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 of the **affected loops** or lexically last **structured**
25 **block sequence** 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

- **distribute** Construct, see [Section 19.7](#)
- **do** Construct, see [Section 19.6.2](#)
- **for** Construct, see [Section 19.6.1](#)
- **loop** Construct, see [Section 19.8](#)
- **private** Clause, see [Section 7.3.3](#)
- **sections** Construct, see [Section 19.3](#)
- **simd** Construct, see [Section 18.4](#)
- **taskloop** Construct, see [Section 20.2](#)

7.3.6 is_device_ptr Clause

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

Arguments

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

Modifiers

Name	Modifies	Type	Properties
directive-name-modifier	<i>all arguments</i>	Keyword: <i>directive-name</i> (a 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 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](#). Each new [list item](#) has the [is-device-ptr attribute](#) and is initialized to the [device address](#) to which the [original list item](#) refers.

1 **Restrictions**

2 Restrictions to the `is_device_ptr` clause are as follows:

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

4 **Cross References**

- 5 • `dispatch` Construct, see [Section 15.7](#)
- 6 • `has_device_addr` Clause, see [Section 7.3.8](#)
- 7 • `target` Construct, see [Section 21.8](#)

8 **7.3.7 use_device_ptr Clause**

Name: <code>use_device_ptr</code>	Properties: all-data-environments, data-environment attribute, data-sharing attribute, device-associated, privatization
-----------------------------------	--

10 **Arguments**

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

12 **Modifiers**

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

14 **Directives**

15 `target_data`

16 **Semantics**

17 Each `list item` in the `use_device_ptr` clause results in a `new list item` that has the
18 `use-device-ptr` attribute and is a `device pointer` that refers to a `device address`. Since the
19 `use_device_ptr` clause is an `all-data-environments` clause, it has this effect even for `minimal`
20 `data environments`. The `device address` is determined as follows. A `list item` is treated as if a
21 `zero-offset assumed-size array` at the `storage location` to which the `list item` points is mapped by a
22 `map` clause on the `construct` with a `map-type` of `storage`. If a `matched candidate` is found for the
23 `assumed-size array` (see [Section 11.3](#)), the `new list item` refers to the `device address` that is the `base`
24 `address` of the `array section` that corresponds to the `assumed-size array` in the `device data`
25 `environment`. Otherwise, the `new list item` refers to the address stored in the `original list item`. All
26 references to the `list item` inside the `structured block` associated with the `construct` are replaced with
27 the `new list item` that is a `private` copy in the associated `data environment` on the `encountering`
28 `device`. Thus, the `use_device_ptr` clause is a `privatization` clause.

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` Construct, see [Section 21.7](#)

7.3.8 has_device_addr Clause

Name: <code>has_device_addr</code>	Properties: data-environment attribute, data-sharing attribute, device-associated, 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: <i>directive-name</i> (a <i>directive name</i>)	<i>unique</i>

Directives

`dispatch`, `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`. Inside the `construct`, the `list items` have the `has-device-addr` attribute. The `list items` may include `array sections`. If the `list item` is a `referencing variable`, the semantics of the `has_device_addr` clause apply to its `referenced pointee`. When the `clause` appears on the `target` construct, if the `device address` of a `list item` is not for the `device` on which the `target` region executes, accessing the `list item` inside the `region` results in `unspecified behavior`.

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 `array base` had been specified in the `clause` instead.

Fortran

- 1
- 2
- 3
- 4

2

- 34

C / C++

- 5

C / C++

Fortran

- 67

- 89

Fortran

O

- 1
- 2

3

4

5

6

7

8

9

20

21

2
3
4
5

1 `construct` are to the `corresponding list item` in the `device data environment`. The `list items` in a
2 `use_device_addr` clause may include `array sections` and `assumed-size arrays`. Since the
3 `use_device_addr` clause is an `all-data-environments` clause, it has this effect even for `minimal`
4 `data environments`.

5 If the `list item` is a `referencing variable`, the semantics of the `use_device_addr` clause apply to
6 its `referenced pointe`. A `private` copy of the `referring pointer` that refers to the corresponding
7 `referenced pointe` is used in place of the original `referring pointer` in the `structured block`.

▼ C / C++ ▼
8 If a `list item` is an `array section` that has a `base pointer`, all references to the `base pointer` inside the
9 `structured block` are replaced with a new pointer that contains the `base address` of the `corresponding`
10 `list item`. This conversion may be elided if no `corresponding list item` is present.
▲ C / C++ ▲

11 **Restrictions**

12 Restrictions to the `use_device_addr` clause are as follows:

- 13 • Each `list item` must have a `corresponding list item` in the `device data environment` or be
14 accessible on the `target device`.
- 15 • If a `list item` is an `array section` or an `array element`, the `array base` must be a `base language`
16 identifier.

17 **Cross References**

- 18 • `target_data` Construct, see [Section 21.7](#)

19 **7.4 saved Modifier**

20 **Modifiers**

Name	Modifies	Type	Properties
<i>saved</i>	<i>list</i>	Keyword: saved	<i>default</i>

22 **Clauses**

23 **`firstprivate`**

24 **Semantics**

25 If the `saved` modifier is present in a `data-sharing attribute clause` that is specified on a `replayable`
26 `construct` then its `original list items` of a `replay execution` are defined by the `saved data environment`
27 of the `replayable construct`. The `saved` modifier has no effect if specified in a `clause` that does not
28 appear on a `replayable construct`.

29 **Cross References**

- 30 • `firstprivate` Clause, see [Section 7.3.4](#)

8 Reduction and Induction Data Control

This chapter presents various **clauses** and **directives** for controlling **reduction** and **induction** data operations. The **reduction clauses** and the **induction clause** are **data-sharing attribute clauses** that can be used to perform **reductions** and **inductions** in parallel. These calculations involve the repeated application of **reduction operations** or **induction operations**. **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. The **linear clause** can be used as a shorthand to express an **induction operation** that corresponds to a linear recurrence. The **scan directive** is used to update a **reduction variable** to successive values of a **scan computation** in a loop.

8.1 OpenMP Reduction and Induction Identifiers

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

C

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

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

C

C++

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

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

C++

Fortran

A **reduction identifier** is either a **base language** identifier, a user-defined operator, an allowed intrinsic procedure name or one of the following operators: **+**, *****, **.and.**, **.or.**, **.eqv.** or **.neqv.**. The intrinsic procedure names that are allowed as **reduction identifiers** are **max**, **min**, **iand**, **ior** and **ieor**.

An **induction identifier** is either a **base language** identifier, a user-defined operator, or one of the following operators: **+** or *****.

Fortran

8.2 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:

- The execution of a **reduction expression** or **induction expression** must not result in the execution of a **construct** or an **OpenMP API routine**.
- A **declare target directive** must be specified for any **procedure** that can be accessed through any **reduction expression** or **induction expression** that respectively corresponds to a **reduction identifier** or an **induction identifier** that is used in a **target region**.

Fortran

- Any generic identifier, defined operation, defined assignment, or specific procedure used in a **reduction expression** or an **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 an **induction expression** must not have any alternate returns appear in the argument list.
- Any **procedure** called in the **region** of a **reduction expression** or an **induction expression** must be pure and must not reference any host-associated or use-associated **variables** nor any **variables** in a common block.

Fortran

8.2.1 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** are **OpenMP identifiers** for special **variables** that refer to storage of the type of the **list item** to which the **reduction** applies. If the **list item** is an array or **array section**, the **intrinsic identifiers** **omp_in** and **omp_out** each refer to an **array element** of that **list item**. Each of these **OpenMP identifiers** denotes one of the values to be combined before executing the **combiner expression**. The **omp_out intrinsic 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.

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

Fortran

Restrictions

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

- The only **variables** allowed in a **combiner expression** are **omp_in** and **omp_out**.

Fortran

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

Fortran

8.2.2 OpenMP Initializer Expressions

If the initialization of the **private** copies of **list items** in a **reduction clause** is not determined *a priori*, the syntax of an **initializer expression** is as follows:

C

```
omp_priv = initializer
```

C

or

C++

```
omp_priv initializer
```

C++

or

C / C++

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

C / C++

or

Fortran

```
omp_priv = expression
```

or

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

Fortran

In the definition of an **initializer expression**, the **intrinsic identifier** `omp_priv` represents a special **variable** that refers to the storage to be initialized. The **intrinsic identifier** `omp_orig` represents a special **variable** that can be used in an **initializer expression** to refer to the storage of the **original list item** to be reduced. The number of times that an **initializer expression** is evaluated and the order of these evaluations are unspecified.

C / C++

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

C / C++

Fortran

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

Fortran

C

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

C

C++

The *a priori* initialization of **private** copies that are created for **reductions** follows the **base language** rules for default initialization.

C++

Fortran

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

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

Fortran

Restrictions

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

- The only **variables** allowed in an **initializer expression** are **omp_priv** and **omp_orig**.
- An **initializer expression** must not modify the **variable omp_orig**.

C

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

C

C++

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

C++

Fortran

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

Fortran

8.2.3 OpenMP Inductor Expressions

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

Fortran

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

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**, that generic name 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

8.2.4 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**, the **intrinsic identifier omp_step** is a special **variable** that refers to the **step expression** and the **intrinsic identifier omp_idx** is a special **variable** that refers to the **collapsed iteration number**.

Restrictions

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

- The only **variables** allowed in a **collector expression** are **omp_step** and **omp_idx**.

8.3 Implicitly Declared OpenMP Reduction Identifiers

C / C++

Table 8.1 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 8.6.

TABLE 8.1: Implicitly Declared C/C++ Reduction Identifiers

Identifier	Initializer	Combiner
+	omp_priv = 0	omp_out += omp_in
*	omp_priv = 1	omp_out *= omp_in
&	omp_priv = ~ 0	omp_out &= omp_in
	omp_priv = 0	omp_out = omp_in
^	omp_priv = 0	omp_out ^= omp_in

table continued on next page

table continued from previous page

Identifier	Initializer	Combiner
<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</code> <i>representable number in the</i> <i>reduction list item type</i>	<code>omp_out = omp_in > omp_out ?</code> <code>omp_in : omp_out</code>
<code>min</code>	<code>omp_priv = Maximal</code> <i>representable number in the</i> <i>reduction list item type</i>	<code>omp_out = omp_in < omp_out ?</code> <code>omp_in : omp_out</code>



Table 8.2 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 8.2: 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</code> <i>representable number in the</i> <i>reduction list item type</i>	<code>omp_out = max(omp_in, omp_out)</code>
<code>min</code>	<code>omp_priv = Maximal</code> <i>representable number in the</i> <i>reduction list item type</i>	<code>omp_out = min(omp_in, omp_out)</code>

table continued on next page

table continued from previous page

Identifier	Initializer	Combiner
iand	omp_priv = All bits on	omp_out = iand(omp_in, omp_out)
ior	omp_priv = 0	omp_out = ior(omp_in, omp_out)
ieor	omp_priv = 0	omp_out = ieor(omp_in, omp_out)

Fortran

8.4 Implicitly Declared OpenMP Induction Identifiers

C / C++

Table 8.3 lists each [induction identifier](#) that is implicitly declared at every scope for arithmetic types and its corresponding [inductor expression](#) and [collector expression](#).

TABLE 8.3: Implicitly Declared C/C++ Induction Identifiers

Identifier	Inductor Expression	Collector Expression
+	omp_var = omp_var + omp_step	omp_step * omp_idx
*	omp_var = omp_var * omp_step	pow(omp_step, omp_idx)

C / C++

Fortran

Table 8.4 lists each [induction identifier](#) that is implicitly declared for numeric types and its corresponding [inductor expression](#) and [collector expression](#).

TABLE 8.4: Implicitly Declared Fortran Induction Identifiers

Identifier	Inductor Expression	Collector Expression
+	omp_var = omp_var + omp_step	omp_step * omp_idx
*	omp_var = omp_var * omp_step	omp_step ** omp_idx

Fortran

8.5 Properties Common to Reduction and induction Clauses

The **list items** that appear in a **reduction clause** or an **induction clause** may include **array sections** and **array elements**.

C++

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

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

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

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

C++

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

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

Fortran

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

Fortran

Restrictions

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

- Any **array element** must be specified at most once in all **list items** on a **directive**.
- For a **reduction identifier** or an **induction identifier** declared in a **declare_reduction** or a **declare_induction** directive, the **directive** must appear before its use in a **reduction clause** or **induction clause**.
- If a **list item** is an **array section**, it must not be a **zero-length array section** and its **array base** must be a **base language** identifier.
- If a **list item** is an **array section** or an **array element**, accesses to the elements of the array outside the specified **array section** or **array element** result in **unspecified behavior**.

C / C++

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

C / C++

Fortran

- The type, type parameters and rank of a **list item** that appears in a **reduction clause** must be valid for the **combiner expression** and the **initializer expression**. The type, type parameters and rank of a **list item** and of the **step expression** that appear in an **induction clause** must be valid for the **inductor expression**.
- A **list item** that appears in a **reduction clause** or an **induction clause** must be definable.
- A procedure pointer must not appear in a **reduction clause** or an **induction clause**.
- A pointer with the **INTENT (IN)** attribute must not appear in a **reduction clause** or an **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 an **inductor expression** must be associated at entry with 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 identifier** 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 identifier** 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 identifier` 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

8.6 Properties Common to All Reduction Clauses

The *clause-specification* of a `reduction clause` has a *clause-argument-specification* that specifies a `variable list` and has a required *reduction-identifier modifier* that specifies the `reduction identifier` to use for 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 8.5](#) as well as one of:

- A default constructor and an assignment operator that accepts a type *T* that can be implicitly constructed from an integer expression, such that the following requirement is valid:

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

- A single-argument constructor that accepts a type *T* that can be implicitly constructed from an integer expression, such that the following requirement is valid:

```
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` have the `reduction attribute` and so are `reduction variables`. These `reduction variables` are initialized with the `initializer` value of the `reduction identifier`. Any copies are combined using the `combiner` associated with the `reduction identifier`.

1 **Execution Model Events**

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

8 **Tool Callbacks**

9 A *thread* dispatches a registered *reduction* callback with *ompt_sync_region_reduction*
10 in its *kind* argument and *ompt_scope_begin* as its *endpoint* argument for each occurrence of a
11 *reduction-begin event* in that *thread*. Similarly, a *thread* dispatches a registered *reduction*
12 callback with *ompt_sync_region_reduction* in its *kind* argument and *ompt_scope_end*
13 as its *endpoint* argument for each occurrence of a *reduction-end event* in that *thread*. These
14 callbacks occur in the context of the *task* that performs the *reduction*.

15 **Restrictions**

16 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++ ▲

23 **Cross References**

- 24
 - **reduction** Callback, see [Section 40.7.6](#)
 - OMPT **scope_endpoint** Type, see [Section 39.27](#)
 - OMPT **sync_region** Type, see [Section 39.33](#)

8.7 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 8.5 and Section 8.6, apply to reduction-scoping clauses.

The number of copies created for each list item and the point at which those copies are initialized are determined by the particular reduction-scoping clause that appears on the construct. The point 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 the update of the original list item that occurs as a result of the reduction, which may occur after execution of the construct on which the reduction-scoping clause appears, for example, due to the use of a 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.

8.8 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 8.5 and Section 8.6, 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.

8.9 reduction-identifier Modifier

Modifiers

Name	Modifies	Type	Properties
reduction-identifier	all arguments	An OpenMP reduction identifier	required, ultimate

Clauses

in_reduction, reduction, task_reduction

Semantics

Reduction clauses use the *reduction-identifier* modifier to specify the reduction identifier for the clause. The reduction identifier determines the initializer expression and combiner expression to use for the reduction.

Cross References

- OpenMP Reduction and Induction Identifiers, see Section 8.1
- `in_reduction` Clause, see Section 8.12
- `reduction` Clause, see Section 8.10
- `task_reduction` Clause, see Section 8.11

8.10 reduction Clause

Name: <code>reduction</code>	Properties: data-environment attribute, data-sharing attribute, original list-item updating, 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>all arguments</i>	An OpenMP reduction identifier	<i>required, ultimate</i>
<i>reduction-modifier</i>	<i>list</i>	Keyword: default , inscan , task	<i>default</i>
<i>original-sharing-modifier</i>	<i>list</i>	Complex, name: original Arguments: <i>sharing</i> Keyword: default , private , shared (<i>default</i>)	<i>default</i>
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: <i>directive-name</i> (a directive name)	<i>unique</i>

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 8.7](#) and [Section 8.8](#). 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 the **clause** appears on a **worksharing construct** and the **original list item** is **private** in the **enclosing context** of that **construct**, the behavior is as if a **shared** copy (initialized with the **initializer** value) specific to the **worksharing region** is updated by combining its value with the values of the **private** copies created by the **clause**; once an **encountering thread** observes that all of those updates are completed, the **original list item** for that **thread** is then updated by combining its value with the value of the **shared** copy.

If the **original-sharing-modifier** is not present, the behavior is as if it were present with the **sharing** argument specified as **default**. If the **sharing** argument is specified as **default**, **original list items** are assumed to be **shared** in the **enclosing context** unless determined not to be **shared** according to the rules specified in [Section 7.1](#). If **shared** or **private** is specified as the **original-sharing-modifier** **sharing** argument, the **original list items** are assumed to be **shared** or **private**, respectively, in the **enclosing context**.

If **reduction-modifier** is not present or the **default reduction-modifier** is present, the behavior is as follows. For **parallel** and **worksharing constructs**, one or more **private** copies of each **list item** are created for each **implicit task**, as if the **private clause** had been used. For the **simd construct**, one or more **private** copies of each **list item** are created for each **SIMD lane**, as if the **private clause** had been used. For the **taskloop construct**, **private** copies are created according to the rules of the **reduction-scoping clause**. For the **teams construct**, one or more **private** copies of each **list item** are created for the **initial task** of each **team** in the **league**, as if the **private clause** had been used. For the **loop construct**, **private** copies are created and used in the **construct** according to the description and restrictions in [Section 7.2](#). At the end of a **region** that corresponds to a **construct** for which the **reduction clause** was specified, the **original list item** is updated by combining its original value with the final value of each of the **private** copies, using the **combiner** of the specified **reduction-identifier**.

If the **inscan reduction-modifier** is present, a **scan computation** is performed over updates to the **list item** performed in each **logical iteration** of the **affected loops** (see [Section 8.17](#)). The **list items** are **privatized** in the **construct** according to the description and restrictions in [Section 7.2](#). At the end of the **region**, each **original list item** is assigned the value described in [Section 8.17](#).

If the **task reduction-modifier** is present for a **parallel** or **worksharing construct**, then each **list item** is **privatized** according to the description and restrictions in [Section 7.2](#), and an unspecified number of additional **private** copies may be created to support **task reductions**. Any copies associated with the **reduction** are initialized before they are accessed by the **tasks** that participate in the **reduction**, which include all **implicit tasks** in the corresponding **region** and all participating **explicit tasks** that specify an **in_reduction clause** (see [Section 8.12](#)). After the end of the **region**, the **original list item** contains the result of the **reduction**.

Restrictions

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

- All restrictions common to all **reduction clauses**, as listed in [Section 8.5](#) and [Section 8.6](#), apply to this **clause**.
- For a given **construct** on which the **clause** appears, the lifetime of all **original list items** must extend at least until after the synchronization point at which the completion of the corresponding **region** by all participants in the **reduction** can be observed by all participants.
- If the **inscan reduction-modifier** is specified on a **reduction clause** that appears on a **worksharing construct** and an **original list item** is **private** in the **enclosing context** of the **construct**, the **private** copies must all have identical values when the **construct** is encountered.
- If the **reduction clause** appears on a **worksharing construct** and the *original-sharing-modifier* specifies **default** as its *sharing* argument, each **original list item** must be **shared** in the **enclosing context** unless it is determined not to be **shared** according to the rules specified in [Section 7.1](#).
- If the **reduction clause** appears on a **worksharing construct** and the *original-sharing-modifier* specifies **shared** or **private** as its *sharing* argument, the **original list items** must be **shared** or **private**, respectively, in the **enclosing context**.
- Each **list item** specified with the **inscan reduction-modifier** must appear as a **list item** in an **inclusive** or **exclusive clause** on a **scan directive** enclosed by the **construct**.
- If the **inscan reduction-modifier** is specified, a **reduction clause** without the **inscan reduction-modifier** must not appear on the same **construct**.
- A **list item** that appears in a **reduction clause** on a **work-distribution construct** for which the corresponding **region** binds to a **teams region** must be **shared** in the **teams region**.
- A **reduction clause** with the **task reduction-modifier** may only appear on a **parallel construct** or a **worksharing construct**, or a **compound construct** for which any of the aforementioned **constructs** is a **constituent construct** and neither **simd** nor **loop** are **constituent constructs**.
- A **reduction clause** with the **inscan reduction-modifier** may only appear on a **worksharing-loop construct** or a **simd construct**, or a **compound construct** for which any of the aforementioned **constructs** is a **constituent construct** and neither **distribute** nor **taskloop** is a **constituent construct**.
- The **inscan reduction-modifier** must not be specified on a **construct** for which the **ordered** or **schedule clause** is specified.
- A **list item** that appears in a **reduction clause** of the innermost enclosing **worksharing construct** or **parallel construct** must not be accessed in an **explicit task** generated by a **construct** unless an **in_reduction clause** with the same **list item** appears on that **construct**.

- The **task** *reduction-modifier* must not appear in a **reduction** clause if the **nowait** clause is specified on the same **construct**.

Fortran

- If the *original-sharing-modifier* for a **reduction** clause on a **worksharing** construct specifies **default** *sharing* and a **list item** in the **clause** either has a base pointer or is a dummy argument without the **VALUE** attribute, the **original list item** must refer to the same object for all **threads** of the **team** that execute the corresponding **region**.

Fortran

C / C++

- If the *original-sharing-modifier* specifies **default** as its *sharing* argument and a **list item** in a **reduction** clause on a **worksharing** construct has a reference type then that **list item** must bind to the same object for all **threads** of the **team**.
- A **variable** of **class type** (or array thereof) that appears in a **reduction** clause with the **inscan** *reduction-modifier* requires an accessible, unambiguous default constructor and copy assignment operator for the **class type**; the number of calls to them while performing the **scan computation** is unspecified.

C / C++

Cross References

- **do** Construct, see [Section 19.6.2](#)
- **for** Construct, see [Section 19.6.1](#)
- List Item Privatization, see [Section 7.2](#)
- **loop** Construct, see [Section 19.8](#)
- **ordered** Clause, see [Section 6.4.6](#)
- **parallel** Construct, see [Section 18.1](#)
- **private** Clause, see [Section 7.3.3](#)
- **scan** Directive, see [Section 8.17](#)
- **schedule** Clause, see [Section 19.6.3](#)
- **scope** Construct, see [Section 19.2](#)
- **sections** Construct, see [Section 19.3](#)
- **simd** Construct, see [Section 18.4](#)
- **taskloop** Construct, see [Section 20.2](#)
- **teams** Construct, see [Section 18.2](#)

8.11 task_reduction Clause

Name: <code>task_reduction</code>	Properties: data-environment attribute, data-sharing attribute, original list-item updating, 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>all arguments</i>	An OpenMP reduction identifier	required, ultimate
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: <i>directive-name</i> (a <i>directive name</i>)	unique

Directives

taskgroup

Semantics

The **task_reduction** clause is a **reduction-scoping clause**, as described in [Section 8.7](#), that specifies a **task reduction**. 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 8.5](#) and [Section 8.6](#), apply to this **clause**.

Cross References

- **taskgroup** Construct, see [Section 23.4](#)

8.12 in_reduction Clause

Name: <code>in_reduction</code>	Properties: data-environment attribute, data-sharing attribute, privatization, reduction participating
--	---

1 **Arguments**

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

3 **Modifiers**

Name	Modifies	Type	Properties
<i>reduction-identifier</i>	<i>all arguments</i>	An OpenMP reduction identifier	required, ultimate
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: <i>directive-name</i> (a <i>directive name</i>)	unique

5 **Directives**

6 **target**, **target_data**, **task**, **taskloop**

7 **Semantics**

8 The **in_reduction** clause is a reduction-participating clause, as described in Section 8.8, that
9 specifies that a **task** participates in a reduction. For a given **list item**, the **in_reduction** clause
10 defines a **task** to be a participant in a **task reduction** that is defined by an enclosing **region** for a
11 matching **list item** that appears in a **task_reduction** clause or a **reduction** clause with the
12 **task reduction-modifier**, where either:

- 13 1. The matching **list item** has the same **storage location** as the **list item** in the **in_reduction**
14 **clause**; or
- 15 2. A **private** copy, derived from the matching **list item**, that is used to perform the **task reduction**
16 has the same **storage location** as the **list item** in the **in_reduction** clause.

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

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

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

25 When specified on the **target_data** directive, the **in_reduction** clause has the
26 **all-data-environments** property.

27 **Restrictions**

28 Restrictions to the **in_reduction** clause are as follows:

- 29 • All restrictions common to all **reduction clauses**, as listed in Section 8.5 and Section 8.6,
30 apply to this **clause**.

- For each `list item`, a matching `list item` must exist that appears in a `task_reduction clause` or a `reduction clause` with the `task reduction-modifier` that is specified on a `construct` that corresponds to a `region` in which the `region` of the participating `task` is `closely nested`. 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` Construct, see [Section 21.8](#)
- `target_data` Construct, see [Section 21.7](#)
- `task` Construct, see [Section 20.1](#)
- `taskloop` Construct, see [Section 20.2](#)

8.13 induction Clause

Name: <code>induction</code>	Properties: data-environment attribute, data-sharing attribute, original list-item updating, 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	<i>required, ultimate</i>
<i>step-modifier</i>	<i>list</i>	Complex, name: step Arguments: <i>induction-step</i> expression of induction-step type (<i>region-invariant</i>)	<i>required</i>
<i>induction-modifier</i>	<i>list</i>	Keyword: relaxed , strict	<i>default</i>
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: <i>directive-name</i> (a <i>directive name</i>)	<i>unique</i>

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 7.3.3, except as otherwise specified. The **new list items** have the **induction** attribute.

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**.

▼ 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 an assignment statement.

If the **original list item** has the **POINTER** attribute, its update occurs as if by pointer assignment.

▲ Fortran ▲

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**.

1 **Restrictions**

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

- 3 • All restrictions listed in [Section 8.5](#) apply to this clause.
- 4 • The *induction-step* must not be an array or array section.
- 5 • If an array section or array element appears as a list item in an **induction** clause on a
- 6 worksharing construct, all threads of the team must specify the same storage location.
- 7 • None of the affected loops of a loop-nest-associated construct that has an **induction**
- 8 clause may be a non-rectangular loop.

9 C / C++

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

15 C / C++

15 **Cross References**

- 16 • **distribute** Construct, see [Section 19.7](#)
- 17 • **do** Construct, see [Section 19.6.2](#)
- 18 • **for** Construct, see [Section 19.6.1](#)
- 19 • List Item Privatization, see [Section 7.2](#)
- 20 • **private** Clause, see [Section 7.3.3](#)
- 21 • **simd** Construct, see [Section 18.4](#)
- 22 • **taskloop** Construct, see [Section 20.2](#)

23 **8.14 linear Clause**

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

25 **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: <i>directive-name</i> (a <i>directive name</i>)	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 7.3.3, except as noted. Additionally, each *new list item* has the *linear* attribute and so is a *linear variable*. 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 one.

When a **linear** clause is specified on a *loop-collapsing construct* and a *list item* is the *loop-iteration variable* of an *affected loop*, the effect is as if that *list item* had appeared in a **lastprivate** clause. Otherwise, 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:

- If a **reduction** clause with the **inscan** modifier also appears on the **construct**, only loop-iteration variables of affected loops may appear as list items in a **linear** clause.
- A *linear-modifier* may be specified as **ref** or **uval** only for **linear** clauses on **declare_simd** directives.
- 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.
- None of the affected loops of a loop-nest-associated construct that has a **linear** clause may be a non-rectangular loop.

C

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

C

C++

- 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.
- If *linear-modifier* is **ref** or **uval**, all list items must be of a reference type.
- If a list item in a **linear** clause on a worksharing construct has a reference type then it must bind to the same object for all threads of the team.
- If a list item in a **linear** clause that is specified on a **declare_simd** directive is of a reference type and *linear-modifier* is not **ref**, the difference between the value of the argument on exit from the function and its value on entry to the function must be the same for all SIMD lanes.

C++

Fortran

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

Fortran

Cross References

- **declare_simd** Directive, see [Section 15.8](#)
- **do** Construct, see [Section 19.6.2](#)
- **for** Construct, see [Section 19.6.1](#)
- **private** Clause, see [Section 7.3.3](#)
- **simd** Construct, see [Section 18.4](#)
- **taskloop** Construct, see [Section 20.2](#)

8.15 declare_reduction Directive

Name: declare_reduction	Association: unassociated
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 **declare_reduction** directive may alternatively be specified with **declare reduction** as the *directive-name*.

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 specified **reduction identifier** and **type-name list** 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 **type-name list**. If the **directive** specifies several types then the behavior is as if a **declare_reduction** directive was specified for each type. The visibility and accessibility of a **user-defined reduction** are the same as those of a **variable** declared at the same location in the program.

C++

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

C++

The **enclosing context** of the **combiner expression** specified by the **combiner** clause and of the **initializer expression** specified by the **initializer** clause is that of the **declare_reduction** directive. The **combiner expression** and the **initializer expression** must be correct in the **base language**, as if they were the body of a **procedure** defined at the same location in the program.

Fortran

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

If the specified **reduction identifier** is the same as the name of a user-defined operator or an extended operator, or the same as a generic name that is one of the allowed intrinsic procedures,

and if the operator or procedure name appears in an accessibility statement in the same module, the accessibility of the corresponding **declare_reduction directive** is determined by the accessibility attribute of the statement.

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

Fortran

Restrictions

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

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

C / C++

- A type name in a **declare_reduction directive** must not be a function type, an array type, a reference type, or a type qualified with **const**, **volatile** or **restrict**.

C / C++

Fortran

- If the length type parameter is specified for a type, it must be a constant, a colon (:) or an asterisk (*).
- If a type with a 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 8.15.1](#)
- OpenMP Combiner Expressions, see [Section 8.2.1](#)
- OpenMP Initializer Expressions, see [Section 8.2.2](#)
- OpenMP Reduction and Induction Identifiers, see [Section 8.1](#)
- **initializer** Clause, see [Section 8.15.2](#)

8.15.1 combiner Clause

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

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: <i>directive-name</i> (a <i>directive name</i>)	<code>unique</code>

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 8.15](#)
- OpenMP Combiner Expressions, see [Section 8.2.1](#)

8.15.2 initializer Clause

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

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: <i>directive-name</i> (a <i>directive name</i>)	<code>unique</code>

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 8.15](#)
- OpenMP Initializer Expressions, see [Section 8.2.2](#)

8.16 `declare_induction` Directive

Name: <code>declare_induction</code> Category: declarative	Association: unassociated Properties: pure
---	---

Arguments

`declare_induction`(*induction-specifier*)

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

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*

where *type-specifier-list* is defined as follows:

type-specifier-list := *type-specifier* | *type-specifier* , *type-specifier-list*
type-specifier := *typename-list-item* | *typename-pair*
typename-pair := (*typename-list-item* , *typename-list-item*)

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

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 *type-specifier-list*, with corresponding [step expressions](#) of the same type if the *type-specifier-list* does not specify a *typename-pair*. If the *type-specifier-list* specifies a *typename-pair* then the [induction identifier](#) can be used in an [induction clause](#) that lists that pair, in which case the [induction variable](#) and `omp_var` must be of the first type specified in the *typename-pair* while the corresponding [step expression](#) and `omp_step` must be of the second type in the *typename-pair*. The type of `omp_idx` is the type used for the [iteration count](#) of the [collapsed iteration space](#) of the [collapse-affected loops](#) of the [construct](#) on which the [induction clause](#) appears.

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

C++

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

C++

The **enclosing context** of the **inductor expression** specified by the **inductor clause** and of the **collector expression** specified by the **collector clause** is that of the **declare_induction directive**. The **inductor expression** and the **collector expression** must be correct in the **base language**, as if they were the body of a **procedure** defined at the same location in the program.

Fortran

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

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

Fortran

Restrictions

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

- An **induction identifier** must not be re-declared in the current scope for the same type or for a type that is compatible according to the **base language** rules.
- A **type-name list item** in the *type-specifier-list* must not declare a new type.

C / C++

- A type name in a **declare_induction directive** must not be a function type, an array type, a reference type, or a type qualified with **const**, **volatile** or **restrict**.

C / C++

Fortran

- A type name in a **declare_induction directive** must not be an enum type or an enumeration type.

Fortran

Cross References

- **collector** Clause, see [Section 8.16.2](#)
- OpenMP Collector Expressions, see [Section 8.2.4](#)
- OpenMP Inductor Expressions, see [Section 8.2.3](#)
- OpenMP Loop-Iteration Spaces and Vectors, see [Section 6.4.3](#)
- OpenMP Reduction and Induction Identifiers, see [Section 8.1](#)
- **inductor** Clause, see [Section 8.16.1](#)

8.16.1 inductor Clause

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

Arguments

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

Modifiers

Name	Modifies	Type	Properties
directive-name-modifier	<i>all arguments</i>	Keyword: <i>directive-name</i> (a 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 8.16](#)
- OpenMP Inductor Expressions, see [Section 8.2.3](#)

8.16.2 collector Clause

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

Arguments

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

Modifiers

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

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 8.16](#)
- OpenMP Collector Expressions, see [Section 8.2.4](#)

8.17 scan Directive

Name: scan Category: <i>subsidiary</i>	Association: <i>separating</i> Properties: <i>pure</i>
--	---

Separated directives

do, **for**, **simd**

Clauses

exclusive, **inclusive**, **init_complete**

Clause set

Properties: <i>unique</i> , <i>required</i> , <i>exclusive</i>	Members: exclusive , inclusive , init_complete
---	--

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 *structured block sequences* that represent different phases of a *scan computation*. The use of **scan** directives results in a *structured block sequence* that serves as an *input phase*, a *structured block sequence* that serves as a *scan phase*, and, optionally, a *structured block sequence* that serves as an *initialization phase*. The optional *initialization phase* begins the *collapsed iteration* by initializing *private variables* that can be used in the *input 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*.

The **clause** that is specified on the **scan directive** determines the phases of the **scan computation** that correspond to the **structured block sequences** that precede and follow the **directive**.

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

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

At the beginning of the **input phase** of each **collapsed iteration**, the **new list item** is either initialized with the value of the **initializer expression** of the **reduction-identifier** specified by the **reduction clause** on the **separated construct** or with the value of the **list item** in the **scan phase** of some **collapsed iteration**. The **update value** of a **new list item** is, for a given **collapsed iteration**, the value the **new list item** would have on completion of its **input phase** if it were initialized with the value of the **initializer expression**.

Let *orig-val* be the value of the **original list item** on entry to the **separated construct**. Let *combiner* be the **combiner expression** for the **reduction-identifier** specified by the **reduction clause** on the **construct**. Let u_i be the **update value** of a **list item** for **collapsed iteration** i . For **list items** that appear in an **inclusive clause** on the **scan directive**, at the beginning of the **scan phase** for **collapsed iteration** i the **new list item** is assigned the result of the operation $\text{PRESUM}_{\text{last}}(\text{combiner}, \text{orig-val}, u_0, \dots, u_i)$. For **list items** that appear in an **exclusive clause** on the **scan directive**, at the beginning of the **scan phase** for **collapsed iteration** $i = 0$ the **list item** is assigned the value *orig-val*, and at the beginning of the **scan phase** for **collapsed iteration** $i > 0$ the **list item** is assigned the result of the operation $\text{PRESUM}_{\text{last}}(\text{combiner}, \text{orig-val}, u_0, \dots, u_{i-1})$.

For **list items** that appear in an **inclusive clause**, at the end of the **separated construct**, the **original list item** is assigned the value of the **private** copy from the last **collapsed iteration** of the **affected loops** of the **separated construct**. For **list items** that appear in an **exclusive clause**, let k be the last **collapsed iteration** of the **affected loops** of the **separated construct**. At the end of the **separated construct**, the **original list item** is assigned the result of the operation $\text{PRESUM}_{\text{last}}(\text{combiner}, \text{orig-val}, u_0, \dots, u_k)$.

Restrictions

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

- The **separated construct** must have at most one **scan directive** with an **inclusive** or **exclusive clause** as a **separating directive**.
- The **separated construct** must have at most one **scan directive** with an **init_complete clause** as a **separating directive**.

- If specified, a **scan** directive with an **init_complete** clause must precede a **scan** directive with an **exclusive** clause that is a subsidiary directive of the same construct.
- The **affected loops** of the **separated construct** must all be perfectly nested loops.
- Each **list item** that appears in the **inclusive** or **exclusive** clause must appear in a **reduction** clause with the **inscan** modifier on the **separated construct**.
- Each **list item** that appears in a **reduction** clause with the **inscan** modifier on the **separated construct** must appear in a **clause** on the **scan** separating directive.
- Cross-iteration dependences across different **collapsed iterations** of the **separated construct** 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 immediately precedes a **scan** directive with an **inclusive** or **exclusive** clause to a statement in the **structured block sequence** that follows that **scan** directive must not exist, except for dependences for the **list items** specified in that **clause**.
- The **private** copy of a **list item** that appears in the **inclusive** or **exclusive** clause must not be modified in the **scan phase**.
- Any **list item** that appears in an **exclusive** clause must not be modified or used in the **initialization phase**.
- Statements in the **initialization phase** must only modify **private variables**. Any **private variables** modified in the **initialization phase** must not be used in the **scan phase**.

Cross References

- **do** Construct, see [Section 19.6.2](#)
- **exclusive** Clause, see [Section 8.17.2](#)
- **for** Construct, see [Section 19.6.1](#)
- **inclusive** Clause, see [Section 8.17.1](#)
- **init_complete** Clause, see [Section 8.17.3](#)
- **reduction** Clause, see [Section 8.10](#)
- **simd** Construct, see [Section 18.4](#)

8.17.1 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: <i>directive-name</i> (a <i>directive name</i>)	unique

Directives

scan

Semantics

The **inclusive** clause is used on a **scan** directive to specify that an inclusive scan computation is performed for each list item of the argument list. The structured block sequence that precedes the directive serves as the input phase of the inclusive scan computation while the structured block sequence that follows the directive serves as the scan phase of the inclusive scan computation. The list items that appear in an **inclusive** clause may include array sections and array elements.

Cross References

- **scan** Directive, see [Section 8.17](#)

8.17.2 exclusive Clause

Name: exclusive	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: <i>directive-name</i> (a <i>directive name</i>)	unique

Directives

scan

Semantics

The **exclusive** clause is used on a **scan** directive to specify an **exclusive scan computation** is performed for each **list item** of the **argument list**. The **structured block sequence** that follows the **directive** serves as the **input phase** of the **exclusive scan computation** while the **structured block sequence** that precedes the **directive** serves as the **scan phase** of the **exclusive scan computation**. The **list items** that appear in an **exclusive** clause may include **array sections** and **array elements**.

Cross References

- **scan** Directive, see [Section 8.17](#)

8.17.3 init_complete Clause

Name: <code>init_complete</code>	Properties: innermost-leaf, unique
---	---

Arguments

Name	Type	Properties
<i>create_init_phase</i>	expression of OpenMP logical type	constant, optional

Modifiers

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

Directives

scan

Semantics

The **init_complete** clause is used on a **scan** directive to demarcate the end of the **initialization phase** of an **exclusive scan computation**. The **structured block sequence** that precedes the **directive** serves as the **initialization phase** of the **exclusive scan computation** while the **structured block sequence** that follows the **directive** serves as the **scan phase** of the **exclusive scan computation**. If *create_init_phase* is not specified, the effect is as if *create_init_phase* evaluates to *true*.

Cross References

- **scan** Directive, see [Section 8.17](#)

9 Static-Lifetime Data Control

This chapter presents various [directives](#) and [clauses](#) for controlling the [data-environment attributes](#) of [variables](#) that have [static storage duration](#).

9.1 threadprivate Directive

Name: <code>threadprivate</code> Category: declarative	Association: explicit Properties: pure
---	---

Arguments

threadprivate (*list*)

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

Semantics

The [threadprivate](#) directive specifies that [variables](#) have the [threadprivate](#) attribute and therefore they 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 [variables](#) with [static storage duration](#) 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, whether it is initialized is unspecified. If it is initialized, its initialization occurs at an unspecified point in the program.

C++

The content of a [threadprivate variable](#) can change across a [task scheduling point](#) if the executing [thread](#) switches to another [task](#) that modifies the [variable](#). For more details on [task](#) scheduling, see [Section 1.2](#) and [Chapter 20](#).

In [parallel](#) regions, references by the [primary thread](#) are to the copy of the [variable](#) of the [thread](#) that encountered the [parallel](#) region.

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

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

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

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

C / C++

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

C / C++

Fortran

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

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

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

Fortran

C++

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

C++

Restrictions

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

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

C / C++

- Each **list item** must be a file-scope, namespace-scope, or static block-scope **variable**.
- No **list item** may have an incomplete type.
- The address of a **threadprivate variable** must not be an address constant.
- If the value of a **variable** referenced in an explicit initializer of a **threadprivate variable** is modified prior to the first reference to any instance of the **threadprivate variable**, the behavior is **unspecified**.

- A **threadprivate** directive for file-scope **variables** must appear outside any definition or declaration, and must lexically precede all references to any of the **variables** in its **argument list**.
- A **threadprivate** directive for namespace-scope **variables** must appear outside any definition or declaration other than the namespace definition itself and must lexically precede all references to any of the **variables** in its **argument list**.
- Each **variable** in the **argument list** of a **threadprivate** directive at file, namespace, or class scope must refer to a **variable** declaration at file, namespace, or class scope that lexically precedes the **directive**.
- A **threadprivate** directive for a static block-scope **variable** must appear in the scope of the **variable** and not in a nested scope. The **directive** must lexically precede all references to any of the **variables** in its **argument list**.
- Each **variable** in the **argument list** of a **threadprivate** directive in block scope must refer to a **variable** declaration in the same scope that lexically precedes the **directive**. The **variable** must have **static storage duration**.
- If a **variable** is specified in a **threadprivate** directive in one **compilation unit**, it must be specified in a **threadprivate** directive in every **compilation unit** in which it is declared.

C / C++

C++

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

C++

Fortran

- Each **list item** must be a named **variable** or a named common block; a named common block must appear between slashes.
- The **list** argument must not include any coarrays or associate names.

- The **threadprivate** directive must appear in the declaration section of a scoping unit in which the common block or **variable** is declared.
- If a **threadprivate** directive that specifies a common block name appears in one **compilation unit**, then such a directive must also appear in every other **compilation unit** that contains a **COMMON** statement that specifies the same name. It must appear after the last such **COMMON** statement in the **compilation unit**.
- If a **threadprivate variable** or a **threadprivate** common block is declared with the **BIND** attribute, the corresponding C entities must also be specified in a **threadprivate directive** in the C program.
- A **variable** may only appear as an argument in a **threadprivate** directive in the scope in which it is declared. It must not be an element of a common block or appear in an **EQUIVALENCE** statement.
- A **variable** that appears as an argument in a **threadprivate** directive must be declared in the scope of a module or have the **SAVE** attribute, either explicitly or implicitly.
- The effect of an access to a **threadprivate variable** in a **DO CONCURRENT** construct is unspecified.

Fortran

Cross References

- **copyin** Clause, see [Section 10.1](#)
- *dyn-var* ICV, see [Table 3.1](#)
- **order** Clause, see [Section 18.3](#)
- Determining the Number of Threads for a **parallel** Region, see [Section 18.1.1](#)

9.2 groupprivate Directive

Name: groupprivate Category: declarative		Association: explicit Properties: pure
Arguments		
groupprivate(list)		
Name	Type	Properties
list	list of variable list item type	default
Clauses		
device_type		

Semantics

The **groupprivate** directive specifies that **list items** have the **groupprivate attribute** and therefore they 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*.
- 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**.
- 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.

C / C++

C++

- 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**.
- A **groupprivate variable** must not have an incomplete type or a reference type.

C++

Fortran

- Each **list item** must be a named **variable** or a named common block; a named common block must appear between slashes.
- The *list* argument must not include any coarrays or associate names.
- The **groupprivate directive** must appear in the declaration section of a scoping unit in which the common block or **variable** is declared.
- 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**.
- 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.
- 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.
- A **variable** that appears as a **list item** in a **groupprivate directive** must be declared in the scope of a module or have the **SAVE** attribute, either explicitly or implicitly.
- The effect of an access to a **groupprivate variable** in a **DO CONCURRENT** construct is **unspecified**.

Fortran

Cross References

- **device_type** Clause, see [Section 21.1](#)
- **local** Clause, see [Section 9.3](#)

9.3 local Clause

Name: <code>local</code>	Properties: <code>data-environment</code> attribute
---------------------------------	--

1 **Arguments**

2

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

3 **Modifiers**

4

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

5 **Directives**

6 `declare_target`

7 **Semantics**

8 The `local` clause specifies that each `list item` has the `device-local` attribute. A reference to a `list`
9 `item` on a given `device` will refer to a copy of the `list item` that is a `device-local variable` and is in
10 `memory` associated with the `device`.

11 **Cross References**

- 12
 - `declare_target` Directive, see [Section 15.9.1](#)

13 **9.4 enter Clause**

14

Name: <code>enter</code>	Properties: data-environment attribute, data-mapping attribute
---------------------------------	---

15 **Arguments**

16

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

17 **Modifiers**

18

Name	Modifies	Type	Properties
<i>automap-modifier</i>	<i>list</i>	Keyword: automap	<i>default</i>
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: <i>directive-name</i> (a directive name)	unique

19 **Directives**

20 `declare_target`

21 **Semantics**

22 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 **devices** to which the **directive** of the **clause** applies.

Fortran

Whether a device-specific version is created is **implementation defined** when the **clause** is in a different **compilation unit** than the definition of the procedure but both are in the same **translation unit**.

Fortran

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, unless otherwise specified.

If a **list item** is a **referencing variable**, the effect of the **enter** clause applies to its **referring pointer**.

Fortran

If a **list item** is an allocatable **variable**, the **automap-modifier** is present, and the **variable** is allocated by an **ALLOCATE** statement or deallocated by a **DEALLOCATE** statement where the **enter** clause is visible, the behavior is as follows:

- Upon allocation due to the **ALLOCATE** statement, the **list item** is mapped to the **device data environment** of the default **device** as if it appeared as a **list item** in a **map** clause on a **target_enter_data** directive; and
- Immediately prior to the deallocation due to the **DEALLOCATE** statement, the **list item** is removed from the **device data environment** of the default **device** as if it appeared as a **list item** in a **map** clause with the **delete-modifier** on a **target_exit_data** directive.

Fortran

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`.

	C / C++	
• The <i>automap-modifier</i> must not be present.	C / C++	
	Fortran	
• If the <i>automap-modifier</i> is present, each <code>list item</code> must be an allocatable <code>variable</code> .	Fortran	

Cross References

- `declare_target` Directive, see [Section 15.9.1](#)

9.5 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: <i>directive-name</i> (a <code>directive name</code>)	<code>unique</code>

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 11.1](#).

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 15.9.1](#)
- Implicit Data-Mapping Attribute Rules, see [Section 11.1](#)

10 Data-Copying Control

This chapter 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`.

10.1 `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: <i>directive-name</i> (a <code>directive name</code>)	<code>unique</code>

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 an assignment statement. If the **original list item** does not have the **POINTER** attribute but 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.

Fortran

Cross References

- **parallel** Construct, see [Section 18.1](#)

10.2 copyprivate Clause

Name: copyprivate	Properties: innermost-leaf, end-clause, data copying
--------------------------	--

1 **Arguments**

2

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

3 **Modifiers**

4

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

5 **Directives**

6 **single**

7 **Semantics**

8 The **copyprivate** clause provides a mechanism to use a **private variable** to broadcast a value
9 from the **data environment** of one **implicit task** to the **data environments** of the other **implicit tasks**
10 that belong to the innermost enclosing **parallel region**. The effect of the **copyprivate** clause on
11 the specified **list items** occurs after the execution of the **structured block** associated with the
12 **construct** on which the **clause** is specified, and before any of the **threads** in the **team** have left the
13 **barrier** at the end of the **construct**. To avoid **data races**, concurrent reads or updates of the **list item**
14 must be synchronized with the update of the **list item** that occurs as a result of the **copyprivate**
15 **clause** if, for example, the **nowait** clause is used to remove the **barrier**.

16 ▼────────────────── C / C++ ───────────────────▼
17 In all other **implicit tasks** that belong to the **parallel region**, each specified **list item** becomes defined
18 with the value of the **corresponding list item** in the **implicit task** associated with the **thread** that
19 executed the **structured block**. For **variables** of non-array type, the definition occurs by copy
20 assignment. For an array of elements of non-array type, each element is copied by copy assignment
21 from an element of the array in the **data environment** of the **implicit task** that is associated with the
22 **thread** that executed the **structured block** to the corresponding element of the array in the **data**
 environment of the other **implicit tasks**.

23 ▲────────────────── C / C++ ───────────────────▲
24 ▼────────────────── C++ ───────────────────▼
 For **class types**, a copy assignment operator is invoked. The order in which copy assignment
 operators for different **variables** of **class type** are called is unspecified.
 ▲────────────────── C++ ───────────────────▲

Fortran

If a **list item** does not have the **POINTER** attribute then, in all other **implicit tasks** that belong to the **parallel region**, the **list item** becomes defined as if by an assignment statement with the value of the **corresponding list item** in the **implicit task** that is associated with the **thread** that executed the **structured block**.

If the **list item** has the **POINTER** attribute then, in all other **implicit tasks** that belong to the **parallel region**, the **list item** receives, as if by pointer assignment, the same association status as the **corresponding list item** in the **implicit task** that is associated with the **thread** that executed the **structured block**.

The order in which any final subroutines for different **variables** of a finalizable type are called is unspecified.

Fortran

Restrictions

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

- All **list items** that appear in a **copyprivate** clause must be either **threadprivate** or **private** in the **enclosing context**.

C++

- A **variable** of **class type** (or array thereof) that appears in a **copyprivate** clause requires an accessible unambiguous copy assignment operator for the **class type**.

C++

Fortran

- A common block that appears in a **copyprivate** clause must be **threadprivate**.
- Pointers with the **INTENT (IN)** attribute 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.

Fortran

Cross References

- List Item Privatization, see [Section 7.2](#)
- **single** Construct, see [Section 19.1](#)
- **threadprivate** Directive, see [Section 9.1](#)

11 Data-Mapping Control

This chapter 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 storage duration** 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.

11.1 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 7.1.1. 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 **compound 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 25.2 specifies this behavior.
- If a **variable** is the **base variable** of a **list item** in an **in_reduction** clause on a **target construct** then it is treated as if the **list item** had appeared in a **map** clause with a *map-type* of **tofrom** and an *always-modifier*.
- If a **defaultmap** clause is present for the category of the **variable** and specifies an implicit behavior other than **default**, the **data-mapping attribute** or **data-sharing attribute** is determined by that **clause**.

C++

- If the **target construct** is within a class non-static member function, and a **variable** is an accessible data member of the object for which the non-static member function is invoked, the **variable** is treated as if the **this[:1]** expression had appeared in a **map clause** with a **map-type** of **tofrom**. Additionally, if the **variable** is of type pointer or reference to pointer, it is also treated as if it is the **array base** of a **zero-offset assumed-size array** that appears in a **map clause** with the **storage map-type**.
- If the **this** keyword is referenced inside a **target construct** within a class non-static member function, it is treated as if the **this[:1]** expression had appeared in a **map clause** with a **map-type** of **tofrom**.

C++

C / C++

- A **variable** that is of type pointer, but is neither a pointer to function nor (for C++) a pointer to a member function, is treated as if it is the **array base** of a **zero-offset assumed-size array** that appears in a **map clause** with the **storage map-type**.

C / C++

C++

- A **variable** that is of type reference to pointer, but is neither a reference to pointer to function nor a reference to a pointer to a member function, is treated as if it is the **array base** of a **zero-offset assumed-size array** that appears in a **map clause** with the **storage map-type**.

C++

Fortran

- If a **compound target construct** is associated with a **DO CONCURRENT** loop, a **variable** that has **REDUCE** or **SHARED** locality in the loop is treated as if it had appeared in a **map clause** with a **map-type** of **tofrom**.

Fortran

- If a **variable** is not a **scalar variable** then it is treated as if it had appeared in a **map clause** with a **map-type** of **tofrom**.

Fortran

- If a **scalar variable** has the **TARGET**, **ALLOCATABLE** or **POINTER** attribute then it is treated as if it had appeared in a **map clause** with a **map-type** of **tofrom**.
- If a **variable** is an assumed-type **variable** then it is treated as if it had appeared in a **map clause** with a **map-type** of **storage**.
- A **procedure pointer** is treated as if it had appeared in a **firstprivate clause**.

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 7.1.1).

11.2 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>)	<i>unique</i>

Clauses

from, **map**, **to**

Semantics

Mapper identifiers can be used to identify uniquely 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 **storage 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

- **declare_mapper** Directive, see [Section 11.5](#)
- **from** Clause, see [Section 12.2](#)
- Data-Motion Control, see [Chapter 12](#)
- **map** Clause, see [Section 11.3](#)
- **to** Clause, see [Section 12.1](#)

11.3 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>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 list of iterator specifier list item type (<i>default</i>)	unique
<i>ref-modifier</i>	<i>all arguments</i>	Keyword: ref_ptee , ref_ptr , ref_ptr_ptee	unique
<i>attach-modifier</i>	<i>all arguments</i>	Complex, name: attach Arguments: attachment-mode Keyword: always , auto , never (<i>default</i>)	unique
<i>map-type</i>	<i>all arguments</i>	Keyword: from , storage , to , tofrom	<i>default</i>
<i>present-modifier</i>	<i>locator-list</i>	Keyword: present	map-type-modifying
<i>always-modifier</i>	<i>locator-list</i>	Keyword: always	map-type-modifying
<i>delete-modifier</i>	<i>locator-list</i>	Keyword: delete	map-type-modifying
<i>close-modifier</i>	<i>locator-list</i>	Keyword: close	map-type-modifying
<i>self-modifier</i>	<i>locator-list</i>	Keyword: self	map-type-modifying
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: <i>directive-name</i> (a <i>directive name</i>)	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. The list items that appear on a **map** clause may include array sections, assumed-size arrays, and structure elements. A list item in a **map** clause may reference any *iterator-identifier* defined in its *iterator* modifier. A list item may appear more than once in the **map** clauses that are specified on the same directive.

Fortran

The list items that appear on a **map** clause may also include assumed-type variables and procedure pointers. For the purposes of the **map** clause, an assumed-type variable list item is treated as if it is an assumed-size array of unspecified type.

Fortran

C / C++

If a list item is a zero-length array section that has a single array subscript, the behavior is as if the list item is an assumed-size array that is instead mapped with the **storage map-type**.

C / C++

Section 11.3.1 describes the various **map** types determined by the *map-type* modifier in a **map** clause, and Section 11.3.2 describes the process of **map-type** decay for determining an output map type.

Section 11.3.3 describes the effect of the *ref-modifier*, when specified, and otherwise describes the behavior of the **map** clause for list items that are referencing variables.

Section 11.3.4 describes the effect of the *attach-modifier*, when specified, and otherwise describes when pointer attachment occurs for list items in a **map** clause.

Section 11.3.5 describes the effect of the **map** clause when an assumed-size array is specified as a list item. If a matched candidate is not found for such list items and the *present-modifier* is specified, runtime error termination is performed. Otherwise, if a matched candidate is not found then no observable behavior that is described for the **map** clause in this section or its subsections occur for the list item.

Section 11.3.6 describes the effect of the **map** clause when an array, array section or structure is specified as a list item.

Section 11.3.7 describes the mapping of storage blocks based on the **map** clauses and their *map-type* and **map-type-modifying** modifiers (i.e., *present-modifier*, *always-modifier*, *delete-modifier*, *close-modifier*, and *self-modifier*) on a given data-mapping construct, as well as the creation of any new list items within the corresponding storage blocks.

Section 11.3.8 describes additional mapping rules for [list items](#) that are implicitly mapped according to [Section 11.1](#).

Execution Model Events

The *target-map* event occurs in a [thread](#) that executes the outermost [region](#) that corresponds to an encountered [device construct](#) with a [map clause](#), after the *target-task-begin* event for the [device construct](#) and before any [mapping operations](#) are performed. The *target-data-op-begin* event occurs before a [thread](#) initiates a data operation on the [target device](#) that is associated with a [map clause](#), in the outermost [region](#) that corresponds to the encountered [construct](#). The *target-data-op-end* event occurs after a [thread](#) initiates a data operation on the [target device](#) that is associated with a [map clause](#), in the outermost [region](#) that corresponds to the encountered [construct](#).

Tool Callbacks

A [thread](#) dispatches one or more registered [target_map_emi](#) callbacks for each occurrence of a *target-map* event in that [thread](#). The [callback](#) occurs in the context of the [target task](#). A [thread](#) dispatches a registered [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 [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](#).

Restrictions

Restrictions to the [map clause](#) are as follows:

- Two [list items](#) of the [map clauses](#) on the same [construct](#) must not share [original storage](#) unless one of the following is true: they are the same [list item](#), one is the [containing structure](#) of the other, at least one is an [assumed-size array](#), or at least one is implicitly mapped due to the [list item](#) also appearing in a [use_device_addr](#) clause.
- If the same [list item](#) appears more than once in [map clauses](#) on the same [construct](#), the [map clauses](#) must specify the same [mapper](#) modifier.
- A [variable](#) that is a [groupprivate variable](#) or a [device-local variable](#) must not appear as a [list item](#) in a [map clause](#).
- If an expression that is used to form a [list item](#) in a [map clause](#) contains an [iterator](#) identifier that is defined by an [iterator modifier](#), the [list item](#) instances that would result from different values of the [iterator](#) must not have the same [containing array](#) and must not have [base pointers](#) that share [original storage](#).
- If multiple [list items](#) are explicitly mapped on the same [construct](#) and have the same [containing array](#) or have [base pointers](#) that share [original storage](#), and if any of the [list items](#) do not have [corresponding list items](#) that are [present](#) in the [device data environment](#) prior to a [task](#) encountering the [construct](#), then the [list items](#) must refer to the same [array elements](#) of either the [containing array](#) or the [implicit array](#) of the [base pointers](#).
- Each [list item](#) must have a [mappable type](#).

- **Handles** for **memory spaces** and **memory allocators** must not appear as **list items** in a **map clause**.

C++

- No type mapped through a reference may contain a reference to its own type, or any references to types that could produce a cycle of references.

C++

C / C++

- A **list item** cannot be a **variable** that is a member of a **structure** of a union type.
- A bit-field cannot appear in a **map clause**.

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** must not be allocated, deallocated or reshaped while the **corresponding list item** has allocated storage.
- **List items** must not be complex part designators.

Fortran

Cross References

- **attach-modifier** Modifier, see [Section 11.3.4](#)
- **declare_mapper** Directive, see [Section 11.5](#)
- Array and Structure Mapping, see [Section 11.3.6](#)
- Array Sections, see [Section 5.2.5](#)
- Assumed-Size Array Mapping, see [Section 11.3.5](#)
- Implicit Data-Mapping, see [Section 11.3.8](#)
- Map Type Decay, see [Section 11.3.2](#)
- Storage Block Mapping, see [Section 11.3.7](#)
- **iterator** Modifier, see [Section 5.2.6](#)
- Mapper Identifiers and **mapper** Modifiers, see [Section 11.2](#)
- *map-type* Modifier, see [Section 11.3.1](#)
- **ref-modifier** Modifier, see [Section 11.3.3](#)
- OMPT **scope_endpoint** Type, see [Section 39.27](#)

- **target** Construct, see [Section 21.8](#)
- **target_data** Construct, see [Section 21.7](#)
- **target_data_op_emi** Callback, see [Section 41.7](#)
- **target_enter_data** Construct, see [Section 21.5](#)
- **target_exit_data** Construct, see [Section 21.6](#)
- **target_map_emi** Callback, see [Section 41.9](#)
- **target_update** Construct, see [Section 21.9](#)

11.3.1 *map-type* Modifier

Modifiers

Name	Modifies	Type	Properties
<i>map-type</i>	<i>all arguments</i>	Keyword: from , storage , to , tofrom	<i>default</i>

Clauses

[map](#)

Additional information

The value **alloc** may be used on [map-entering constructs](#) and the value **release** may be used on [map-exiting constructs](#) with identical meaning to the value **storage**.

Semantics

The *map-type* modifier determines the type of [mapping operations](#) that are performed as a result of the [clause](#) on which it appears. All [mapping operations](#) update the reference count of [corresponding storage](#) in a [device data environment](#), which may entail creation or removal of that storage. The **storage** *map-type* never includes an assignment operation. If the *map-type* is **to**, **from**, or **tofrom**, the *map-type* is an [assigning map type](#) and may include an assignment operation to or from the [target device](#).

The *map-type* is a [map-entering map type](#) if it is **to**, **tofrom**, or **storage**. The *map-type* is a [map-exiting map type](#) if it is **from**, **tofrom**, or **storage**. If the *map-type* is a [map-entering map type](#), the [clause](#) on which the *map-type* appears is a [map-entering clause](#). If the *map-type* is a [map-exiting map type](#), the [clause](#) on which the *map-type* appears is a [map-exiting clause](#).

When a *map-type* is not specified for a [clause](#) on which it may be specified, the *map-type* defaults to **storage** if the [delete-modifier](#) is present on the [clause](#) or if the [list item](#) for which the *map-type* is not specified is an [assumed-size array](#). Otherwise, the *map-type* defaults to **tofrom** if a *map-type* is not specified for a [clause](#) on which it may be specified, unless otherwise specified.

TABLE 11.1: Map-Type Decay of Map Type Combinations

	storage	to	from	tofrom
storage	storage	storage	storage	storage
to	storage	to	storage	to
from	storage	storage	from	from
tofrom	storage	to	from	tofrom

Fortran

When a *map-type* is not specified for a *clause* on which it may be specified, the *map-type* defaults to **storage** if the *list item* for which the *map-type* is not specified is an assumed-type *variable*.

Fortran

Restrictions

Restrictions to the *map-type modifier* are as follows:

- If the *clause* on which the *map-type* appears is specified on a *construct* that is *map-entering* but not *map-exiting*, the *map-type* must be *map-entering*.
- If the *clause* on which the *map-type* appears is specified on a *construct* that is *map-exiting* but not *map-entering*, the *map-type* must be *map-exiting*.

Cross References

- **map** Clause, see [Section 11.3](#)

11.3.2 Map Type Decay

Map-type decay is a process that derives an *output map type* from a given *input map type* according to an *underlying map type*. This process is defined by Table 11.1, where the *output map type* is shown at the row and column that corresponds to the *underlying map type* and *input map type*, respectively. When *map-type decay* determines the *map-type modifier* to apply for a **map** clause on a *data-mapping constituent directive* of a *composite construct*, the *input map type* is given by the *map-type modifier* specified by the **map** clause on the *composite construct* and the *underlying map type* is respectively **to** or **from** for a *map-entering constituent directive* or a *map-exiting constituent directive*. When *map-type decay* is applied by an invoked *mapper*, the *underlying map type* is given by the *map-type modifier* of the **map** clause specified by the *mapper* and the *input map type* is given by the *map-type modifier* of the **map** clause that invokes the *mapper*.

11.3.3 ref-modifier Modifier

Modifiers

Name	Modifies	Type	Properties
<i>ref-modifier</i>	<i>all arguments</i>	Keyword: ref_ptee , ref_ptr , ref_ptr_ptee	unique

Clauses

map

Semantics

The *ref-modifier* for a given **map** clause indicates how to interpret the identity of a **list item** argument of that **clause**. If the **ref_ptr** or **ref_ptr_ptee** *ref-modifier* is specified, the semantics of the **clause** apply to the **referring pointer** of the **referencing variable**. If the **ref_ptee** or **ref_ptr_ptee** *ref-modifier* is specified and a **referenced pointee** of the **referencing variable** exists, the semantics of the **clause** apply to the **referenced pointee**.

If the *ref-modifier* is not specified and a **list item** is a **referencing variable**, unless otherwise specified the effect of the **map** clause is applied to its **referring pointer** and, if a **referenced pointee** exists, its **referenced pointee**.

C++

If a **list item** is a reference and it does not have a **containing structure** then the **map** clause is applied only to its **referenced pointee**.

C++

Fortran

If a **list item** in a **map** clause is an associated pointer that is **attach-ineligible** or the pointer is the **base pointer** of another **list item** in a **map** clause on the same **construct** then the effect of the **map** clause does not apply to its pointer target.

Fortran

For the purposes of the **map** clause when applied to a **list item** that has a **base referencing variable**, the storage of the **referring pointer** of the **referencing variable** does not change for the lifetime of that **referencing variable**. When the **referencing variable** is passed as an argument to a **procedure**, the storage of the **referring pointer** for the argument inside the **procedure** may differ from the storage of the **referring pointer** at the call site.

Restrictions

Restrictions to the *ref-modifier* are as follows:

- A **list item** that appears in a **clause** with the *ref-modifier* must be a **referencing variable**.
- A **list item** that appears in a **clause** for which the *ref-modifier* is specified must have a **containing structure**.

C / C++

C / C++

Cross References

- **map** Clause, see [Section 11.3](#)

11.3.4 attach-modifier Modifier

Modifiers

Name	Modifies	Type	Properties
<i>attach-modifier</i>	<i>all arguments</i>	Complex, name: attach Arguments: <i>attachment-mode</i> Keyword: always , auto , never (<i>default</i>)	unique

Clauses

map

Semantics

The *attach-modifier* for a given **map** clause determines whether a **pointer attachment** occurs for the **base pointer** or **base referring pointer** of any of its **list items** on a **map-entering construct**. If **pointer attachment** occurs, the corresponding **base pointer** or **base referring pointer** becomes an **attached pointer** to the **corresponding list item** in the **device data environment** via **corresponding pointer initialization**. Any expression that uses the pointer to refer to some part of the **list item** will, if it occurs in a **target region**, instead refer to that part of the **corresponding list item**.

If the *attachment-mode* of the *attach-modifier* is **never** or if the pointer is **attach-ineligible**, **pointer attachment** will not occur. Otherwise, **pointer attachment** will occur if a **corresponding pointer** to the **base pointer** or **base referring pointer** of the **list item** exists in the **device data environment** after all **mappable storage blocks** specified by the **map clauses** on the **construct** are mapped, if:

- The *attachment-mode* is **always**; or
- The *attachment-mode* is **auto** and either a **new list item** or the **corresponding pointer** is created in the **device data environment** on entry to the **region**.

If the *attach-modifier* is not present on a **map clause** on a **map-entering construct**, the behavior is as if it is specified with the **auto attachment-mode** for any **list item** that has a **base pointer** or **base referring pointer**.

Fortran

If a **list item** is a **procedure pointer**, it is **attach-ineligible**.

Whether **pointer attachment** occurs for components of a derived type that are mapped as a result of a **predefined default mapper** is covered in [Section 11.3.6](#).

Fortran

Restrictions

Restrictions to the *attach-modifier* are as follows:

- A **map clause** with an *attach-modifier* can only be specified on a **map-entering construct** or a **declare_mapper directive**.
- A **list item** that appears in a **map clause** with the *attach-modifier* must have a **base pointer** or **base referring pointer**.

Cross References

- Array and Structure Mapping, see [Section 11.3.6](#)
- **map** Clause, see [Section 11.3](#)

11.3.5 Assumed-Size Array Mapping

For supporting the mapping of **list items** in a **map** clause that are **assumed-size arrays**, a set of **matchable candidates** is tracked as follows. When a **map** clause **list item** that is not an **assumed-size array** is mapped on a **map-entering construct** and **corresponding storage** is created in the **device data environment** on entry to the **region**, the **list item** becomes a **matchable candidate** with an associated **starting address**, **ending address**, and **base address** that define its **mapped address range** and **extended address range**. The current set of **matchable candidates** consists of any **map** clause **list item** on the **construct** that is a **matchable candidate** and all **matchable candidates** that were previously mapped and are still mapped.

A **list item** that is an **assumed-size array** is treated as if a **substitute array section**, with an **array base**, lower bound and length determined as follows, is substituted in its place if a **matched candidate** from the set of **matchable candidates** is found. If the **assumed-size array** is an **array section**, the **array base** of the **substitute array section** is the same as for the **assumed-size array**; otherwise, the **array base** is the **assumed-size array**. If the **mapped address range** of a **matchable candidate** includes the first **storage location** of the **assumed-size array**, it is a **matched candidate**. If a **matchable candidate** does not exist for which the **mapped address range** includes the first **storage location** of the **assumed-size array** then a **matchable candidate** is a **matched candidate** if its **extended address range** includes the first **storage location** of the **assumed-size array**. If multiple **matched candidates** exist, an arbitrary one of them is the found **matched candidate**. The **substitute array section** is treated as if its storage is identical to the storage of the found **matched candidate**.

Restrictions

Additional restrictions to the **map** clause for mapping of **assumed-size arrays** are as follows:

- If a **list item** is an **assumed-size array**, multiple **matched candidates** must not exist unless they are subobjects of the same **containing structure**.
- If a **list item** is an **assumed-size array**, the **map-type** must be **storage**.

11.3.6 Array and Structure Mapping

If a **list item** is an **array** or **array section**, the **array elements** become implicit **list items** with the same **modifiers** (including the **map-type**) specified in the **clause**. If the **array** or **array section** is implicitly mapped and **corresponding storage** exists in the **device data environment** prior to a **task** encountering the **construct** on which the **map** clause appears, only those **array elements** that have **corresponding storage** are implicitly mapped.

If a *mapper modifier* is not present, the behavior is as if a *mapper modifier* was specified with the **default** parameter. The map behavior of a *list item* in a **map clause** is modified by a visible *user-defined mapper* (see Section 11.5) if the *mapper-identifier* of the *mapper modifier* is defined for a *base language* type that matches the type of the *list item*. Otherwise, the *predefined default mapper* for the type of the *list item* applies. The effect of the *mapper modifier* is to remove the *list item* from the **map clause** and to apply the *clauses* specified in the declared *mapper* to the *construct* on which the **map clause** appears. In the *clauses* applied by the *mapper*, references to *var* are replaced with references to the *list item* and the *map-type* is replaced with the *output map type* that is determined according to the rules of *map-type decay*. If any *modifier* with the *map-type-modifying property* appears in the **map clause** then the effect is as if that *modifier* appears in each **map clause** specified in the declared *mapper*.

Fortran

If a component of a derived type *list item* is a **map clause list item** that results from the *predefined default mapper* for that derived type, and if the derived type component is not an explicit *list item* or the *array base* of an explicit *list item* in a **map clause** on the *construct* then:

- If it has the **POINTER** attribute, it is *attach-ineligible*; and
- If it has the **ALLOCATABLE** attribute and an allocated allocation status, and it is *present* in the *device data environment* when the *construct* is encountered, the **map clause** may treat its allocation status as if it is unallocated if the corresponding component does not have allocated storage.

Fortran

C++

If a *list item* has a closure type that is associated with a lambda expression, it is mapped as if it has a *structure* type. For each *variable* that is captured by reference by the lambda expression, the behavior is as if the closure type contains a non-static data member that is a reference to that *variable* unless otherwise specified. If a *variable* that is captured by reference is a reference that binds to an object with *static storage duration*, a corresponding non-static data member might not exist in the closure type. For the *corresponding list item* of closure type, references in the body of the lambda expression to a *variable* that is captured by reference refer to the *corresponding storage* of the *variable* in the *device data environment*. For each pointer, that is not a function pointer, that is captured by the lambda expression, the behavior is as if the pointer or, if a corresponding pointer member exists, the corresponding pointer member of the closure object is the *base pointer* of a *zero-offset assumed-size array* that appears as a *list item* in a **map clause** with the *storage map-type*.

If the **this** pointer is captured by a lambda expression in class scope, and a *variable* of the associated closure type is later mapped explicitly or implicitly with its full static type, the behavior is as if the object to which **this** points is also mapped as an *array section*, of length one, for which the *base pointer* is the non-static data member that corresponds to the **this** pointer in the closure object.

C++

Restrictions

Additional restrictions to the **map** clause for mapping of arrays, **array sections**, and **structures** are as follows:

- If a **list item** is an array or an **array section**, it must specify contiguous storage.
- If a **list item** is an element of a **structure**, and a different element of the **structure** has a **corresponding list item** in the **device data environment** prior to a **task** encountering the **construct** associated with the **map** clause, then the **list item** must also have a **corresponding list item** in the **device data environment** prior to the **task** encountering the **construct**.
- If a **mapper modifier** appears in a **map** clause, the type on which the specified **mapper** operates must match the type of the **list items** in the **clause**.

C++

- If a **list item** has a polymorphic **class type** and its static type does not match its dynamic type, the behavior is unspecified if the **map** clause is specified on a **map-entering construct** and a **corresponding list item** is not **present** in the **device data environment** prior to a **task** encountering the **construct**.
- If a given **variable** is captured by reference by the associated lambda expression of a **list item** that has a closure type and that **variable** is a reference that binds to a **variable** with **static storage duration**, the **variable** to which it binds must appear in an **enter** clause or a **link** clause on a **declare target directive** and must have **corresponding storage** in the **device data environment** prior to a **task** encountering the **construct**.

C++

Fortran

- If a **list item** has polymorphic type, the behavior is unspecified.
- 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.

Fortran

Cross References

- **declare_mapper** Directive, see [Section 11.5](#)

11.3.7 Storage Block Mapping

The **effective map clause set** of a **data-mapping construct** is the set of all **map clauses** that apply to that **construct**, including implicit **map clauses** and **map clauses** applied by **mappers**. The **effective map clause set** of a **construct** determines the set of **mappable storage blocks** for that **construct**. All **map clause list items** that share storage or have the same **containing structure** or **containing array** result in a single **mappable storage block** that contains the storage of the **list items**, unless otherwise

specified. The storage for each other **map clause list item** becomes a distinct **mappable storage block**. If a **list item** is a **referencing variable** that has a **containing structure**, the behavior is as if only the storage for its **referring pointer** is part of that **structure**. In general, if a **list item** is a **referencing variable** then the storage for its **referring pointer** and its **referenced pointee** occupy distinct **mappable storage blocks**.

For each **mappable storage block** that is determined by the **effective map clause set** of a **map-entering construct**, on entry to the **region** the following sequence of steps occurs as if performed as a single **atomic operation**:

1. If a **corresponding storage block** is not **present** in the **device data environment** then:
 - a) A **corresponding storage block**, which may share storage with the **original storage block**, is created in the **device data environment** of the **target device**. The *close-modifier* is a hint that the **corresponding storage block** should be created in a storage resource that is close to the **target device**, while the *self-modifier* specifies a **self map** that requires that the **corresponding storage block** is identical to the **original storage block**.
 - b) The **corresponding storage block** receives a reference count that is initialized to zero. This reference count also applies to any part of the **corresponding storage block**.
2. The reference count of the **corresponding storage block** is incremented by one.
3. For each **map clause list item** in the **effective map clause set** that is contained by the **mappable storage block**:
 - a) If the reference count of the **corresponding storage block** is one, a **new list item** with language-specific attributes derived from the **original list item** is created in the **corresponding storage block**. The reference count of the **new list item** is always equal to the reference count of its storage.
 - b) If the reference count of the **corresponding list item** is one or if the *always-modifier* is specified, and if the **map type** is **to**, the **corresponding list item** is updated as if the **list item** appeared in a **to clause** on a **target_update** directive.

For each **mappable storage block** that is determined by the **effective map clause set** of a **map-exiting construct**, and for which **corresponding storage** is **present** in the **device data environment**, on exit from the **region** the following sequence of steps occurs as if performed as a single **atomic operation**:

1. For each **map clause list item** in the **effective map clause set** that is contained by the **mappable storage block**:
 - a) If the reference count of the **corresponding list item** is one or if the *always-modifier* or *delete-modifier* is specified, and if the **map type** is **from**, the **original list item** is updated as if the **list item** appeared in a **from clause** on a **target_update** directive.
2. If the *delete-modifier* is not present and the reference count of the **corresponding storage block** is finite then the reference count is decremented by one.

3. If the *delete-modifier* is present and the reference count of the **corresponding storage block** is finite then the reference count is set to zero.

4. If the reference count of the **corresponding storage block** is zero, all storage to which that reference count applies is removed from the **device data environment**.

If the effect of the **map clauses** on a **construct** would assign the same value to an **original list item** or to a **corresponding list item** more than once then an implementation is allowed to ignore the additional assignments of the same value.

In all cases on entry to or exit from the **region**, concurrent reads or updates of any part of the **original list item** or **corresponding list item** must be synchronized with any update of the **list item** that occurs as a result of the **map clause** to avoid **data races**.

The **original list item** and **corresponding list item** may share storage such that writes to either item by one **task** followed by a read or write of the other **list item** by another **task** without intervening synchronization can result in **data races**. They are guaranteed to share storage if the **mapping operation** is a **self map** resulting from the *self-modifier*, if the **map clause** appears on a **data-mapping construct** for which the **target device** is the **encountering device**, or if the **corresponding list item** has an **attached pointer** that shares storage with its **original pointer**.

C / C++

If a **new list item** is created then the **new list item** will have the same static type as the **original list item**, and language-specific attributes of the **new list item**, including size and alignment, are determined by that type.

C / C++

C++

If **corresponding storage** that differs from the **original storage** is created in a **device data environment**, all **new list items** that are created in that **corresponding storage** are default initialized. Default initialization for **new list items** of **class type**, including their data members, is performed as if with an implicitly-declared default constructor and as if non-static data member initializers are ignored.

C++

Fortran

If a **new list item** is created then the **new list item** will have the same type, type parameter, and rank as the **original list item**. The **new list item** inherits all default values for the type parameters from the **original list item**.

Fortran

If a **map clause** has the *present-modifier* and on entry to the **region** the **corresponding list item** is not **present** in the **device data environment** then **runtime error termination** is performed.

If a **map-entering clause** specifies a **self map** for a **list item** then **runtime error termination** is performed if any of the following is true:

- 1 • The **original list item** is not **accessible** and cannot be made **accessible** from the **target device**;
- 2 • The **corresponding list item** is **present** prior to a **task** encountering the **construct** on which the
- 3 **clause** appears, and the **corresponding storage** differs from the **original storage**; or
- 4 • The **list item** is a pointer that would be assigned a different value as a result of **pointer**
- 5 **attachment**.

6 **Restrictions**

7 Additional restrictions to the **map clause** for mapping of **storage blocks** are as follows:

- 8 • If any part of the **original storage** of a **list item** that is explicitly mapped by a **map clause** has
- 9 **corresponding storage** in the **device data environment** prior to a **task** encountering the
- 10 **construct** associated with the **map clause**, all of the **original storage** must have **corresponding**
- 11 **storage** in the **device data environment** prior to the **task** encountering the **construct**.
- 12 • If a **list item** in a **map clause** has **corresponding storage** in the **device data environment**, all
- 13 **corresponding storage** must correspond to a single **mappable storage block** that was
- 14 previously mapped.
- 15 • If a **list item** appears in a **map clause** with the *self-modifier*, any other **list item** in a **map**
- 16 **clause** on the same **construct** that has the same **base variable** or **base pointer** must also be
- 17 specified with the *self-modifier*.

18 **11.3.8 Implicit Data-Mapping**

19 If a single contiguous part of the **original storage** of a **list item** that results from an **implicitly**

20 **determined data-mapping attribute** has **corresponding storage** in the **device data environment** prior

21 to a **task** encountering the **construct** on which the **map clause** appears, only that part of the **original**

22 **storage** will have **corresponding storage** in the **device data environment** as a result of the **map clause**.

23 If a **list item** with an **implicitly determined data-mapping attribute** does not have any **corresponding**

24 **storage** in the **device data environment** prior to a **task** encountering the **construct** associated with the

25 **map clause**, and one or more contiguous parts of the **original storage** are either **list items** or **base**

26 **pointers to list items** that are explicitly mapped on the **construct**, only those parts of the **original**

27 **storage** will have **corresponding storage** in the **device data environment** as a result of the **map**

28 **clauses** on the **construct**.

29 **11.4 defaultmap Clause**

30 Name: defaultmap	Properties: unique, post-modified
---------------------------	-----------------------------------

Arguments

Name	Type	Properties
<i>implicit-behavior</i>	Keyword: default , firstprivate , from , none , present , private , self , storage , to , tofrom	<i>default</i>

Modifiers

Name	Modifies	Type	Properties
<i>variable-category</i>	<i>implicit-behavior</i>	Keyword: aggregate , all , allocatable , pointer , scalar	<i>default</i>
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: <i>directive-name</i> (a <i>directive name</i>)	unique

Directives

target

Additional information

The value **alloc** may also be specified as *implicit-behavior* with identical meaning to the value **storage**.

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 11.1. 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 <i>variable-category</i> specifies non-pointer scalar variables.		
▲	C / C++	▲
▼	Fortran	▼
The scalar <i>variable-category</i> specifies non-pointer and non-allocatable scalar variables. The allocatable <i>variable-category</i> specifies variables with the ALLOCATABLE attribute.		
▲	Fortran	▲

1 The **pointer** *variable-category* specifies *variables* of pointer type. The **aggregate**
2 *variable-category* specifies *aggregate variables*. Finally, the **all** *variable-category* specifies all
3 *variables*.

4 If *implicit-behavior* corresponds to a *map-type*, the attribute is a *data-mapping attribute* determined
5 by an implicit **map** clause with the specified *map-type*. If *implicit-behavior* is **private**, the
6 attribute is a *data-sharing attribute* of **private**. If *implicit-behavior* is **firstprivate**, the
7 attribute is a *data-sharing attribute* of **firstprivate**. If *implicit-behavior* is **present**, the attribute is
8 a *data-mapping attribute* determined by an implicit **map** clause with a *map-type* of **storage** and
9 the *present-modifier*. If *implicit-behavior* is **self**, the attribute is a *data-mapping attribute*
10 determined by an implicit **map** clause with a *map-type* of **storage** and the *self-modifier*. If
11 *implicit-behavior* is **none** then no *implicitly determined data-mapping attributes* or *implicitly*
12 *determined data-sharing attributes* are defined for *variables* in *variable-category*, except for
13 *variables* that appear in the **enter** or **link** clause of a **declare_target** directive. If
14 *implicit-behavior* is **default** then the *clause* has no effect.

15 **Restrictions**

16 Restrictions to the **defaultmap** clause are as follows:

- 17 • A given *variable-category* may be specified in at most one **defaultmap** clause on a
18 *construct*.
- 19 • If a **defaultmap** clause specifies the **all** *variable-category*, no other **defaultmap**
20 *clause* may appear on the *construct*.
- 21 • If *implicit-behavior* is **none**, each *variable* that is specified by *variable-category* and is
22 referenced in the *construct* but does not have a *predetermined data-sharing attribute* and does
23 not appear in an **enter** or **link** clause on a **declare_target** directive must be
24 explicitly listed in a *data-environment attribute clause* on the *construct*.

25 C / C++
 • The specified *variable-category* must not be **allocatable**.
 C / C++

26 **Cross References**

- 27 • Implicit Data-Mapping Attribute Rules, see [Section 11.1](#)
- 28 • **target** Construct, see [Section 21.8](#)

29 **11.5 declare_mapper Directive**

30	Name: <code>declare_mapper</code> Category: <code>declarative</code>	Association: <code>unassociated</code> Properties: <code>pure</code>
----	---	---

1 **Arguments**

2 **declare_mapper** (*mapper-specifier*)

3

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

4 **Clauses**

5 **map**

6 **Additional information**

7 The **declare_mapper** directive may alternatively be specified with **declare mapper** as the
8 *directive-name*.

9 **Semantics**

10 User-defined **mappers** can be defined using the **declare_mapper** directive. The
11 *mapper-specifier* argument declares the **mapper** using the following syntax:

12

	C / C++
	Fortran

13

	Fortran
--	---------

14 where *mapper-identifier* is a **mapper** identifier, *type* is a type that is permitted in a **type-name list**,
15 and *var* is a **base language** identifier.

16 The *type* and an optional *mapper-identifier* uniquely identify the **mapper** for use in a **map clause** or
17 **data-motion clause** later in the **OpenMP program**.

18 If *mapper-identifier* is not specified, the behavior is as if *mapper-identifier* is **default**.

19 The **variable** declared by *var* is available for use in all **map clauses** on the **directive**, and no part of
20 the **variable** to be mapped is mapped by default.

21 The effect that a **user-defined mapper** has on either a **map clause** that maps a **list item** of the given
22 **base language** type or a **data-motion clause** that invokes the **mapper** and updates a **list item** of the
23 given **base language** type is to replace the map or update with a set of **map clauses** or updates
24 derived from the **map clauses** specified by the **mapper**, as described in **Section 11.3** and **Chapter 12**.

25 A **list item** in a **map clause** that appears on a **declare_mapper** directive may include **array**
26 **sections**.

27 All **map clauses** that are introduced by a **mapper** are further subject to **mappers** that are in scope,
28 except a **map clause** with **list item** *var* maps *var* without invoking a **mapper**.

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.
- 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, must 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 a **struct** or **class** type, it must not be derived from any virtual base class.

C++

Fortran

- *type* must not be an intrinsic type, a parameterized derived type, an enum type, or an enumeration type.

Fortran

Cross References

- **map** Clause, see [Section 11.3](#)

12 Data-Motion Control

This chapter describes **data-motion clauses** that specify data movement between **devices** in a **device set** that is specified by the **construct** on which the **clause** appears, where one of the **devices** in the set is the **encountering device** and the remaining **devices** are **target devices** of the **construct**. 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 an **iterator modifier** on the **clause**. The **list items** may include **array sections** with **stride** expressions.

C / C++

The **list items** may use **shape-operators**.

C / C++

If a **list item** is an array or **array section** then it is treated as if it is replaced by each of its **array elements** in the **clause**.

If the **mapper modifier** is not specified, the behavior is as if the **modifier** was specified with the **default mapper identifier**. The effect of a **data-motion clause** on a **list item** is modified by a visible **user-defined mapper** if 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-modifier** 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 **present-modifier** only applies to those **mapper list items** that share storage with the **original list item**.

If a **list item** is a **referencing variable** then the effect of the **data-motion clause** is applied only to its **referenced pointee** and only if the **referenced pointee** exists.

Fortran

If a **list item** is an associated **procedure** pointer, the **corresponding list item** on the **device** is associated with the target **procedure** of the **host device**.

Fortran

C / C++

On exit from the associated **region**, if the **corresponding list item** is an **attached pointer**, the **original list item** will have the value it had on entry to the **region** and the **corresponding list item** will have the value it had on entry to the **region**.

C / C++

For each **list item** that is not an **attached pointer**, the value of the **assigned list item** is assigned the value of the other **list item**. To avoid **data races**, concurrent reads or updates of the **assigned list item** must be synchronized with the update of an **assigned list item** that occurs as a result of a **data-motion clause**.

Restrictions

Restrictions to **data-motion clauses** are as follows:

- Each **list item** of *locator-list* must have a **mappable type**.
- If an array appears as a **list item** in a **data-motion clause** and it has **corresponding storage** in the **device data environment**, the **corresponding storage** must correspond to a single **mappable storage block** that was previously mapped.
- If a **list item** in a **data-motion clause** has **corresponding storage** in the **device data environment**, all **corresponding storage** must correspond to a single **mappable storage block** that was previously mapped.
- If a **mapper modifier** appears in a **data-motion clause**, the specified **mapper** must operate on a type that matches either the type or **array element** type of each **list item** in the **clause**.

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**.

Fortran

12.1 to Clause

Name: to	Properties: data-motion attribute
-----------------	--

Arguments

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

1 **Modifiers**


Name	Modifies	Type	Properties
<i>present-modifier</i>	<i>locator-list</i>	Keyword: present	<i>default</i>
<i>mapper</i>	<i>locator-list</i>	Complex, name: mapper Arguments: mapper-identifier OpenMP identifier (<i>default</i>)	<i>unique</i>
<i>iterator</i>	<i>locator-list</i>	Complex, name: iterator Arguments: iterator-specifier list of iter- ator specifier list item type (<i>default</i>)	<i>unique</i>
<i>directive-name- modifier</i>	<i>all arguments</i>	Keyword: <i>directive-name</i> (a <i>directive name</i>)	<i>unique</i>



3 **Directives**

4 **target_update**

5 **Semantics**

6 The **to** clause is a *data-motion clause* that specifies data movement to the *target devices* from the
7 *encountering device* so the *corresponding list items* are the *assigned list items* and the *compatible*
8 *map types* are **to** and **tofrom**.

9 
A list item for which a *mapper* does not exist is ignored if it has *static storage duration* and has the
10 **constexpr** specifier.

11 

A list item for which a *mapper* does not exist is ignored if it has *static storage duration* and either it
12 has the **constexpr** specifier or it is a non-mutable member of a *structure* that has the
13 **constexpr** specifier.

14 **Cross References**

- **iterator** Modifier, see [Section 5.2.6](#)
- **target_update** Construct, see [Section 21.9](#)

12.2 from Clause

Name: <code>from</code>	Properties: <code>data-motion</code> attribute
--------------------------------	---

Arguments

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

Modifiers

Name	Modifies	Type	Properties
<i>present-modifier</i>	<i>locator-list</i>	Keyword: present	<i>default</i>
<i>mapper</i>	<i>locator-list</i>	Complex, name: mapper Arguments: <i>mapper-identifier</i> OpenMP identifier (<i>default</i>)	<i>unique</i>
<i>iterator</i>	<i>locator-list</i>	Complex, name: iterator Arguments: <i>iterator-specifier</i> list of iterator specifier list item type (<i>default</i>)	<i>unique</i>
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: <i>directive-name</i> (a <i>directive name</i>)	<i>unique</i>

Directives

`target_update`

Semantics

The **from** clause is a `data-motion` clause that specifies data 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** or **constexpr** 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++

Fortran

A list item for which a [mapper](#) does not exist is ignored if it has the **INTENT (IN)** attribute and does not have the **POINTER** attribute.

Fortran

Cross References

- **iterator** Modifier, see [Section 5.2.6](#)
- **target_update** Construct, see [Section 21.9](#)

13 Memory Management

This chapter defines `directives`, `clauses` and related concepts for managing `memory` used by OpenMP programs.

13.1 Memory Spaces

OpenMP `memory spaces` represent storage resources where `variables` can be stored and retrieved. Table 13.1 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 13.1: 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 <code>variables</code> with <code>constant</code> 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.

13.2 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 contiguous [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](#).

If an [allocator](#) is used to allocate [memory](#) for a [variable](#) with [static storage duration](#) that is not a [local static variable](#) then the [task](#) that requested the allocation is unspecified. If an [allocator](#) is used to allocate [memory](#) for a [local static variable](#) then the [task](#) that requested the allocation is considered to be the [current task](#) of the first [thread](#) that executes code in which the [variable](#) is visible.

The behavior of the allocation process can be affected by the [allocator traits](#) that the user specifies. Table 13.2 shows the allowed [allocator traits](#), their possible values and the default value of each [trait](#).

TABLE 13.2: Allocator Traits

Allocator Trait	Allowed Values	Default Value
<code>sync_hint</code>	<code>contended</code> , <code>uncontended</code> , <code>serialized</code> , <code>private</code>	<code>contended</code>
<code>alignment</code>	Non-negative integer powers of 2	1 byte
<code>access</code>	<code>all</code> , <code>memspace</code> , <code>device</code> , <code>cgroup</code> , <code>pteam</code> , <code>thread</code>	<code>memspace</code>
<code>pool_size</code>	Any positive integer	Implementation de- fined
<code>fallback</code>	<code>default_mem_fb</code> , <code>null_fb</code> , <code>abort_fb</code> , <code>allocator_fb</code>	See below
<code>fb_data</code>	An allocator handle	(none)
<code>pinned</code>	true , false	false
<code>partition</code>	<code>environment</code> , <code>nearest</code> , <code>blocked</code> , <code>interleaved</code> , <code>partitioner</code>	<code>environment</code>
<code>pin_device</code>	Conforming device number	(none)
<code>preferred_device</code>	Conforming device number	(none)
<code>target_access</code>	<code>single</code> , <code>multiple</code>	<code>single</code>
<code>atomic_scope</code>	<code>all</code> , <code>device</code>	<code>device</code>

table continued on next page

table continued from previous page

Allocator Trait	Allowed Values	Default Value
part_size	Positive integer value	Implementation defined
partitioner	A memory partitioner handle	(none)
partitioner_arg	An integer value	0

The **sync_hint trait** describes the expected manner in which multiple threads may use the allocator. The values and their descriptions are:

- **contended**: high contention is expected on the allocator; that is, many tasks are expected to request allocations simultaneously;
- **uncontended**: low contention is expected on the allocator; that is, few tasks are expected to request allocations simultaneously;
- **serialized**: one task at a time will request allocations with the allocator. Requesting two allocations simultaneously when specifying **serialized** results in **unspecified behavior**; and
- **private**: the same thread will execute all tasks that request allocations with the allocator. Requesting an allocation from tasks that different threads execute, simultaneously or not, when specifying **private** results in **unspecified behavior**.

Allocated memory will be byte aligned to at least the value specified for the **alignment trait** of the allocator. Some directives and routines can specify additional requirements on alignment beyond those described in this section.

The **access trait** defines the *access group* of tasks that may access memory that is allocated by a memory allocator. If the value is **all**, the access group consists of all tasks that execute on all available devices. If the value is **memspace**, the access group consists of all tasks that execute on all devices that are associated with the allocator. If the value is **device**, the access group consists of all tasks that execute on the device where the allocation was requested. If the value is **cgroup**, the access group consists of all tasks in the same contention group as the task that requested the allocation. If the value is **pteam**, the access group consists of all current team tasks of the innermost enclosing parallel region in which the allocation was requested. If the value is **thread**, the access group consists of all tasks that are executed by the same thread that executed the allocation request. Memory returned by the allocator will be memory accessible by all tasks in the same access group as the task that requested the allocation. Attempts to access this memory from a task that is not in same access group results in **unspecified behavior**.

The total amount of storage in bytes that an allocator can use for allocation requests from tasks in the same access group is limited by the **pool_size trait**. Requests that would result in using more storage than **pool_size** will not be fulfilled by the allocator.

The **fallback trait** specifies how the **memory allocator** behaves when it cannot fulfill an allocation request. If the **fallback trait** is set to **null_fb**, the **allocator** returns the value zero if it fails to allocate the **memory**. If the **fallback trait** is set to **abort_fb**, the behavior is as if an **error directive** for which *sev-level* is **fatal** and *action-time* is **execution** is encountered if the allocation fails. If the **fallback trait** is set to **allocator_fb** then when an allocation fails the request will be delegated to the **allocator** specified in the **fb_data trait**. If the **fallback trait** is set to **default_mem_fb** then when an allocation fails another allocation will be tried in **omp_default_mem_space**, which assumes all **allocator traits** to be set to their default values except for **fallback trait**, which will be set to **null_fb**. The default value for the **fallback trait** is **null_fb** for any **allocator** that is associated with a **target memory space**. Otherwise, the default value is **default_mem_fb**.

All **memory** that is allocated with an **allocator** for which the **pinned trait** is specified as **true** must remain in the same storage resource at the same location for its entire lifetime. If **pin_device** is also specified then the allocation must be allocated in that **device**.

The **partition trait** describes the partitioning of allocated **memory** over the storage resources represented by the **memory space** associated with the **allocator**. The partitioning will be done in parts with a minimum size that is **implementation defined**. The values are:

- **environment**: the placement of allocated **memory** is determined by the execution environment;
- **nearest**: allocated **memory** is placed in the storage resource that is nearest to the **thread** that requests the allocation;
- **blocked**: allocated **memory** is partitioned into parts of approximately the same size with at most one part per storage resource; and
- **interleaved**: allocated **memory** parts are distributed in a round-robin fashion across the storage resources such that the size of each part is the value of the **part_size trait** except possibly the last part, which can be smaller.
- **partitioner**: the number of **memory** parts and how they are distributed across the storage are defined by the **memory partition** object created by the **memory partitioner** specified by the **partitioner trait**.

The **part_size trait** specifies the size of the parts allocated over the storage resources for some of the **memory partition trait** policies. The actual value of the **trait** might be rounded up to an **implementation defined** value to comply with hardware restrictions of the storage resources.

If the **preferred_device trait** is specified then storage resources of the specified **device** are preferred to fulfill the allocation.

If the value of the **target_access trait** is **single** then data from this **allocator** cannot be accessed on two different **devices** unless, for any given **host device** access, the entry and exit of the **target region** in which any accesses occur either both precede or both follow the **host device** access in **happens-before order**. Additionally, for any two **target regions** that may access data

from this `allocator` and execute on distinct `devices`, the entry and exit of one of the `regions` must precede those of the other in `happens-before` order. If the value of the `target_access_trait` is `multiple` then accesses of data from this `allocator` from different `devices` may be arbitrarily interleaved, provided that synchronization ensures `data races` do not occur.

If the value of the `atomic_scope_trait` is `all` then all `storage locations` of data from this `allocator` have an `atomic scope` that consists of all `threads` on the devices associated with the `allocator`. If the value is `device` then all `storage locations` have an `atomic scope` that consists of all `threads` on the `device` on which the `atomic operation` is performed.

Table 13.3 shows the list of predefined `memory allocators` and their `associated memory spaces`. The predefined `memory allocators` have default values for their `allocator traits` unless otherwise specified.

TABLE 13.3: 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`.

- The `omp_cgroup_mem_alloc`, `omp_pteam_mem_alloc`, and `omp_thread_mem_alloc` predefined [memory allocators](#) must not be used to allocate a [variable](#) with [static storage duration](#) unless the [variable](#) is a [local static variable](#).

13.3 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: <i>directive-name</i> (a 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 13.5](#)
- Memory Allocators, see [Section 13.2](#)

13.4 allocator Clause

Name: <code>allocator</code>	Properties: <code>ICV-defaulted</code> , <code>unique</code>
-------------------------------------	---

Arguments

Name	Type	Properties
<i>allocator</i>	expression of <code>allocator_</code> -handle type	<i>default</i>

Modifiers

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

Directives

`allocators`

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 **`directives`** 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*.

Restrictions

Restrictions to the **`allocator`** clause are as follows:

- The **`memory allocator`** specified by *allocator* at the **`directive`** on which the **`allocator`** clause appears must be the same as the **`memory allocator`** specified by *allocator* at the declaration of the **`variables`** to be allocated.

Cross References

- **`allocators`** Construct, see [Section 13.7](#)
- Memory Allocators, see [Section 13.2](#)
- *def-allocator-var ICV*, see [Table 3.1](#)

13.5 allocate Directive

Name: <code>allocate</code> Category: <code>declarative</code>	Association: <code>explicit</code> Properties: <code>pure</code>
---	---

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 its **list items** and must follow the 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** and alignment.
- 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 be a coarray or have a coarray as an ultimate component, or 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](#) and alignment 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 13.3](#)
- `allocator` Clause, see [Section 13.4](#)
- Memory Allocators, see [Section 13.2](#)

13.6 allocate Clause

Name: `allocate`

Properties: [all-privatizing](#), [target-consistent](#)

Arguments

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

Modifiers

Name	Modifies	Type	Properties
<i>allocator-simple-modifier</i>	<i>list</i>	expression of OpenMP allocator_handle type	exclusive, unique
<i>allocator-complex-modifier</i>	<i>list</i>	Complex, name: allocator Arguments: allocator expression of allocator_handle type (<i>default</i>)	unique
<i>align-modifier</i>	<i>list</i>	Complex, name: align Arguments: alignment expression of integer type (<i>constant</i> , <i>positive</i>)	unique
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: <i>directive-name</i> (a <i>directive name</i>)	unique

Directives

allocators, distribute, do, for, parallel, scope, sections, single, target, target_data, task, taskgroup, taskloop, teams

Semantics

The **allocate** clause specifies the **memory allocator** to be used to obtain storage for a **variable list**. 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**. On a **target** constructs, the **allocate** clause only affects allocations that happen on the **target device**. If the *allocator-simple-modifier* is specified, the behavior is as if the *allocator-complex-modifier* is instead specified with *allocator-simple-modifier* as its *allocator* argument. The *allocator-complex-modifier* and *align-modifier* have the same syntax and semantics for the **allocate** clause as the **allocator** and **align** clauses have for the **allocate directive**. For allocations that arise from this **clause**, the **null_fb** value of the fallback **allocator trait** behaves as if the **abort_fb** value had been specified.

For an **allocate** clauses specified on a **target** construct, the **allocator** for the **target device** is used. If a predefined **allocator** is specified then the predefined **allocator** for the **target device** is used. Otherwise, the specified **allocator variable** must also be a **list item** in a **uses_allocators** clause on the **construct** and the **allocator** created for its **corresponding list item** on the **target device** is used.

Restrictions

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

- For any **list item** that is specified in the **allocate** clause on a **directive** other than the **allocators** directive, a **data-sharing attribute clause** that may create a **private** copy of that **list item** must be specified on the same **directive**.

- For **task**, **taskloop** or **target** directives, allocation requests to **memory allocators** with the **access trait** set to **thread** result in **unspecified behavior**.
- **allocate** clauses that appear in a **target region** must specify an *allocator-simple-modifier* or *allocator-complex-modifier* unless a **requires** directive with the **dynamic_allocators** clause is present in the same **compilation unit**.
- **allocate** clauses that appear on a **target construct** must specify an *allocator-simple-modifier* or *allocator-complex-modifier* with an **allocator** that is either a predefined **memory allocator** or an **allocator** that was specified in a **uses_allocators** clause on the **construct**.
- **allocate** clauses that appear on a **teams construct** that is an **immediately nested construct** of a **target construct** must specify an *allocator-simple-modifier* or *allocator-complex-modifier* with a predefined **memory allocator**.

Cross References

- **align** Clause, see [Section 13.3](#)
- **allocator** Clause, see [Section 13.4](#)
- **allocators** Construct, see [Section 13.7](#)
- **distribute** Construct, see [Section 19.7](#)
- **do** Construct, see [Section 19.6.2](#)
- **for** Construct, see [Section 19.6.1](#)
- **Memory Allocators**, see [Section 13.2](#)
- **parallel** Construct, see [Section 18.1](#)
- **scope** Construct, see [Section 19.2](#)
- **sections** Construct, see [Section 19.3](#)
- **single** Construct, see [Section 19.1](#)
- **target** Construct, see [Section 21.8](#)
- **target_data** Construct, see [Section 21.7](#)
- **task** Construct, see [Section 20.1](#)
- **taskgroup** Construct, see [Section 23.4](#)
- **taskloop** Construct, see [Section 20.2](#)
- **teams** Construct, see [Section 18.2](#)

13.7 allocators Construct

Name: <code>allocators</code> Category: <code>executable</code>	Association: <code>block : allocator</code> Properties: <code>default</code>
--	---

Clauses

`allocate`

Semantics

The `allocators` construct specifies that if a `variable` that is to be allocated by the associated `allocate-stmt`, 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. The `list items` that appear in an `allocate` clause may include `structure elements`.

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`.
- A `list item` must not be a coarray or have a coarray as an ultimate component.

Cross References

- `allocate` Clause, see [Section 13.6](#)
- Memory Allocators, see [Section 13.2](#)
- OpenMP Allocator Structured Blocks, see [Section 6.3.1](#)

13.8 uses_allocators Clause

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

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>)	<i>default</i>
<i>traits-array</i>	<i>allocator</i>	Complex, name: traits Arguments: traits variable of alloctrait array type (<i>default</i>)	<i>default</i>
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: <i>directive-name</i> (a <i>directive name</i>)	<i>unique</i>

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. The *clause* has no effect for an *allocator* argument value of **omp_null_allocator**. If *allocator* is an identifier that matches the name of a predefined *allocator* (see Table 13.3), that predefined *allocator* will be available for use in the *region*. Otherwise, the effect is as if *allocator* is specified on a **private** clause. The resulting *corresponding list 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**.

More than one *clause-argument-specification* may be specified.

Restrictions

- The *allocator* expression must be a *base language* identifier.
- If *allocator* is an identifier that matches the name of a predefined *allocator*, no *modifiers* may be specified.
- If *allocator* is not the name of a predefined *allocator* and is not **omp_null_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*.

1
2

3

4
5
6
7
8
9
10
11
12
13
14

15
16
17
18

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

- OpenMP **allocator_handle** Type, see [Section 26.8.2](#)
- OpenMP **alloctrait** Type, see [Section 26.8.3](#)
- Memory Allocators, see [Section 13.2](#)
- Memory Spaces, see [Section 13.1](#)
- OpenMP **memspace_handle** Type, see [Section 26.8.12](#)
- omp_destroy_allocator** Routine, see [Section 33.7](#)
- omp_init_allocator** Routine, see [Section 33.6](#)
- target** Construct, see [Section 21.8](#)

13.9 dyn_groupprivate Clause

Name: dyn_groupprivate		Properties: target-consistent
Arguments		
Name	Type	Properties
<i>size</i>	expression of integer type	non-negative

Modifiers

Name	Modifies	Type	Properties
<i>access-group</i>	<i>size</i>	Keyword: cgroup	unique
<i>fallback-modifier</i>	<i>size</i>	Complex, name: fallback Arguments: <i>fallback-mode</i> Keyword: abort , default_mem , null (<i>default</i>)	unique
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: <i>directive-name</i> (a <i>directive name</i>)	unique

Directives

target, teams

Semantics

The **dyn_groupprivate** clause specifies the instantiation of a **dynamic groupprivate block** of *size* bytes in the **region** that corresponds to the **groupprivate-instantiating construct** on which it appears. The instantiated **dynamic groupprivate block** has the **groupprivate attribute**, and each **access group** of **tasks** receives one of the **groupprivate** copies. Each copy is uninitialized upon creation. The lifetime of a **groupprivate** copy is limited to the lifetime of **all tasks** in the **access group**.

The *access-group* modifier specifies the **access group type** for the **dynamic groupprivate block**. If **cgroup** is specified, **tasks** are grouped based on the **contention group** to which they belong. If *access-group* is not specified then the effect is as if **cgroup** is specified.

Each **groupprivate** copy uses **memory** from the **groupprivate memory space** associated with the **memory allocator** of the **access group type**, if the requested size falls within the **groupprivate effective limit** of that **memory space**. Two limits, in bytes, are associated with the request of **memory** from that **groupprivate memory space** for the specified **access group type** and the **region** corresponding to the **groupprivate-instantiating construct** – the **groupprivate absolute limit** and the **groupprivate effective limit**. The **groupprivate absolute limit** for a **groupprivate memory space** specified by an **access group type** and **device** can be obtained using the **omp_get_groupprivate_limit** routine. The **groupprivate effective limit** is less than or equal to the **groupprivate absolute limit**.

The *size* argument specifies the number of bytes of the instantiated **dynamic groupprivate block**. The *fallback-modifier* specifies the behavior when *size* exceeds the **groupprivate effective limit** for the **access group type**, and thus, the implementation is unable to instantiate a **dynamic groupprivate block** from the **groupprivate memory space** for an **access group** of **tasks**. If the *fallback-mode* of the *fallback-modifier* is **abort**, **runtime error termination** is performed. If the *fallback-mode* is **null**, the **dynamic groupprivate block** is not instantiated. If the *fallback-mode* is **default_mem**, the implementation uses **memory** from an **implementation defined memory space** for that **access group** of **tasks**. If the *fallback-modifier* is not specified then the effect is as if the *fallback-mode* is

specified as **default_mem**.

When the *size* argument evaluates to zero, a [dynamic groupprivate block](#) is not instantiated. On [groupprivate-instantiating constructs](#), if the [dyn_groupprivate clause](#) is not specified then the effect is as if the [clause](#) appears with a *size* argument that evaluates to zero.

The [dynamic-groupprivate-information routines](#) provide information regarding the [dynamic groupprivate blocks](#). Any [task](#) in an [access group](#) can obtain the size of and a pointer to the [groupprivate](#) copy of the [storage block](#) by using those [routines](#), as defined in [Section 33.13](#).

Restrictions

Restrictions to the [dyn_groupprivate clause](#) are as follows:

- A given [access-group](#) must not be specified in more than one [dyn_groupprivate clause](#) on a [construct](#).
- A [task](#) can only access the [groupprivate](#) copy of a [dynamic groupprivate block](#) for the [access groups](#) to which it belongs.

Cross References

- Groupprivate Memory Routines, see [Section 33.13](#)
- **groupprivate** Directive, see [Section 9.2](#)
- **target** Construct, see [Section 21.8](#)
- **teams** Construct, see [Section 18.2](#)

14 SIMD Data Control

This chapter presents [clauses](#) that control [data-environment attributes](#) that are relevant for [SIMD](#) execution.

14.1 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: <i>directive-name</i> (a 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](#).

Restrictions

Restrictions to the [uniform clause](#) are as follows:

- Only [named parameter list items](#) can be specified in the *parameter-list*.

Cross References

- `declare_simd` Directive, see [Section 15.8](#)

14.2 aligned Clause

Name: <code>aligned</code>	Properties: data-environment attribute , post-modified
-----------------------------------	---

1 **Arguments**

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

3 **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: <i>directive-name</i> (a <i>directive name</i>)	unique

5 **Directives**

6 **declare_simd, simd**

7 **Semantics**

8  C / C++

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

10  C / C++

11  Fortran

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

13  Fortran

14 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.

15 **Restrictions**

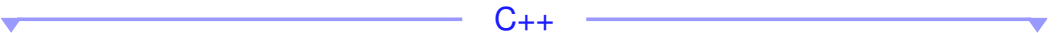
16 Restrictions to the **aligned** clause are as follows:

- 17 • If the **clause** appears on a **declare_simd** directive, each **list item** must be a **named**
18 **parameter list item** of the associated **procedure**.

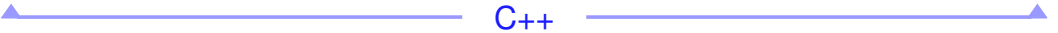
19  C

- 20 • The type of each **list item** must be an array or pointer type.

21  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 15.8](#)
- `simd` Construct, see [Section 18.4](#)

1

Part III

2

General Directives and Clauses

15 Variant Directives

This chapter defines **directives** and related concepts to support the seamless adaption of **OpenMP programs** to **OpenMP contexts**.

15.1 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 15.2**.

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**. **Compound constructs** are added to the set as their **leaf 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.

The *simd* trait is a [clause-list trait](#) that is defined with [properties](#) that match the [clauses](#) that can be specified on the [declare_simd](#) directive with the same names and semantics. The *simd* trait defines at least the *simdden* [property](#) and one of the *inbranch* or *notinbranch* [properties](#). Traits in the [construct trait set](#) other than *simd* are [non-property traits](#).

The [device trait set](#) includes [traits](#) that define the characteristics of the [device](#) that the compiler determines will be the [current device](#) during program execution at a given point in the [OpenMP program](#). A [trait](#) in the [device trait set](#) is considered to be active at program points that fall outside a defined [procedure](#) if it defines a characteristic of some [available device](#), including the [host device](#). For each [target device](#) that the implementation supports, a [target device trait set](#) exists that defines the characteristics of that [device](#). At least the following [traits](#) must be defined for the [device trait set](#) and all [target device trait sets](#):

- The *kind(kind-list)* [name-list trait](#) specifies the general kind of the [device](#). Each member of *kind-list* is a *kind-name*, for which the following values are defined:
 - *host*, which specifies that the [device](#) is the [host device](#);
 - *nohost*, which specifies that the [device](#) is not the [host device](#); and
 - the values defined in the [OpenMP Additional Definitions document](#).
- The *isa(isa-list)* [name-list trait](#) specifies the Instruction Set Architectures supported by the [device](#). Each member of *isa-list* is an *isa-name*, for which the accepted values are [implementation defined](#).
- The *arch(arch-list)* [name-list trait](#) specifies the architectures supported by the [device](#). Each member of *arch-list* is an *arch-name*, for which the accepted values are [implementation defined](#).

The [target device trait set](#) also defines the following [traits](#):

- The *device_num* [trait](#) specifies the *device number* of the [device](#).
- The *uid* [trait](#) specifies a unique identifier string of the [device](#), for which the accepted values are [implementation defined](#).

The [implementation trait set](#) includes [traits](#) that describe the functionality supported by the OpenMP implementation at that point in the [OpenMP program](#). At least the following [traits](#) can be defined:

- The *vendor(vendor-list)* [name-list trait](#), which specifies the vendor identifiers of the implementation. Each member of *vendor-list* is a *vendor-name*, for which the defined values are in the [OpenMP Additional Definitions document](#).
- The *extension(extension-list)* [name-list trait](#), which specifies vendor-specific extensions to the OpenMP specification. Each member of *extension-list* is an *extension-name*, for which the accepted values are [implementation defined](#).
- A *requires(requires-list)* [clause-list trait](#), for which the [properties](#) are the [clauses](#) that have been supplied to the [requires](#) directive prior to the program point as well as

1 **implementation defined** implicit requirements.

2 Implementations can define additional **traits** in the **device trait set**, **target device trait set** and
3 **implementation trait set**; these **traits** are **extension traits**.

4 The **dynamic trait set** includes **traits** that define the dynamic **properties** of an **OpenMP program** at a
5 point in its execution. The *data state trait* in the **dynamic trait set** refers to the complete data state of
6 the **OpenMP program** that may be accessed at runtime.

7 15.2 Context Selectors

8 **Context selectors** are used to define the **properties** that can match an **OpenMP context**. OpenMP
9 defines different **trait selector sets**, each of which contains different **trait selectors**.

10 The syntax for a **context selector** is *context-selector-specification* as described in the following
11 grammar:

```
12       context-selector-specification :  
13                trait-set-selector [ , trait-set-selector [ , ... ] ]  
14  
15       trait-set-selector :  
16                trait-set-selector-name = { trait-selector [ , trait-selector [ , ... ] ] }  
17  
18       trait-selector :  
19                trait-selector-name [ ( [ trait-score : ] trait-property [ , trait-property [ , ... ] ] ) ]  
20  
21       trait-property :  
22                trait-property-name  
23                trait-property-clause  
24                trait-property-expression  
25                trait-property-extension  
26  
27       trait-property-clause :  
28                clause  
29  
30       trait-property-name :  
31                identifier  
32                string-literal  
33  
34       trait-property-expression  
35                scalar-expression (for C/C++)  
36                scalar-logical-expression (for Fortran)  
37                scalar-integer-expression (for Fortran)  
38  
39       trait-score :
```

```

1      score (score-expression)
2
3      trait-property-extension :
4          trait-property-name
5          identifier (trait-property-extension [, trait-property-extension [, ...]])
6          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 corresponding string literal (for C/C++) and the corresponding character literal constant (for Fortran) representation are considered representations of the same value.

For **trait selectors** that correspond to **clause-list traits**, each *trait-property* should be *trait-property-clause*. The syntax is the same as for the matching **clause**.

The **construct selector set** defines the **traits** in the **construct trait set** that should be active in the **OpenMP context**. Each **trait selector** that can be defined in the **construct selector set** is the *directive-name* of a **context-matching construct**. Each *trait-property* of the **simd trait selector** is a *trait-property-clause*. The syntax is the same as for a valid **clause** of the **declare_simd** directive and the restrictions on the **clauses** from that **directive** apply. The **construct selector set** is an ordered list c_1, \dots, c_N .

The **device selector set** and **implementation selector set** define the **traits** that should be active in the corresponding **trait set** of the **OpenMP context**. The **target_device selector set** defines the **traits** that should be active in the **target device trait set** for the **device** that the specified **device_num trait selector** identifies. The same **traits** that are defined in the corresponding **trait sets** can be used as **trait selectors** with the same **properties**. The **kind trait selector** of the **device selector set** and **target_device selector set** can also specify the value **any**, which is as if no **kind trait selector** was specified. If a **device_num trait selector** does not appear in the **target_device selector set** then a **device_num trait selector** that specifies the value of the *default-device-var* ICV is implied. For the **device_num trait selector** of the **target_device selector set**, a single *trait-property-expression* must be specified. The **device_num trait selector** can be **true** only if that *trait-property-expression* evaluates to a **conforming device number** other than **omp_invalid_device**. For the **atomic_default_mem_order trait selector** of the **implementation selector set**, a single *trait-property* must be specified as an identifier equal to one of the valid arguments to the **atomic_default_mem_order clause** on the **requires directive**. For the **requires trait selector** of the **implementation selector set**, each *trait-property* is a *trait-property-clause*. The syntax is the same as for a valid **clause** of the **requires directive** and the restrictions on the **clauses** from that **directive** apply.

The **user selector set** defines the **condition trait selector** that provides additional user-defined conditions. The **condition trait selector** contains a single *trait-property-expression* that must evaluate to **true** for the **trait selector** to be **true** and otherwise **false**. Any **non-constant trait-property-expression** that is evaluated to determine the suitability of a variant is evaluated according to the *data state trait* in the **dynamic trait set** of the **OpenMP context**. The **user selector set** is dynamic if the **condition trait selector** is present and the expression in the **condition**

`trait selector` is not a `constant` expression; otherwise, it is static.

All parts of a `context selector` define the static part of the `context selector` except the following parts, which define the dynamic part of the `context selector`:

- Its `user selector set` if it is dynamic; and
- Its `target_device selector set`.

For the `match` clause of a `declare_variant` directive, any argument of the `base function` that is referenced in an expression that appears in the `context selector` is treated as a reference to the expression that is passed into that argument at the call to the `base function`. Otherwise, a `variable` or `procedure` reference in an expression that appears in a `context selector` is a reference to the `variable` or `procedure` of that name that is visible at the location of the `directive` on which the `context selector` appears.

▼ C++ ▼

Each occurrence of the `this` pointer in an expression in a `context selector` that appears in the `match` clause of a `declare_variant` directive is treated as an expression that is the address of the object on which the associated `base function` is invoked.

▲ C++ ▲

Implementations can allow further `trait selectors` to be specified. Each specified *trait-property* for these *implementation defined trait selectors* should be a *trait-property-extension*. Implementations can ignore specified `trait selectors` that are not those described in this section.

Restrictions

Restrictions to `context selectors` are as follows:

- Each *trait-property* may only be specified once in a `trait selector` other than those in the `construct selector set`.
- Each *trait-set-selector-name* may only be specified once in a `context selector`.
- Each *trait-selector-name* may only be specified once in a `trait selector set`.
- A *trait-score* cannot be specified in `traits` from the `construct selector set`, the `device selector set` or the `target_device selector sets`.
- A *score-expression* must be a `non-negative constant` integer expression.
- The expression of a `device_num` `trait` must evaluate to a `conforming device number`.
- A `variable` or `procedure` that is referenced in an expression that appears in a `context selector` must be visible at the location of the `directive` on which the `context selector` appears unless the `directive` is a `declare_variant` directive and the `variable` is an argument of the associated `base function`.
- If *trait-property any* is specified in the `kind` *trait-selector* of the `device selector set` or the `target_device selector sets`, no other *trait-property* may be specified in the same `selector set`.

- For a *trait-selector* that corresponds to a **name-list trait**, at least one *trait-property* must be specified.
- For a *trait-selector* that corresponds to a **non-property trait**, no *trait-property* may be specified.
- For the **requires trait selector** of the **implementation selector set**, at least one *trait-property* must be specified.

15.3 Matching and Scoring Context Selectors

A **compatible context selector** for an **OpenMP context** satisfies the following conditions:

- 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**; and
- **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**;

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
6. A **context selector** that is a strict subset of another **compatible context selector** has a score of zero. For other **context selectors**, the final score is the sum of the values of all specified **trait selectors** plus 1.

15.4 Metadirectives

A **metadirective** is a **directive** that can specify multiple **directive variants** of which one may be conditionally selected to replace the **metadirective** based on the **enclosing context**. A **metadirective** is replaced by a **nothing directive** or one of the **directive variants** specified by the **when clauses** or the **otherwise clause**. If no **otherwise clause** is specified the effect is as if one was specified without an associated **directive variant**.

The **OpenMP context** for a given **metadirective** is defined according to **Section 15.1**. The order of **clauses** that appear on a **metadirective** is significant and, if specified, **otherwise** must be the last **clause** specified on a **metadirective**.

Replacement candidates for a **metadirective** are ordered according to the following rules in decreasing precedence:

- A **candidate** is before another one if the score associated with the **context selector** of the corresponding **when clause** is higher.
- A **candidate** that was explicitly specified is before one that was implicitly specified.
- **Candidates** are ordered according to the order in which they lexically appear on the **metadirective**.

The list of **dynamic replacement candidates** is the prefix of the sorted list of **replacement candidates** up to and including the first **candidate** for which the corresponding **when** or **otherwise clause** has a **static context selector**. The first **dynamic replacement candidate** for which the corresponding **when** or **otherwise clause** has a **compatible context selector**, according to the matching rules defined in **Section 15.3**, replaces the **metadirective**.

Restrictions

Restrictions to **metadirectives** are as follows:

- Replacement of the **metadirective** with the **directive variant** associated with any of the **dynamic replacement candidates** must result in a **conforming program**.

- Insertion of user code at the location of a **metadirective** must be allowed if the first **dynamic replacement candidate** does not have a **static context selector**.
- If the list of **dynamic replacement candidates** has multiple items then all items must be **executable directives**.

Fortran

- A **metadirective** that appears in the specification part of a subprogram must follow all **variant-generating 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

15.4.1 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: <i>directive-name</i> (a 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 15.3**. 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**, see [Section 15.4.4](#)
- Context Selectors, see [Section 15.2](#)
- **metadirective**, see [Section 15.4.3](#)
- **nothing** Directive, see [Section 16.7](#)

15.4.2 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: <i>directive-name</i> (a 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

- **begin metadirective**, see [Section 15.4.4](#)
- **metadirective**, see [Section 15.4.3](#)
- **when** Clause, see [Section 15.4.1](#)

15.4.3 metadirective

Name: metadirective Category: meta	Association: unassociated Properties: pure
---	---

Clauses

[otherwise](#), [when](#)

Semantics

The [metadirective](#) specifies [metadirective](#) semantics.

Cross References

- Metadirectives, see [Section 15.4](#)
- **otherwise** Clause, see [Section 15.4.2](#)
- **when** Clause, see [Section 15.4.1](#)

15.4.4 begin metadirective

Name: begin metadirective Category: meta	Association: delimited Properties: pure
---	--

Clauses

[otherwise](#), [when](#)

Semantics

The [begin metadirective](#) is a [metadirective](#) that is a [delimited directive](#) and 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 required paired [end directive](#) is implicitly replaced by the [end directive](#) of the [directive variant](#) to demarcate the statements that are associated with the [directive variant](#). If the [nothing directive](#) is selected to replace the [begin metadirective directive](#), the [end directive](#) 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

- Metadirectives, see [Section 15.4](#)
- **nothing** Directive, see [Section 16.7](#)
- **otherwise** Clause, see [Section 15.4.2](#)
- **when** Clause, see [Section 15.4.1](#)

15.5 Semantic Requirement Set

The [semantic requirement set](#) of each [task](#) is a logical set of elements that can be added to or removed from the set by different [directives](#) in the scope of the [task region](#), as well as affect the semantics of those [directives](#).

A [directive](#) can add the following elements to the set:

- *depend*, which specifies that a [construct](#) requires enforcement of the synchronization relationship expressed by the [depend clause](#);
- *nowait*, which specifies that a [construct](#) is asynchronous;
- *is_device_ptr(list-item)*, which specifies that the *list-item* is a [device pointer](#) in a [construct](#);
- *has_device_addr(list-item)*, which specifies that the *list-item* has a [device address](#) in a [construct](#); and
- *interop(list-item)*, which specifies that the *list-item* is a user-provided [interoperability object](#) to be used in a [construct](#). The order in which the *interop* elements are added is relevant.

If an implementation supports the [unified_address](#) requirement then:

- Adding an *is_device_ptr* element for a [list item](#) also adds a *has_device_addr* element for any [data entity](#) for which the [list item](#) is a [base pointer](#); and
- Adding a *has_device_addr* element for a [list item](#) that has a [base pointer](#) also adds an *is_device_ptr* element for that [base pointer](#) if the [base pointer](#) is an identifier.

The following [directives](#) may add elements to the set:

- **dispatch**.

The following [directives](#) may remove elements from the set:

- **declare_variant**

Cross References

- **dispatch** Construct, see [Section 15.7](#)
- Declare Variant Directives, see [Section 15.6](#)

15.6 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 15.1.

For a function variant to be a replacement candidate to be called instead of the base function, its declare variant directive for the base function must be visible at the call site and the static part of its associated context selector must be compatible with the OpenMP context of the call according to the matching rules defined in Section 15.3. In addition, if the base function is called from a non-host device, the declare variant directive must not specify an **append_args** clause or an **adjust_args** clause with a **need_device_ptr** or **need_device_addr** *adjust-op*.

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++

Fortran

- The characteristic of the `function variant` must be compatible with the characteristic of the `base function` after the `implementation defined` transformation for its `OpenMP context`.

Fortran

Cross References

- Context Selectors, see [Section 15.2](#)
- OpenMP Contexts, see [Section 15.1](#)

15.6.1 match Clause

Name: match		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: <i>directive-name</i> (a 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 15.6.5](#)
- `declare_variant` Directive, see [Section 15.6.4](#)
- Context Selectors, see [Section 15.2](#)

15.6.2 `adjust_args` Clause

Name: <code>adjust_args</code>	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: <code>need_device_addr</code> , <code>need_device_ptr</code> , <code>nothing</code>	<i>required</i>
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: <i>directive-name</i> (a <i>directive name</i>)	<i>unique</i>

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 in the context of a *function-dispatch structured block*. 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*. Any argument specified in the *clause* that does not exist at a given *function* call site is ignored.

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* element for a given argument in its *semantic requirement set* (as added by the *dispatch construct* that encloses the call to the *base function*), 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*, except that if the argument cannot be converted into a *device pointer* then **NULL** is passed as the argument.

If the *adjust-op* modifier is **need_device_addr**, the arguments are replaced with references to the corresponding objects in the *device data environment* of the default *device* if they do not already have *device addresses*. If the *current task* has a *has_device_addr* element for a given argument in its *semantic requirement set*, as added by the *dispatch construct* that encloses the call to the *base function*, the argument is not adjusted. Otherwise, the argument is converted in the same manner that a *use_device_addr* *clause* on a *target_data* *construct* replaces references to the *list items*.

Restrictions

- If the **need_device_addr** *adjust-op* modifier is present and the *has-device-addr* element does not exist for a specified argument in the *semantic requirement set* of the *current task*, all restrictions that apply to a *list item* in a *use_device_addr* *clause* also apply to the corresponding argument that is passed by the call.

C

- If the **need_device_ptr** *adjust-op* modifier is present, each *list item* that appears in the *clause* that refers to a specific named argument in the declaration of the *function variant* must be of pointer type.
- The **need_device_addr** *adjust-op* modifier must not be specified in the *clause*.

C

C++

- If the **need_device_ptr** *adjust-op* modifier is present, each *list item* that appears in the *clause* that refers to a specific named argument in the declaration of the *function variant* must be of pointer type or reference to pointer type.
- If the **need_device_addr** *adjust-op* modifier is present, each *list item* that appears in the *clause* must refer to an argument in the declaration of the *function variant* that has a reference type.

C++

Fortran

- If the `need_device_ptr` *adjust-op* modifier is present, each `list item` that appears in the `clause` must refer to a `C pointer` dummy argument in the declaration of the `function variant`.
- If the `need_device_addr` *adjust-op* modifier is present, each `list item` that appears in the `clause` must refer to a dummy argument in the declaration of the `function variant` that does not have the `VALUE` attribute.
- If the `need_device_addr` *adjust-op* modifier is present, the corresponding actual argument for each specified argument must be contiguous.

Fortran

Cross References

- `declare_variant` Directive, see [Section 15.6.4](#)
- `use_device_addr` Clause, see [Section 7.3.9](#)
- `use_device_ptr` Clause, see [Section 7.3.7](#)

15.6.3 `append_args` Clause

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

Arguments

Name	Type	Properties
<i>append-op-list</i>	list of OpenMP operation list item type	<i>default</i>

Modifiers

Name	Modifies	Type	Properties
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: <i>directive-name</i> (a directive name)	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 in the context of a [function-dispatch structured block](#). The arguments are formed according to each specified `list item` in *append-op-list*, in the order those `list items` appear. The arguments are passed to the `function variant` after any named arguments of the `base function` in the same order in which they are formed. If the `base function` is variadic, the formed 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.

For each **interop** operation specified, an argument is formed and appended as follows. If the **semantic requirement set** contains one or more *interop* elements, the first of those elements that was added to the set is removed and the associated **interoperability object** of that removed element is appended as an argument. Otherwise, the **interop** operation constructs an argument of **interop OpenMP type** using the **semantic 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 **semantic requirement set** contains one or more elements (as added by the **dispatch** construct) that correspond to **clauses** for an **interop** construct of *interop-type*, the behavior is as if the corresponding **clauses** are specified on the **interop** construct and those elements are removed from the **semantic requirement set**.

Any appended arguments that were not obtained from the *interop* elements of the **semantic requirement set** are 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

- **declare_variant** Directive, see [Section 15.6.4](#)
- **destroy** Clause, see [Section 5.8](#)
- OpenMP Operations, see [Section 5.2.3](#)
- Semantic Requirement Set, see [Section 15.5](#)
- **init** Clause, see [Section 5.7](#)
- **interop** Construct, see [Section 22.1](#)

15.6.4 declare_variant Directive

Name: <code>declare_variant</code>	Association: declaration
Category: declarative	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	default

Clauses

`adjust_args`, `append_args`, `match`

Additional information

The `declare_variant directive` may alternatively be specified with `declare variant` as the *directive-name*.

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

The restrictions to the `declare_variant directive` are as follows:

C / C++

- If *base-name* is specified, it must match the name used in the associated declaration, if any declaration is associated.

C / C++

C++

- If an expression in the `context selector` that appears in a `match clause` references the `this` pointer, the `base function` must be a non-static member function.

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 15.6.2](#)
- **append_args** Clause, see [Section 15.6.3](#)
- Declare Variant Directives, see [Section 15.6](#)
- **match** Clause, see [Section 15.6.1](#)

15.6.5 begin declare_variant Directive

Name: begin declare_variant Category: declarative	Association: delimited Properties: default
---	---

Clauses

[match](#)

Additional information

The **begin declare_variant** directive may alternatively be specified with **begin declare variant** as the *directive-name*.

Semantics

The **begin declare_variant** directive associates the **context selector** in the **match** clause with each procedure definition in the delimited code region formed by the **directive** and its paired **end directive**. The delimited code region is a **declaration sequence**. For the purpose of call resolution, each function definition that appears in the delimited code region is a **function variant** for an assumed **base function**, with the same name and a compatible prototype, that is declared elsewhere without an associated **declare variant directive**.

If a **declare variant directive** appears between a **begin declare_variant** directive and its paired **end directive**, the **effective context selectors** of the outer **directive** are appended to the **context selector** of the inner **directive** to form the **effective context selector** of the inner **directive**. If a *trait-set-selector* is present on both **directives**, the *trait-selector* list of the outer **directive** is appended to the *trait-selector* list of the inner **directive** after equivalent *trait-selectors* have been

removed from the outer list. Restrictions that apply to explicitly specified **context selectors** also apply to **effective context selectors** constructed through this process.

The symbol name of a function definition that appears between a **begin declare_variant directive** and its paired **end directive** is determined through the **base language** rules after the name of the **function** has been augmented with a string that is determined according to the **effective context selector** of the **begin declare_variant directive**. The symbol names of two definitions of a **function** are considered to be equal if and only if their **effective context selectors** are equivalent.

If the **context selector** of a **begin declare_variant directive** contains **traits** in the **device trait set** or **implementation trait set** that are known never to be compatible with an **OpenMP context** during the current compilation, the **preprocessed code** that follows the **begin declare_variant directive** up to its paired **end directive** is elided.

Any expressions in the **match clause** are interpreted at the location of the **directive**.

Restrictions

The restrictions to **begin declare_variant directive** are as follows:

- **match clause** must not contain a **simd trait selector**.
- Two **begin declare_variant directives** and their paired **end directives** must either encompass disjoint source ranges or be perfectly nested.

▼ C++ ▼

- A **match clause** must not contain a **dynamic context selector** that references the **this** pointer.

▲ C++ ▲

▼ Fortran ▼

- The delimited code region must be one of the following:
 - one or more internal subprograms;
 - one or more module subprograms; or
 - one or more statements in a specification part.
- Procedure definitions in the delimited code region must use the same argument keywords as in the **base function** declaration, if any.
- If the delimited code region is in the specification part, the **context selector** may only contain **traits** in the **device trait set** or **implementation trait set**.
- If the delimited code region is one or more internal subprograms or module subprograms, any variable referenced in the expressions that appear in a context selector of the directive must be accessible as if by host association.

▲ Fortran ▲

1 **Cross References**

- 2 • Declare Variant Directives, see [Section 15.6](#)
3 • `match` Clause, see [Section 15.6.1](#)

4 **15.7 dispatch Construct**

5

Name: <code>dispatch</code> Category: <code>executable</code>	Association: <code>block</code> : <code>function-dispatch</code> Properties: <code>context-matching</code>
--	---

6 **Clauses**

7 `depend`, `device`, `has_device_addr`, `interop`, `is_device_ptr`, `nocontext`,
8 `novariants`, `nowait`

9 **Binding**

10 The `binding task set` for a `dispatch` region is the `generating task`. The `dispatch` region binds
11 to the `region` of the `generating task`.

12 **Semantics**

13 The `dispatch` construct controls whether `variant substitution` occurs for `target-call` in the
14 associated `function-dispatch structured block`. The `dispatch` construct may also modify the
15 `semantic requirement set` of elements that affect the arguments of the `function variant` if `variant`
16 `substitution` occurs (see [Section 15.6.2](#) and [Section 15.6.3](#)).

17 Elements added to the `semantic requirement set` by the `dispatch` construct can be removed by
18 the effect of `declare variant directives` (see [Section 15.5](#)) before the `dispatch` region is executed.
19 If one or more `depend` clauses are present on the `dispatch` construct, they are added as `depend`
20 elements of the `semantic requirement set`. If a `nowait` clause is present on the `dispatch`
21 construct the `nowait` element is added to the `semantic requirement set`. For each `list item` specified
22 in an `is_device_ptr` clause, an `is_device_ptr` element for that `list item` is added to the `semantic`
23 `requirement set`. For each `list item` specified in a `has_device_addr` clause, a `has_device_addr`
24 element for that `list item` is added to the `semantic requirement set`. For each `list item` specified in an
25 `interop` clause, an `interop` element for that `list item` is added to the `semantic requirement set` in
26 the same order that they were specified on the `directive`.

27 If the `dispatch` directive adds one or more `depend` element to the `semantic requirement set`, and
28 those element are not removed by the effect of a `declare variant directive`, the behavior is as if those
29 elements were applied as `depend` clauses to a `taskwait` construct that is executed before the
30 `dispatch` region is executed.

31 The addition of the `nowait` and `interop` elements to the `semantic requirement set` by the `dispatch`
32 `directive` has no effect on the `dispatch` construct apart from the effect it may have on the
33 arguments that are passed when calling a `function variant`.

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 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 23.9.5](#)
- **device** Clause, see [Section 21.2](#)
- OpenMP Function Dispatch Structured Blocks, see [Section 6.3.2](#)
- Semantic Requirement Set, see [Section 15.5](#)
- **has_device_addr** Clause, see [Section 7.3.8](#)
- **interop** Clause, see [Section 15.7.1](#)
- **is_device_ptr** Clause, see [Section 7.3.6](#)
- **nocontext** Clause, see [Section 15.7.3](#)
- **novariants** Clause, see [Section 15.7.2](#)
- **nowait** Clause, see [Section 23.6](#)
- **taskwait** Construct, see [Section 23.5](#)

15.7.1 interop Clause

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

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: <i>directive-name</i> (a <i>directive name</i>)	<code>unique</code>

Directives

dispatch

Semantics

The **interop** clause specifies interoperability objects to be added to the semantic requirement set of the encountering task. They are added to the semantic requirement set 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** Construct, see [Section 15.7](#)

15.7.2 novariants Clause

Name: novariants	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: <i>directive-name</i> (a <i>directive name</i>)	<i>unique</i>

Directives

dispatch

Semantics

If *do-not-use-variant* evaluates to *true*, no *function variant* is selected for the *target-call* of the **dispatch** region associated with the **novariants** clause even if one would be selected normally. The use of a *variable* in *do-not-use-variant* causes an implicit reference to the *variable* in all enclosing constructs. *do-not-use-variant* is evaluated in the enclosing context.

Cross References

- **dispatch** Construct, see [Section 15.7](#)

15.7.3 nocontext Clause

Name: nocontext	Properties: unique
------------------------	------------------------------------

Arguments

Name	Type	Properties
<i>do-not-update-context</i>	expression of OpenMP logical type	default

Modifiers

Name	Modifies	Type	Properties
directive-name-modifier	<i>all arguments</i>	Keyword: <i>directive-name</i> (a 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](#) Construct, see [Section 15.7](#)

15.8 declare_simd Directive

Name: declare_simd	Association: declaration
Category: declarative	Properties: pure , variant-generating

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](#)

1 **Additional information**

2 The `declare_simd` directive may alternatively be specified with `declare simd` as the
3 *directive-name*.

4 **Semantics**

5 The association of one or more `declare_simd` directives with a `procedure` declaration or
6 definition enables the creation of corresponding `SIMD` versions of the associated `procedure` that
7 can be used to process multiple arguments from a single invocation in a `SIMD loop` concurrently.

8 If a `SIMD` version is created and the `simdlen` clause is not specified, the number of concurrent
9 arguments for the function is *implementation defined*.

10 For purposes of the `linear` clause, any integer-typed parameter that is specified in a `uniform`
11 *clause* on the *directive* is considered to be constant and so may be used in a *step-complex-modifier*
12 as *linear-step*.

▼ C / C++ ▼

13 The expressions that appear in the `clauses` of each *directive* are evaluated in the scope of the
14 arguments of the `procedure` declaration or definition.

▲ C / C++ ▲

▼ C++ ▼

15 The special `this` pointer can be used as if it was one of the arguments to the `procedure` in any of
16 the `linear`, `aligned`, or `uniform` clauses.

▲ C++ ▲

17 **Restrictions**

18 Restrictions to the `declare_simd` directive are as follows:

- 19 • The `procedure` body must be a `structured block`.
- 20 • The execution of the `procedure`, when called from a `SIMD loop`, must not result in the
21 execution of any `constructs` except for `atomic` constructs and `ordered` constructs on
22 which the `simd` clause is specified.
- 23 • The execution of the `procedure` must not have any side effects that would alter its execution
24 for concurrent iterations of a `SIMD chunk`.

▼ C / C++ ▼

- 25 • If a `declare_simd` directive is specified for a declaration of a `procedure` then the
26 definition of the `procedure` must have a `declare_simd` directive with identical `clauses`
27 with identical arguments and `modifiers`.
- 28 • The `procedure` must not contain calls to the `longjmp` or `setjmp` functions.

▲ C / C++ ▲

C++

- The `procedure` must 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 `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` then the definition of the `procedure` must contain a `declare_simd` directive with identical `clauses` with identical arguments and `modifiers`.
- `Procedures` pointers may not be used to access versions created by the `declare_simd` directive.

Fortran

Cross References

- `aligned` Clause, see [Section 14.2](#)
- `linear` Clause, see [Section 8.14](#)
- `simdlen` Clause, see [Section 18.4.3](#)
- `uniform` Clause, see [Section 14.1](#)

15.8.1 *branch* Clauses

Clause groups

Properties: <code>exclusive</code> , <code>unique</code>	Members: Clauses <code>inbranch</code> , <code>notinbranch</code>
---	---

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 15.8](#)

15.8.1.1 inbranch Clause

Name: <code>inbranch</code>	Properties: <i>unique</i>
------------------------------------	----------------------------------

Arguments

Name	Type	Properties
<i>inbranch</i>	expression of OpenMP logical type	<i>constant, optional</i>

Modifiers

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

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 15.8](#)

15.8.1.2 notinbranch Clause

Name: <code>notinbranch</code>	Properties: <i>unique</i>
---------------------------------------	----------------------------------

Arguments

Name	Type	Properties
<i>notinbranch</i>	expression of OpenMP logical type	<i>constant, optional</i>

1 **Modifiers**

2

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

3 **Directives**

4 **declare_simd**

5 **Semantics**

6 If *notinbranch* evaluates to *true*, the **notinbranch** clause specifies that the *procedure* will never
7 be called from inside a conditional statement of the calling context. If *notinbranch* evaluates to
8 *false*, the *procedure* may be called from inside a conditional statement. If *notinbranch* is not
9 specified, the effect is as if *notinbranch* evaluates to *true*.

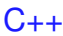
10 **Cross References**

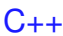
- 11
 - **declare_simd** Directive, see [Section 15.8](#)

12 **15.9 Declare Target Directives**

13 **Declare target directives** apply to *procedures* and/or *variables* to ensure that they can be executed or
14 accessed on a *device*. *Variables* are either replicated as *device-local variables* for each *device*
15 through a **local** clause, are mapped for all *device* executions through an **enter** clause, or are
16 mapped for specific *device* executions through a **link** clause. An implementation may generate
17 different versions of a *procedure* to be used for **target regions** that execute on different *devices*.
18 Whether it generates different versions, and whether it calls a different version in a **target region**
19 from the version that it calls outside a **target region**, are *implementation defined*.

20 To facilitate *device* usage, OpenMP defines rules that implicitly specify **declare target directives** for
21 *procedures* and *variables*. The remainder of this section defines those rules as well as restrictions
22 that apply to all **declare target directives**.

23  If a *variable* with *static storage duration* has the **constexpr** specifier and is not a *groupprivate*
24 *variable* then the *variable* is treated as if it had appeared as a *list item* in an **enter** clause on a
25 **declare target directive**.

26  If a *variable* with *static storage duration* that is not a *device-local variable* (including that it is not a
27 *groupprivate variable*) is declared in a *device procedure* then the *variable* is treated as if it had
28 appeared as a *list item* in an **enter** clause on a **declare target directive**.

29 If a *procedure* is referenced outside of any *reverse-offload region* in a *procedure* that appears as a
30 *list item* in an **enter** clause on a *non-host declare target directive* then the name of the referenced
31 *procedure* is treated as if it had appeared in an **enter** clause on a **declare target directive**.

C / C++

If a **variable** with **static storage duration** or a function (except *lambda* for C++) is referenced in the **initializer expression list** of a **variable** with **static storage duration** that appears as a **list item** in an **enter** or **local** clause on a **declare target directive** then the name of the referenced **variable** or **procedure** is treated as if it had appeared in an **enter** clause on a **declare target directive**.

C / C++

Fortran

If a **declare_target** directive has a **device_type** clause then any enclosed internal **procedure** cannot contain any **declare_target** directives. The enclosing **device_type** clause implicitly applies to internal **procedures**.

Fortran

A reference to a **device-local** **variable** that has **static storage duration** inside a **device** **procedure** is replaced with a reference to the copy of the **variable** for the **device**. Otherwise, a reference to a **variable** that has **static storage duration** in a **device** **procedure** is replaced with a reference to a corresponding **variable** in the **device data environment**. If the corresponding **variable** does not exist or the **variable** does not appear in an **enter** or **link** clause on a **declare target directive**, the behavior is unspecified.

Execution Model Events

The *target-global-data-op* **event** occurs when an **original list item** is associated with a **corresponding list item** on a **device** as a result of a **declare target directive**; the **event** occurs before the first access to the **corresponding list item**.

Tool Callbacks

A **thread** dispatches a registered **target_data_op_emi** callback with **ompt_scope_beginend** as its *endpoint* argument for each occurrence of a *target-global-data-op* **event** in that **thread**.

Restrictions

Restrictions to any **declare target directive** are as follows:

- The same **list item** must not explicitly appear in both an **enter** clause on one **declare target directive** and a **link** or **local** clause on another **declare target directive**.
- The same **list item** must not explicitly appear in both a **link** clause on one **declare target directive** and a **local** clause on another **declare target directive**.
- If a **variable** appears in an **enter** clause on a **declare target directive**, its initializer must not refer to a **variable** that appears in a **link** clause on a **declare target directive**.

Cross References

- `begin declare_target` Directive, see [Section 15.9.2](#)
- `declare_target` Directive, see [Section 15.9.1](#)
- `enter` Clause, see [Section 9.4](#)
- `link` Clause, see [Section 9.5](#)
- OMPT `scope_endpoint` Type, see [Section 39.27](#)
- `target` Construct, see [Section 21.8](#)
- `target_data_op_emi` Callback, see [Section 41.7](#)

15.9.1 declare_target Directive

Name: <code>declare_target</code> Category: declarative	Association: explicit Properties: declare-target , device , pure , variant-generating
--	---

Arguments

`declare_target` (*extended-list*)

Name	Type	Properties
<i>extended-list</i>	list of extended list item type	optional

Clauses

[device_type](#), [enter](#), [indirect](#), [link](#), [local](#)

Additional information

The [declare_target directive](#) may alternatively be specified with `declare target` as the *directive-name*.

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 *list* argument of an implicit [enter clause](#) and any [list items](#) that are [groupprivate variables](#) appear in the *list* argument of an implicit [local clause](#).

If neither the *extended-list* argument nor a [data-environment attribute clause](#) is specified then the [directive](#) is a [declaration-associated directive](#). The effect is as if the name of the associated [procedure](#) appears as a [list item](#) in an [enter clause](#) of a [declare target directive](#) that otherwise specifies the same set of [clauses](#).

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++

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** is not a **declaration-associated directive** and an *extended-list* argument is not specified, a **data-environment attribute clause** must be present.
- A **variable** for which **nohost** is specified must not appear in a **link clause**.
- A **groupprivate variable** must not appear in any **enter clauses** or **link clauses**.

C / C++

- If the **directive** is not a **declaration-associated directive**, it must appear at the same scope as the declaration of every **list item** in its *extended-list* or in its **data-environment attribute clauses**.

C / C++

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 the **directive** is a **declaration-associated directive**, the **directive** must appear in the specification part of a subroutine subprogram, function subprogram or interface body.
- If a **list item** is a **procedure** name that is not declared via a **procedure** declaration statement, the **directive** must be in the specification part of the subprogram or interface body of that **procedure**.
- If a **list item** in *extended-list* is a **variable**, the **directive** must appear in the specification part in which the **variable** is declared.
- If a **declare_target** directive is specified for a **procedure** that has an explicit interface then the definition of the **procedure** must contain a **declare_target** directive with identical **clauses** with identical arguments and **modifiers**.
- If an external **procedure** is a type-bound **procedure** of a derived type and the **directive** is specified in the definition of the external **procedure**, it must appear in the interface block that is accessible to the derived-type definition.

- If any `procedure` is declared via a `procedure` declaration statement that is not in the type-bound `procedure` part of a derived-type definition, any `declare_target` directive with the `procedure` name must appear in the same specification part.
- If a `declare_target` directive that specifies a common block name appears in one program unit, then such a `directive` must also appear in every other program unit that contains a `COMMON` statement that specifies the same name, after the last such `COMMON` statement in the program unit.
- If a `list item` is declared with the `BIND` attribute, the corresponding C entities must also be specified in a `declare_target` directive in the C program.
- A `variable` can only appear in a `declare_target` directive in the scope in which it is declared. It must not be an element of a common block or appear in an `EQUIVALENCE` statement.

Fortran

Cross References

- `device_type` Clause, see [Section 21.1](#)
- `enter` Clause, see [Section 9.4](#)
- Declare Target Directives, see [Section 15.9](#)
- `indirect` Clause, see [Section 15.9.3](#)
- `link` Clause, see [Section 9.5](#)
- `local` Clause, see [Section 9.3](#)

C / C++

15.9.2 begin declare_target Directive

Name: <code>begin declare_target</code> Category: <code>declarative</code>	Association: <code>delimited</code> Properties: <code>declare-target</code> , <code>device</code> , <code>variant-generating</code>
---	--

Clauses

`device_type`, `indirect`

Additional information

The `begin declare_target` directive may alternatively be specified with `begin declare target` as the *directive-name*.

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 delimited code region is a **declaration sequence**.

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:

C++

- The function names of overloaded functions or template functions may only be specified within an implicit *extended-list*.
- If a *lambda declaration and definition* appears between a **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 21.1](#)
- **enter** Clause, see [Section 9.4](#)
- **Declare Target Directives**, see [Section 15.9](#)
- **indirect** Clause, see [Section 15.9.3](#)

C / C++

15.9.3 indirect Clause

Name: indirect	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: <i>directive-name</i> (a <i>directive name</i>)	unique

Directives

begin declare_target, declare_target

Semantics

If *invoked-by-fptr* evaluates to *true*, any *procedures* 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 on that *procedure*.

C / C++

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 15.9.2](#)
- **declare_target** Directive, see [Section 15.9.1](#)

16 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 is an [executable directive](#).

16.1 error Directive

Name: <code>error</code> Category: utility	Association: unassociated 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 the *action-time* of the [at clause](#) 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 [error callback](#) for each occurrence of a *runtime-error event* in the context of the [encountering task](#).

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 16.2](#)
- [error](#) Callback, see [Section 40.2](#)

- **message** Clause, see [Section 16.3](#)
- **severity** Clause, see [Section 16.4](#)

16.2 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: <i>directive-name</i> (a 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 16.1](#)

16.3 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: <i>directive-name</i> (a <i>directive name</i>)	<i>unique</i>

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 16.1](#)
- **parallel** Construct, see [Section 18.1](#)

16.4 severity Clause

Name: severity	Properties: <i>unique</i>
------------------------------	----------------------------------

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: <i>directive-name</i> (a <i>directive name</i>)	<i>unique</i>

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 16.1](#)
- **parallel** Construct, see [Section 18.1](#)

16.5 requires Directive

Name: requires	Association: unassociated
Category: informational	Properties: default

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** must not appear lexically after the specification of a **context selector** in which any **clause** of that **requires directive** is used, and it must not appear lexically after any code that depends on such a **context selector**.
- The **requires directive** must only appear at file scope.
- The **requires directive** must only appear at file or namespace scope.
- Any **requires directive** that specifies a **device global requirement clause** must appear lexically before any **device constructs** or **device procedures**.

- 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.

16.5.1 *requirement* Clauses

Clause groups

Properties: required, unique

Members:

Clauses

atomic_default_mem_order,
device_safesync,
dynamic_allocators,
reverse_offload,
self_maps, **unified_address**,
unified_shared_memory

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 16.5](#)

16.5.1.1 **atomic_default_mem_order** Clause

Name: **atomic_default_mem_order**

Properties: unique

1 **Arguments**

2

Name	Type	Properties
<i>memory-order</i>	Keyword: acq_rel , acquire , relaxed , release , seq_cst	<i>default</i>

3 **Modifiers**

4

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

5 **Directives**

6 **requires**

7 **Semantics**

8 The **atomic_default_mem_order** clause specifies the default memory ordering behavior for
9 **atomic constructs** that an implementation must provide. The effect is as if its argument appears as
10 a clause on any **atomic construct** that does not specify a *memory-order* clause.

11 **Restrictions**

12 Restrictions to the **atomic_default_mem_order** clause are as follows:

- 13
- All **requires directives** in the same **compilation unit** that specify the

14 **atomic_default_mem_order** requirement must specify the same argument.

 - Any **directive** that specifies the **atomic_default_mem_order** clause must not appear

15 lexically after any **atomic construct** on which a *memory-order* clause is not specified.

16

17 **Cross References**

- 18
- **atomic** Construct, see [Section 23.8.5](#)

19

 - *memory-order* Clauses, see [Section 23.8.1](#)

20

 - **requires** Directive, see [Section 16.5](#)

21 **16.5.1.2 dynamic_allocators Clause**

22

Name: dynamic_allocators	Properties: unique
---------------------------------	---------------------------

23 **Arguments**

24

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: <i>directive-name</i> (a <i>directive name</i>)	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 13.6](#)
- *def-allocator-var* ICV, see [Table 3.1](#)
- **omp_destroy_allocator** Routine, see [Section 33.7](#)
- **omp_init_allocator** Routine, see [Section 33.6](#)
- **requires** Directive, see [Section 16.5](#)
- **target** Construct, see [Section 21.8](#)
- **uses_allocators** Clause, see [Section 13.8](#)

16.5.1.3 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: <i>directive-name</i> (a <i>directive name</i>)	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** *device-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*.

Cross References

- **device** Clause, see [Section 21.2](#)
- **requires** Directive, see [Section 16.5](#)
- **target** Construct, see [Section 21.8](#)

16.5.1.4 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: <i>directive-name</i> (a <i>directive name</i>)	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*.

Cross References

- `is_device_ptr` Clause, see Section 7.3.6
- `requires` Directive, see Section 16.5
- `target` Construct, see Section 21.8

16.5.1.5 `unified_shared_memory` Clause

Name: <code>unified_shared_memory</code>	Properties: <code>unique</code> , device global requirement
--	---

Arguments

Name	Type	Properties
<i>required</i>	expression of OpenMP logical type	<code>constant</code> , <code>optional</code>

Modifiers

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

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 16.5.1.4). 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`.

Cross References

- `enter` Clause, see [Section 9.4](#)
- `requires` Directive, see [Section 16.5](#)
- `unified_address` Clause, see [Section 16.5.1.4](#)

16.5.1.6 `self_maps` Clause

Name: <code>self_maps</code>	Properties: <code>unique</code> , <code>device global requirement</code>
------------------------------	--

Arguments

Name	Type	Properties
<code>required</code>	expression of OpenMP logical type	<code>constant</code> , <code>optional</code>

Modifiers

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

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`. If `required` is not specified, the effect is as if `required` evaluates to `true`.

Cross References

- **enter** Clause, see [Section 9.4](#)
- **requires** Directive, see [Section 16.5](#)
- **unified_shared_memory** Clause, see [Section 16.5.1.5](#)

16.5.1.7 device_safesync Clause

Name: device_safesync	Properties: unique , device global requirement
------------------------------	---

Arguments

Name	Type	Properties
<i>required</i>	expression of OpenMP logical type	constant , optional

Modifiers

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

Directives

[requires](#)

Semantics

If *required* evaluates to [true](#), the [device_safesync](#) clause indicates that any two [synchronizing divergent threads](#) in a [team](#) that execute on a [non-host device](#) must be able to make progress, unless indicated otherwise by the use of a [safesync](#) clause. If *required* is not specified, the effect is as if *required* evaluates to [true](#).

Cross References

- **requires** Directive, see [Section 16.5](#)
- **safesync** Clause, see [Section 18.1.5](#)

16.6 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 specified by the [assumption directive](#) do not hold at runtime, the behavior is [unspecified](#).

16.6.1 assumption Clauses

Clause groups

Properties: required , unique	Members: Clauses absent , contains , holds , no_omp , no_omp_constructs , no_omp_routines , no_parallelism
--	--

Directives

[assume](#), [assumes](#), [begin assumes](#)

Semantics

The [assumption clause group](#) defines a [clause set](#) that indicates 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 one of its [constituent 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 [directive](#) that is a [declarative directive](#), an [informational directive](#), or a [metadirective](#).

Cross References

- [assume](#) Directive, see [Section 16.6.3](#)
- [assumes](#) Directive, see [Section 16.6.2](#)
- [begin assumes](#) Directive, see [Section 16.6.4](#)

16.6.1.1 absent Clause

Name: absent	Properties: unique
-------------------------------------	---

Arguments

Name	Type	Properties
directive-name-list	list of directive-name list item type	default

Modifiers

Name	Modifies	Type	Properties
directive-name-modifier	all arguments	Keyword: directive-name (a directive name)	unique

Directives

[assume](#), [assumes](#), [begin assumes](#)

Semantics

The [absent](#) clause specifies that the program guarantees that no [construct](#) that matches a *directive-name* [list item](#) is encountered in the [assumption scope](#).

Cross References

- **assume** Directive, see [Section 16.6.3](#)
- **assumes** Directive, see [Section 16.6.2](#)
- **begin assumes** Directive, see [Section 16.6.4](#)

16.6.1.2 contains Clause

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

Arguments

Name	Type	Properties
<i>directive-name-list</i>	list of <i>directive-name</i> list item type	default

Modifiers

Name	Modifies	Type	Properties
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: <i>directive-name</i> (a <i>directive name</i>)	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 16.6.3](#)
- **assumes** Directive, see [Section 16.6.2](#)
- **begin assumes** Directive, see [Section 16.6.4](#)

16.6.1.3 holds Clause

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

1 **Arguments**

2

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

3 **Modifiers**

4

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

5 **Directives**

6 ***assume***, ***assumes***, ***begin assumes***

7 **Semantics**

8 When the ***holds*** clause appears on an **assumption directive**, the program guarantees that the listed
9 expression evaluates to *true* in the **assumption scope**. The effect of the **clause** does not include any
10 evaluation of the expression that affects the behavior of the program.

11 **Cross References**

- 12 • ***assume*** Directive, see [Section 16.6.3](#)
13 • ***assumes*** Directive, see [Section 16.6.2](#)
14 • ***begin assumes*** Directive, see [Section 16.6.4](#)

15 **16.6.1.4 no_openmp Clause**

16

Name: <i>no_openmp</i>	Properties: <i>unique</i>
-------------------------------	----------------------------------

17 **Arguments**

18

Name	Type	Properties
<i>can_assume</i>	expression of OpenMP logical type	<i>constant</i> , <i>optional</i>

19 **Modifiers**

20

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

21 **Directives**

22 ***assume***, ***assumes***, ***begin assumes***

23 **Semantics**

24 If *can_assume* evaluates to *true*, the ***no_openmp*** clause implies the ***no_openmp_constructs***
25 **clause** and the ***no_openmp_routines*** clause. If *can_assume* is not specified, the effect is as if
26 *can_assume* evaluates to *true*.

The `no_omp` 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`.

Cross References

- `assume` Directive, see [Section 16.6.3](#)
- `assumes` Directive, see [Section 16.6.2](#)
- `begin assumes` Directive, see [Section 16.6.4](#)

16.6.1.5 no_omp_constructs Clause

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

Arguments

Name	Type	Properties
<code>can_assume</code>	expression of OpenMP logical type	<code>constant</code> , <code>optional</code>

Modifiers

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

Directives

`assume`, `assumes`, `begin assumes`

Semantics

If `can_assume` evaluates to `true`, the `no_omp_constructs` clause guarantees that no `constructs` are encountered in the `assumption scope`. If `can_assume` is not specified, the effect is as if `can_assume` evaluates to `true`.

Cross References

- `assume` Directive, see [Section 16.6.3](#)
- `assumes` Directive, see [Section 16.6.2](#)
- `begin assumes` Directive, see [Section 16.6.4](#)

16.6.1.6 no_omp_routines Clause

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

1 **Arguments**

2

Name	Type	Properties
<i>can_assume</i>	expression of OpenMP logical type	constant, optional

3 **Modifiers**

4

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

5 **Directives**

6 **assume, assumes, begin assumes**

7 **Semantics**

8 If *can_assume* evaluates to *true*, the **no_openmp_routines** clause guarantees that no OpenMP
9 API routines are executed in the **assumption scope**. If *can_assume* is not specified, the effect is as if
10 *can_assume* evaluates to *true*.

11 **Cross References**

- 12
 - **assume** Directive, see [Section 16.6.3](#)
 - **assumes** Directive, see [Section 16.6.2](#)
 - **begin assumes** Directive, see [Section 16.6.4](#)

15 **16.6.1.7 no_parallelism Clause**

16

Name: no_parallelism	Properties: unique
------------------------------------	---------------------------

17 **Arguments**

18

Name	Type	Properties
<i>can_assume</i>	expression of OpenMP logical type	constant, optional

19 **Modifiers**

20

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

21 **Directives**

22 **assume, assumes, begin assumes**

23 **Semantics**

24 If *can_assume* evaluates to *true*, the **no_parallelism** clause guarantees that no
25 **parallelism-generating constructs** will be encountered in the **assumption scope**. If *can_assume* is
26 not specified, the effect is as if *can_assume* evaluates to *true*.

Cross References

- **assume** Directive, see [Section 16.6.3](#)
- **assumes** Directive, see [Section 16.6.2](#)
- **begin assumes** Directive, see [Section 16.6.4](#)

16.6.2 assumes Directive

Name: assumes Category: informational	Association: unassociated 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 must only appear at file scope.

C

C++

- The **assumes** directive must only appear at file or namespace scope.

C++

Fortran

- The **assumes** directive must only appear in the specification part of a module or subprogram, after all **USE** statements, **IMPORT** statements, and **IMPLICIT** statements.

Fortran

16.6.3 assume Directive

Name: assume Category: informational	Association: block Properties: pure
--	--

1 **Clause groups**

2 *assumption*

3 **Semantics**

4 The *assumption scope* of the **assume** directive is the corresponding *region* and any *nested region*
5 of that *region*.

▼ C / C++ ▼

6 **16.6.4 begin assumes Directive**

Name: begin assumes	Association: <i>delimited</i>
Category: <i>informational</i>	Properties: <i>default</i>

8 **Clause groups**

9 *assumption*

10 **Semantics**

11 The *assumption scope* of the **begin assumes** directive is the code that is executed and reached
12 from any of the declared functions in the delimited code region. The delimited code region is a
13 *declaration sequence*.

▲ C / C++ ▲

14 **16.7 nothing Directive**

Name: nothing	Association: <i>unassociated</i>
Category: <i>utility</i>	Properties: <i>pure, loop-transforming</i>

16 **Clauses**

17 **apply**

18 **Loop Modifiers for the apply Clause**

<i>loop-modifier</i>	Number of Generated Loops	Description
identity (<i>default</i>)	1	the copy of the <i>transformation-affected loop</i>

21 **Semantics**

22 The **nothing** directive has no effect on the execution of the *OpenMP program* unless otherwise
23 specified by the **apply** clause.

24 If the **nothing** directive immediately precedes a *canonical loop nest* then it forms a
25 *loop-transforming construct*. It is associated with the outermost loop and generates one loop that
26 has the same *logical iterations* in the same order as the *transformation-affected loop*.

Restrictions

- The [apply](#) clause can be specified if and only if the [nothing](#) directive forms a [loop-transforming construct](#).

Cross References

- [apply](#) Clause, see [Section 17.1](#)
- Loop-Transforming Constructs, see [Chapter 17](#)

17 Loop-Transforming Constructs

A **loop-transforming construct** replaces itself, including its **associated loop nest** (see [Section 6.4.1](#)) or **associated loop sequence** (see [Section 6.4.2](#)), 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** from its associated **canonical loop sequence**. The **canonical loop nests** that precede or follow the **affected loop nests** in the associated **canonical loop sequence** will respectively precede or follow, in the generated **canonical loop sequence**, the **generated loop nest** or **generated loop sequence** that replaces the **affected loop nests**.

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 **transformation-affected loop** was not associated with any **directive**. After the execution of the **loop-transforming construct**, the **loop-iteration variables** of any of its **transformation-affected 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**.
- A **generated loop** of a **loop-transforming construct** must not be a **doacross-affected loop**.
- The arguments of any **clauses** on a **loop-transforming construct** must not refer to **loop-iteration variables** of surrounding loops in the same **canonical loop nest**.
- The *lb* and *ub* expressions of an **affected loop** (see [Section 6.4.1](#)) may only reference the **loop-iteration variable** of an enclosing loop affected by a **loop-transforming construct** if that **loop-transforming construct** has the **nonrectangular-compatible property**.

- A [generated loop](#) of a [loop-transforming construct](#) may only be a [non-rectangular affected loop](#) of an enclosing [loop-nest-associated directive](#) if that [loop-transforming construct](#) has the [nonrectangular-compatible property](#).

Cross References

- Canonical Loop Nest Form, see [Section 6.4.1](#)

17.1 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>	Complex, Keyword: flattened, fused, grid, identity, interchanged, intratile, offsets, reversed, split, unrolled Arguments: <i>indices</i> list of expression of integer type (<i>optional</i>)	<i>optional</i>
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: <i>directive-name</i> (a <i>directive name</i>)	<i>unique</i>

Directives

[fuse](#), [interchange](#), [nothing](#), [reverse](#), [split](#), [stripe](#), [tile](#), [unroll](#)

Semantics

The [apply](#) clause applies [loop-nest-associated constructs](#), specified by the *applied-directives* list, to [generated loops](#) of a [loop-transforming construct](#). The *loop-modifier* specifies to which [generated loops](#) the [directives](#) are applied. If the [loop-transforming construct](#) generates a [canonical loop sequence](#), the [generated loops](#) to which the [directives](#) are applied are the outermost loops of each [generated loop nest](#). 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. Each of the *indices* in the argument of the *loop-modifier* specifies the position of the [generated loop](#) to which the respective *applied-directives* item is applied.

If the *loop-modifier* is specified with no argument, the behavior is as if the list 1, 2, ..., m is specified, where m is the number of *generated loops* according to the specification of the *loop-modifier* keyword. 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* with the argument 1, 2, ..., m 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:

- Each *list item* in the *applied-directives* list of any **apply** clause must be **nothing** or the *directive-specification* of a *loop-nest-associated construct*.
- The *loop-transforming construct* on which the **apply** clause is specified must either have the *generally-composable property* or every *list item* in the *applied-directives* list of any **apply** clause must be the *directive-specification* of a *loop-transforming directive*.
- Every *list item* in the *applied-directives* list of any **apply** clause that is specified on a *loop-transforming construct* that is itself specified as a *list item* in the *applied-directives* list of another **apply** clause must be the *directive-specification* of a *loop-transforming directive*.
- For a given *loop-modifier* keyword, every *indices list item* may appear at most once in any **apply** clause on the *directive*.
- Every *indices list item* must be a *positive constant* less than or equal to m , the number of *generated loops* according to the specification of the *loop-modifier* keyword.
- The *list items* in *indices* must be in ascending order.
- If a *directive* does not define a default *loop-modifier* keyword, a *loop-modifier* is required.

Cross References

- **fuse** Construct, see [Section 17.4](#)
- **interchange** Construct, see [Section 17.5](#)
- **metadirective**, see [Section 15.4.3](#)
- **nothing** Directive, see [Section 16.7](#)
- **reverse** Construct, see [Section 17.6](#)
- **split** Construct, see [Section 17.7](#)
- **stripe** Construct, see [Section 17.8](#)
- **tile** Construct, see [Section 17.9](#)
- **unroll** Construct, see [Section 17.10](#)

17.2 sizes Clause

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

Arguments

Name	Type	Properties
<i>size-list</i>	list of OpenMP integer expression type	<code>positive</code>

Modifiers

Name	Modifies	Type	Properties
<i>sizes-selector</i>	<i>size-list</i>	Keyword: <code>grid</code> , <code>tile</code>	<code>default</code>
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: <i>directive-name</i> (a <code>directive name</code>)	<code>unique</code>

Directives

`stripe`, `tile`

Semantics

For a given `loop-transforming directive` on which the `clause` appears, the `sizes clause` specifies the manner in which the `logical iteration space` of the affected `canonical loop nest` is subdivided into m -dimensional grid cells that are relevant to the loop transformation, where m is the number of `list items` in *size-list*. Specifically, each `list item` in *size-list* specifies the size of the grid cells along the corresponding dimension. `List items` in *size-list* are not required to be unique.

For a given `loop-transforming directive` on which the `clause` appears, the `sizes clause` specifies the manner in which the `logical iteration space` of the `affected loops` are subdivided into an m -dimensional grid of m -dimensional `tiles` that are relevant to the `loop-transforming construct`, where m is the number of `list items` in *size-list*. If *sizes-selector* is `tile`, each `list item` in *size-list* specifies the size of a `tile` along the corresponding dimension. If *sizes-selector* is `grid`, each `list item` in *size-list* specifies the number of `tiles` along the corresponding dimension. If *sizes-selector* is not specified, the behavior is as if it is specified as `tile`. The partition and iteration order of the `tiles` in the grid is defined by the `construct` on which the `sizes clause` appears.

Restrictions

Restrictions to the `sizes clause` are as follows:

- The `loop nest depth` of the `associated loop nest` of the `loop-transforming construct` on which the `clause` is specified must be greater than or equal to m .

Cross References

- `stripe` Construct, see [Section 17.8](#)
- `tile` Construct, see [Section 17.9](#)

17.3 flatten Construct

Name: <code>flatten</code> Category: <code>executable</code>	Association: <code>loop nest</code> Properties: <code>generally-composable</code> , <code>loop-transforming</code> , <code>order-concurrent-nestable</code> , <code>pure</code> , <code>simdizable</code> , <code>teams-nestable</code>
---	---

Clauses
`apply`, `depth`

Loop Modifiers for the `apply` Clause

<i>loop-modifier</i>	Number of Generated Loops	Description
<code>flattened</code> (<i>default</i>)	1	the <code>flattened loop</code>

Semantics
The `flatten` construct combines the loops of a `canonical loop nest` into a single loop, the `flattened loop`. The number of `flatten-affected loops` is specified by the `depth` clause.
The iterations of the `flatten-affected loops` are flattened into one larger `logical iteration space` that is the `flattened iteration space`. The particular integer type used to compute the `iteration count` for the `flattened 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 `affected loop`. If any of the `transformation-affected loops` is a `non-rectangular loop`, the `flattened iteration space` may include additional empty `logical iterations`. The number of empty `logical iterations` is `implementation defined`. The `flattened loop` iterates over the `flattened iteration spaces`.

Restrictions
Restrictions to the `flatten` construct are as follows:

- The `flatten-affected loops` must be `perfectly nested loops`.

Cross References

- `apply` Clause, see [Section 17.1](#)
- `depth` Clause, see [Section 6.4.7](#)

17.4 fuse Construct

Name: <code>fuse</code> Category: <code>executable</code>	Association: <code>loop sequence</code> Properties: <code>loop-transforming</code> , <code>order-concurrent-nestable</code> , <code>pure</code> , <code>simdizable</code> , <code>teams-nestable</code>
--	--

Clauses

apply, **depth**, **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**.

The **depth** clause specifies the number of **affected loops** of each **affected loop nest**. If the **depth** clause is not present, the effect is as if a **depth** clause with a *depth-expr* equal to one was specified.

The **construct** generates a **canonical loop sequence** that consists of a **canonical loop nest** with at least *depth-expr* loops. The **affected loops** of the **affected loop nests** in the **canonical loop sequence** are fused from outermost to innermost. At each level of fusion the **generated loop** has the maximum number of **logical iterations** of any of the **affected loops** at that level. The *i*th **logical iteration** of the **generated loop** executes the *i*th **logical iteration** of each **affected loop**, in the order of their **affected loop nest** in the **associated loop sequence**. If an **affected loop** does not have an *i*th **logical iteration**, that loop is treated as if it has an *i*th **logical iteration** that is empty.

Restrictions

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

- The **transformation-affected loops** of any **affected loop nest** must be **perfectly nested**.
- No **transformation-affected loop** of any **affected loop nest** may be a **non-rectangular loop**.

Cross References

- **apply** Clause, see [Section 17.1](#)
- **looprange** Clause, see [Section 6.4.8](#)
- **depth** Clause, see [Section 6.4.7](#)

17.5 interchange Construct

Name: interchange Category: executable	Association: loop nest Properties: loop-transforming , nonrectangular-compatible , order- concurrent-nestable , pure , simdizable , teams-nestable
---	--

Clauses
apply, permutation

Loop Modifiers for the apply Clause

loop-modifier	Number of Generated Loops	Description
interchanged (de-fault)	n	the generated loops, in the new order

Semantics

The **interchange** construct has n transformation-affected 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 transformation-affected loops, from outermost to innermost. The original transformation-affected 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:

- No transformation-affected loop may be a non-rectangular loop.
- The transformation-affected loops must be perfectly nested loops.

Cross References

- **apply** Clause, see Section 17.1
- **permutation** Clause, see Section 17.5.1

17.5.1 permutation Clause

Name: permutation	Properties: unique
-------------------	--------------------

Arguments

Name	Type	Properties
permutation-list	list of OpenMP integer expression type	constant, positive

Modifiers

Name	Modifies	Type	Properties
directive-name-modifier	all arguments	Keyword: directive-name (a directive name)	unique

Directives

interchange

1 **Semantics**

2 The **permutation** clause specifies a list of n **positive constant** expressions of integer **OpenMP**
3 **type**.

4 **Restrictions**

5 Restrictions to the **permutation** clause are as follows:

- 6 • Every integer from 1 to n must appear exactly once in *permutation-list*.
7 • n must be at least 2.

8 **Cross References**

- 9 • **interchange** Construct, see [Section 17.5](#)

10 **17.6 reverse Construct**

Name: reverse Category: executable	Association: loop nest Properties: generally-composable , loop-transforming , order-concurrent- nestable , pure , simdizable , teams- nestable
---	---

12 **Clauses**

13 **apply**

14 **Loop Modifiers for the apply Clause**

<i>loop-modifier</i>	Number of Generated Loops	Description
reversed (<i>default</i>)	1	the reversed loop

17 **Semantics**

18 The **reverse** construct has one **transformation-affected loop**, the outermost loop, where
19 $0, 1, \dots, n - 2, n - 1$ are the **logical iteration** numbers of that loop. The **construct** transforms that
20 loop into a loop in which iterations occur in the order $n - 1, n - 2, \dots, 1, 0$.

21 **Cross References**

- 22 • **apply** Clause, see [Section 17.1](#)

23 **17.7 split Construct**

Name: split Category: executable	Association: loop nest Properties: generally-composable , loop-transforming , order-concurrent- nestable , pure , simdizable , teams- nestable
---	---

Clauses
apply, counts

Loop Modifiers for the apply Clause

<i>loop-modifier</i>	Number of Generated Loop Nests	Description
split	m	the loops of each logical iteration space partition

Semantics

The **split** loop-transforming construct implements index-set splitting, which partitions a logical iteration space into a sequence of smaller logical iteration spaces. It has one transformation-affected loop and generates a canonical loop sequence with m loop nests where m is the number of list items in the *count-list* argument of the **counts** clause. Let n be the number of logical iterations of the affected loop and c_1, \dots, c_m be the list items of the *count-list* argument. Let the k^{th} list item be the list item with the intrinsic identifier **omp_fill**. c_k is defined as

$$c_k = \max(0, n - \sum_{\substack{t=1 \\ t \neq k}}^m c_t)$$

Each generated loop in the sequence contains a copy of the loop body of the affected loop. The i^{th} generated loop executes the next c_i logical iterations except any logical iteration beyond the n original logical iterations.

Restrictions

The following restrictions apply to the **split** construct:

- Exactly one list item in the **counts** clause must be the intrinsic identifier **omp_fill**.

Cross References

- apply** Clause, see Section 17.1
- counts** Clause, see Section 17.7.1

17.7.1 counts Clause

Name: counts	Properties: unique, required
---------------------	-------------------------------------

Arguments

Name	Type	Properties
<i>count-list</i>	list of OpenMP integer expression type	non-negative

Modifiers

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

Directives

split

Semantics

For a given **loop-transforming directive** on which the **clause** appears, the **counts clause** specifies the manner in which the **logical iteration space** of the **transformation-affected loop** is subdivided into n partitions, where m is the number of **list items** in *count-list* and where each partition is associated with a **generated loop** of the **directive**. Specifically, each **list item** in *count-list* specifies the **iteration count** of one of the **generated loops**. **List items** in *count-list* are not required to be unique.

Cross References

- **split** Construct, see [Section 17.7](#)

17.8 stripe Construct

Name: stripe Category: executable	Association: loop nest Properties: loop-transforming , order-concurrent-nestable , pure , simdizable , teams-nestable
--	--

Clauses

apply, **sizes**

Loop Modifiers for the **apply** Clause

<i>loop-modifier</i>	Number of Generated Loops	Description
offsets	m	the offsetting loops o_1, \dots, o_m
grid	m	the grid loops g_1, \dots, g_m

Semantics

The **stripe** construct has m **transformation-affected loops**, where m is the number of **list items** in the *size-list* argument of the **sizes clause**, which consists of the **list items** s_1, \dots, s_m . The **construct** has the effect of **striping** the execution order of the **logical iterations** across the grid cells of the **logical iteration space** that result from the **sizes clause**. Let ℓ_1, \dots, ℓ_m be the **transformation-affected loops**, from outermost to innermost, which the **construct** replaces with a **canonical loop nest** that consists of $2m$ **perfectly nested loops**. Let $o_1, \dots, o_m, g_1, \dots, g_m$ be the **generated loops**, from outermost to innermost. The loops o_1, \dots, o_m are the **offsetting loops** and the loops g_1, \dots, g_m are the **grid loops**.

Let n_1, \dots, n_m be number of **logical iterations** of each **affected loop** and $O = \{G_{\alpha_1, \dots, \alpha_m} \mid \forall k \in \{1, \dots, m\} : 0 \leq \alpha_k < s_k\}$ the **logical iteration vector space** of the **offsetting loops**. The **logical iteration** (i_1, \dots, i_m) is executed in the **logical iteration space** of $G_{i_1 \bmod s_1, \dots, i_m \bmod s_m}$.

The **offsetting loops** iterate over all $G_{\alpha_1, \dots, \alpha_m}$ in **lexicographic order** of their indices and the **grid loops** iterate over the **logical iteration space** in the **lexicographic order** of the corresponding **logical iteration vectors**.

If an **offsetting loop** and a **grid loop** that are generated from the same **stripe construct** are **affected loops** of the same **loop-nest-associated construct**, the **grid loops** may execute additional empty **logical iterations**. The number of empty **logical iterations** is **implementation defined**.

Restrictions

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

- The **transformation-affected loops** must be **perfectly nested loops**.
- No **transformation-affected loops** may be a **non-rectangular loop**.
- The **sizes clause** must not specify the **grid sizes-selector**.

Cross References

- **apply** Clause, see [Section 17.1](#)
- Consistent Loop Schedules, see [Section 6.4.4](#)
- **sizes** Clause, see [Section 17.2](#)

17.9 tile Construct

Name: tile Category: executable	Association: loop nest Properties: loop-transforming , order-concurrent-nestable , pure , simdizable , teams-nestable
--	--

Clauses

apply, **sizes**

Loop Modifiers for the **apply** Clause

<i>loop-modifier</i>	Number of Generated Loops	Description
grid	m	the grid loops g_1, \dots, g_m
intratile	m	the tile loops t_1, \dots, t_m

Semantics

The **tile** construct has m transformation-affected loops, where m is the number of list items in the *size-list* argument of the **sizes** clause, which consists of list items s_1, \dots, s_m . Let ℓ_1, \dots, ℓ_m be the transformation-affected loops, from outermost to innermost, which the construct replaces with a canonical loop nest that consists of $2m$ perfectly nested loops. Let $g_1, \dots, g_m, t_1, \dots, t_m$ be the generated loops, from outermost to innermost. The loops g_1, \dots, g_m are the grid loops and the loops t_1, \dots, t_m are the tile loops.

Let Ω be the logical iteration vector space of the transformation-affected loops and let $G = \{T_{\alpha_1, \dots, \alpha_m} \mid T_{\alpha_1, \dots, \alpha_m} \neq \emptyset, (\alpha_1, \dots, \alpha_m) \in \mathbb{N}^m\}$ be the grid that consists of the logical iteration spaces of all tiles with at least one logical iteration.

If the *sizes-selector* of the **sizes** clause is **tile**, then define

$T_{\alpha_1, \dots, \alpha_m} = \{(i_1, \dots, i_m) \in \Omega \mid \forall k \in \{1, \dots, m\} : s_k \alpha_k \leq i_k < s_k \alpha_k + s_k\}$. Tiles that contain $\prod_{k=1}^m s_k$ iterations are complete tile. Otherwise, they are partial tiles.

If the *sizes-selector* of the **sizes** clause is **grid**, then define

$T_{\alpha_1, \dots, \alpha_m} = \{(i_1, \dots, i_m) \in \Omega \mid \forall k \in \{1, \dots, m\} : s_k \alpha_k \leq i_k < s_k \alpha_k + s_k\}$ where for each $d \in \{1, \dots, m\}$, $0 = u_d^0, u_d^1, \dots, u_d^{s_d} = n_d$ are implementation defined sequences where $\lfloor n_d/s_d \rfloor \leq u_d^k - u_d^{k-1} \leq \lceil n_d/s_d \rceil$ for $1 \leq k \leq s_d$.

The grid loops iterate over all tiles $\{T_{\alpha_1, \dots, \alpha_m} \in G\}$ in lexicographic order with respect to their indices $(\alpha_1, \dots, \alpha_m)$ and the tile loops iterate over the iterations in $T_{\alpha_1, \dots, \alpha_m}$ 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.

If a grid loop and a tile loop that are generated from the same tile construct are affected loops of the same loop-nest-associated construct, the tile loops may execute additional empty logical iterations. The number of empty logical iterations is implementation defined.

Restrictions

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

- The transformation-affected loops must be perfectly nested loops.
- No transformation-affected loops may be a non-rectangular loop.
- The implementation defined sequences $u_d^0, u_d^1, \dots, u_d^{s_d}$ must conform to the consistent schedules requirement.

Cross References

- **apply** Clause, see Section 17.1
- Consistent Loop Schedules, see Section 6.4.4
- **sizes** Clause, see Section 17.2

17.10 unroll Construct

Name: <code>unroll</code> Category: <code>executable</code>	Association: <code>loop nest</code> Properties: <code>generally-composable</code> , <code>loop-transforming</code> , <code>order-concurrent-nestable</code> , <code>pure</code> , <code>simdizable</code> , <code>teams-nestable</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
<code>unrolled</code> (<i>default</i>)	1	the grid loop g_1 of the tiling step

Semantics

The `unroll` construct has one `transformation-affected 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.

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 17.1](#)
- `full` Clause, see [Section 17.10.1](#)
- `partial` Clause, see [Section 17.10.2](#)

17.10.1 full Clause

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

Arguments

Name	Type	Properties
<i>fully_unroll</i>	expression of OpenMP logical type	<code>constant</code> , <code>optional</code>

Modifiers

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

Directives

unroll

Semantics

If *fully_unroll* evaluates to *true*, the **full clause** specifies that the **transformation-affected 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* **affected 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 **transformation-affected loop** must be **constant**.

Cross References

- **unroll** Construct, see [Section 17.10](#)

17.10.2 partial Clause

Name: partial	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: <i>directive-name</i> (a <i>directive name</i>)	unique

Directives

unroll

Semantics

The **partial clause** specifies that the **transformation-affected loop** is first tiled with a **tile** size of *unroll-factor*. Then, the generated **tile loop** is fully unrolled. If the **partial clause** is used without an *unroll-factor* argument then *unroll-factor* is an **implementation defined positive** integer.

Cross References

- **unroll** Construct, see [Section 17.10](#)

18 Parallelism Generation and Control

This chapter defines `constructs` for generating and controlling parallelism.

18.1 `parallel` Construct

Name: <code>parallel</code> Category: <code>executable</code>	Association: <code>block</code> Properties: <code>cancellable</code> , <code>context-matching</code> , <code>order-concurrent-nestable</code> , <code>parallelism-generating</code> , <code>team-generating</code> , <code>teams-nestable</code> , <code>thread-limiting</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`. 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` is `region-invariant` and, so, does not change for the duration of that `parallel region`.

Within a `parallel region`, `thread numbers` uniquely identify each `thread`. `Thread numbers` are consecutive `non-negative` integers ranging from zero for the `primary thread` up to one less than the number of `threads` in the `team`. A `thread` may obtain its own `thread number` by a call to the `omp_get_thread_num` library routine.

A set of `implicit tasks`, equal in number to the number of `threads` in the `team`, is generated by the `encountering thread`. The `structured block` of the `parallel` construct determines the code that will be executed in each `implicit task`. Each `task` is assigned to a different `thread` in the `team` and becomes a `tied`. The `task region` of the `task` that the `encountering thread` is executing is suspended

and each **thread** in the **team** executes its **implicit task**. Each **thread** can execute a path of statements that is different from that of the other **threads**.

The implementation may cause any **thread** to suspend execution of its **implicit task** at a **task scheduling point**, and to switch to execution of any **explicit task** generated by any of the **threads** in the **team**, before eventually resuming execution of the **implicit task**.

An **implicit barrier** occurs at the end of a **parallel region**. After the end of a **parallel region**, only the **primary thread** of the **team** resumes execution of the enclosing **task region**.

If a **thread** in a **team** that is executing a **parallel region** encounters another **parallel directive**, it forms a new **team** and becomes the **primary thread** of that new **team**.

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

Unless a **requires directive** is specified on which the **device_safesync** clause appears, if the **parallel** construct is encountered on a **non-host device** and the **safesync** clause is not present then the behavior is as if the **safesync** clause appears on the **directive** with a *width* value that is **implementation defined**.

Execution Model Events

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

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

If a new **native thread** is created for the **team** that executes the **parallel 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 *implicit-task-begin event*.

Events associated with **implicit barriers** occur at the end of a **parallel region**. Section 23.3.2 describes **events** associated with **implicit barriers**.

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

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

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

Tool Callbacks

A **thread** dispatches a registered **parallel_begin** callback for each occurrence of a *parallel-begin event* in that **thread**. The **callback** occurs in the **task** that encounters the **parallel** construct. In the dispatched **callback**, (*flags & omp_parallel_team*) evaluates to *true*.

A **thread** dispatches a registered **implicit_task** callback with **ompt_scope_begin** as its *endpoint* argument for each occurrence of an *implicit-task-begin event* in that **thread**. Similarly, a **thread** dispatches a registered **implicit_task** callback with **ompt_scope_end** as its *endpoint* argument for each occurrence of an *implicit-task-end event* in that **thread**. The **callbacks** occur in the context of the **implicit task**. In the dispatched **callback**, (*flags & ompt_task_implicit*) evaluates to *true*.

A **thread** dispatches a registered **parallel_end** callback for each occurrence of a *parallel-end event* in that **thread**. The **callback** occurs in the **task** that encounters the **parallel** construct.

A **thread** dispatches a registered **thread_begin** callback for any *native-thread-begin event* in that **thread**. The **callback** occurs in the context of the **thread**.

A **thread** dispatches a registered **thread_end** callback for any *native-thread-end event* in that **thread**. The **callback** occurs in the context of the **thread**.

Cross References

- **allocate** Clause, see [Section 13.6](#)
- **copyin** Clause, see [Section 10.1](#)
- **default** Clause, see [Section 7.3.1](#)
- **firstprivate** Clause, see [Section 7.3.4](#)
- **if** Clause, see [Section 5.6](#)
- **implicit_task** Callback, see [Section 40.5.3](#)
- **message** Clause, see [Section 16.3](#)
- **num_threads** Clause, see [Section 18.1.2](#)
- **omp_get_thread_num** Routine, see [Section 27.3](#)
- Determining the Number of Threads for a **parallel** Region, see [Section 18.1.1](#)
- **parallel_begin** Callback, see [Section 40.3.1](#)
- **parallel_end** Callback, see [Section 40.3.2](#)
- OMPT **parallel_flag** Type, see [Section 39.22](#)
- **private** Clause, see [Section 7.3.3](#)
- **proc_bind** Clause, see [Section 18.1.4](#)
- **reduction** Clause, see [Section 8.10](#)

Algorithm 18.1 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** is specified **then let** *IfClauseValue* be the value of *if-expression*;
else let *IfClauseValue* = **true**;
if a **num_threads clause** is specified **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(*thread-limit-var* - *ThreadsBusy*,
structured-thread-limit-var - *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**

- 1 • **safesync** Clause, see [Section 18.1.5](#)
- 2 • OMPT **scope_endpoint** Type, see [Section 39.27](#)
- 3 • **severity** Clause, see [Section 16.4](#)
- 4 • **shared** Clause, see [Section 7.3.2](#)
- 5 • OMPT **task_flag** Type, see [Section 39.37](#)
- 6 • **thread_begin** Callback, see [Section 40.1.3](#)
- 7 • **thread_end** Callback, see [Section 40.1.4](#)

18.1.1 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**. When a **thread** encounters a **parallel** construct, the number of **threads** is determined according to Algorithm 18.1.

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.

Cross References

- *dyn-var* ICV, see Table 3.1
- *max-active-levels-var* ICV, see Table 3.1
- *nthreads-var* ICV, see Table 3.1
- *thread-limit-var* ICV, see Table 3.1
- **if** Clause, see Section 5.6
- **num_threads** Clause, see Section 18.1.2
- **parallel** Construct, see Section 18.1

18.1.2 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: <i>directive-name</i> (a directive name)	unique

Directives

`parallel`

Semantics

The `num_threads` clause specifies the desired number of `threads` to execute a `parallel` region. Algorithm 18.1 determines the number of `threads` that execute the `parallel` region. If *prescriptiveness* is specified as `strict` and an implementation determines that Algorithm 18.1 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 18.1 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 16.2](#)
- `error` Directive, see [Section 16.1](#)
- `message` Clause, see [Section 16.3](#)
- `parallel` Construct, see [Section 18.1](#)

18.1.3 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.

TABLE 18.1: Affinity-related Symbols used in this Section

Symbol	Symbol Description
L	the value of the <i>thread-limit-var</i> ICV
NG	the total number of place-assignment groups
g_i	the i^{th} place-assignment group
P	the number of places in the input place partition
T	the number of threads in the team
AT	$\lceil T/NG \rceil$ ("above-thread" count)
BT	$\lfloor T/NG \rfloor$ ("below-thread" count)
ET	$T \bmod NG$ ("excess-thread" count)

The *place-assignment-var* ICV is a list of L place numbers, where L is the value of the *thread-limit-var* ICV, that defines the place assignment of threads that participate in the execution of tasks bound to a given team. Any such thread corresponds to a position in the list, meaning it will be assigned to the place given by the place number at that position. If a thread is an assigned thread of the team with thread number i , it corresponds to position i in the *place-assignment-var* list. If a thread is a free-agent thread, it corresponds to the first position for which another thread has not yet been assigned to the associated place. If another thread is already assigned to the place associated with that position, the place to which the free-agent thread is assigned is implementation defined.

Each thread affinity policy determines how threads are assigned to places. A policy assigns each place in the input place partition to one of NG place-assignment groups, g_0, \dots, g_{NG-1} ; additionally, it assigns each position from the *place-assignment-var* ICV to one of these groups. In a given group, the place number of each place is then assigned to a *place-assignment-var* position, in round robin fashion, starting with the first place. Threads are thus assigned to places according to the resulting *place-assignment-var* of the policy.

Under the **primary** thread affinity policy, $NG = 1$ and place-assignment group g_0 is assigned the place to which the encountering thread is assigned, and all positions of *place-assignment-var* are assigned to the same group. Thus, the corresponding threads of all positions of the *place-assignment-var* ICV are assigned to the same place as the primary thread.

For the **close** and **spread** thread affinity policies, let P be the number of places in the input place partition and let T be the number of assigned threads in the team. The following paragraphs describe how places in the input place partition are subdivided into place-assignment groups for these policies. A general description of how positions in *place-assignment-var* are assigned to these places, and thus how place assignment for threads under the policies is determined, then

follows these descriptions.

The **close thread affinity** policy distributes assignment of **places** evenly across a **team** of **threads**, while ensuring **threads** with consecutive numbers are assigned to the same **place** or adjacent **places**. Each **place** in the **input place partition** is assigned to one **place-assignment group** (so, $NG = P$). **Place-assignment group** g_0 is assigned the **place** to which the **encountering thread** is assigned. The **place** assigned to group g_i is then the next **place** in the **place partition** of the one assigned to group g_{i-1} , with wrap around with respect to the **input place partition**.

The **spread thread affinity** policy creates a sparse distribution for a **team** of T **threads** among the P **places** of the **input place partition**. A sparse distribution is achieved by first subdividing the **input place partition** into T subpartitions if $T \leq P$ (in which case $NG = T$), or P subpartitions if $T > P$ (in which case $NG = P$). The subpartitions are determined as follows:

- $T \leq P$: The **input place partition** is split into T subpartitions, where each subpartition contains $\lfloor P/T \rfloor$ or $\lceil P/T \rceil$ consecutive **places**; if $P \bmod T$ is not zero, which subpartitions contain $\lceil P/T \rceil$ **places** is **implementation defined**;
- $T > P$: The **input place partition** is split into P subpartitions, each with a single **place**.

In either case, the **places** from each subpartition are assigned to a **place-assignment group** that corresponds to the subpartition. The subpartition that corresponds to group g_0 is the one that includes the **place** on which the **encountering thread** is executing. The subpartition that corresponds to group g_i is the one that includes the next **place** to those in the subpartition corresponding to group g_{i-1} , with wrap around with respect to the **input place partition**. For a given **implicit task** and corresponding **place-assignment-var** position to its **assigned thread**, the **place-partition-var ICV** of the **implicit task** is set to the subpartition that corresponds to the group that includes the position. Thus, the subpartitioning is not only a mechanism for achieving a sparse distribution, it also defines a subset of **places** for a **thread** to use when creating a nested **parallel region**.

Let AT equal $\lceil T/NG \rceil$, BT equal $\lfloor T/NG \rfloor$, and ET equal $T \bmod NG$. The **close** and the **spread thread affinity** policies assign the positions of the **place-assignment-var ICV** to **place-assignment groups** as follows.

- For positions from 0 up to $T - 1$: The positions are partitioned into NG sets of consecutive positions, ET of which have AT positions and $NG - ET$ of which have only BT positions (when ET is not zero, which sets have which count is **implementation defined** unless the **thread affinity** policy is **close** and $T < P$, in which case the first T groups are assigned the sets with AT positions). The sets are assigned to each group, with the first set, starting at position 0, assigned to group g_0 , and with each successive set i , starting at the position immediately after the last position in the set assigned to group g_{i-1} , assigned to the next group g_i ;
- If $ET \neq 0$, for the positions from T up to $(AT * NG) - 1$: Each of these positions is assigned to a group g_i that received only BT positions in the above step, such that each such g_i is then assigned AT positions (which positions are assigned to which group is **implementation defined**);

- For the remaining positions from $AT * NG$ up to L : Each position is assigned to a group in round robin fashion, starting with the first group g_0 .

The determination of whether the [thread affinity](#) request can be fulfilled is [implementation defined](#). If it cannot be fulfilled, then the affinity of [threads](#) in the [team](#) is [implementation defined](#).

Note – Wrap around is needed if the end of a [place partition](#) is reached before all [thread](#) assignments are done. For example, wrap around may be needed in the case of **close** and $T \leq P$, if the [primary thread](#) is assigned to a [place](#) other than the first [place](#) in the [place partition](#). In this case, [thread](#) 1 is assigned to the [place](#) after the [place](#) of the [primary thread](#), thread 2 is assigned to the [place](#) after that, and so on. The end of the [place partition](#) may be reached before all [threads](#) are assigned. In this case, assignment of [threads](#) is resumed with the first [place](#) in the [place partition](#).

Cross References

- *bind-var* ICV, see [Table 3.1](#)
- *place-assignment-var* ICV, see [Table 3.1](#)
- *place-partition-var* ICV, see [Table 3.1](#)
- *thread-limit-var* ICV, see [Table 3.1](#)
- **parallel** Construct, see [Section 18.1](#)
- **proc_bind** Clause, see [Section 18.1.4](#)

18.1.4 proc_bind Clause

Name: proc_bind	Properties: unique
-------------------------------	---

Arguments

Name	Type	Properties
<i>affinity-policy</i>	Keyword: close , primary , spread	default

Modifiers

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

Directives

[parallel](#)

Semantics

The `proc_bind` clause specifies the mapping of threads to places within the input place partition. The effect of the possible values for *affinity-policy* are described in Section 18.1.3

Cross References

- Controlling OpenMP Thread Affinity, see Section 18.1.3
- `parallel` Construct, see Section 18.1

18.1.5 safesync Clause

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

Arguments

Name	Type	Properties
<i>width</i>	expression of integer type	<code>positive, optional</code>

Modifiers

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

Directives

`parallel`

Semantics

The `safesync` clause determines whether two synchronizing threads in a team can make progress (see Section 1.2). The 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 may execute in the same progress unit. If the *width* argument is not specified, the behavior is as if the *width* argument is one.

Restrictions

Restrictions to the `safesync` clause are as follows:

- The *width* argument must be a `safesync-compatible` expression.

Cross References

- `parallel` Construct, see Section 18.1

18.2 teams Construct

Name: <code>teams</code> Category: <code>executable</code>	Association: <code>block</code> Properties: <code>parallelism-generating</code> , <code>team-generating</code> , <code>thread-limiting</code> , <code>context-matching</code> , <code>groupprivate-instantiating</code>
---	--

Clauses

`allocate`, `default`, `dyn_groupprivate`, `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` are `region-invariant`, thus do not change for the duration of the `teams` region. Within a `teams` region, `initial team` numbers uniquely identify each `initial team`. `Initial teams` numbers are consecutive `non-negative` integers 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.

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 **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**. In the dispatched **callback**, *(flags & ompt_parallel_league)* evaluates to *true*.

A **thread** dispatches a registered **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 **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**. In the dispatched **callback**, *(flags & ompt_task_initial)* and *(flags & ompt_task_implicit)* evaluate to *true*.

A **thread** dispatches a registered **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**.

A **thread** dispatches a registered **thread_begin callback** for each *native-thread-begin event* in that **thread**. The **callback** occurs in the context of the **thread**.

A **thread** dispatches a registered **thread_end callback** for each *native-thread-end event* in that **thread**. The **callback** occurs in the context of the **thread**.

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 a **strictly nested region** of 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**.

- For a **teams** construct that is an immediately nested construct of a **target** construct, the bounds expressions of any **array sections** and the index expressions of any array elements used in any **clause** on the **construct**, as well as all expressions of any **target-consistent clauses** on the **construct**, must be **target-consistent expressions**.
- Only **regions** that are generated by **teams-nestable constructs** or **teams-nestable routines** may be **strictly nested regions** of **teams regions**.

Cross References

- **allocate** Clause, see [Section 13.6](#)
- **default** Clause, see [Section 7.3.1](#)
- **distribute** Construct, see [Section 19.7](#)
- **dyn_groupprivate** Clause, see [Section 13.9](#)
- **firstprivate** Clause, see [Section 7.3.4](#)
- *default-device-var* ICV, see [Table 3.1](#)
- *place-partition-var* ICV, see [Table 3.1](#)
- **if** Clause, see [Section 5.6](#)
- **implicit_task** Callback, see [Section 40.5.3](#)
- **num_teams** Clause, see [Section 18.2.1](#)
- **omp_get_num_teams** Routine, see [Section 28.1](#)
- **omp_get_team_num** Routine, see [Section 28.3](#)
- **parallel** Construct, see [Section 18.1](#)
- **parallel_begin** Callback, see [Section 40.3.1](#)
- **parallel_end** Callback, see [Section 40.3.2](#)
- OMPT **parallel_flag** Type, see [Section 39.22](#)
- **private** Clause, see [Section 7.3.3](#)
- **reduction** Clause, see [Section 8.10](#)
- OMPT **scope_endpoint** Type, see [Section 39.27](#)
- **shared** Clause, see [Section 7.3.2](#)
- **target** Construct, see [Section 21.8](#)
- OMPT **task_flag** Type, see [Section 39.37](#)
- **thread_begin** Callback, see [Section 40.1.3](#)

- **thread_end** Callback, see [Section 40.1.4](#)
- **thread_limit** Clause, see [Section 21.3](#)

18.2.1 num_teams Clause

Name: <code>num_teams</code>	Properties: target-consistent , 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: <i>directive-name</i> (a 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 *ntteams-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 *ntteams-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

- *ntteams-var* ICV, see [Table 3.1](#)
- **teams** Construct, see [Section 18.2](#)

18.3 order Clause

Name: order	Properties: schedule-specification , unique
--------------------	---

Arguments

Name	Type	Properties
<i>ordering</i>	Keyword: concurrent	default

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: <i>directive-name</i> (a 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 6.4.4](#)). 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 [routines](#) for which a call may be nested inside a [region](#) that corresponds to a [construct](#) on which the [order](#) clause is specified with **concurrent** as the *ordering* argument are [order-concurrent-nestable routines](#).
- Only [regions](#) that correspond to [order-concurrent-nestable constructs](#) or [order-concurrent-nestable routines](#) may be [strictly nested regions](#) of [regions](#) that correspond to [constructs](#) on which the [order](#) clause is specified with **concurrent** as the *ordering* argument.
- 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.

1 **Cross References**

- 2 • **distribute** Construct, see [Section 19.7](#)
- 3 • **do** Construct, see [Section 19.6.2](#)
- 4 • **for** Construct, see [Section 19.6.1](#)
- 5 • **loop** Construct, see [Section 19.8](#)
- 6 • **simd** Construct, see [Section 18.4](#)

7 **18.4 simd Construct**

8

Name: <code>simd</code> Category: executable	Association: loop nest Properties: context-matching , order-concurrent-nestable , parallelism-generating , pure , simdizable
---	---

9 **Separating directives**

10 [scan](#)

11 **Clauses**

12 [aligned](#), [collapse](#), [if](#), [induction](#), [lastprivate](#), [linear](#), [nontemporal](#), [order](#),
13 [private](#), [reduction](#), [safelen](#), [simdlen](#)

14 **Binding**

15 A [simd region](#) binds to the [current task region](#). The [binding thread set](#) of the [simd region](#) is the
16 [current team](#).

17 **Semantics**

18 The [simd construct](#) enables the execution of multiple [collapsed iterations](#) concurrently by using
19 [SIMD instructions](#). The number of [collapsed iterations](#) that are executed concurrently at any given
20 time is [implementation defined](#). Each concurrent iteration will be executed by a different [SIMD](#)
21 [lane](#). Each set of concurrent iterations is a [SIMD chunk](#). Lexical forward dependences in the
22 iterations of the original loop must be preserved within each [SIMD chunk](#), unless an [order clause](#)
23 that specifies **concurrent** is present.

24 When an [if clause](#) is present with an *if-expression* that evaluates to *false*, the preferred number of
25 iterations to be executed concurrently is one, regardless of whether a [simdlen clause](#) is specified.

26 **Restrictions**

27 Restrictions to the [simd construct](#) are as follows:

- 28 • If both [simdlen](#) and [safelen clauses](#) are specified, the value of the [simdlen length](#)
29 must be less than or equal to the value of the [safelen length](#).
- 30 • 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++

C++

- No exceptions can be raised in the **simd** region.
- The only random access **iterator** types that are allowed for the **collapse-affected** loops are pointer types.

C++

Cross References

- **aligned** Clause, see [Section 14.2](#)
- **collapse** Clause, see [Section 6.4.5](#)
- **if** Clause, see [Section 5.6](#)
- **induction** Clause, see [Section 8.13](#)
- **lastprivate** Clause, see [Section 7.3.5](#)
- **linear** Clause, see [Section 8.14](#)
- **nontemporal** Clause, see [Section 18.4.1](#)
- **order** Clause, see [Section 18.3](#)
- **private** Clause, see [Section 7.3.3](#)
- **reduction** Clause, see [Section 8.10](#)
- **safelen** Clause, see [Section 18.4.2](#)
- **scan** Directive, see [Section 8.17](#)
- **simdlen** Clause, see [Section 18.4.3](#)

18.4.1 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: <i>directive-name</i> (a <i>directive name</i>)	<i>unique</i>

Directives

simd

Semantics

The *nontemporal* clause specifies that accesses to the *storage locations* to which the *list items* refer have low temporal locality across the *logical 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* Construct, see [Section 18.4](#)

18.4.2 safelen Clause

Name: <i>safelen</i>	Properties: <i>unique</i>
-----------------------------	----------------------------------

Arguments

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

Modifiers

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

Directives

simd

Semantics

The *safelen* clause specifies that no two concurrent *logical iterations* within a *SIMD chunk* can have a distance in the *collapsed iteration space* that is greater than or equal to the *length* argument.

Cross References

- *simd* Construct, see [Section 18.4](#)

18.4.3 simdlen Clause

Name: simdlen	Properties: unique
----------------------	---------------------------

Arguments

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

Modifiers

Name	Modifies	Type	Properties
<i>scaled-modifier</i>	<i>length</i>	Complex, name: scaled Arguments: type type name (<i>default</i>) divisor expression of integer type (constant , positive , optional)	<i>default</i>
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: <i>directive-name</i> (a 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** directive, if a **SIMD** version of the associated **procedure** is created, *length* corresponds to the number of concurrent arguments of the **procedure**.

If the *scaled-modifier* is present, the preferred number of **collapsed iterations** or the number of concurrent arguments is given by the *length* argument multiplied by a scaling factor that is computed as follows. Let $VL(type)$ be an **implementation defined** value, possibly determined at runtime, equal to some number of **SIMD lanes** that can be used to process *type* data elements in parallel, and let *d* be the value of *divisor*, if specified, and otherwise one. The scaling factor is then equal to $\max(1, \lfloor VL(type)/d \rfloor)$.

Restrictions

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

	C / C++
• In a <i>scaled-modifier</i> , <i>type</i> must be an arithmetic type.	
	C / C++
	Fortran
• In a <i>scaled-modifier</i> , <i>type</i> must be a numeric intrinsic type.	
	Fortran

Cross References

- `declare_simd` Directive, see [Section 15.8](#)
- `simd` Construct, see [Section 18.4](#)

18.5 masked Construct

Name: <code>masked</code>	Association: <code>block</code>
Category: <code>executable</code>	Properties: <code>thread-limiting</code> , <code>thread-selecting</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`.

Tool Callbacks

A `thread` dispatches a registered `masked callback` with `ompt_scope_begin` as its `endpoint` argument for each occurrence of a `masked-begin event` in that `thread`. Similarly, a `thread` dispatches a registered `masked callback` with `ompt_scope_end` as its `endpoint` argument for each occurrence of a `masked-end event` in that `thread`. These `callbacks` occur in the context of the `task` executed by the `encountering thread`.

Cross References

- `filter` Clause, see [Section 18.5.1](#)
- `masked` Callback, see [Section 40.3.3](#)

- OMPT `scope_endpoint` Type, see [Section 39.27](#)

18.5.1 filter Clause

Name: filter	Properties: unique
---------------------	------------------------------------

Arguments

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

Modifiers

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

Directives

[masked](#)

Semantics

If *thread_num* specifies the [thread number](#) of the [encountering thread](#) in the [current team](#) then the [filter clause](#) selects the [encountering thread](#). If the [filter clause](#) is not specified, the effect is as if the [clause](#) is specified with *thread_num* equal to zero, so that the [filter clause](#) selects the [primary thread](#). The use of a [variable](#) in a *thread_num* argument expression causes an implicit reference to the [variable](#) in all enclosing [constructs](#).

Cross References

- [masked](#) Construct, see [Section 18.5](#)

19 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 thread 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**.

For **loop-nest-associated constructs**, the **loop schedule** is determined by a **schedule specification** for the **construct**, which is defined by **schedule-specification clauses** and (where applicable) the *run-sched-var* **ICV**. **OpenMP programs** can only depend on which **thread** executes a particular **collapsed iteration** if the **construct** specifies a **reproducible schedule**. Schedule reproducibility also determines whether **constructs** with the same **schedule specification** will have **consistent schedules** (see [Section 6.4.4](#)).

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**.

- 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**.

19.1 single Construct

Name: single	Association: block
Category: executable	Properties: work-distribution, team-executed, partitioned, worksharing, thread-limiting, thread-selecting

Clauses

allocate, copyprivate, firstprivate, nowait, private

Clause set

Properties: exclusive	Members: copyprivate, nowait
-------------------------------------	--

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**.

1 **Tool Callbacks**

2 A **thread** dispatches a registered **work** callback with **ompt_scope_begin** as its *endpoint*
3 argument for each occurrence of a *single-begin event* in that **thread**. Similarly, a **thread** dispatches a
4 registered **work** callback with **ompt_scope_end** as its *endpoint* argument for each occurrence
5 of a *single-end event* in that **thread**. For each of these **callbacks**, the *work_type* argument is
6 **ompt_work_single_executor** if the **thread** executes the **structured block** associated with the
7 **single region**; otherwise, the *work_type* argument is **ompt_work_single_other**.

8 **Cross References**

- 9 • **allocate** Clause, see [Section 13.6](#)
- 10 • **copyprivate** Clause, see [Section 10.2](#)
- 11 • **firstprivate** Clause, see [Section 7.3.4](#)
- 12 • **nowait** Clause, see [Section 23.6](#)
- 13 • **private** Clause, see [Section 7.3.3](#)
- 14 • OMPT **scope_endpoint** Type, see [Section 39.27](#)
- 15 • **work** Callback, see [Section 40.4.1](#)
- 16 • OMPT **work** Type, see [Section 39.41](#)

17 **19.2 scope Construct**

18 Name: scope Category: executable	Association: block Properties: work-distribution, team- executed, worksharing, thread-limiting
--	---

19 **Clauses**

20 **[allocate](#), [firstprivate](#), [nowait](#), [private](#), [reduction](#)**

21 **Binding**

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

26 **Semantics**

27 The **scope construct** specifies that all **threads** in a **team** execute the associated **structured block** and
28 any additionally specified **OpenMP operations**. An **implicit barrier** occurs at the end of a **scope**
29 **region** if the **nowait** clause does not specify otherwise.

1 **Execution Model Events**

2 The *scope-begin event* occurs after an **implicit task** encounters a **scope construct** but before the
3 **task** starts to execute the **structured block** of the **scope region**. The *scope-end event* occurs after
4 an **implicit task** finishes execution of a **scope region** but before it resumes execution of the
5 enclosing **region**.

6 **Tool Callbacks**

7 A **thread** dispatches a registered **work callback** with **ompt_scope_begin** as its *endpoint*
8 argument and **ompt_work_scope** as its *work_type* argument for each occurrence of a
9 *scope-begin event* in that **thread**. Similarly, a **thread** dispatches a registered **work callback** with
10 **ompt_scope_end** as its *endpoint* argument and **ompt_work_scope** as its *work_type*
11 argument for each occurrence of a *scope-end event* in that **thread**. The **callbacks** occur in the
12 context of the **implicit task**.

13 **Cross References**

- 14 • **allocate** Clause, see [Section 13.6](#)
15 • **firstprivate** Clause, see [Section 7.3.4](#)
16 • **nowait** Clause, see [Section 23.6](#)
17 • **private** Clause, see [Section 7.3.3](#)
18 • **reduction** Clause, see [Section 8.10](#)
19 • OMPT **scope_endpoint** Type, see [Section 39.27](#)
20 • **work** Callback, see [Section 40.4.1](#)
21 • OMPT **work** Type, see [Section 39.41](#)

22 **19.3 sections Construct**

23

Name: sections Category: executable	Association: block Properties: work-distribution, team- executed, partitioned, worksharing, thread-limiting, cancellable
--	---

24 **Separating directives**
25 **[section](#)**

26 **Clauses**
27 **[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 **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 **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**.

Cross References

- **allocate** Clause, see [Section 13.6](#)
- **firstprivate** Clause, see [Section 7.3.4](#)
- **lastprivate** Clause, see [Section 7.3.5](#)
- **nowait** Clause, see [Section 23.6](#)
- **private** Clause, see [Section 7.3.3](#)
- **reduction** Clause, see [Section 8.10](#)
- OMPT **scope_endpoint** Type, see [Section 39.27](#)
- **section** Directive, see [Section 19.3.1](#)

- **work** Callback, see [Section 40.4.1](#)
- OMPT **work** Type, see [Section 39.41](#)

19.3.1 section Directive

Name: section Category: subsidiary	Association: separating Properties: default
--	--

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 [dispatch callback](#) for each occurrence of a *section-begin* event in that [thread](#). The [callback](#) occurs in the context of the [implicit task](#).

Cross References

- **dispatch** Callback, see [Section 40.4.2](#)
- **sections** Construct, see [Section 19.3](#)



Fortran

19.4 workshare Construct

Name: workshare Category: executable	Association: block Properties: work-distribution , team-executed , partitioned , worksharing
--	---

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.

Semantics

The **workshare** construct divides the execution of the associated **structured block** into separate **units of work** and causes the **threads** of the **team** to share the work such that each **unit of work** is executed only once by one **thread**, in the context of its **implicit task**. An **implicit barrier** occurs at the end of a **workshare region** if a **nowait** clause does not specify otherwise.

An implementation of the **workshare** construct must insert any synchronization 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 following statements, and the evaluation of the right hand side of an assignment must appear to complete prior to the effects of assigning to the left hand side.

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

- For array expressions within each statement, including transformational array intrinsic functions that compute scalar values from arrays:
 - Evaluation of each element of the array expression, including any references to elemental functions, is a **unit of work**.
 - Evaluation of transformational array intrinsic functions may be 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**.
- For **WHERE** statements or constructs, evaluation of the mask expression and the masked assignments are each a **unit of work**.
- For **FORALL** statements or constructs, evaluation of the mask expression, expressions occurring in the specification of the iteration space, and the masked assignments are each a **unit of work**.
- For **atomic** constructs, **critical** constructs, and **parallel** constructs, the **construct** is a **unit of work**. A new **team** executes the statements contained in a **parallel** construct.
- If none of the rules above apply to a portion of a statement in the **structured block**, then that portion is a **unit of work**.

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

The **units of work** are assigned to the **threads** that execute a **workshare 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**.

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

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

Execution Model Events

The *workshare-begin event* occurs after an **implicit task** encounters a **workshare construct** but before the **task** starts to execute the **structured block** of the **workshare region**. The *workshare-end event* occurs after an **implicit task** finishes execution of a **workshare region** but before it resumes execution of the **enclosing context**.

Tool Callbacks

A **thread** dispatches a registered **work callback** with **ompt_scope_begin** as its *endpoint* argument and **ompt_work_workshare** as its *work_type* argument for each occurrence of a *workshare-begin event* in that **thread**. Similarly, a **thread** dispatches a registered **work callback** with **ompt_scope_end** as its *endpoint* argument and **ompt_work_workshare** as its *work_type* argument for each occurrence of a *workshare-end event* in that **thread**. The **callbacks** occur in the context of the **implicit task**.

Restrictions

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

- The only OpenMP **constructs** that may be **closely nested constructs** of a **workshare construct** are the **atomic**, **critical**, and **parallel** constructs.
- **Base language** statements that are encountered inside a **workshare construct** but that are not enclosed within a **parallel** or **atomic construct** that is nested inside the **workshare construct** must consist of only the following:
 - array assignments;
 - scalar assignments;
 - **FORALL** statements;
 - **FORALL** constructs;
 - **WHERE** statements;
 - **WHERE** constructs; and
 - **BLOCK** constructs that are **strictly structured blocks** associated with **directives**.
- All array assignments, scalar assignments, and masked array assignments that are encountered inside a **workshare construct** but are not nested inside a **parallel construct** that is nested inside the **workshare construct** must be **intrinsic assignments**.

- The **construct** must not contain any user-defined function calls unless either the function is pure and elemental or the function call is contained inside a **parallel construct** that is nested inside the **workshare construct**.

Cross References

- **atomic** Construct, see [Section 23.8.5](#)
- **critical** Construct, see [Section 23.2](#)
- **nowait** Clause, see [Section 23.6](#)
- **parallel** Construct, see [Section 18.1](#)
- OMPT **scope_endpoint** Type, see [Section 39.27](#)
- **work** Callback, see [Section 40.4.1](#)
- OMPT **work** Type, see [Section 39.41](#)



19.5 workdistribute Construct

Name: workdistribute Category: executable	Association: block Properties: work-distribution, partitioned
--	--

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 **workdistribute 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 **workdistribute 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 **workdistribute 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 **workdistribute 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 *workdistribute-begin event* occurs after an **initial task** encounters a **workdistribute construct** but before the **task** starts to execute the **structured block** of the **workdistribute region**. The *workdistribute-end event* occurs after an **initial task** finishes execution of a **workdistribute region** but before it resumes execution of the **enclosing context**.

Tool Callbacks

A **thread** dispatches a registered **work callback** with **ompt_scope_begin** as its *endpoint* argument and **ompt_work_workdistribute** as its *work_type* argument for each occurrence of a *workdistribute-begin event* in that **thread**. Similarly, a **thread** dispatches a registered **work callback** with **ompt_scope_end** as its *endpoint* argument and **ompt_work_workdistribute** as its *work_type* argument for each occurrence of a *workdistribute-end event* in that **thread**. The **callbacks** occur in the context of the **implicit task**.

Restrictions

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

- The **workdistribute** construct must be a **closely nested construct** inside a **teams construct**.
- No **explicit region** may be nested inside a **workdistribute region**.
- Base language statements that are encountered inside a **workdistribute** 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 [workdistribute 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 [workdistribute region](#), the behavior is unspecified.

Cross References

- OMPT [scope_endpoint](#) Type, see [Section 39.27](#)
- [target](#) Construct, see [Section 21.8](#)
- [teams](#) Construct, see [Section 18.2](#)
- [work](#) Callback, see [Section 40.4.1](#)
- OMPT [work](#) Type, see [Section 39.41](#)

Fortran

19.6 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 [collapsed 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 the [assigned threads](#) of 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 [collapse-affected loop](#) has the value that it would have if the [collapse-affected loops](#) were executed sequentially.

The [loop schedule](#) 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`.

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 *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 *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 *work_type* argument indicates the *schedule type* as shown in Table 19.1.

A thread dispatches a registered *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.

TABLE 19.1: *work* OMPT types for Worksharing-Loop

Value of <i>work_type</i>	If determined schedule is
<code>ompt_work_loop</code>	unknown at runtime
<code>ompt_work_loop_static</code>	<code>static</code>
<code>ompt_work_loop_dynamic</code>	<code>dynamic</code>
<code>ompt_work_loop_guided</code>	<code>guided</code>
<code>ompt_work_loop_other</code>	implementation defined

Restrictions

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

- The collapsed 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

- `dispatch` Callback, see Section 40.4.2

- *run-sched-var* ICV, see [Table 3.1](#)
- **nowait** Clause, see [Section 23.6](#)
- **order** Clause, see [Section 18.3](#)
- **schedule** Clause, see [Section 19.6.3](#)
- OMPT **scope_endpoint** Type, see [Section 39.27](#)
- **work** Callback, see [Section 40.4.1](#)
- OMPT **work** Type, see [Section 39.41](#)

C / C++

19.6.1 for Construct

Name: <code>for</code> Category: executable	Association: loop nest Properties: work-distribution , team-executed , partitioned , SIMD-partitionable , worksharing , worksharing-loop , cancellable , context-matching
--	---

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 13.6](#)
- **collapse** Clause, see [Section 6.4.5](#)
- **firstprivate** Clause, see [Section 7.3.4](#)
- Worksharing-Loop Constructs, see [Section 19.6](#)
- **induction** Clause, see [Section 8.13](#)
- **lastprivate** Clause, see [Section 7.3.5](#)
- **linear** Clause, see [Section 8.14](#)
- **nowait** Clause, see [Section 23.6](#)

- **order** Clause, see [Section 18.3](#)
- **ordered** Clause, see [Section 6.4.6](#)
- **private** Clause, see [Section 7.3.3](#)
- **reduction** Clause, see [Section 8.10](#)
- **scan** Directive, see [Section 8.17](#)
- **schedule** Clause, see [Section 19.6.3](#)



19.6.2 do Construct

Name: <code>do</code> Category: <code>executable</code>	Association: <code>loop nest</code> Properties: <code>work-distribution</code> , <code>team-executed</code> , <code>partitioned</code> , <code>SIMD-partitionable</code> , <code>worksharing</code> , <code>worksharing-loop</code> , <code>cancellable</code> , <code>context-</code> <code>matching</code>
--	---

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 13.6](#)
- **collapse** Clause, see [Section 6.4.5](#)
- **firstprivate** Clause, see [Section 7.3.4](#)
- Worksharing-Loop Constructs, see [Section 19.6](#)
- **induction** Clause, see [Section 8.13](#)
- **lastprivate** Clause, see [Section 7.3.5](#)
- **linear** Clause, see [Section 8.14](#)
- **nowait** Clause, see [Section 23.6](#)

- **order** Clause, see [Section 18.3](#)
- **ordered** Clause, see [Section 6.4.6](#)
- **private** Clause, see [Section 7.3.3](#)
- **reduction** Clause, see [Section 8.10](#)
- **scan** Directive, see [Section 8.17](#)
- **schedule** Clause, see [Section 19.6.3](#)

Fortran

19.6.3 schedule Clause

Name: <code>schedule</code>	Properties: <code>schedule-specification</code> , <code>unique</code>
------------------------------------	--

Arguments

Name	Type	Properties
<i>kind</i>	Keyword: auto , dynamic , guided , runtime , static	<i>default</i>
<i>chunk_size</i>	expression of integer type	<code>ultimate</code> , <code>optional</code> , <code>posi- tive</code> , <code>region-invariant</code>

Modifiers

Name	Modifies	Type	Properties
<i>ordering-modifier</i>	<i>kind</i>	Keyword: monotonic , nonmonotonic	<code>unique</code>
<i>chunk-modifier</i>	<i>kind</i>	Keyword: simd	<code>unique</code>
<i>directive-name- modifier</i>	<i>all arguments</i>	Keyword: <i>directive-name</i> (a <code>directive name</code>)	<code>unique</code>

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*.

If the *kind* argument is **dynamic**, each *thread* executes a *chunk*, then requests another *chunk*, until no *chunks* remain to be assigned. Each *chunk* contains *chunk_size collapsed iterations*, except for the *chunk* that contains the sequentially last iteration, which may have fewer iterations. If *chunk_size* is not specified, it defaults to 1.

If the *kind* argument is **guided**, each *thread* executes a *chunk*, then requests another *chunk*, until no *chunks* remain to be assigned. For a *chunk_size* of 1, the size of each *chunk* is proportional to the number of unassigned *collapsed iterations* divided by the number of *threads* in the *team*, decreasing to 1. For a *chunk_size* with value $k > 1$, the size of each *chunk* is determined in the same way, with the restriction that the *chunks* do not contain fewer than k *collapsed iterations* (except for the *chunk* that contains the sequentially last iteration, which may have fewer than k iterations). If *chunk_size* is not specified, it defaults to 1.

If the *kind* argument is **auto**, the decision regarding scheduling is *implementation defined*. If the *schedule clause* is not specified on a *worksharing-loop construct* then the effect is as if the *schedule* clause was specified with **auto** as its *kind* argument.

If the *kind* argument is **runtime**, the decision regarding scheduling is deferred until runtime, and the behavior is as if the *clause* specifies *kind*, *chunk-size* and *ordering-modifier* as set in the *run-sched-var* ICV. If the *schedule clause* explicitly specifies any *modifiers* then they override any corresponding *modifiers* that are specified in the *run-sched-var* ICV.

If the **simd chunk-modifier** is specified and the *canonical loop nest* is associated with a **SIMD construct**, $new_chunk_size = \lceil chunk_size / simd_width \rceil * simd_width$ is the *chunk_size* for all *chunks* except the first and last *chunks*, where *simd_width* is an *implementation defined* value. The first *chunk* will have at least *new_chunk_size collapsed iterations* except if it is also the last *chunk*. The last *chunk* may have fewer *collapsed iterations* than *new_chunk_size*. If the **simd chunk-modifier** is specified and the *canonical loop nest* is not associated with a **SIMD construct**, the *modifier* is ignored.

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

A *compliant implementation* of the **guided schedule type** with a *chunk_size* value of k would

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

If the **monotonic ordering-modifier** is specified then each thread executes the chunks that it is assigned in increasing collapsed iteration order. When the **nonmonotonic ordering-modifier** is specified then chunks may be assigned to threads in any order and the behavior of an application that depends on any execution order of the chunks is unspecified. If an **ordering-modifier** is not specified, the effect is as if the **monotonic ordering-modifier** is specified if the *kind* argument is **static** or an **ordered clause** is specified on the **construct**; otherwise, the effect is as if the **nonmonotonic ordering-modifier** is specified.

Restrictions

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

- The **schedule clause** cannot be specified if any of the collapse-affected loops is a non-rectangular loop.
- The value of the *chunk_size* expression must be the same for all threads in the team.
- If **runtime** or **auto** is specified for *kind*, *chunk_size* must not be specified.
- The **nonmonotonic ordering-modifier** cannot be specified if an **ordered clause** is specified on the same **construct**.

Cross References

- **do Construct**, see [Section 19.6.2](#)
- **for Construct**, see [Section 19.6.1](#)
- *run-sched-var* ICV, see [Table 3.1](#)
- **ordered Clause**, see [Section 6.4.6](#)

19.7 distribute Construct

Name: distribute	Association: loop nest
Category: executable	Properties: SIMD-partitionable, teams-nestable, work-distribution, partitioned

Clauses

allocate, **collapse**, **dist_schedule**, **firstprivate**, **induction**, **lastprivate**, **order**, **private**

Binding

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

Semantics

The **distribute** construct specifies that the **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 **loop schedule** is **implementation defined**.

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*.

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 **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 **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**.

A **thread** dispatches a registered **dispatch callback** for each occurrence of a *distribute-chunk-begin* event in that **thread**. The **callback** occurs in the context of the **initial task**.

Restrictions

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

- The **collapsed 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 13.6](#)
- **collapse** Clause, see [Section 6.4.5](#)
- **dispatch** Callback, see [Section 40.4.2](#)
- **dist_schedule** Clause, see [Section 19.7.1](#)
- **firstprivate** Clause, see [Section 7.3.4](#)
- Consistent Loop Schedules, see [Section 6.4.4](#)
- **induction** Clause, see [Section 8.13](#)
- **lastprivate** Clause, see [Section 7.3.5](#)
- **order** Clause, see [Section 18.3](#)
- **private** Clause, see [Section 7.3.3](#)
- OMPT **scope_endpoint** Type, see [Section 39.27](#)
- **teams** Construct, see [Section 18.2](#)
- **work** Callback, see [Section 40.4.1](#)
- OMPT **work** Type, see [Section 39.41](#)

19.7.1 dist_schedule Clause

Name: <code>dist_schedule</code>	Properties: schedule-specification , unique
---	--

Arguments

Name	Type	Properties
<i>kind</i>	Keyword: static	default
<i>chunk_size</i>	expression of integer type	ultimate , optional , posi- tive , region-invariant

Modifiers

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

1 **Directives**
2 **distribute**

3 **Semantics**

4 The **dist_schedule** clause specifies how **collapsed iterations** of a **distribute** construct are
5 divided into **chunks**, and how these **chunks** are distributed among the **teams** of the **league**. If
6 *chunk_size* is not specified, the **collapsed iteration space** is divided into **chunks** that are
7 approximately equal in size, and at most one **chunk** is distributed to each **initial team** of the **league**.
8 If the *chunk_size* argument is specified, **collapsed iterations** are divided into **chunks** of *chunk_size*
9 iterations. The *chunk_size* expression is evaluated using the **original list items** of any **variables** that
10 become **private variables** in the **distribute construct**. Whether, in what order, or how many
11 times, any side effects of the evaluation of this expression occur is unspecified. The use of a
12 **variable** in a **dist_schedule** clause expression of a **distribute construct** causes an implicit
13 reference to the **variable** in all enclosing **constructs**. These **chunks** are assigned to the **initial teams**
14 of the **league** in a round-robin fashion in the order of their **team number**.

15 **Restrictions**

16 Restrictions to the **dist_schedule** clause are as follows:

- 17
 - 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 **collapse-affected loops** is a
19 **non-rectangular loop**.

20 **Cross References**

- 21
 - **distribute** Construct, see [Section 19.7](#)

22 **19.8 loop Construct**

23	Name: loop Category: executable	Association: loop nest Properties: order-concurrent-nestable , partitioned , simdizable , team-executed , teams-nestable , work-distribution , worksharing
----	--	--

24 **Clauses**

25 **bind**, **collapse**, **lastprivate**, **order**, **private**, **reduction**

26 **Binding**

27 The **bind** clause determines the **binding region**, which determines the **binding thread set**.

28 **Semantics**

29 A **loop construct** specifies that the **collapsed iterations** execute in the context of the **binding thread**
30 **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**.

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**. Further, the **loop construct** has the **SIMDizable property** if and only if its **binding region** is not defined.

Fortran

The **collapse-affected loop** may be a **DO CONCURRENT** loop.

Fortran

Restrictions

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

- A **list item** must not appear in a **lastprivate clause** unless it is the **loop-iteration variable** of an **affected 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 **collapse-affected loop** is a **DO CONCURRENT** loop, neither the **data-sharing attribute clauses** nor the **collapse clause** may be specified.
- If a **variable** is accessed in more than one iteration of a **DO CONCURRENT** loop that is associated with a **loop construct** and at least one of the accesses modifies the **variable**, the **variable** must have locality specified in the **DO CONCURRENT** loop.

Fortran

Cross References

- **bind** Clause, see [Section 19.8.1](#)
- **collapse** Clause, see [Section 6.4.5](#)
- Consistent Loop Schedules, see [Section 6.4.4](#)
- **lastprivate** Clause, see [Section 7.3.5](#)
- **order** Clause, see [Section 18.3](#)
- **private** Clause, see [Section 7.3.3](#)
- **reduction** Clause, see [Section 8.10](#)
- **teams** Construct, see [Section 18.2](#)

19.8.1 bind Clause

Name: bind	Properties: unique
-------------------	------------------------------------

Arguments

Name	Type	Properties
<i>binding</i>	Keyword: parallel , teams , thread	default

Modifiers

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

1 threads that are executing that region. If the binding region is not defined, then the binding thread
2 set is the encountering thread.

3 Restrictions

4 Restrictions to the bind clause are as follows:

- 5 • If teams is specified as *binding* then the corresponding loop region must be a strictly
6 nested region of a teams region.
- 7 • If teams is specified as *binding* and the corresponding loop region executes on a non-host
8 device then the behavior of a reduction clause that appears on the corresponding loop
9 construct is unspecified if the construct is not nested inside a teams construct.
- 10 • If parallel is specified as *binding*, the behavior is unspecified if the corresponding loop
11 region is a closely nested region of a simd region.

12 Cross References

- 13 • loop Construct, see Section 19.8

20 Tasking Constructs

This chapter defines [directives](#) and concepts related to [explicit tasks](#).

20.1 task Construct

Name: <code>task</code> Category: executable	Association: block Properties: parallelism-generating , thread-limiting , task-generating
---	---

Clauses

[affinity](#), [allocate](#), [default](#), [depend](#), [detach](#), [final](#), [firstprivate](#), [if](#),
[in_reduction](#), [mergeable](#), [priority](#), [private](#), [replayable](#), [shared](#),
[threadset](#), [transparent](#), [untied](#)

Clause set

Properties: exclusive	Members: detach , mergeable
--	--

Binding

The [binding thread set](#) of the [task region](#) is the set of [threads](#) specified in the [threadset](#) clause. A [task region](#) binds to the innermost enclosing [parallel region](#).

Semantics

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

The [encountering thread](#) may immediately execute the [task](#), or defer its execution. In the latter case, any [thread](#) of the current [binding thread set](#) may be assigned the [task](#). [Task completion](#) of the [task](#) can be guaranteed using [task synchronization constructs](#) and [clauses](#). If a [task construct](#) is encountered during execution of an outer [task](#), the [generated task region](#) that corresponds to this [construct](#) is not a part of the outer [task region](#) unless the [generated task](#) is an [included task](#).

A [detachable task](#) is completed when the execution of its associated [structured block](#) is completed and the [allow-completion event](#) is fulfilled. If no [detach](#) clause is present on a [task construct](#), the [generated task](#) is completed when the execution of its associated [structured block](#) is completed.

A **thread** that encounters a **task scheduling point** within the **task region** may temporarily suspend the **task region**.

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

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

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

Execution Model Events

The *task-create event* occurs when a **thread** encounters a **task-generating construct**. The **event** occurs after the **task** is initialized but before its execution begins and before the encountering thread resumes execution of any **task**.

Tool Callbacks

A **thread** dispatches a registered **task_create callback** for each occurrence of a *task-create event* in the context of the **encountering task**. The *flags* argument of this **callback** indicates the **task** types shown in Table 20.1.

TABLE 20.1: **task_create** Callback Flags Evaluation

Operation	Evaluates to <i>true</i>
(<i>flags</i> & ompt_task_explicit)	Always in the dispatched callback
(<i>flags</i> & ompt_task_importing)	If the task is an importing task
(<i>flags</i> & ompt_task_exporting)	If the task is an exporting task
(<i>flags</i> & ompt_task_underrferred)	If the task is an underrferred task
(<i>flags</i> & ompt_task_final)	If the task is a final task
(<i>flags</i> & ompt_task_untied)	If the task is an untied task

table continued on next page

table continued from previous page

Operation	Evaluates to <i>true</i>
<code>(flags & omp_task_mergeable)</code>	If the task is a mergeable task
<code>(flags & omp_task_merged)</code>	If the task is a merged task

Cross References

- **affinity** Clause, see [Section 20.10](#)
- **allocate** Clause, see [Section 13.6](#)
- **default** Clause, see [Section 7.3.1](#)
- **depend** Clause, see [Section 23.9.5](#)
- **detach** Clause, see [Section 20.11](#)
- **final** Clause, see [Section 20.7](#)
- **firstprivate** Clause, see [Section 7.3.4](#)
- Task Scheduling, see [Section 20.14](#)
- **if** Clause, see [Section 5.6](#)
- **in_reduction** Clause, see [Section 8.12](#)
- **mergeable** Clause, see [Section 20.5](#)
- **omp_fulfill_event** Routine, see [Section 29.2.1](#)
- **priority** Clause, see [Section 20.9](#)
- **private** Clause, see [Section 7.3.3](#)
- **replayable** Clause, see [Section 20.6](#)
- **shared** Clause, see [Section 7.3.2](#)
- **task_create** Callback, see [Section 40.5.1](#)
- OMPT **task_flag** Type, see [Section 39.37](#)
- **threadset** Clause, see [Section 20.8](#)
- **transparent** Clause, see [Section 23.9.6](#)
- **untied** Clause, see [Section 20.4](#)

20.2 taskloop Construct

Name: taskloop Category: executable	Association: loop nest Properties: parallelism-generating, SIMD-partitionable, task-generating
--	---

Subsidiary directives

task_iteration

Clauses

allocate, collapse, default, final, firstprivate, grainsize, if, in_reduction, induction, lastprivate, mergeable, nogroup, num_tasks, priority, private, reduction, replayable, shared, threadset, transparent, untied

Clause set

synchronization-clause

Properties: exclusive	Members: nogroup, reduction
------------------------------	------------------------------------

Clause set

granularity-clause

Properties: exclusive	Members: grainsize, num_tasks
------------------------------	--------------------------------------

Binding

The binding thread set of the taskloop region is the 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 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. Tasks created by a taskloop directive can be affected by task_iteration directives that are subsidiary directives of that taskloop directive. If a task_iteration directive on which a depend clause appears is a subsidiary directive of the taskloop construct then the behavior is as if the order of the creation of the generated tasks is in increasing collapsed iteration order with respect to their assigned chunks. Otherwise, the order of the creation of the generated tasks is unspecified and programs that rely on the 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 items** is present. Thus, the **generated tasks** are participants of a **reduction** previously defined by a **reduction-scoping** clause.

If a **threadset** clause is present, the behavior is as if each **generated task** was defined by a **task** construct on which a **threadset** clause with the same set of **threads** is present. Thus, the **binding thread set** of the generated **tasks** is the same as that of the **taskloop** region.

If a **transparent** clause is present, the behavior is as if each **generated task** was defined by a **task** construct on which a **transparent** clause with the same *impex-type* argument is present.

If no clause from the *granularity-clause* clause set is present, the number of loop **tasks** generated and the number of **logical iterations** assigned to these **tasks** is **implementation defined**.

When an **if** clause is present and the **if** clause expression evaluates to *false*, **underrun tasks** are generated. The use of a **variable** in an **if** clause expression causes an implicit reference to the **variable** in all enclosing **constructs**.

▼ C++ ▼

For **firstprivate** variables of class type, the number of invocations of copy constructors that perform the initialization is **implementation defined**.

▲ C++ ▲

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

Execution Model Events

The *taskloop-begin* event occurs upon entering the **taskloop** region. A *taskloop-begin* will precede any *task-create* events for the generated **tasks**. The *taskloop-end* event occurs upon completion of the **taskloop** region.

Events for an implicit **taskgroup** region that surrounds the **taskloop** region are the same as for the **taskgroup** construct.

The *taskloop-iteration-begin* event occurs at the beginning of each *logical-iteration* of a **taskloop** region before an **explicit task** executes the **logical iteration**. The *taskloop-chunk-begin* event occurs before an **explicit task** executes any of its associated **logical iterations** in a **taskloop** region.

Tool Callbacks

A **thread** dispatches a registered **work** callback for each occurrence of a *taskloop-begin* and *taskloop-end* event in that **thread**. The callback occurs in the context of the **encountering task**. The callback receives **ompt_scope_begin** or **ompt_scope_end** as its *endpoint* argument, as appropriate, and **ompt_work_taskloop** as its *work_type* argument.

A **thread** dispatches a registered **dispatch** callback for each occurrence of a *taskloop-iteration-begin* or *taskloop-chunk-begin* event in that **thread**. The callback binds to the **explicit task** executing the **logical iterations**.

Restrictions

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

- The *reduction-modifier* must be **default**.
- The **conditional** *lastprivate-modifier* must not be specified.
- If the **taskloop** construct is associated with a **task_iteration** directive, none of the **taskloop-affected loops** may be the **generated loop** of a **loop-transforming construct**.

Cross References

- **allocate** Clause, see [Section 13.6](#)
- **collapse** Clause, see [Section 6.4.5](#)
- **default** Clause, see [Section 7.3.1](#)
- **dispatch** Callback, see [Section 40.4.2](#)
- **final** Clause, see [Section 20.7](#)
- **firstprivate** Clause, see [Section 7.3.4](#)
- Canonical Loop Nest Form, see [Section 6.4.1](#)
- **grainsize** Clause, see [Section 20.2.1](#)
- **if** Clause, see [Section 5.6](#)
- **in_reduction** Clause, see [Section 8.12](#)
- **induction** Clause, see [Section 8.13](#)
- **lastprivate** Clause, see [Section 7.3.5](#)
- **mergeable** Clause, see [Section 20.5](#)
- **nogroup** Clause, see [Section 23.7](#)
- **num_tasks** Clause, see [Section 20.2.2](#)
- **priority** Clause, see [Section 20.9](#)

- **private** Clause, see [Section 7.3.3](#)
- **reduction** Clause, see [Section 8.10](#)
- **replayable** Clause, see [Section 20.6](#)
- OMPT **scope_endpoint** Type, see [Section 39.27](#)
- **shared** Clause, see [Section 7.3.2](#)
- **task** Construct, see [Section 20.1](#)
- **task_iteration** Directive, see [Section 20.2.3](#)
- **taskgroup** Construct, see [Section 23.4](#)
- **threadset** Clause, see [Section 20.8](#)
- **transparent** Clause, see [Section 23.9.6](#)
- **untied** Clause, see [Section 20.4](#)
- **work** Callback, see [Section 40.4.1](#)
- OMPT **work** Type, see [Section 39.41](#)

20.2.1 grainsize Clause

Name: grainsize	Properties: taskgraph-altering , 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: <i>directive-name</i> (a 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 `collapse-affected loops` may be `non-rectangular loops`.

Cross References

- `taskloop` Construct, see [Section 20.2](#)

20.2.2 num_tasks Clause

Name: <code>num_tasks</code>	Properties: <code>taskgraph-altering</code> , <code>unique</code>
-------------------------------------	--

Arguments

Name	Type	Properties
<i>num-tasks</i>	expression of integer type	<code>positive</code>

Modifiers

Name	Modifies	Type	Properties
<i>prescriptiveness</i>	<i>num-tasks</i>	Keyword: <code>strict</code>	<code>unique</code>
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: <i>directive-name</i> (a <code>directive name</code>)	<code>unique</code>

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 `collapse-affected loops` may be `non-rectangular loops`.

1 **Cross References**

- 2 • **taskloop** Construct, see [Section 20.2](#)

3 **20.2.3 task_iteration Directive**

4

Name: <code>task_iteration</code> Category: <code>subsidiary</code>	Association: <code>unassociated</code> Properties: <code>default</code>
--	--

5 **Enclosing directives**

6 `taskloop`

7 **Clauses**

8 `affinity`, `depend`, `if`

9 **Semantics**

10 The `task_iteration` directive is a `subsidiary directive` that controls the per-iteration
11 `task-execution` attributes of the `generated tasks` of its associated `taskloop construct`, which is the
12 innermost enclosing `taskloop construct`, as described below.

13 For each `task-inherited clause` specified on the `task_iteration` directive, the behavior is as if
14 each `task` generated by the enclosing `taskloop construct` is specified with a corresponding `clause`
15 that has the same *clause-specification*, but adjusted as follows. These `clauses` are instantiated for
16 each instance of the `loop-iteration variables` for which the *if-expression* of the `if clause` evaluates
17 to *true*. If an `if clause` is not specified on the `task_iteration` directive, the behavior is as if
18 the *if-expression* evaluates to *true*.

19 **Restrictions**

20 The restrictions to the `task_iteration` directive are as follows:

- 21 • Each `task_iteration` directive must appear in the `loop body` of one of the
22 `taskloop-affected loops` and must precede all statements and `directives` (except other
23 `task_iteration` directives) in that `loop body`.
- 24 • If a `task_iteration` directive appears in the `loop body` of one of the
25 `taskloop-affected loops`, no `intervening code` may occur between any two
26 `collapse-affected loops` of the `taskloop-affected loops`.

27 **Cross References**

- 28 • **affinity** Clause, see [Section 20.10](#)
- 29 • **depend** Clause, see [Section 23.9.5](#)
- 30 • **if** Clause, see [Section 5.6](#)
- 31 • **iterator** Modifier, see [Section 5.2.6](#)
- 32 • **task** Construct, see [Section 20.1](#)

- **taskloop** Construct, see [Section 20.2](#)

20.3 taskgraph Construct

Name: taskgraph	Association: block
Category: executable	Properties: default

Clauses

[graph_id](#), [graph_reset](#), [if](#), [nogroup](#)

Binding

The [binding thread set](#) of a **taskgraph** region is all [threads](#) on the [current device](#). The [binding task set](#) of a **taskgraph** region is all [tasks](#) of the [current team](#) that are generated in the [region](#).

Semantics

When a [thread](#) encounters a **taskgraph** construct, a **taskgraph** region is generated for which execution entails one of the following:

- Execution of the [structured block](#) associated with the [construct](#), while optionally creating a [taskgraph record](#) of all encountered [replayable constructs](#) and the sequence in which they are encountered; or
- A [replay execution](#) of the last [matching taskgraph record](#) of the [construct](#).

If a **taskloop** construct is encountered in the **taskgraph** region, the behavior is as if each [task](#) that it generates is instead generated by a **task** construct. If a [task-generating construct](#) is encountered in the **taskgraph** construct as part of its corresponding [region](#), then it is a [replayable construct](#) of the [region](#) unless otherwise specified by the [replayable clause](#). If a [depend clause](#) with a [depobj task-dependence-type](#) is present on a [replayable construct](#) then for each listed [depend object](#) the behavior is as if a [depend clause](#) with the [dependence type](#) and [locator list item](#) represented by the [depend object](#) is instead present on the [construct](#). Whether a [task-generating construct](#) that is encountered as part of the **taskgraph** region, but not in the **taskgraph** construct, is a [replayable construct](#) of the [region](#) is unspecified, unless the [replayable clause](#) is present on that [construct](#). For the purposes of the **taskgraph** region, a **taskwait** construct on which the [depend clause](#) appears is a [task-generating construct](#).

A [taskgraph record](#) contains a record of the following:

- The [graph-id-value](#) specified in the [graph_id clause](#) upon encountering the [construct](#);
- The sequence of encountered [replayable constructs](#) in the **taskgraph** region, along with their [subsidiary directives](#); and
- For each [replayable construct](#), a [saved data environment](#).

A [clause](#) or [modifier](#) argument for a [replayable construct](#) is recorded after evaluating all expressions that compose the argument and substituting the resulting values for those expressions. Additionally,

if a **clause** argument or a **modifier** argument specification requires a **locator list item** or a **variable list item**, then:

- For a **locator list item** of a **taskgraph-altering clause**, only the **storage locations** are recorded;
- Otherwise, the identifier that designates the **base variable** or **base pointer** of the **list item** is recorded along with any values that are needed to reconstruct the **list item**.

The **saved data environment** of each **replayable construct** in the **taskgraph record** includes copies of all **variables** that do not have **static storage duration** and that are **firstprivate** in the **replayable construct**, with values that are captured from the **enclosing data environment** when the **construct** is encountered. Additionally, it includes copies of all **variables** that have **static storage duration** and that appear in a **firstprivate clause** that has the **saved modifier** on the **construct**. Finally, it includes references to any other **variables** that have **static storage duration**, exist in the **enclosing data environment** of the **replayable construct**, and do not exist in the **enclosing data environment** of the **taskgraph construct**.

The **taskgraph record** becomes a **finalized taskgraph record** on exit from the **taskgraph region** in which it is created. An implementation may create a **finalized taskgraph record** prior to the first execution of the **taskgraph region**, if it can guarantee that the contents of the record would match the record that would have been created during an execution of the **region**. In this case, a **replay execution** of that **taskgraph record** may occur upon first encountering the **taskgraph construct**.

If the **graph_id clause** is not present, an existing **finalized taskgraph record** that was generated for the **construct** when encountered on the same **device** is the **matching taskgraph record**. Otherwise, an existing **finalized taskgraph record** that was generated for the **construct** when encountered on the same **device** is the **matching taskgraph record** if the *graph-id-value* specified in the **graph_id clause** matches the value in the **graph_id clause** that was saved in the record.

Each **finalized taskgraph record** has an associated *replay count* that is initialized to zero. If the **graph_reset clause** is not present or its argument evaluates to *false*, the **encountering task** of the **taskgraph region** is not a **final task**, and a **matching taskgraph record** exists, the **matching taskgraph record** is replayed and its replay count is incremented by one. A **replay execution** of a **taskgraph record** has the effect of encountering the recorded **replayable constructs**, with their recorded **clause** and **modifier** arguments unless otherwise specified, in their recorded sequence and implies all semantics defined for those **constructs** except as otherwise specified in this section. A **replay execution** does not entail execution of any code that is part of both the **taskgraph region** and the **encountering task region**. Any changes from when the **matching taskgraph record** was created to the arguments or **modifiers** of a **taskgraph-altering clause** that appears on a **replayable construct** does not alter the behavior of a **replay execution** of that **taskgraph record**. The replay count is decremented by one once all **tasks** that are generated by the **replayable constructs** have completed.

If completion of a **taskgraph region** results in a new **finalized taskgraph record** when a **matching taskgraph record** already exists, the behavior is as if the new record replaces the old record, with the old record being discarded once its replay count reaches zero.

When executing a **replayable construct** during a **replay execution**, unless otherwise specified by a **saved modifier** on a **data-environment attribute clause**, its **enclosing data environment** (inclusive of ICVs with **data environment ICV scope**) is the **enclosing data environment** of the **taskgraph construct**. If a **variable** does not exist in the **enclosing data environment** of the **taskgraph construct** then the **saved data environment** in the **taskgraph record** is used as the **enclosing data environment** for that **variable**. If the **replayable construct** permits an **ICV-defaulted clause** and the **clause** is not present, in a **replay execution** of the **construct** the **ICV** in the **enclosing data environment** of the **taskgraph construct** determines the value of the **clause** argument.

If the **if clause** is present and its argument evaluates to *false*, execution of the **taskgraph region** will not create a **taskgraph record** or entail replaying a **matching taskgraph record** of the **construct**.

If the **nogroup clause** is not present, the **taskgraph region** executes as if enclosed by a **taskgroup region**.

Whether **foreign tasks** are recorded or not in a **taskgraph record** and the manner in which they are executed during a **replay execution** if they are recorded is **implementation defined**.

Execution Model Events

Events for the implicit **taskgroup region** that surrounds the **taskgraph region** when no **nogroup clause** is specified are the same as for the **taskgroup construct**.

The events that occur during a **replay execution** of a **taskgraph region** is unspecified.

Tool Callbacks

Callbacks associated with **events** for the **taskgroup region** are the same as for the **taskgroup construct** as defined in [Section 23.4](#).

Restrictions

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

- **Task-generating constructs** are the only **constructs** that may be encountered as part of the **taskgraph region**.
- A **taskgraph construct** must not be encountered in a **final task region**.
- A **replayable construct** that generates an **importing** or **exporting transparent task**, a **detachable task**, or an **undelayed task** must not be encountered in a **taskgraph region**.
- Any **variable** referenced in a **replayable construct** that does not have **static storage duration** and that does not exist in the **enclosing data environment** of the **taskgraph construct** must be a **private-only** or **firstprivate variable** in the **replayable construct**.
- A **list item** of a **clause** on a **replayable construct** that accepts a **locator list** and is not a **taskgraph-altering clause** must have a **base variable** or **base pointer**.
- Any **variable** that appears in an expression of a **variable list item** or **locator list item** for a **clause** on a **replayable construct** and does not designate the **base variable** or **base pointer** of

that [list item](#) must be listed in a [data-environment attribute clause](#) with the *saved* modifier on that [construct](#).

- If a [construct](#) that permits the [nogroup](#) clause is encountered in a [taskgraph](#) region then the [nogroup](#) clause must be specified with the *do_not_synchronize* argument evaluating to *true*.
- If a [task](#) generated by one [construct](#) in a [taskgraph](#) region is an [antecedent task](#) of a [task](#) generated by another [construct](#) in the [region](#), either both [constructs](#) must be [replayable](#) or both must be not [replayable](#).

Cross References

- [graph_id](#) Clause, see [Section 20.3.1](#)
- [graph_reset](#) Clause, see [Section 20.3.2](#)
- [if](#) Clause, see [Section 5.6](#)
- [nogroup](#) Clause, see [Section 23.7](#)
- [task](#) Construct, see [Section 20.1](#)
- [taskgroup](#) Construct, see [Section 23.4](#)

20.3.1 [graph_id](#) Clause

Name: graph_id	Properties: unique
--------------------------------	------------------------------------

Arguments

Name	Type	Properties
<i>graph-id-value</i>	expression of OpenMP integer type	<i>default</i>

Directives

[taskgraph](#)

Semantics

The [graph_id](#) clause specifies the *graph-id-value* that identifies a [taskgraph](#) record. At most, one [matching taskgraph record](#) exists for a given *graph-id-value*.

Cross References

- [taskgraph](#) Construct, see [Section 20.3](#)

20.3.2 [graph_reset](#) Clause

Name: graph_reset	Properties: unique
-----------------------------------	------------------------------------

Arguments

Name	Type	Properties
<i>graph-reset-expression</i>	expression of OpenMP logical type	optional

Directives

taskgraph

Semantics

If *graph-reset-expression* evaluates to *true*, any existing **matching taskgraph record** is discarded if a replay of the record is not in progress (i.e., if its replay count equals zero). If the replay count is non-zero, the **matching taskgraph record** is not replayed and instead the **structured block** associated with the **taskgraph construct** is executed; in this case, the **matching taskgraph record** is discarded once its replay count reaches zero. If *graph-reset-expression* is not specified, the effect is as if *graph-reset-expression* evaluates to *true*.

Cross References

- **taskgraph** Construct, see [Section 20.3](#)

20.4 untied Clause

Name: untied	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: <i>directive-name</i> (a 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** Construct, see [Section 20.1](#)
- **taskloop** Construct, see [Section 20.2](#)

20.5 mergeable Clause

Name: mergeable	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: <i>directive-name</i> (a directive name)	unique

Directives

[target_data](#), [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*. If the [generated task](#) is a [mergeable task](#) that is also an [underrferred task](#), the implementation may generate a [merged task](#) instead.

Cross References

- **target_data** Construct, see [Section 21.7](#)
- **task** Construct, see [Section 20.1](#)
- **taskloop** Construct, see [Section 20.2](#)

20.6 replayable Clause

Name: replayable	Properties: unique
-------------------------	------------------------------------

1 **Arguments**

2

Name	Type	Properties
<i>replayable-expression</i>	expression of OpenMP logical type	constant, optional

3 **Modifiers**

4

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

5 **Directives**

6 **target**, **target_enter_data**, **target_exit_data**, **target_update**, **task**,
7 **taskloop**, **taskwait**

8 **Semantics**

9 If *replayable-expression* evaluates to *true*, the **replayable** clause specifies that the **construct** on
10 which it appears is a **replayable construct**. If *replayable-expression* evaluates to *false*, the
11 **replayable** clause specifies that the **construct** on which it appears is not a **replayable construct**.
12 If *replayable-expression* is not specified, the effect is as if *replayable-expression* evaluates to *true*.

13 **Cross References**

- 14 • **target** Construct, see [Section 21.8](#)
15 • **target_enter_data** Construct, see [Section 21.5](#)
16 • **target_exit_data** Construct, see [Section 21.6](#)
17 • **target_update** Construct, see [Section 21.9](#)
18 • **task** Construct, see [Section 20.1](#)
19 • **taskloop** Construct, see [Section 20.2](#)
20 • **taskwait** Construct, see [Section 23.5](#)

21 **20.7 final Clause**

22

Name: final	Properties: unique
---------------------------	---------------------------

23 **Arguments**

24

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: <i>directive-name</i> (a <i>directive name</i>)	<i>unique</i>

Directives

task, taskloop

Semantics

The **final** clause specifies that **tasks** generated by the **construct** on which it appears are **final tasks** if the *finalize* expression evaluates to *true*. All **task-generating constructs** on which the **final** clause may be specified 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. If *finalize* is not specified, the effect is as if *finalize* evaluates to *true*.

Cross References

- **task** Construct, see [Section 20.1](#)
- **taskloop** Construct, see [Section 20.2](#)

20.8 threadset Clause

Name: threadset	Properties: <i>unique</i>
------------------------	----------------------------------

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: <i>directive-name</i> (a <i>directive name</i>)	<i>unique</i>

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** Construct, see [Section 20.1](#)
- **taskloop** Construct, see [Section 20.2](#)

20.9 priority Clause

Name: priority	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: <i>directive-name</i> (a directive name)	unique

Directives

target, **target_data**, **target_enter_data**, **target_exit_data**, **target_update**, **task**, **taskgraph**, **taskloop**

Semantics

The **priority** clause specifies, in the *priority-value* argument, a **task priority** for the **construct** on which it appears . 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 **task 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 **task priorities** may have **unspecified behavior**.

Cross References

- *max-task-priority-var* **ICV**, see [Table 3.1](#)
- **target** Construct, see [Section 21.8](#)
- **target_data** Construct, see [Section 21.7](#)
- **target_enter_data** Construct, see [Section 21.5](#)
- **target_exit_data** Construct, see [Section 21.6](#)

- **target_update** Construct, see [Section 21.9](#)
- **task** Construct, see [Section 20.1](#)
- **taskgraph** Construct, see [Section 20.3](#)
- **taskloop** Construct, see [Section 20.2](#)

20.10 affinity Clause

Name: <code>affinity</code>	Properties: <code>task-inherited</code>
------------------------------------	--

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 list of iterator specifier list item type (<i>default</i>)	unique
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: <i>directive-name</i> (a directive name)	unique

Directives

`target_data`, `task`, `task_iteration`

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

- `iterator` Modifier, see [Section 5.2.6](#)
- `target_data` Construct, see [Section 21.7](#)
- `task` Construct, see [Section 20.1](#)
- `task_iteration` Directive, see [Section 20.2.3](#)

20.11 detach Clause

Name: <code>detach</code>	Properties: data-sharing attribute, innermost-leaf, privatization, unique
---------------------------	---

Arguments

Name	Type	Properties
<i>event-handle</i>	variable of <code>event_handle</code> type	<i>default</i>

Modifiers

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

Directives

`target_data`, `task`

Semantics

The `detach` clause specifies that the `task` generated by the `construct` on which it appears is a detachable task. The clause provides a superset of the functionality provided by the `private` clause. 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 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 any `data-environment attribute clause` on the same `construct`.
- A `variable` that is part of an `aggregate variable` cannot appear in a `detach` clause.

- *event-handle* must not have the **POINTER** attribute.
- If *event-handle* has the **ALLOCATABLE** attribute, the allocation status must be allocated when the **task construct** is encountered, and the allocation status must not be changed, either explicitly or implicitly, in the **task region**.

Cross References

- OpenMP **event_handle** Type, see [Section 26.6.1](#)
- **target_data** Construct, see [Section 21.7](#)
- **task** Construct, see [Section 20.1](#)

20.12 taskyield Construct

Name: taskyield Category: executable	Association: unassociated Properties: default
---	--

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 20.14](#)

20.13 Initial Task

Execution Model Events

While no **events** are associated with the **implicit parallel region** in each **initial thread**, several **events** are associated with **initial tasks**. The *initial-thread-begin event* occurs in an **initial thread** after the OpenMP runtime invokes the **OMPT-tool initializer** but before the **initial thread** begins to execute the first **explicit region** in the **initial task**. The *initial-task-begin event* occurs after an *initial-thread-begin event* but before the first **explicit region** in the **initial task** begins to execute. The *initial-task-end event* occurs before an *initial-thread-end event* but after the last **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 **OMPT-tool finalizer**.

Tool Callbacks

A **thread** dispatches a registered **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** receives **ompt_thread_initial** as its *thread_type* argument.

A **thread** dispatches a registered **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 **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**. In the dispatched **callback**, $(flags \ \& \ \mathbf{ompt_task_initial})$ and $(flags \ \& \ \mathbf{ompt_task_implicit})$ evaluate to *true*.

A **thread** dispatches a registered **thread_end** callback for the *initial-thread-end event* in that **thread**. The **callback** occurs in the context of the **thread**. The **implicit parallel region** does not dispatch a **parallel_end** callback; however, the **implicit parallel region** can be finalized within this **thread_end** callback.

Cross References

- **implicit_task** Callback, see [Section 40.5.3](#)
- **parallel_end** Callback, see [Section 40.3.2](#)
- OMPT **scope_endpoint** Type, see [Section 39.27](#)
- OMPT **task_flag** Type, see [Section 39.37](#)
- OMPT **thread** Type, see [Section 39.39](#)
- **thread_begin** Callback, see [Section 40.1.3](#)
- **thread_end** Callback, see [Section 40.1.4](#)

20.14 Task Scheduling

Whenever a **thread** reaches a **task scheduling point**, it may begin or resume execution of a **task** from its **schedulable task set**. An **idle thread** is treated as if it is always at a **task scheduling point**. For other **threads**, **task scheduling points** are implied at the following locations:

- During the generation of an **explicit task**;
- The point immediately following the generation of an **explicit task**;
- After the point of completion of the **structured block** associated with a **task**;
- In a **taskyield** region;
- In a **taskwait** region;
- At the end of a **taskgroup** region;

- At the beginning and end of a **taskgraph** region;
- In an **implicit barrier** region;
- In an explicit **barrier** region;
- During the generation of a **target** region;
- The point immediately following the generation of a **target** region;
- In a **target_update** region;
- In a **target_enter_data** region;
- In a **target_exit_data** region;
- In each instance of any **memory-copying routine**; and
- In each instance of any **memory-setting routine**.

When a **thread** encounters a **task scheduling point** it may do one of the following, subject to the **task** scheduling constraints specified below:

- Begin execution of a **tied task** in its **schedulable task set**;
- Resume the suspended **task region** of any **task** to which it is **tied**;
- Begin execution of an **untied task** in its **schedulable task set**; or
- Resume the suspended **task region** of any **untied task** in its **schedulable task set**.

If more than one of the above choices is available, which one is chosen is unspecified.

Task Scheduling Constraints are as follows:

1. If any suspended **tasks** are **tied** to the **thread** and are not suspended in a **barrier region**, a new explicit **tied task** may be scheduled only if it is a **descendent task** of all of those suspended **tasks**. Otherwise, any new explicit **tied task** may be scheduled.
2. A **dependent task** shall not start its execution until its **task dependences** are fulfilled.
3. A **task** shall not be scheduled while another **task** has been scheduled but has not yet completed, if they are **mutually exclusive tasks**.
4. A **task** shall not start or resume execution on an **unassigned thread** if it would result in the total number of **free-agent threads** in the **OpenMP thread pool** exceeding *free-agent-thread-limit-var*.

Task scheduling points dynamically divide **task regions** into **subtasks**. Each **subtask** is executed uninterrupted from start to end. Different **subtasks** of the same **task region** are executed in the order in which they are encountered. In the absence of **task synchronization constructs**, the order in which a **thread** executes **subtasks** of different **tasks** in its **schedulable task set** is unspecified.

A program must behave correctly and consistently with all conceivable scheduling sequences that are compatible with the rules above. A program that relies on any other assumption about `task` scheduling is a `non-conforming program`.

Note – For example, if `threadprivate memory` is accessed (explicitly in the source code or implicitly in calls to library `procedures`) in one `subtask` of a `task region`, its value cannot be assumed to be preserved into the next `subtask` of the same `task region` if another `schedulable task` exists that modifies it.

As another example, if different `subtasks` of a `task region` invoke a `lock-acquiring routine` and its corresponding `lock-releasing routine`, no invocation of a `lock-acquiring routine` for the same `lock` should be made in any `subtask` of another `task` that the executing `thread` may schedule. Otherwise, deadlock is possible. A similar situation can occur when a `critical region` spans multiple `subtasks` of a `task` and another `schedulable task` contains a `critical region` with the same name.

Execution Model Events

The `task-schedule event` occurs in a `thread` when the `thread` switches `tasks` at a `task scheduling point`; no `event` occurs when switching to or from a `merged task`.

Tool Callbacks

A `thread` dispatches a registered `task_schedule callback` for each occurrence of a `task-schedule event` in the context of the `task` that begins or resumes. The `prior_task_status` argument is used to indicate the cause for suspending the prior `task`. This cause may be the completion of the prior `task region`, the encountering of a `taskyield construct`, or the encountering of an active `cancellation point`.

Cross References

- `task_schedule` Callback, see [Section 40.5.2](#)

21 Device Directives and Clauses

This chapter defines constructs and concepts related to device execution.

21.1 device_type Clause

Name: device_type	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: <i>directive-name</i> (a directive name)	unique

Directives

begin declare_target, **declare_target**, **groupprivate**, **target**

Semantics

If the **device_type** clause appears on a **declarative directive**, the *device-type-description* argument specifies the type of **devices** for which a version of the **procedure** or **variable** should be made available. If the **device_type** clause appears on a **target construct**, the argument specifies the type of **devices** for which the implementation should support execution of the corresponding **target region**.

The **host** *device-type-description* specifies the **host device**. The **nohost** *device-type-description* specifies any supported **non-host device**. The **any** *device-type-description* specifies any **supported device**. If the **device_type** clause is not specified, the behavior is as if the **device_type** clause appears with **any** specified.

If the **device_type** clause specifies the **host device** on a **target construct** for which the **target device** is a **non-host device**, the corresponding **region** executes on the **host device**. Otherwise, if the **devices** specified by the **device_type** clause does not include the **target device** then **runtime error termination** is performed.

Cross References

- `begin declare_target` Directive, see [Section 15.9.2](#)
- `declare_target` Directive, see [Section 15.9.1](#)
- `groupprivate` Directive, see [Section 9.2](#)
- `target` Construct, see [Section 21.8](#)

21.2 device Clause

Name: <code>device</code>	Properties: <code>ICV-defaulted</code> , <code>unique</code>
----------------------------------	---

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: <code>ancestor</code> , <code>device_num</code>	<i>default</i>
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: <i>directive-name</i> (a directive name)	<code>unique</code>

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 with `device_num` as *device-modifier* and with a *device-description* that evaluates to the value of the *default-device-var* ICV.

1 **Restrictions**

- 2 • The **ancestor** *device-modifier* must not appear on the **device** clause on any **directive**
- 3 other than the **target** construct.
- 4 • If the **ancestor** *device-modifier* is specified, the *device-description* must evaluate to 1 and
- 5 a **requires** directive with the **reverse_offload** clause must be specified;
- 6 • If the **device_num** *device-modifier* is specified and *target-offload-var* is not **mandatory**,
- 7 *device-description* must evaluate to a **conforming device number**.

8 **Cross References**

- 9 • **dispatch** Construct, see [Section 15.7](#)
- 10 • *target-offload-var* ICV, see [Table 3.1](#)
- 11 • **interop** Construct, see [Section 22.1](#)
- 12 • **target** Construct, see [Section 21.8](#)
- 13 • **target_data** Construct, see [Section 21.7](#)
- 14 • **target_enter_data** Construct, see [Section 21.5](#)
- 15 • **target_exit_data** Construct, see [Section 21.6](#)
- 16 • **target_update** Construct, see [Section 21.9](#)

17 **21.3 thread_limit Clause**

18

Name: <code>thread_limit</code>	Properties: ICV-modifying, target-consistent, unique
--	---

19 **Arguments**

20

Name	Type	Properties
<i>threadlim</i>	expression of integer type	positive

21 **Modifiers**

22

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

23 **Directives**

24 **target, teams**

Semantics

As described in [Section 3.4](#), 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 on a [teams construct](#) 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](#) Construct, see [Section 21.8](#)
- [teams](#) Construct, see [Section 18.2](#)

21.4 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 [device_initialize callback](#) for each occurrence of a [device-initialize event](#) in that [thread](#). A [thread](#) dispatches a registered [device_load callback](#) for each occurrence of a [device-load event](#) in that [thread](#). A [thread](#) dispatches a registered [device_unload callback](#) for each occurrence of a [device-unload event](#) in that [thread](#). A [thread](#) dispatches a registered [device_finalize callback](#) for each occurrence of a [device-finalize event](#) in that [thread](#).

Restrictions

Restrictions to OpenMP [device](#) initialization are as follows:

- No [thread](#) may offload execution of a [construct](#) to a [device](#) until a dispatched [device_initialize callback](#) completes.
- No [thread](#) may offload execution of a [construct](#) to a [device](#) after a dispatched [device_finalize callback](#) occurs.

Cross References

- [device_finalize](#) Callback, see [Section 41.2](#)

- `device_initialize` Callback, see [Section 41.1](#)
- `device_load` Callback, see [Section 41.3](#)
- `device_unload` Callback, see [Section 41.4](#)

21.5 `target_enter_data` Construct

Name: <code>target_enter_data</code> Category: <code>executable</code>	Association: <code>unassociated</code> Properties: <code>parallelism-generating</code> , <code>task-generating</code> , <code>device</code> , <code>device-affecting</code> , <code>data-mapping</code> , <code>map-entering</code>
---	---

Clauses

`depend`, `device`, `if`, `map`, `nowait`, `priority`, `replayable`

Additional information

The `target_enter_data` directive may alternatively be specified with `target enter data` as the *directive-name*.

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` in the `map` clause 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`.

Assignment operations associated with mapping a `variable` (see [Section 11.3](#)) 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 20.1](#).

The *target-enter-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-enter-data-begin event* is a *target-task-begin event*. The *target-enter-data-end event* occurs after all other **events** associated with the **target_enter_data** construct.

Tool Callbacks

Callbacks associated with **events** for **target tasks** are the same as for the **task construct** defined in [Section 20.1](#); (*flags & ompt_target_target*) always evaluates to *true* in the dispatched **callback**.

A **thread** dispatches a registered **target_emi** callback with **ompt_scope_begin** as its *endpoint* argument and **ompt_target_enter_data** or **ompt_target_enter_data_nowait** if the **nowait** clause is present as its *kind* argument for each occurrence of a *target-enter-data-begin event* in that **thread** in the context of the **target task** on the **host device**. Similarly, a **thread** dispatches a registered **target_emi** callback with **ompt_scope_end** as its *endpoint* argument and **ompt_target_enter_data** or **ompt_target_enter_data_nowait** if the **nowait** clause is present as its *kind* argument for each occurrence of a *target-enter-data-end event* in that **thread** in the context of the **target task** on the **host device**.

Restrictions

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

- At least one **map clause** must appear on the **directive**.
- All **map clauses** must be **map-entering clauses**.

Cross References

- **depend** Clause, see [Section 23.9.5](#)
- **device** Clause, see [Section 21.2](#)
- **if** Clause, see [Section 5.6](#)
- **map** Clause, see [Section 11.3](#)
- **nowait** Clause, see [Section 23.6](#)
- **priority** Clause, see [Section 20.9](#)
- **replayable** Clause, see [Section 20.6](#)
- OMPT **scope_endpoint** Type, see [Section 39.27](#)
- OMPT **target** Type, see [Section 39.34](#)
- **target_emi** Callback, see [Section 41.8](#)
- **task** Construct, see [Section 20.1](#)

- OMPT `task_flag` Type, see [Section 39.37](#)

21.6 target_exit_data Construct

Name: <code>target_exit_data</code> Category: <code>executable</code>	Association: <code>unassociated</code> Properties: <code>parallelism-generating</code> , <code>task-generating</code> , <code>device</code> , <code>device-affecting</code> , <code>data-mapping</code> , <code>map-exiting</code>
--	--

Clauses

`depend`, `device`, `if`, `map`, `nowait`, `priority`, `replayable`

Additional information

The `target_exit_data` directive may alternatively be specified with `target exit data` as the *directive-name*.

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 11.3](#)) 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 20.1](#).

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 20.1; (*flags & ompt_target_target*) always evaluates to *true* in the dispatched **callback**.

A **thread** dispatches a registered **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 **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**.

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 23.9.5
- **device** Clause, see Section 21.2
- **if** Clause, see Section 5.6
- **map** Clause, see Section 11.3
- **nowait** Clause, see Section 23.6
- **priority** Clause, see Section 20.9
- **replayable** Clause, see Section 20.6
- OMPT **scope_endpoint** Type, see Section 39.27
- OMPT **target** Type, see Section 39.34
- **target_emi** Callback, see Section 41.8
- **task** Construct, see Section 20.1
- OMPT **task_flag** Type, see Section 39.37

21.7 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>parallelism-generating</code> , <code>sharing-task</code> , <code>simple-composite</code> , <code>task-generating</code>
---	--

Clauses

`affinity`, `allocate`, `default`, `depend`, `detach`, `device`, `firstprivate`, `if`, `in_reduction`, `map`, `mergeable`, `nogroup`, `nowait`, `priority`, `private`, `shared`, `transparent`, `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>
--	---

Additional information

The `target_data` directive may alternatively be specified with `target data` as the *directive-name*.

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 is a `composite directive` that provides a superset of the functionality provided by the `target_enter_data` and `target_exit_data` directives. The functionality added by the `target_data` directive is the inclusion of a `task region` for which `data-sharing attributes` may be specified. The effect of a `target_data` directive is equivalent to that of specifying three `constituent directives`, as described in the following, except expressions in all `clauses` are evaluated when the `target_data` construct is encountered.

The first `constituent directive` is a `target_enter_data` directive that is specified in the same code location as the `target_data` directive. The second `constituent directive` is a `task` directive that is specified immediately after the `target_enter_data` directive and that is associated with the `structured block` associated with the `target_data` directive. This `task` directive generates a `sharing task`. The third `constituent directive` is a `target_exit_data` directive that is specified immediately following the `structured block` that is associated with the `target_data` directive.

Since each `constituent directive` is a `task-generating construct`, the `target_data` directive generates three `tasks`. The `task` that is generated by the constituent `target_exit_data` directive

is a **dependent task** of the **task** that is generated by the constituent **task directive**, which is a **dependent task** of the **task** that is generated by the constituent **target_enter_data directive**.

When an **if clause** is present on a **target_data construct**, the effect is as if the **clause** is present only on the constituent **data-mapping constructs**.

When a **nowait clause** is present on a **target_data construct**, the effect is as if the **clause** is present on the constituent **data-mapping constructs**. In addition, the **task** associated with the **structured block** may be deferred unless otherwise specified. If the **nowait clause** is not present, all **tasks** associated with the **constituent directives** are **included tasks** and, in addition, the **task** associated with the **structured block** is a **merged task**.

If the **transparent clause** is not specified then the effect is as if a **transparent clause** is specified such that *impex-type* evaluates to **omp_impex**. If the **mergeable clause** is not specified then the effect is as if a **mergeable clause** is specified such that *can_merge* evaluates to **true**.

When a **map clause** is present on a **target_data construct**, the effect is as if the **clause** is present on the constituent **data-mapping constructs** with substituted *map-type modifiers* that are determined according to the rules of **map-type decay**.

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**.

If the **nogroup clause** is not present, the **target_data construct** executes as if the **structured block** of the constituent **task** were enclosed in a **taskgroup region**. If the **nogroup clause** is present, no implicit **taskgroup region** is created.

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 21.5](#), followed by the same **events** that are associated with a **task construct**, as described in [Section 20.1](#).

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 21.6](#).

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 21.5](#), followed by the same **tool callbacks** that are dispatched when encountering a **task construct**, as described in [Section 20.1](#).

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 21.6](#).

Cross References

- **affinity** Clause, see [Section 20.10](#)
- **allocate** Clause, see [Section 13.6](#)
- **default** Clause, see [Section 7.3.1](#)
- **depend** Clause, see [Section 23.9.5](#)
- **detach** Clause, see [Section 20.11](#)
- **device** Clause, see [Section 21.2](#)
- **firstprivate** Clause, see [Section 7.3.4](#)
- **if** Clause, see [Section 5.6](#)
- **in_reduction** Clause, see [Section 8.12](#)
- **map** Clause, see [Section 11.3](#)
- **mergeable** Clause, see [Section 20.5](#)
- **nogroup** Clause, see [Section 23.7](#)
- **nowait** Clause, see [Section 23.6](#)
- **priority** Clause, see [Section 20.9](#)
- **private** Clause, see [Section 7.3.3](#)
- **shared** Clause, see [Section 7.3.2](#)
- **target_enter_data** Construct, see [Section 21.5](#)
- **target_exit_data** Construct, see [Section 21.6](#)
- **task** Construct, see [Section 20.1](#)
- **transparent** Clause, see [Section 23.9.6](#)
- **use_device_addr** Clause, see [Section 7.3.9](#)
- **use_device_ptr** Clause, see [Section 7.3.7](#)

21.8 target Construct

Name: <code>target</code> Category: <code>executable</code>	Association: <code>block</code> Properties: <code>parallelism-generating</code> , <code>team-generating</code> , <code>thread-limiting</code> , <code>exception-aborting</code> , <code>task-generating</code> , <code>device</code> , <code>device-affecting</code> , <code>data-mapping</code> , <code>map-entering</code> , <code>map-exiting</code> , <code>context-</code> <code>matching</code> , <code>groupprivate-instantiating</code>
--	---

Clauses

`allocate`, `default`, `defaultmap`, `depend`, `device`, `device_type`,
`dyn_groupprivate`, `firstprivate`, `has_device_addr`, `if`, `in_reduction`,
`is_device_ptr`, `map`, `nowait`, `priority`, `private`, `replayable`, `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 generates a `target task` that encloses a `target region` to be executed on a `device`. If a `depend` clause is present, it is associated with the `target task`. The `device` and `device_type` clauses determine the `device` on which to execute the `target task region`. If the `nowait` clause is present, execution of the `target tasks` may be deferred. If the `nowait` clause is not present, the `target task` is an `included task`. The effect of any `map` clauses occur on entry to and exit from the generated `target region`, as specified in [Section 11.3](#).

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 11.3](#)) 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`.

If a **procedure** is explicitly or implicitly referenced in a **target construct** that does not specify a **device clause** in which the **ancestor device-modifier** appears then that **procedure** is treated as if its name had appeared in an **enter clause** on a **declare target directive**.

If a **variable** with **static storage duration** is declared in a **target construct** that does not specify a **device clause** in which the **ancestor device-modifier** appears then the named **variable** is treated as if it had appeared in an **enter clause** on a **declare target directive** if it is not a **groupprivate variable** and otherwise as if it had appeared in a **local clause** on a **declare target directive**.

If a **list item** in a **map clause** has a **base pointer** that is **predetermined firstprivate** or a **base referencing variable** for which the **referring pointer** is **predetermined firstprivate** (see Section 7.1.1), and on entry to the **target region** the **list item** is mapped, the **firstprivate** pointer is updated via corresponding pointer initialization.

Fortran

When an internal procedure is called in a **target region**, any references to **variables** that are host associated in the **procedure** have **unspecified behavior**.

Fortran

Execution Model Events

Events associated with a **target task** are the same as for the **task construct** defined in Section 20.1. **Events** associated with the **initial task** that executes the **target region** are defined in Section 20.13. The **target-submit-begin event** occurs prior to initiating creation of an **initial task** on a **target device** for a **target region**. The **target-submit-end event** occurs after initiating creation of an **initial task** on a **target device** for a **target region**. The **target-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-begin event** is a **target-task-begin event**. The **target-end event** occurs after the **target-submit-begin**, **target-submit-end** and **target-begin events** associated with the **target construct** and any **events** associated with **map clauses** on the **construct**. If the **nowait clause** is not present, the **target-end event** also occurs after all **events** associated with the **target task** and **initial task** but before the **thread** resumes execution of the **encountering task**.

Tool Callbacks

Callbacks associated with **events** for **target tasks** are the same as for the **task construct** defined in Section 20.1; **(flags & omp_target)** always evaluates to **true** in the dispatched **callback**.

A **thread** dispatches a registered **target_emi callback** with **ompt_scope_begin** as its **endpoint** argument and **ompt_target** or **ompt_target_nowait** if the **nowait clause** is present as its **kind** argument for each occurrence of a **target-begin event** in that **thread** in the context of the **target task** on the **host device**. Similarly, a **thread** dispatches a registered **target_emi callback** with **ompt_scope_end** as its **endpoint** argument and **ompt_target** or **ompt_target_nowait** if the **nowait clause** is present as its **kind** argument for each occurrence of a **target-end event** in that **thread** in the context of the **target task** on the **host device**.

A **thread** dispatches a registered **target_submit_emi callback** with **ompt_scope_begin** as its **endpoint** argument for each occurrence of a **target-submit-begin event** in that **thread**. Similarly, a

`thread` dispatches a registered `target_submit_emi` callback with `ompt_scope_end` as its *endpoint* argument for each occurrence of a *target-submit-end* event in that `thread`. These callbacks occur in the context of the `target` task.

Restrictions

Restrictions to the `target` construct are as follows:

- Device-affecting constructs, other than `target` constructs for which the `ancestor device-modifier` is specified, must not be encountered during execution of a `target` region.
- The result of an `omp_set_default_device`, `omp_get_default_device`, or `omp_get_num_devices` routine called within a `target` region is unspecified.
- The effect of an access to a `threadprivate` variable in a `target` region is unspecified.
- If a `list` item in a `map` clause is a `structure` element, any other element of that `structure` that is referenced in the `target` construct must also appear as a `list` item in a `map` clause.
- A `list` item in a `map` clause that is specified on a `target` construct must have a `base` variable or `base` pointer.
- A `list` item in a `data-sharing` attribute clause that is specified on a `target` construct must not have the same `base` variable as a `list` item in a `map` clause on the construct.
- A `variable` referenced in a `target` region but not the `target` construct that is not declared in the `target` region must appear in a `declare target` directive.
- 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.
- If a `device` clause is specified with the `ancestor device-modifier`, whether a `storage` block on the encountering device that has no corresponding `storage` on the specified device may be mapped is *implementation defined*.
- `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`. This restriction does not apply to predefined `memory allocators` in `allocate` clauses on the `target` directive.
- 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, `pointer association queries`, 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 disassociated upon entry to the **target region**, the **list item** must be disassociated upon exit from the **region**.
- If the association status of a **list item** with the **POINTER** attribute 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**.
- 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 13.6](#)
- **default** Clause, see [Section 7.3.1](#)
- **defaultmap** Clause, see [Section 11.4](#)
- **depend** Clause, see [Section 23.9.5](#)
- **device** Clause, see [Section 21.2](#)
- **device_type** Clause, see [Section 21.1](#)
- **dyn_groupprivate** Clause, see [Section 13.9](#)
- **firstprivate** Clause, see [Section 7.3.4](#)
- **has_device_addr** Clause, see [Section 7.3.8](#)
- **if** Clause, see [Section 5.6](#)
- **in_reduction** Clause, see [Section 8.12](#)
- **is_device_ptr** Clause, see [Section 7.3.6](#)
- **map** Clause, see [Section 11.3](#)
- **nowait** Clause, see [Section 23.6](#)
- **priority** Clause, see [Section 20.9](#)
- **private** Clause, see [Section 7.3.3](#)
- **replayable** Clause, see [Section 20.6](#)
- OMPT **scope_endpoint** Type, see [Section 39.27](#)
- OMPT **target** Type, see [Section 39.34](#)
- **target_data** Construct, see [Section 21.7](#)
- **target_emi** Callback, see [Section 41.8](#)
- **target_submit_emi** Callback, see [Section 41.10](#)
- **task** Construct, see [Section 20.1](#)
- OMPT **task_flag** Type, see [Section 39.37](#)
- **thread_limit** Clause, see [Section 21.3](#)
- **uses_allocators** Clause, see [Section 13.8](#)

21.9 target_update Construct

Name: <code>target_update</code> Category: <code>executable</code>	Association: <code>unassociated</code> Properties: <code>parallelism-generating</code> , <code>task-generating</code> , <code>device</code> , <code>device-affecting</code>
---	---

Clauses

`depend`, `device`, `from`, `if`, `nowait`, `priority`, `replayable`, `to`

Clause set

Properties: <code>required</code>	Members: <code>from</code> , <code>to</code>
--	---

Additional information

The `target_update` directive may alternatively be specified with `target update` as the *directive-name*.

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 20.1.

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 20.1; (*flags & ompt_task_target*) always evaluates to *true* in the dispatched callback.

A **thread** dispatches a registered **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 **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**.

A **thread** dispatches a registered **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 **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 occur in the context of the **target** task.

Cross References

- **depend** Clause, see Section 23.9.5
- **device** Clause, see Section 21.2
- **from** Clause, see Section 12.2
- **if** Clause, see Section 5.6
- **nowait** Clause, see Section 23.6
- **priority** Clause, see Section 20.9
- **replayable** Clause, see Section 20.6
- OMPT **scope_endpoint** Type, see Section 39.27
- OMPT **target** Type, see Section 39.34
- **target_data_op_emi** Callback, see Section 41.7
- **target_emi** Callback, see Section 41.8
- **task** Construct, see Section 20.1
- OMPT **task_flag** Type, see Section 39.37
- **to** Clause, see Section 12.1

22 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 [function variant](#) and the [interoperability routines](#).

Cross References

- Interoperability Routines, see [Chapter 32](#)

22.1 interop Construct

Name: interop Category: executable	Association: unassociated 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](#).

The [interop-type](#) set for an **init** clause is the set of specified [interop-type](#) modifiers. For any other [action-clause](#) and the [interoperability object](#) that its argument specifies, the [interop-type](#) set is the set of [modifiers](#) that were specified by the **init** clause that initialized that [interoperability object](#).

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](#). If the [interop-type](#) set does not include **targetsync**, the **nowait** clause has no effect. 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**, 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 between the **encountering thread** and either **foreign tasks** or OpenMP **tasks** by the **interop** construct.

Restrictions

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

- A **depend** clause must only appear on the **directive** if the *interop-type* includes **targetsync**.
- An **interoperability** object must not be specified in more than one *action-clause* that appears on the **interop** construct.

Cross References

- **depend** Clause, see [Section 23.9.5](#)
- **destroy** Clause, see [Section 5.8](#)
- **device** Clause, see [Section 21.2](#)
- **init** Clause, see [Section 5.7](#)
- **nowait** Clause, see [Section 23.6](#)
- **use** Clause, see [Section 22.1.2](#)

22.1.1 OpenMP Foreign Runtime Identifiers

Allowed values for **foreign runtime identifiers** include the names (as **string literals**) and integer values that the **OpenMP Additional Definitions document** specifies and the corresponding **omp_ifr_name** values of the **interop_fr** OpenMP type. **Implementation defined** values for **foreign runtime identifiers** may also be supported.

22.1.2 use Clause

Name: use		Properties: <i>default</i>
Arguments		
Name	Type	Properties
<i>interop-var</i>	variable of interop OpenMP type	<i>default</i>

Modifiers

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

Directives

interop

Semantics

The *use clause* specifies the *interop-var* that is used for the effects of the *directive* on which the *clause* appears. However, *interop-var* is not initialized, destroyed or otherwise modified. The *interop-type* set is inferred based on the *interop-type modifiers* used to initialize *interop-var*.

Restrictions

- The state of *interop-var* must be *initialized*.

Cross References

- *interop* Construct, see [Section 22.1](#)

22.1.3 *prefer-type* Modifier

Modifiers

Name	Modifies	Type	Properties
<i>prefer-type</i>	<i>init-var</i>	Complex, name: <i>prefer_type</i> Arguments: <i>prefer-type-specification</i> list of preference specification list item type (<i>default</i>)	<i>complex, unique</i>

Clauses

init

Semantics

The *prefer-type modifier* specifies a set of preferences to be used to initialize an *interoperability object*. Each *preference specification list item* specified in the *prefer-type-specification* argument is a *preference specification* that has the following syntax:

```
preference-specification :  
    {preference-selector[, preference-selector[, ...]]}  
    foreign-runtime-identifier  
  
preference-selector :  
    fr (foreign-runtime-identifier)
```

```

attr(preference-property-extension[, preference-property-extension[, ...]])

preference-property-extension:
    ext-string-literal

```

Where *foreign-runtime-identifier* is a [foreign runtime identifier](#) and an [implementation defined ext-string-literal](#) is a [string literal](#) that must start with the `omp_x_` prefix and must not include any commas (i.e., instances of the character `,`).

The `fr` *preference-selector* specifies a [foreign runtime environment](#) identified by its [foreign runtime identifier](#). The `attr` *preference-selector* specifies a preference for the attributes specified as its arguments.

If a *preference-specification* is a *foreign-runtime-identifier*, it is equivalent to specifying a *preference-specification* that uses the `fr` *preference-selector* and the [foreign runtime identifier](#) as its argument.

The [interoperability object](#) specified by the *init-var* argument of the `init` clause is initialized based on the first supported [preference specification](#), if any, in left-to-right order. If the implementation does not support any of the specified [preference specifications](#), *init-var* is initialized based on an [implementation defined preference specification](#).

Restrictions

Restrictions to the [prefer-type modifier](#) are as follows:

- At most one `fr` *preference-selector* may be specified for each *preference-specification*.

Cross References

- `init` Clause, see [Section 5.7](#)

23 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 1.3.4](#) and [Section 1.3.6](#) describe synchronization through the use of synchronizing [flushes](#) and [atomic operations](#). [Section 23.8.7](#) defines the behavior of synchronizing [flushes](#) that are implied at various other locations in an [OpenMP program](#).

23.1 hint Clause

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

Arguments

Name	Type	Properties
<i>hint-expr</i>	expression of <code>sync_hint</code> type	constant

Modifiers

Name	Modifies	Type	Properties
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: <i>directive-name</i> (a 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 `omp_sync_hint_none` had been specified as *hint-expr*.

Restrictions

- *hint-expr* must evaluate to a valid [synchronization hint](#).

Cross References

- **atomic** Construct, see [Section 23.8.5](#)
- **critical** Construct, see [Section 23.2](#)
- OpenMP **sync_hint** Type, see [Section 26.9.5](#)

23.2 critical Construct

Name: critical Category: executable	Association: block Properties: mutual-exclusion , thread-limiting , thread-exclusive
---	---

Arguments

critical (<i>name</i>)		
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 *mutex_acquire* callback for each occurrence of a *critical-acquiring event* in that *thread*. A *thread* dispatches a registered *mutex_acquired* callback for each occurrence of a *critical-acquired event* in that *thread*. A *thread* dispatches a registered *mutex_released* callback for each occurrence of a *critical-released event* in that *thread*. These *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.
- A *critical* region must not be nested (closely or otherwise) inside a *critical* region with the same *name*. This restriction is not sufficient to prevent deadlock.

Fortran

- If a *name* is specified on a *critical directive* and a paired *end directive* is specified, the same *name* must also be specified on the *end directive*.
- If no *name* appears on the *critical directive* and a paired *end directive* is specified, no *name* can appear on the *end directive*.

Fortran

Cross References

- *hint* Clause, see [Section 23.1](#)
- OMPT *mutex* Type, see [Section 39.20](#)
- *mutex_acquire* Callback, see [Section 40.7.8](#)
- *mutex_acquired* Callback, see [Section 40.7.12](#)
- *mutex_released* Callback, see [Section 40.7.13](#)
- OpenMP *sync_hint* Type, see [Section 26.9.5](#)

23.3 Barriers

23.3.1 barrier Construct

Name: <code>barrier</code> Category: <code>executable</code>	Association: <code>unassociated</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 a waiting interval in a `barrier region`. The `explicit-barrier-wait-end event` occurs when a `task` ends a waiting interval and resumes execution in a `barrier region`. The `explicit-barrier-end event` occurs in each `thread` that encounters the `barrier` construct after the `barrier` synchronization on exit from the `barrier region`. A `cancellation event` occurs if `cancellation` is activated at an implicit `cancellation point` in a `barrier region`.

Tool Callbacks

A `thread` dispatches a registered `sync_region` callback with `ompt_sync_region_barrier_explicit` as its `kind` argument and `ompt_scope_begin` as its `endpoint` argument for each occurrence of an `explicit-barrier-begin event`. Similarly, a `thread` dispatches a registered `sync_region` callback with `ompt_sync_region_barrier_explicit` as its `kind` argument and `ompt_scope_end` as its `endpoint` argument for each occurrence of an `explicit-barrier-end event`. These `callbacks` occur in the context of the `task` that encountered the `barrier` construct.

A `thread` dispatches a registered `sync_region_wait` callback with `ompt_sync_region_barrier_explicit` as its `kind` argument and `ompt_scope_begin` as its `endpoint` argument for each occurrence of an `explicit-barrier-wait-begin event`. Similarly, a `thread` dispatches a registered `sync_region_wait` callback with `ompt_sync_region_barrier_explicit` as its `kind` argument and `ompt_scope_end` as its `endpoint` argument for each occurrence of an `explicit-barrier-wait-end event`. These `callbacks` occur in the context of the `task` that encountered the `barrier` construct.

A `thread` dispatches a registered `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*.

Restrictions

Restrictions to the *barrier construct* are as follows:

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

Cross References

- `cancel` Callback, see [Section 40.6](#)
- OMPT `cancel_flag` Type, see [Section 39.7](#)
- OMPT `scope_endpoint` Type, see [Section 39.27](#)
- `sync_region` Callback, see [Section 40.7.4](#)
- OMPT `sync_region` Type, see [Section 39.33](#)
- `sync_region_wait` Callback, see [Section 40.7.5](#)

23.3.2 Implicit Barriers

This section describes the *OMPT events* and *tool callbacks* associated with *implicit barriers*, which occur at the end of various *regions* as defined in the description of the *constructs* to which they correspond. *Implicit barriers* are *task scheduling points*. For a description of *task scheduling points*, associated *events*, and *tool callbacks*, see [Section 20.14](#).

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 a waiting interval in an *implicit barrier region*. The *implicit-barrier-wait-end event* occurs when a *task* ends a waiting interval 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 *sync_region* callback for each *implicit-barrier-begin* and *implicit-barrier-end event*. Similarly, a *thread* dispatches a registered *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*.

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 `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`.

Restrictions

Restrictions to `implicit barriers` are as follows:

- If a `thread` is in the `ompt_state_wait_barrier_implicit_parallel` state, a call to `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 `get_parallel_info` when a `thread` is in the `ompt_state_wait_barrier_implicit_parallel` state results in `unspecified behavior`.

Cross References

- `cancel` Callback, see [Section 40.6](#)
- OMPT `cancel_flag` Type, see [Section 39.7](#)
- `get_parallel_info` Entry Point, see [Section 42.14](#)
- OMPT `scope_endpoint` Type, see [Section 39.27](#)
- OMPT `state` Type, see [Section 39.31](#)
- `sync_region` Callback, see [Section 40.7.4](#)
- OMPT `sync_region` Type, see [Section 39.33](#)
- `sync_region_wait` Callback, see [Section 40.7.5](#)

23.3.3 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` as the *kind* argument when they are dispatched.

23.4 taskgroup Construct

Name: taskgroup Category: executable	Association: block Properties: cancellable
---	---

Clauses

allocate, **task_reduction**

Binding

The **binding task set** of a **taskgroup region** is all **tasks** of the **current team** that are generated in the **region**. A **taskgroup region** binds to the innermost enclosing **parallel region**.

Semantics

The **taskgroup construct** specifies a wait on completion of the **taskgroup set** associated with the **taskgroup region**. When a **thread** encounters a **taskgroup construct**, it starts executing the **region**. An implicit **task scheduling point** occurs at the end of the **taskgroup region**. The **current task** is suspended at the **task scheduling point** until all **tasks** in the **taskgroup set** complete execution.

Execution Model Events

The *taskgroup-begin event* occurs in each **thread** that encounters the **taskgroup construct** on entry to the **taskgroup region**. The *taskgroup-wait-begin event* occurs when a **task** begins a waiting interval in a **taskgroup region**. The *taskgroup-wait-end event* occurs when a **task** ends a waiting interval 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 **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 **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**.

A **thread** dispatches a registered **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 **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**.

Cross References

- `allocate` Clause, see [Section 13.6](#)
- Task Scheduling, see [Section 20.14](#)
- OMPT `scope_endpoint` Type, see [Section 39.27](#)
- `sync_region` Callback, see [Section 40.7.4](#)
- OMPT `sync_region` Type, see [Section 39.33](#)
- `sync_region_wait` Callback, see [Section 40.7.5](#)
- `task_reduction` Clause, see [Section 8.11](#)

23.5 taskwait Construct

Name: <code>taskwait</code> Category: executable	Association: unassociated Properties: default
---	--

Clauses

[depend](#), [nowait](#), [replayable](#)

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 a waiting interval in a *region* that corresponds to a *taskwait construct* with no *depend clause*. The *taskwait-wait-end event* occurs when a *task* ends a waiting interval 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 the *dependent task* and before any subsequently *generated task* that depends on the *dependent task* starts its execution.

Tool Callbacks

A *thread* dispatches a registered *sync_region callback* with *ompt_sync_region_taskwait* as its *kind* argument and *ompt_scope_begin* as its *endpoint* argument for each occurrence of a *taskwait-begin event* in the *task* that encounters the *taskwait construct*. Similarly, a *thread* dispatches a registered *sync_region callback* with *ompt_sync_region_taskwait* as its *kind* argument and *ompt_scope_end* as its *endpoint* argument for each occurrence of a *taskwait-end event* in the *task* that encounters the *taskwait construct*. These *callbacks* occur in the *task* that encounters the *taskwait construct*.

A *thread* dispatches a registered *sync_region_wait callback* with *ompt_sync_region_taskwait* as its *kind* argument and *ompt_scope_begin* as its *endpoint* argument for each occurrence of a *taskwait-wait-begin event*. Similarly, a *thread* dispatches a registered *sync_region_wait callback* with *ompt_sync_region_taskwait* as its *kind* argument and *ompt_scope_end* as its *endpoint* argument for each occurrence of a *taskwait-wait-end event*. These *callbacks* occur in the context of the *task* that encounters the *taskwait construct*.

A *thread* dispatches a registered *task_create callback* for each occurrence of a *taskwait-init event* in the context of the *encountering task*. In the dispatched *callback*, (*flags & ompt_task_taskwait*) always evaluates to *true*. If the *nowait clause* is not present, (*flags & ompt_task_underrred*) also evaluates to *true*.

A *thread* dispatches a registered *task_schedule callback* for each occurrence of a *taskwait-complete event*. This *callback* has *ompt_taskwait_complete* as its *prior_task_status* argument.

Restrictions

Restrictions to the *taskwait construct* are as follows:

- The *mutexinoutset task-dependence-type* may not appear in a *depend clause* on a *taskwait construct*.

- If the *task-dependence-type* of a **depend** clause is **depobj** then the **depend** objects may not represent dependences of the **mutexinoutset** dependence type.
- The **nowait** clause may only appear on a **taskwait** directive if the **depend** clause is present.
- The **replayable** clause may only appear on a **taskwait** directive if the **depend** clause is present.

Cross References

- **depend** Clause, see [Section 23.9.5](#)
- **nowait** Clause, see [Section 23.6](#)
- **replayable** Clause, see [Section 20.6](#)
- OMPT **scope_endpoint** Type, see [Section 39.27](#)
- **sync_region** Callback, see [Section 40.7.4](#)
- OMPT **sync_region** Type, see [Section 39.33](#)
- **sync_region_wait** Callback, see [Section 40.7.5](#)
- **task** Construct, see [Section 20.1](#)
- OMPT **task_flag** Type, see [Section 39.37](#)
- **task_schedule** Callback, see [Section 40.5.2](#)
- OMPT **task_status** Type, see [Section 39.38](#)

23.6 nowait Clause

Name: nowait	Properties: outermost-leaf, unique, end-clause
----------------------------	---

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: <i>directive-name</i> (a directive name)	unique

Directives

`dispatch`, `do`, `for`, `interop`, `scope`, `sections`, `single`, `target`, `target_data`, `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 a **semantic 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 **directives** that generate a **semantic 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** Construct, see [Section 15.7](#)
- **do** Construct, see [Section 19.6.2](#)
- **for** Construct, see [Section 19.6.1](#)
- **interop** Construct, see [Section 22.1](#)
- **scope** Construct, see [Section 19.2](#)
- **sections** Construct, see [Section 19.3](#)
- **single** Construct, see [Section 19.1](#)
- **target** Construct, see [Section 21.8](#)
- **target_data** Construct, see [Section 21.7](#)

- **target_enter_data** Construct, see [Section 21.5](#)
- **target_exit_data** Construct, see [Section 21.6](#)
- **target_update** Construct, see [Section 21.9](#)
- **taskwait** Construct, see [Section 23.5](#)
- **workshare** Construct, see [Section 19.4](#)

23.7 nogroup Clause

Name: <code>nogroup</code>	Properties: <code>outermost-leaf</code> , <code>unique</code>
-----------------------------------	--

Arguments

Name	Type	Properties
<code>do_not_synchronize</code>	expression of OpenMP logical type	<code>optional</code>

Modifiers

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

Directives

`target_data`, `taskgraph`, `taskloop`

Semantics

If `do_not_synchronize` evaluates to `true`, the `nogroup clause` overrides any implicit `taskgroup` that would otherwise enclose the `construct`. If `do_not_synchronize` evaluates to `false`, the effect is as if the `nogroup clause` is not specified on the `directive`. If `do_not_synchronize` is not specified, the effect is as if `do_not_synchronize` evaluates to `true`.

Cross References

- **target_data** Construct, see [Section 21.7](#)
- **taskgraph** Construct, see [Section 20.3](#)
- **taskloop** Construct, see [Section 20.2](#)

23.8 OpenMP Memory Ordering

This sections describes [constructs](#) and [clauses](#) that support ordering of [memory](#) operations.

23.8.1 *memory-order* Clauses

Clause groups

Properties: exclusive , unique	Members: Clauses acq_rel , acquire , relaxed , release , seq_cst
---	---

Directives

[atomic](#), [flush](#)

Semantics

The [memory-order clause group](#) defines a set of [clauses](#) that indicate the [memory](#) ordering requirements for the visibility of the effects of the [constructs](#) on which they may be specified.

Cross References

- **atomic** Construct, see [Section 23.8.5](#)
- **flush** Construct, see [Section 23.8.6](#)
- OpenMP Memory Consistency, see [Section 1.3.6](#)

23.8.1.1 [acq_rel](#) Clause

Name: acq_rel	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: <i>directive-name</i> (a directive name)	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** Construct, see [Section 23.8.5](#)
- **flush** Construct, see [Section 23.8.6](#)
- OpenMP Memory Consistency, see [Section 1.3.6](#)

23.8.1.2 acquire Clause

Name: acquire	Properties: unique
-----------------------------	---

Arguments

Name	Type	Properties
<i>use_semantics</i>	expression of OpenMP logical type	constant , optional

Modifiers

Name	Modifies	Type	Properties
directive-name-modifier	<i>all arguments</i>	Keyword: <i>directive-name</i> (a directive name)	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** Construct, see [Section 23.8.5](#)
- **flush** Construct, see [Section 23.8.6](#)
- OpenMP Memory Consistency, see [Section 1.3.6](#)

23.8.1.3 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
directive-name-modifier	<i>all arguments</i>	Keyword: <i>directive-name</i> (a 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** Construct, see [Section 23.8.5](#)
- **flush** Construct, see [Section 23.8.6](#)
- OpenMP Memory Consistency, see [Section 1.3.6](#)

23.8.1.4 release Clause

Name: release	Properties: unique
----------------------	------------------------------------

Arguments

Name	Type	Properties
<i>use_semantics</i>	expression of OpenMP logical type	constant, optional

Modifiers

Name	Modifies	Type	Properties
directive-name-modifier	<i>all arguments</i>	Keyword: <i>directive-name</i> (a 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** Construct, see [Section 23.8.5](#)
- **flush** Construct, see [Section 23.8.6](#)
- OpenMP Memory Consistency, see [Section 1.3.6](#)

23.8.1.5 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
directive-name-modifier	<i>all arguments</i>	Keyword: <i>directive-name</i> (a 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** Construct, see [Section 23.8.5](#)
- **flush** Construct, see [Section 23.8.6](#)
- OpenMP Memory Consistency, see [Section 1.3.6](#)

23.8.2 atomic Clauses

Clause groups

Properties: exclusive , unique	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](#) Construct, see [Section 23.8.5](#)

23.8.2.1 read Clause

Name: read	Properties: innermost-leaf , unique
-----------------------------------	--

Arguments

Name	Type	Properties
use_semantics	expression of OpenMP logical type	constant , optional

Modifiers

Name	Modifies	Type	Properties
directive-name-modifier	all arguments	Keyword: directive-name (a 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](#) Construct, see [Section 23.8.5](#)

23.8.2.2 update Clause

Name: <code>update</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: <i>directive-name</i> (a 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 clause** is not specified. If *use_semantics* is not specified, the effect is as if *use_semantics* evaluates to *true*.

Cross References

- **atomic** Construct, see [Section 23.8.5](#)

23.8.2.3 write Clause

Name: <code>write</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: <i>directive-name</i> (a directive name)	unique

Directives

atomic

1 **Semantics**

2 If *use_semantics* evaluates to *true*, the **write clause** specifies that the **atomic construct** has
3 **atomic write** semantics, which write the value of the **shared variable** atomically. If *use_semantics*
4 evaluates to *false*, the effect is as if the **write clause** is not specified. If *use_semantics* is not
5 specified, the effect is as if *use_semantics* evaluates to *true*.

6 **Cross References**

- 7
 - **atomic** Construct, see [Section 23.8.5](#)

8 **23.8.3 extended-atomic Clauses**

9 **Clause groups**

10

Properties: <i>unique</i>	Members: Clauses <i>capture, compare, fail, weak</i>
----------------------------------	--

11 **Directives**

12 ***atomic***

13 **Semantics**

14 The *extended-atomic clause group* defines a set of **clauses** that extend the atomicity semantics
15 specified by members of the *atomic clause group*.

16 **Restrictions**

17 Restrictions to the *extended-atomic clause group* are as follows:

- 18
 - The **compare clause** may not be specified such that *use_semantics* evaluates to *false* if the
19 **weak clause** is specified such that *use_semantics* evaluates to *true*.

20 **Cross References**

- 21
 - **atomic** Construct, see [Section 23.8.5](#)
 - *atomic* Clauses, see [Section 23.8.2](#)

23 **23.8.3.1 capture Clause**

24

Name: <i>capture</i>	Properties: <i>innermost-leaf, unique</i>
-----------------------------	--

25 **Arguments**

26

Name	Type	Properties
<i>use_semantics</i>	expression of OpenMP logical type	<i>constant, optional</i>

Modifiers

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

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** Construct, see [Section 23.8.5](#)

23.8.3.2 compare Clause

Name: compare	Properties: innermost-leaf, <i>unique</i>
-----------------------------	--

Arguments

Name	Type	Properties
<i>use_semantics</i>	expression of OpenMP logical type	<i>constant, optional</i>

Modifiers

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

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** Construct, see [Section 23.8.5](#)

23.8.3.3 fail Clause

Name: fail	Properties: innermost-leaf, unique
-------------------	------------------------------------

Arguments

Name	Type	Properties
<i>memorder</i>	Keyword: acquire , relaxed , seq_cst	<i>default</i>

Modifiers

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

- **atomic** Construct, see [Section 23.8.5](#)
- *memory-order* Clauses, see [Section 23.8.1](#)

23.8.3.4 weak Clause

Name: weak	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: <i>directive-name</i> (a <i>directive name</i>)	<i>unique</i>

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** Construct, see [Section 23.8.5](#)

23.8.4 memscope Clause

Name: memscope	Properties: <i>unique</i>
------------------------------	----------------------------------

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: <i>directive-name</i> (a <i>directive name</i>)	<i>unique</i>

Directives

atomic, **flush**

1 **Semantics**

2 The **memscope** clause determines the **binding thread set** of the **region** that corresponds to the

3 **construct** on which it is specified.

4 If the *scope-specifier* is **device**, the **binding thread set** consists of all **threads** on the **device**. If the

5 *scope-specifier* is **cgroup**, the **binding thread set** consists of all **threads** that are executing **tasks** in

6 the **contention group**. If the *scope-specifier* is **all**, the **binding thread set** consists of **all threads** on

7 all **devices**.

8 Unless otherwise stated, the **thread-set** of any **flushes** that are performed in an **atomic** or **flush**

9 **region** is the same as the **binding thread set** of the **region**, as determined by the **memscope** clause.

10 **Restrictions**

11 The restrictions for the **memscope** clause are as follows:

- 12 • The **binding thread set** defined by the *scope-specifier* of the **memscope** clause on an
- 13 **atomic** **construct** must be a subset of the **atomic scope** of the atomically accessed **memory**.
- 14 • The **binding thread set** defined by the *scope-specifier* of the **memscope** clause on an
- 15 **atomic** **construct** must be a subset of all **threads** that are executing **tasks** in the **contention**
- 16 **group** if the size of the atomically accessed **storage location** is not 8, 16, 32, or 64 bits.

17 **Cross References**

- 18 • **atomic** Construct, see [Section 23.8.5](#)
- 19 • **flush** Construct, see [Section 23.8.6](#)

20 **23.8.5 atomic Construct**

21 Name: atomic	Association: block : atomic
Category: executable	Properties: mutual-exclusion , order- concurrent-nestable , simdizable

22 **Clause groups**

23 ***atomic***, ***extended-atomic***, ***memory-order***

24 **Clauses**

25 ***hint***, ***memscope***

26 **Binding**

27 The **memscope** clause determines the **binding thread set** for an **atomic** **region**. If the **memscope**

28 **clause** is not present, the behavior is as if the **memscope** clause appeared on the **construct** with the

29 **device** *scope-specifier*.

Semantics

This section refers to the symbols defined for **atomic structured blocks**. 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. The read and write of the **storage location** designated by x are performed as a single atomic operation. The evaluation of $expr$ or $expr$ -list need not be atomic with respect to the read and 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 **atomic structured blocks** of the **atomic** construct. The read and write of the **storage location** designated by x are performed as a single atomic operation. 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 and 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 **atomic 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 . The read and write of the **storage location** designated by x are performed as a single atomic operation. 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 and 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++

If a *memory-order clause* is present, or implicitly provided by a **requires directive**, it specifies the effective memory ordering. Otherwise the effect is as if the *relaxed memory-order clause* is specified.

The **atomic construct** may be used to enforce memory consistency between *threads*, based on the guarantees provided by Section 1.3.6. A *strong flush* on the *storage location* designated by x is performed on entry to and exit from the *atomic operation*, ensuring that the set of all *atomic operations* applied to the same *storage location* in a race-free program has a total completion order. If the **write** or **update clause** is specified, the *atomic operation* is not an *atomic conditional update* for which the comparison fails, and the effective memory ordering is **release**, **acq_rel**, or **seq_cst**, the *strong flush* on entry to the *atomic operation* is also a *release flush*. If the **read** or **update clause** is specified and the effective memory ordering is **acquire**, **acq_rel**, or **seq_cst** then the *strong flush* on exit from the *atomic operation* is also an *acquire flush*. Therefore, if the effective memory ordering is not **relaxed**, *release flushes* and/or *acquire flushes* are implied and permit synchronization between the *threads* without the use of explicit **flush directives**.

For all forms of the **atomic construct**, any combination of two or more of these **atomic constructs** enforces mutually exclusive access to the *storage locations* designated by x among *threads* in the *binding thread set*. To avoid *data races*, all accesses of the *storage locations* designated by x that could potentially occur in parallel must be protected with an **atomic construct**.

atomic regions do not guarantee exclusive access with respect to any accesses outside of **atomic regions** to the same *storage location* x even if those accesses occur during a **critical** or **ordered region**, while a **lock** is owned by the executing *task*, or during the execution of a **reduction clause**.

However, other OpenMP synchronization can ensure the desired exclusive access. For example, a **barrier** that follows a series of *atomic updates* to x guarantees that subsequent accesses do not form a *data race* with the atomic accesses.

A **compliant implementation** may enforce exclusive access between **atomic regions** that update different *storage locations*. The circumstances under which this occurs are **implementation defined**.

If the *storage location* designated by x is not size-aligned (that is, if the byte alignment of x is not a multiple of the size of x), then the behavior of the **atomic region** is **implementation defined**.

Execution Model Events

The *atomic-acquiring event* occurs in the *thread* that encounters the **atomic construct** on entry to the **atomic region** before initiating synchronization for the *region*. The *atomic-acquired event* occurs in the *thread* that encounters the **atomic construct** after it enters the *region*, but before it executes the *atomic structured block* of the **atomic region**. The *atomic-released event* occurs in the *thread* that encounters the **atomic construct** after it completes any synchronization on exit from the **atomic region**.

Tool Callbacks

A **thread** dispatches a registered **mutex_acquire** callback for each occurrence of an *atomic-acquiring event* in that **thread**. A **thread** dispatches a registered **mutex_acquired** callback for each occurrence of an *atomic-acquired event* in that **thread**. A **thread** dispatches a registered **mutex_released** callback with **ompt_mutex_atomic** as the *kind* argument if practical, although a less specific *kind* may be used, for each occurrence of an *atomic-released event* in that **thread**. These **callbacks** occurs in the **task** that encounters the **atomic** construct.

Restrictions

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

- **Constructs** may not be encountered during execution of an **atomic** region.
- If a **capture** or **compare** clause is specified, the *atomic* clause must be **update**.
- If an **update** clause is specified, an **update structured block** or **capture structured block** must be associated with the **construct**.
- If a **capture** clause is specified but the **compare** clause is not specified, an **update-capture structured block** must be associated with the **construct**.
- If both **capture** and **compare** clauses are specified, a **conditional-update-capture structured block** must be associated with the **construct**.
- If a **compare** clause is specified but the **capture** clause is not specified, a **conditional-update structured block** must be associated with the **construct**.
- If a **write** clause is specified, a **write structured block** must be associated with the **construct**.
- If a **read** clause is specified, a **read structured block** must be associated with the **construct**.
- If the *atomic* clause is **read** then the *memory-order* clause must not be **release**.
- If the *atomic* clause is **write** then the *memory-order* clause must not be **acquire**.
- The **weak** clause may only appear if the resulting **atomic operation** is an **atomic conditional update** for which the comparison tests for equality.

▼ C / C++ ▼

- All atomic accesses to the **storage locations** designated by *x* throughout the **OpenMP program** are required to have a compatible type.
- The **fail** clause may only appear if the resulting **atomic operation** is an **atomic conditional 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

- **barrier** Construct, see [Section 23.3.1](#)
- **critical** Construct, see [Section 23.2](#)
- **flush** Construct, see [Section 23.8.6](#)
- Lock Routines, see [Chapter 34](#)
- OpenMP Atomic Structured Blocks, see [Section 6.3.3](#)
- **hint** Clause, see [Section 23.1](#)
- **memscope** Clause, see [Section 23.8.4](#)
- OMPT **mutex** Type, see [Section 39.20](#)
- **mutex_acquire** Callback, see [Section 40.7.8](#)
- **mutex_acquired** Callback, see [Section 40.7.12](#)
- **mutex_released** Callback, see [Section 40.7.13](#)
- **ordered** Construct, see [Section 23.10](#)
- **requires** Directive, see [Section 16.5](#)

23.8.6 flush Construct

Name: flush	Association: unassociated
Category: executable	Properties: default

Arguments

flush (*list*)

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

Clause groups

[memory-order](#)

Clauses

`memscope`

Binding

The `memscope` clause determines the binding thread set for a `flush` 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 `flush` construct executes the flush OpenMP operation. This operation makes the temporary view of the memory of a thread consistent with the memory and enforces an order on the memory operations of the variables explicitly specified or implied. Execution of a `flush` region affects the memory and it affects the temporary view of the memory of the encountering thread. It does not affect the temporary view of other threads. Other threads in the thread-set must themselves execute a flush in order to be guaranteed to observe the effects of the flush of the encountering thread. See the memory model description in Section 1.3 and the `memscope` clause description in Section 23.8.4 for more details on thread-sets.

If neither a *memory-order* clause nor a *list* argument appears on a `flush` construct then the behavior is as if the *memory-order* clause is `seq_cst`.

A `flush` construct with the `seq_cst` clause, executed on a given thread, operates as if all storage locations that are accessible to the thread are flushed by a strong flush; that is, the flush has the strong flush property. A `flush` construct with a *list* applies a strong flush to the items in the *list*, and the flush does not complete until the operation is complete for all specified list items. An implementation may implement a `flush` construct with a *list* by ignoring the *list* and treating it the same as a `flush` construct with the `seq_cst` clause.

If no list items are specified, the flush operation has the release flush property and/or the acquire flush property:

- If the *memory-order* clause is `seq_cst` or `acq_rel`, the flush is both a release flush and an acquire flush.
- If the *memory-order* clause is `release`, the flush is a release flush.
- If the *memory-order* clause is `acquire`, the flush is an acquire flush.

C / C++

If a pointer is present in the *list*, the pointer itself is flushed, not the storage locations to which the pointer refers.

A `flush` construct without a *list* corresponds to a call to `atomic_thread_fence`, where the argument is given by the identifier that results from prefixing `memory_order_` to the *memory-order* clause name.

For a `flush` construct without a *list*, the generated flush region implicitly performs the corresponding call to `atomic_thread_fence`. The behavior of an explicit call to

atomic_thread_fence that occurs in an **OpenMP program** and does not have the argument **memory_order_consume** is as if the call is replaced by its corresponding **flush construct**.

▲ C / C++ ▲

▼ Fortran ▼

If the **list item** or a subobject of the **list item** has the **POINTER** attribute, the allocation or association status of the **POINTER** item is flushed, but the pointer target is not. If the **list item** is of type **C_PTR**, the **variable** is flushed, but the **storage location** that corresponds to that address is not flushed. If the **list item** or the subobject of the **list item** has the **ALLOCATABLE** attribute and has an allocation status of allocated, the allocated **variable** is flushed; otherwise the allocation status is flushed.

▲ Fortran ▲

Execution Model Events

The **flush event** occurs in a **thread** that encounters the **flush construct**.

Tool Callbacks

A **thread** dispatches a registered **flush callback** for each occurrence of a **flush event** in that **thread**.

Restrictions

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

- If a **memory-order clause** is specified, the **list** argument must not be specified.
- The **memory-order clause** must not be **relaxed**.

Cross References

- **flush** Callback, see [Section 40.7.15](#)
- **memscope** Clause, see [Section 23.8.4](#)

23.8.7 Implicit Flushes

Flushes implied when executing an **atomic region** are described in [Section 23.8.5](#).

A **flush region** that corresponds to a **flush directive** with the **release clause** present is implied at the following locations:

- During a **barrier region**;
- At entry to a **parallel region**;
- At entry to a **teams region**;
- At exit from a **critical region**;
- During an **omp_unset_lock region**;
- During an **omp_unset_nest_lock region**;

- During an **omp_fulfill_event** region;
- Immediately before every **task scheduling point**;
- At exit from the **task region** of each **implicit task**;
- At exit from an **ordered region**, if a **threads** clause or a **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**.

The synchronization behavior of **implicit flushes** is as follows:

- When a **thread** executes an **atomic region** for which the corresponding **construct** has the **release**, **acq_rel**, or **seq_cst** clause and specifies an **atomic operation** that starts a given **release sequence**, the **release flush** that is performed on entry to the **atomic operation** synchronizes with an **acquire flush** that is performed by a different **thread** and has an associated **atomic operation** that reads a value written by a modification in the **release sequence**.
- When a **thread** executes an **atomic region** for which the corresponding **construct** has the **acquire**, **acq_rel**, or **seq_cst** clause and specifies an **atomic operation** that reads a value written by a given modification, a **release flush** that is performed by a different **thread** and has an associated **release sequence** that contains that modification synchronizes with the **acquire flush** that is performed on exit from the **atomic operation**.
- When a **thread** executes a **critical region** that has a given *name*, the behavior is as if the **release flush** performed on exit from the **region** synchronizes with the **acquire flush** performed on entry to the next **critical region** with the same *name* that is performed by a different **thread**, if it exists.
- When a **team** executes a **barrier region**, the behavior is as if the **release flush** performed by each **thread** within the **region**, and the **release flush** performed by any other **thread** upon completion of the associated **structured block** of an **explicit task** bound to the binding **parallel region** of the **region** or upon fulfilling the *allow-completion event* for any such **task** that is **detachable**, synchronizes with the **acquire flush** performed by all other **threads** within the **region**.
- When a **thread** executes a **taskwait region** that does not result in the creation of a **dependent task** and the **task** that encounters the corresponding **taskwait construct** has at least one **child task**, the behavior is as if each **thread** that executes a **child task** that is generated before the **taskwait region** performs a **release flush** upon completion of the associated **structured block** of the **child task** that synchronizes with an **acquire flush** performed in the **taskwait region**. If the **child task** is a **detachable task**, the **thread** that fulfills its *allow-completion event* performs a **release flush** upon fulfilling the *event* that synchronizes with the **acquire flush** performed in the **taskwait region**.
- When a **thread** executes a **taskgroup region**, the behavior is as if each **thread** that executes a remaining **descendent task** performs a **release flush** upon completion of the associated **structured block** of the **descendent task** that synchronizes with an **acquire flush** performed on

exit from the **taskgroup** region. If the **descendent task** is a **detachable task**, the **thread** that fulfills its *allow-completion event* performs a **release flush** upon fulfilling the **event** that synchronizes with the **acquire flush** performed in the **taskgroup** region.

- When a **thread** executes an **ordered region** that does not arise from a stand-alone **ordered directive**, the behavior is as if the **release flush** performed on exit from the **region** synchronizes with the **acquire flush** performed on entry to an **ordered region** encountered in the next **collapsed iteration** to be executed by a different **thread**, if it exists.
- When a **thread** executes an **ordered region** that arises from a stand-alone **ordered directive**, the behavior is as if the **release flush** performed in the **ordered region** from a given source **doacross iteration** synchronizes with the **acquire flush** performed in all **ordered regions** executed by a different **thread** that are waiting for dependences on that **doacross iteration** to be satisfied.
- When a **team** begins execution of a **parallel region**, the behavior is as if the **release flush** performed by the **primary thread** on entry to the **parallel region** synchronizes with the **acquire flush** performed on entry to each **implicit task** that is assigned to a different **thread**.
- When an **initial thread** begins execution of a **target region** that is generated by a different **thread** from a **target task**, the behavior is as if the **release flush** performed by the generating **thread** in the **target task** synchronizes with the **acquire flush** performed by the **initial thread** on entry to its **initial task region**.
- When an **initial thread** completes execution of a **target region** that is generated by a different **thread** from a **target task**, the behavior is as if the **release flush** performed by the **initial thread** on exit from its **initial task region** synchronizes with the **acquire flush** performed by the generating **thread** in the **target task**.
- When a **thread** encounters a **teams construct**, the behavior is as if the **release flush** performed by the **thread** on entry to the **teams region** synchronizes with the **acquire flush** performed on entry to each **initial task** that is executed by a different **initial thread** that participates in the execution of the **teams region**.
- When a **thread** that encounters a **teams construct** reaches the end of the **teams region**, the behavior is as if the **release flush** performed by each different participating **initial thread** at exit from its **initial task** synchronizes with the **acquire flush** performed by the **thread** at exit from the **teams region**.
- When a **task** generates an **explicit task** that begins execution on a different **thread**, the behavior is as if the **thread** that is executing the **generating task** performs a **release flush** that synchronizes with the **acquire flush** performed by the **thread** that begins to execute the **explicit task**.
- When an **untied task** suspends execution on a **thread** and resumes execution on a different **thread**, the behavior is as if the **thread** on which the **task region** was suspended performs a **release flush** that synchronizes with the **acquire flush** performed by the **thread** on which the **task region** is resumed.

- When an **underrferred task** completes execution on a given **thread** that is different from the **thread** on which its **generating task** is suspended, the behavior is as if a **release flush** performed by the **thread** that completes execution of the associated **structured block** of the **underrferred task** **synchronizes with** an **acquire flush** performed by the **thread** that resumes execution of the **generating task**.
- When a **dependent task** with one or more **antecedent tasks** begins execution on a given **thread**, the behavior is as if each **release flush** performed by a different **thread** on completion of the associated **structured block** of an **antecedent task** **synchronizes with** the **acquire flush** performed by the **thread** that begins to execute the **dependent task**. If the **antecedent task** is a **detachable task**, the **thread** that fulfills its *allow-completion event* performs a **release flush** upon fulfilling the **event** that **synchronizes with** the **acquire flush** performed when the **dependent task** begins to execute.
- When a **task** begins execution on a given **thread** and it is mutually exclusive with respect to another **dependence-compatible task** that is executed by a different **thread**, the behavior is as if each **release flush** performed on completion of the **dependence-compatible task** **synchronizes with** the **acquire flush** performed by the **thread** that begins to execute the **task**.
- When a **thread** executes a **cancel region**, the *cancel-var* ICV is *true*, and **cancellation** is not already activated for the specified **region**, the behavior is as if the **release flush** performed during the **cancel region** **synchronizes with** the **acquire flush** performed by a different **thread** immediately before a **cancellation point** in which that **thread** observes **cancellation** was activated for the **region**.
- When a **thread** executes an **omp_unset_lock** region that causes the specified **lock** to be unset, the behavior is as if a **release flush** is performed during the **omp_unset_lock** region that **synchronizes with** an **acquire flush** that is performed during the next **omp_set_lock** or **omp_test_lock** region to be executed by a different **thread** that causes the specified **lock** to be set.
- When a **thread** executes an **omp_unset_nest_lock** region that causes the specified **nestable lock** to be unset, the behavior is as if a **release flush** is performed during the **omp_unset_nest_lock** region that **synchronizes with** an **acquire flush** that is performed during the next **omp_set_nest_lock** or **omp_test_nest_lock** region to be executed by a different **thread** that causes the specified **nestable lock** to be set.

23.9 OpenMP Dependences

This section describes **constructs** and **clauses** in OpenMP that support the specification and enforcement of **dependences**. OpenMP supports two kinds of **dependences**: **task dependences**, which enforce orderings between **dependence-compatible tasks**; and **doacross dependences**, which enforce orderings between **doacross iterations** of a loop.

23.9.1 task-dependence-type Modifier

Modifiers

Name	Modifies	Type	Properties
<i>task-dependence-type</i>	<i>all arguments</i>	Keyword: depobj , in , inout , inoutset , mutexinoutset , out	unique

Clauses

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 23.9.5.

Cross References

- **depend** Clause, see Section 23.9.5
- **update** Clause, see Section 23.9.4

23.9.2 Depend Objects

Depend objects are OpenMP objects that can be used to supply user-computed dependences to **depend** clauses. Depend objects must be accessed only through the **depobj** construct, the **depend** clause and the asynchronous device routines; 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.

23.9.3 depobj Construct

Name: depobj Category: executable	Association: unassociated Properties: default
--	--

Clauses

destroy, **init**, **update**

Clause set

Properties: required	Members: destroy , init , update
------------------------------------	--

Additional information

The **depobj** construct may alternatively be specified with a **directive** argument *depend-object* that is a depend object. If this syntax is used, the **init** clause must not be specified and instead the **depend** clause may be specified to initialize *depend-object* to represent a given dependence type and locator list item. With this syntax the **update** clause is only permitted to specify the

task-dependence-type as if it is the sole argument of the [clause](#), with the effect being that the specified [dependence](#) type applies to *depend-object*. With this syntax, any *update-var* or *destroy-var* that is specified in an [update](#) or [destroy](#) [clause](#) must be the same as *depend-object*. Finally, with this syntax only one [clause](#) may be specified and it must be [depend](#), [update](#), or [destroy](#).

Binding

The [binding thread set](#) for a [depobj](#) region is the [encountering thread](#).

Semantics

The [depobj](#) construct initializes, updates or destroys [depend objects](#). If an [init](#) [clause](#) is specified, the state of the specified [depend object](#) is set to *initialized* and the [depend object](#) is set to represent the specified [dependence](#) type and [locator list item](#). If an [update](#) [clause](#) is specified, the specified [depend object](#) is updated to represent the new [dependence](#) type. If a [destroy](#) [clause](#) is specified, the specified [depend object](#) is set to *uninitialized*.

Cross References

- [destroy](#) Clause, see [Section 5.8](#)
- [init](#) Clause, see [Section 5.7](#)
- [update](#) Clause, see [Section 23.9.4](#)

23.9.4 update Clause

Name: update	Properties: innermost-leaf , unique
------------------------------	---

Arguments

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

Modifiers

Name	Modifies	Type	Properties
<i>task-dependence-type</i>	<i>all arguments</i>	Keyword: depobj , in , inout , inoutset , mutexinoutset , out	unique
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: <i>directive-name</i> (a directive name)	unique

Directives

[depobj](#)

Semantics

The [update](#) [clause](#) sets the [dependence](#) type of *update-var* to *task-dependence-type*.

Restrictions

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

- *task-dependence-type* must not be **depobj**.
- The state of *update-var* must be *initialized*.
- If the *locator list item* represented by *update-var* is the **omp_all_memory** reserved locator, *task-dependence-type* must be either **out** or **inout**.

Cross References

- **depobj** Construct, see [Section 23.9.3](#)
- *task-dependence-type* Modifier, see [Section 23.9.1](#)

23.9.5 depend Clause

Name: depend	Properties: taskgraph-altering , task-inherited
----------------------------	--

Arguments

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

Modifiers

Name	Modifies	Type	Properties
<i>task-dependence-type</i>	<i>all arguments</i>	Keyword: depobj , in , inout , inoutset , mutexinoutset , out	unique
<i>iterator</i>	<i>locator-list</i>	Complex, name: iterator Arguments: <i>iterator-specifier</i> list of <i>iterator specifier list item</i> type (<i>default</i>)	unique
<i>directive-name-modifier</i>	<i>all arguments</i>	Keyword: <i>directive-name</i> (a <i>directive name</i>)	unique

Directives

[dispatch](#), [interop](#), [target](#), [target_data](#), [target_enter_data](#), [target_exit_data](#), [target_update](#), [task](#), [task_iteration](#), [taskwait](#)

Semantics

The **depend** clause enforces additional constraints on the scheduling of *tasks*. These constraints establish [dependences](#) only between two [dependence-compatible tasks](#): the [antecedent task](#) and the [dependent task](#). The scheduling constraints are transitive so that the [antecedent task](#) must complete execution before any of its [successor tasks](#) execute. Similarly, the [dependent task](#) cannot start

execution before all of its predecessor tasks complete execution. Task dependences are derived from the *task-dependence-type* and the *list items* in the *locator-list* argument.

One task, *A*, is a preceding dependence-compatible task of another task, *B*, if one of the following is true:

- *A* is a previously generated sibling task of *B*;
- *A* is a preceding dependence-compatible task of an importing task for which *B* is a child task;
- *A* is a child task of an exporting task that is a predecessor task of *B*;
- *A* is a child task of an undeferred exporting task that is a previously generated sibling task of *B*.

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 preceding dependence-compatible task was generated then the generated task will be a dependent task of that preceding dependence-compatible task.

For the **out task-dependence-type** and **inout task-dependence-type**, if the storage location of at least one of the list items matches the storage location of a list item appearing in a **depend** clause with an **in**, **out**, **inout**, **mutexinoutset**, or **inoutset task-dependence-type** on a construct from which a preceding dependence-compatible task was generated then the generated task will be a dependent task of that preceding dependence-compatible task.

For the **mutexinoutset task-dependence-type**, if the storage location of at least one of the list items matches the storage location of a list item appearing in a **depend** clause with an **in**, **out**, **inout**, or **inoutset task-dependence-type** on a construct from which a preceding dependence-compatible task was generated then the generated task will be a dependent task of that preceding dependence-compatible task.

If a list item appearing in a **depend** clause with a **mutexinoutset task-dependence-type** on a task-generating construct matches a list item appearing in a **depend** clause with a **mutexinoutset task-dependence-type** on a different task-generating construct, and both constructs generate dependence-compatible tasks, the dependence-compatible tasks will be mutually exclusive tasks.

For the **inoutset task-dependence-type**, if the storage location of at least one of the list items matches the storage location of a list item appearing in a **depend** clause with an **in**, **out**, **inout**, or **mutexinoutset task-dependence-type** on a construct from which a preceding dependence-compatible task was generated then the generated task will be a dependent task of that preceding dependence-compatible task.

When the *task-dependence-type* is **depobj**, the behavior is as if the *dependence* type and *locator*

`list item` that each specified `depend object list item` represents was specified by `depend clauses` on the current `construct`.

The `list items` that appear in the `depend clause` may reference any *iterator-identifier* defined in its *iterator* modifier.

The `list items` that appear in the `depend clause` may include `array sections` or the `omp_all_memory` reserved locator.

C / C++

The `list items` that appear in a `depend clause` may use `shape-operators`.

C / C++

Note – The enforced `task dependence` establishes a synchronization of `memory` accesses performed by a `dependent task` with respect to accesses performed by the `antecedent tasks`. However, the programmer must properly synchronize with respect to other concurrent accesses that occur outside of those `tasks`.

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 `generated 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 `dependences callback` for each occurrence of the *task-dependences event* to announce its `dependences` with respect to the `list items` in the `depend clause`. A `thread` dispatches the `task_dependence callback` for a *task-dependence event* to report a `dependence` between a `antecedent task` (`src_task_data`) and a `dependent task` (`sink_task_data`).

Restrictions

Restrictions to the `depend clause` are as follows:

- `List items`, other than `reserved locators`, used in `depend clauses` of the same `task` or `dependence-compatible 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 **depend** OpenMP type that correspond to **depend objects** in the *initialized* state.
- **List items** that are expressions of the **depend** OpenMP 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.
- If a **locator list item** has the **ALLOCATABLE** attribute and its allocation status is unallocated, the behavior is **unspecified**.
- If a **locator list item** has the **POINTER** attribute and its association status is disassociated or undefined, the behavior is **unspecified**.

Fortran

C / C++

- A bit-field cannot appear in a **depend** clause.

C / C++

Cross References

- **dependences** Callback, see [Section 40.7.1](#)
- **dispatch** Construct, see [Section 15.7](#)
- Array Sections, see [Section 5.2.5](#)
- Array Shaping, see [Section 5.2.4](#)
- **interop** Construct, see [Section 22.1](#)
- **iterator** Modifier, see [Section 5.2.6](#)
- *task-dependence-type* Modifier, see [Section 23.9.1](#)
- **target** Construct, see [Section 21.8](#)
- **target_data** Construct, see [Section 21.7](#)
- **target_enter_data** Construct, see [Section 21.5](#)
- **target_exit_data** Construct, see [Section 21.6](#)
- **target_update** Construct, see [Section 21.9](#)
- **task** Construct, see [Section 20.1](#)
- **task_dependence** Callback, see [Section 40.7.2](#)
- **task_iteration** Directive, see [Section 20.2.3](#)
- **taskwait** Construct, see [Section 23.5](#)

23.9.6 transparent Clause

Name: transparent	Properties: unique
--------------------------	---------------------------

Arguments

Name	Type	Properties
<i>impex-type</i>	expression of impex OpenMP type	optional

Modifiers

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

Directives

target_data, **task**, **taskloop**

Semantics

The **transparent** clause controls the **task dependence** importing and exporting characteristics of any **generated tasks** of the **construct** on which it appears. If *impex-type* evaluates to **omp_not_impex** then the **generated tasks** are neither **importing tasks** nor **exporting tasks** and so are not **transparent tasks**. Otherwise the **clause** extends the set of **dependence-compatible tasks** of any **child task** of any of the **generated tasks** as follows. If *impex-type* evaluates to **omp_import** then the **generated tasks** are **importing tasks**. If *impex-type* evaluates to **omp_export** then the **generated tasks** are **exporting tasks**. If *impex-type* evaluates to **omp_impex** then the **generated tasks** are both **importing tasks** and **exporting tasks**.

The use of a **variable** in an *impex-type* expression causes an implicit reference to the **variable** in all enclosing **constructs**. The *impex-type* expression is evaluated in the context outside of the **construct** on which the **clause** appears. If *impex-type* is not specified, the effect is as if *impex-type* evaluates to **omp_impex**.

Cross References

- **depend** Clause, see [Section 23.9.5](#)
- **target_data** Construct, see [Section 21.7](#)
- **task** Construct, see [Section 20.1](#)
- **taskloop** Construct, see [Section 20.2](#)

23.9.7 doacross Clause

Name: doacross	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: <i>directive-name</i> (a <i>directive name</i>)	unique

Directives

ordered

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 6.4.3).

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** that 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

- OpenMP Loop-Iteration Spaces and Vectors, see [Section 6.4.3](#)
- **ordered** Clause, see [Section 6.4.6](#)
- Stand-alone **ordered** Construct, see [Section 23.10.1](#)

23.10 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-released event* occurs in the *task* that encounters the *ordered construct* after it completes any synchronization on exit from the *region*.

Tool Callbacks

A *thread* dispatches a registered *mutex_acquire* callback for each occurrence of an *ordered-acquiring event* in that *thread*. A *thread* dispatches a registered *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*. These *callback* occur in the *task* that encounters the *construct*.

1 **Restrictions**

- 2 • The **construct** that corresponds to the **binding region** of an **ordered region** must specify an
- 3 **ordered clause** if one of its **constituent constructs** is a **worksharing-loop construct**.
- 4 • The **construct** that corresponds to the **binding region** of an **ordered region** must not specify
- 5 a **reduction clause** with the **inscan** modifier.
- 6 • The **region** of a block-associated **ordered construct** must not have a **binding region** that
- 7 corresponds to a **construct** in which a stand-alone **ordered construct** is closely nested.
- 8 • An **ordered region** that corresponds to an **ordered construct** with the **threads** or
- 9 **doacross clause** may not be closely nested inside a **critical**, **ordered**, **loop**, **task**,
- 10 or **taskloop region** (see [Section 23.10](#)).
- 11 • The **doacross-affected loops** of a **doacross loop nest** must be **perfectly nested loops**.
- 12 • The **construct** that corresponds to the **binding region** of an **ordered region** must not specify
- 13 a **linear clause**.

14 • The **doacross-affected loops** of a **doacross loop nest** must not be range-based **for** loops.

C++

15 **Cross References**

- 16 • OMPT **mutex** Type, see [Section 39.20](#)
- 17 • **mutex_acquire** Callback, see [Section 40.7.8](#)
- 18 • **mutex_released** Callback, see [Section 40.7.13](#)

19 **23.10.1 Stand-alone ordered Construct**

20

Name: ordered	Association: unassociated
Category: executable	Properties: mutual-exclusion

21 **Clauses**

22 **doacross**

23 **Binding**

24 The **binding thread set** for a stand-alone **ordered region** is the **current team**. A stand-alone

25 **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 the n [doacross-affected loops](#). 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 [dependences callback](#) with all vector entries listed as [ompt_dependence_type_sink](#) in the *deps* argument for each occurrence of a *doacross-sink event* in that [thread](#). A [thread](#) dispatches a registered [dependences callback](#) with all vector entries listed as [ompt_dependence_type_source](#) in the *deps* argument for each occurrence of a *doacross-source event* in that [thread](#).

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](#).
- The [construct](#) must be closely nested inside a [worksharing-loop construct](#).

Cross References

- OMPT [dependence_type](#) Type, see [Section 39.10](#)
- **dependences** Callback, see [Section 40.7.1](#)
- **doacross** Clause, see [Section 23.9.7](#)
- Worksharing-Loop Constructs, see [Section 19.6](#)

23.10.2 Block-associated ordered Construct

Name: `ordered`
Category: `executable`

Association: `block`
Properties: `mutual-exclusion`, `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 region that corresponds to a construct for which a `worksharing-loop` construct or `simd` construct is a `constituent construct`.

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.

Execution Model Events

The *ordered-acquired* event occurs in the `task` that encounters the `ordered` construct after it enters the `region`, but before it executes the associated `structured block`.

Tool Callbacks

A `thread` dispatches a registered `mutex_acquired` callback for each occurrence of an *ordered-acquired* event in that `thread`. This `callback` occurs in the `task` that encounters the `construct`.

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 correspond to a `construct` for which the `simd` construct is a `leaf construct`.

- If the **threads parallelization-level** clause is specified, the **binding region** must correspond to a **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 **compound 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 an argument value equal to the number of **collapse-affected loops** must appear on the **construct** that corresponds to the **binding region**, if the **binding region** is not a **simd region**.

Cross References

- *parallelization-level* Clauses, see [Section 23.10.3](#)
- Worksharing-Loop Constructs, see [Section 19.6](#)
- **mutex_acquired** Callback, see [Section 40.7.12](#)
- **ordered** Clause, see [Section 6.4.6](#)
- **simd** Construct, see [Section 18.4](#)

23.10.3 *parallelization-level* Clauses

Clause groups

Properties: unique	Members: Clauses simd, threads
----------------------------------	---

Directives

ordered

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

- Block-associated **ordered** Construct, see [Section 23.10.2](#)

23.10.3.1 threads Clause

Name: threads	Properties: innermost-leaf, unique
-----------------------------	--

1 **Arguments**

2

Name	Type	Properties
<i>apply-to-threads</i>	expression of OpenMP logical type	constant, optional

3 **Modifiers**

4

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

5 **Directives**

6 **ordered**

7 **Semantics**

8 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*.

11 **Cross References**

- 12
 - Block-associated **ordered** Construct, see [Section 23.10.2](#)

13 **23.10.3.2 simd Clause**

14

Name: simd	Properties: innermost-leaf, unique
--------------------------	---

15 **Arguments**

16

Name	Type	Properties
<i>apply-to-simd</i>	expression of OpenMP logical type	constant, optional

17 **Modifiers**

18

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

19 **Directives**

20 **ordered**

21 **Semantics**

22 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*.

25 **Cross References**

- 26
 - Block-associated **ordered** Construct, see [Section 23.10.2](#)

24 Cancellation Constructs

This chapter defines constructs related to cancellation of OpenMP regions.

24.1 *cancel-directive-name* Clauses

Clause groups

Properties: exclusive, required, unique	Members: Clauses do, for, parallel, sections, taskgroup
--	---

Modifiers

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

Directives

cancel, cancellation_point

Semantics

For each directive that has the cancellable property (i.e., the directive may subject to cancellation and 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 OpenMP 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.

1 **Cross References**

- 2 • **cancel** Construct, see [Section 24.2](#)
- 3 • **cancellation_point** Construct, see [Section 24.3](#)
- 4 • **do** Construct, see [Section 19.6.2](#)
- 5 • **for** Construct, see [Section 19.6.1](#)
- 6 • **parallel** Construct, see [Section 18.1](#)
- 7 • **sections** Construct, see [Section 19.3](#)
- 8 • **taskgroup** Construct, see [Section 23.4](#)

9 **24.2 cancel Construct**

10

Name: <code>cancel</code> Category: <code>executable</code>	Association: <code>unassociated</code> Properties: <code>default</code>
--	--

11 **Clause groups**

12 *`cancel-directive-name`*

13 **Clauses**

14 `if`

15 **Binding**

16 The `binding thread set` of the `cancel` region is the `current team`. The `binding region` of the
17 `cancel` region is the innermost enclosing `region` of the type that corresponds to
18 *`cancel-directive-name`*.

19 **Semantics**

20 The `cancel` construct activates `cancellation` of the innermost enclosing `region` of the type
21 specified by *`cancel-directive-name`*, which must be the *directive-name* of a `cancellable construct`.
22 `Cancellation` of the `binding region` is activated only if the *`cancel-var ICV`* is `true`, in which case the
23 `cancel` construct causes the `encountering task` to continue execution at the end of the `binding`
24 `region` if *`cancel-directive-name`* is not `taskgroup`. If the *`cancel-var ICV`* is `true` and
25 *`cancel-directive-name`* is `taskgroup`, the `encountering task` continues execution at the end of the
26 `current task region`. If the *`cancel-var ICV`* is `false`, the `cancel` construct is ignored.

27 `Threads` check for active `cancellation` only at `cancellation points` that are implied at the following
28 locations:

- 29 • `cancel` regions;
- 30 • `cancellation_point` regions;
- 31 • `barrier` regions;

- at the end of a **worksharing-loop construct** with a **nowait** clause and for which the same **list item** appears in both **firstprivate** and **lastprivate** clauses; and
- **implicit barrier regions**.

When a **thread** reaches one of the above **cancellation points** and if the *cancel-var* ICV is *true*, then:

- If the **thread** is at a **cancel** or **cancellation_point** region and *cancel-directive-name* is not **taskgroup**, the **thread** continues execution at the end of the canceled **region** if **cancellation** has been activated for the innermost enclosing **region** of the type specified.
- If the **thread** is at a **cancel** or **cancellation_point** region and *cancel-directive-name* is **taskgroup**, the **encountering task** checks for active **cancellation** of all of the **taskgroup sets** to which the **encountering task** belongs, and continues execution at the end of the **current task region** if **cancellation** has been activated for any of the **taskgroup sets**.
- If the **encountering task** is at a **barrier region** or at the end of a **worksharing-loop construct** with a **nowait** clause and for which the same **list item** appears in both **firstprivate** and **lastprivate** clauses, the **encountering task** checks for active **cancellation** of the innermost enclosing **parallel region**. If **cancellation** has been activated, then the **encountering task** continues execution at the end of the canceled **region**.

When **cancellation** of **tasks** is activated through a **cancel** construct with **taskgroup** for *cancel-directive-name*, the **tasks** that belong to the **taskgroup set** of the innermost enclosing **taskgroup region** will be canceled; that **taskgroup set** is then the **canceled taskgroup set** corresponding to that **cancel region**. The **task** that encountered that **construct** continues execution at the end of its **task region**, which implies completion of that **task**. Any **task** that belongs to the **canceled taskgroup set** and has already begun execution must run to completion or until a **cancellation point** is reached. Upon reaching a **cancellation point** and if **cancellation** is active, the **task** continues execution at the end of its **task region**, which implies the completion of the **task**. Any **task** that belongs to the **canceled taskgroup set** and that has not begun execution or that has not yet been fulfilled through an **event** variable may be discarded, which implies its completion.

When **cancellation** of **tasks** is activated through a **cancel** construct with *cancel-directive-name* other than **taskgroup**, each **thread** of the **binding thread set** resumes execution at the end of the canceled **region** if a **cancellation point** is encountered. If the canceled **region** is a **parallel region**, any **tasks** that have been created by a **task** or a **taskloop** construct and their **descendent tasks** are canceled according to the above **taskgroup cancellation** semantics. If the canceled **region** is not a **parallel region**, no **task** cancellation occurs.

C++

The usual C++ rules for object destruction are followed when **cancellation** is performed.

C++

Fortran

All **private** objects or subobjects with the **ALLOCATABLE** attribute that are allocated inside the canceled **construct** are deallocated.

Fortran

If the canceled **construct** specifies an **original list-item updating clause**, the final values of the **list items** that appear in those **clauses** are **undefined**.

When an **if clause** is present on a **cancel construct** and *if-expression* evaluates to *false*, the **cancel construct** does not activate **cancellation**. The **cancellation point** associated with the **cancel construct** is always encountered regardless of the value of *if-expression*.

Note – The programmer is responsible for releasing locks and other synchronization data structures that might cause a deadlock when a **cancel construct** is encountered and blocked **threads** cannot be canceled. The programmer is also responsible for ensuring proper synchronizations to avoid deadlocks that might arise from **cancellation** of **regions** that contain **synchronization constructs**.

Execution Model Events

If a **task** encounters a **cancel construct** that will activate **cancellation** then a *cancel event* occurs. A *discarded-task event* occurs for any discarded **tasks**.

Tool Callbacks

A **thread** dispatches a registered **cancel callback** for each occurrence of a *cancel event* in the context of the **encountering task**. (*flags & omp_t_cancel_activated*) always evaluates to *true* in the dispatched **callback**; (*flags & omp_t_cancel_parallel*) evaluates to *true* in the dispatched **callback** if *cancel-directive-name* is **parallel**; (*flags & omp_t_cancel_sections*) evaluates to *true* in the dispatched **callback** if *cancel-directive-name* is **sections**; (*flags & omp_t_cancel_loop*) evaluates to *true* in the dispatched **callback** if *cancel-directive-name* is **for** or **do**; and (*flags & omp_t_cancel_taskgroup*) evaluates to *true* in the dispatched **callback** if *cancel-directive-name* is **taskgroup**.

A **thread** dispatches a registered **cancel callback** with its *task_data* argument pointing to the **data** object associated with the discarded **task** and with **omp_t_cancel_discarded_task** as its *flags* argument for each occurrence of a *discarded-task event*. The **callback** occurs in the context of the **task** that discards the **task**.

Restrictions

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

- The behavior for concurrent **cancellation** of a **region** and a **region** nested within it is **unspecified**.
- If *cancel-directive-name* is **taskgroup**, the **cancel construct** must be a **closely nested construct** of a **task** or a **taskloop construct** and the **cancel region** must be a **closely nested region** of a **taskgroup region**.
- If *cancel-directive-name* is not **taskgroup**, the **cancel construct** must be a **closely nested construct** of a **construct** that matches *cancel-directive-name*.

- A **worksharing construct** that is canceled must not have a **nowait** clause or a **reduction** clause with a **user-defined reduction** that uses **omp_orig** in the *initializer-expr* of the corresponding **declare_reduction** directive.
- A **worksharing-loop construct** that is canceled must not have an **ordered** clause or a **reduction** clause with the **inscan** *reduction-modifier*.
- When **cancellation** is active for a **parallel region**, a **thread** in the **team** that binds to that **region** must not be executing or encounter a **worksharing construct** with an **ordered** clause, a **reduction** clause with the **inscan** *reduction-modifier* or a **reduction** clause with a **user-defined reduction** that uses **omp_orig** in the *initializer-expr* of the corresponding **declare_reduction** directive.
- During execution of a **construct** that may be subject to **cancellation**, a **thread** must not encounter an orphaned **cancellation point**. That is, a **cancellation point** must only be encountered within that **construct** and must not be encountered elsewhere in its **region**.

Cross References

- **barrier** Construct, see [Section 23.3.1](#)
- **cancel** Callback, see [Section 40.6](#)
- OMPT **cancel_flag** Type, see [Section 39.7](#)
- **cancellation_point** Construct, see [Section 24.3](#)
- OMPT **data** Type, see [Section 39.8](#)
- **declare_reduction** Directive, see [Section 8.15](#)
- **firstprivate** Clause, see [Section 7.3.4](#)
- *cancel-var* ICV, see [Table 3.1](#)
- **if** Clause, see [Section 5.6](#)
- **nowait** Clause, see [Section 23.6](#)
- **omp_get_cancellation** Routine, see [Section 36.1](#)
- **ordered** Clause, see [Section 6.4.6](#)
- **private** Clause, see [Section 7.3.3](#)
- **reduction** Clause, see [Section 8.10](#)
- **task** Construct, see [Section 20.1](#)

24.3 cancellation_point Construct

Name: <code>cancellation_point</code> Category: <code>executable</code>	Association: <code>unassociated</code> Properties: <code>default</code>
--	--

Clause groups

cancel-directive-name

Additional information

The `cancellation_point` directive may alternatively be specified with `cancellation_point` as the *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 24.2](#).

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 3.1](#)
- `omp_get_cancellation` Routine, see [Section 36.1](#)

25 Composition of Constructs

This chapter defines rules and mechanisms for nesting [regions](#) and for combining [constructs](#).

25.1 Compound Directive Names

Unless explicitly specified otherwise, the *directive-name* of a [compound directive](#) concatenates two or more [directive names](#), with an intervening separating character, the [directive-name separator](#) between each of them. Each [directive name](#), as well as any concatenation of consecutive [directive names](#) and their [directive-name separator](#), is a [constituent-directive name](#). Any [constituent-directive name](#) that is not itself a [compound-directive name](#) is a [leaf-directive name](#).

Let *directive-name-A* refer to the first [leaf-directive name](#) or [simple-composite-directive name](#) that appears in a [compound-directive name](#), and let *directive-name-B* refer to the [constituent-directive name](#) that forms the remainder of the [compound-directive name](#). If the [construct](#) named by *directive-name-B* can be immediately nested inside the [construct](#) named by *directive-name-A*, the [compound-directive name](#) is a [combined-directive name](#), the name of [combined directive](#). Otherwise, the [compound-directive name](#) is a [composite-directive name](#). The syntax for a [compound-directive name](#) is `<compound-directive-name>`, as described in the following grammar:

```
<compound-directive-name> :  
    <combined-directive-name>  
    <composite-directive-name>  
  
<combined-directive-name> :  
    <directive-name-A-for-combined-name><separator><directive-name-B-for-combined-name>  
    >  
  
<directive-name-A-for-combined-name> :  
    <parallelism-generating-directive-name>  
    <thread-selecting-directive-name>  
  
<directive-name-B-for-combined-name> :  
    <parallelism-generating-directive-name>  
    <thread-selecting-directive-name>  
    <partitioned-directive-name>  
  
<composite-directive-name> :
```

```

1  <distributed-loop-composite-directive-name>
2  <simd-partitionable-composite-directive-name>
3  <simple-composite-directive-name>
4
5  <distributed-loop-composite-directive-name> :
6      <distribute-directive-name><separator><parallel-loop-combined-directive-name>
7
8  <parallel-loop-combined-directive-name> :
9      <parallel-directive-name><separator><worksharing-loop-directive-name>
10
11 <simd-partitionable-composite-directive-name> :
12 <simd-partitionable-directive-name><separator><simd-directive-name>

```

where:

- <parallelism-generating-directive-name> is the name of a [parallelism-generating construct](#) or a [compound directive](#) for which *directive-name-A* is the name of a [parallelism-generating construct](#);
- <thread-selecting-directive-name> is the name of a [thread-selecting construct](#) or a [compound directive](#) for which *directive-name-A* is the name of a [thread-selecting construct](#);
- <partitioned-directive-name> is the name of a [partitioned construct](#) or a [composite directive](#) for which *directive-name-A* is the name of a [partitioned construct](#);
- <simple-composite-directive-name> is the name of a [simple-composite directive](#);
- <distribute-directive-name> is [distribute](#);
- <parallel-directive-name> is [parallel](#);
- <worksharing-loop-directive-name> is an instance of <partitioned-directive-name> for which the [partitioned construct](#) is a [worksharing-loop construct](#);
- <simd-partitionable-directive-name> is the name of a [SIMD-partitionable construct](#);
- <simd-directive-name> is [simd](#);

	C / C++	
• <separator>, the directive-name separator , is white space .		
	C / C++	
	Fortran	
• <separator>, the directive-name separator , is white space or a plus sign (i.e., '+').		
	Fortran	

The section that defines any [simple-composite directive](#), such as those that combine a series of [directives](#) into one [directive](#), also specifies its [leaf directives](#). Unless otherwise specified, those [leaf directives](#) may be specified by their [leaf-directive names](#) in a [directive-name-modifier](#).

Restrictions

Restrictions to **compound-directive names** are as follows:

- Any given instance of a **compound-directive name** must use the same character for all instances of *<separator>*.
- **Leaf-directive names** and **simple-composite-directive names** that include spaces are not permitted in a **compound-directive name**; they must instead be specified with an underscore replacing each space in the **directive name**.
- The **leaf-directive names** of a given **compound-directive name** must be unique.
- The **construct** corresponding to *<directive-name-B-for-combined-name>* must be permitted to be immediately nested inside the **construct** corresponding to *<directive-name-A-for-combined-name>*.
- If the first **leaf-directive name** or **simple-composite-directive name** of *<directive-name-B-for-combined-name>* is the name of a **worksharing construct** or a **thread-selecting construct** then *<directive-name-A-for-combined-name>* must be **parallel**.
- If *<directive-name-A-for-combined-name>* and the first **leaf-directive name** or **simple-composite-directive name** of *<directive-name-B-for-combined-name>* are the names of **task-generating constructs** then their respective **explicit task regions** must not bind to the same **parallel region**.
- The **compound construct** named by a given **compound-directive name** must have at most one **constituent construct** that is a **map-entering construct**.
- The **compound construct** named by a given **compound-directive name** must have at most one **constituent construct** that is a **map-exiting construct**.

Fortran

- If a **directive name** is ambiguous due to the use of optional intervening spaces between **leaf-directive names**, the **directive-name separator** must be a plus sign.

Fortran

Cross References

- **distribute** Construct, see [Section 19.7](#)
- **parallel** Construct, see [Section 18.1](#)
- **simd** Construct, see [Section 18.4](#)

25.2 Clauses on Compound Constructs

This section specifies the handling of **clauses** on **compound constructs** and the handling of implicit **clauses** that arise from any **variable** with **predetermined data-sharing attributes** on more than one **leaf construct**. For any **clause** for which a **directive-name-modifier** is specified, the effect of the **modifier** is applied prior to any of the rules that are specified in this section. Some **clauses** are permitted only on a single **leaf construct** of the **compound construct**, in which case the effect is as if the **clause** is applied to that specific **construct**. Other **clauses** that are permitted on more than one **leaf construct** have the effect as if they are applied to a subset of those **constructs**, as detailed in this section. Unless otherwise specified, the effect of a **clause** on a **compound directive** is as if it is applied to all **leaf constructs** that permit it (i.e., it has the default **all-constituents property**).

Unless otherwise specified, certain **clause properties** determine how each **clause** with those **properties** applies to any **constituent directives** of a **compound directive** on which it appears. Regardless of any specified **directive-name-modifier**, the effect of any **clause** with the **once-for-all-constituents property** on a **compound construct** is as if it is applied once to the **compound construct** regardless of how many **constituent constructs** to which they may apply.

The effect of any **clause** with the **all-privatizing property** on a **compound directive** is as if it is applied to all **leaf constructs** that permit the **clause** and to which a **data-sharing attribute clause** that may create a **private** copy of the same **list item** is applied. Unless otherwise specified, the effect of any **clause** with the **innermost-leaf property** on a **compound construct** is as if it is applied only to the innermost **leaf construct** that permits it. Unless otherwise specified, the effect of any **clause** with the **outermost-leaf property** on a **compound construct** is as if it is applied only to the outermost **leaf construct** that permits it.

The effect of the **firstprivate** **clause** is as if it is applied to one or more **leaf constructs** as follows:

- To the **distribute** **construct** if it is among the **constituent constructs**;
- To the **teams** **construct** if it is among the **constituent constructs** and the **distribute** **construct** is not;
- To a **worksharing** **construct** that accepts the **clause** if one is among the **constituent constructs**;
- To the **taskloop** **construct** if it is among the **constituent constructs**;
- To the **parallel** **construct** if it is among the **constituent construct** and neither a **taskloop** **construct** nor a **worksharing** **construct** that accepts the **clause** is among them;
- To the **target** **construct** if it is among the **constituent constructs** and the same **list item** neither appears in a **lastprivate** **clause** nor is the **base variable** or **base pointer** of a **list item** that appears in a **map** **clause**.

If the **parallel** **construct** is among the **constituent constructs** and the effect is not as if the **firstprivate** **clause** is applied to it by the above rules, then the effect is as if the **shared** **clause** with the same **list item** is applied to the **parallel** **construct**. If the **teams** **construct** is

among the **constituent constructs** and the effect is not as if the **firstprivate** clause is applied to it by the above rules, then the effect is as if the **shared** clause with the same **list item** is applied to the **teams** construct.

The effect of the **lastprivate** clause is as if it is applied to all **leaf constructs** that permit the **clause**. If the **parallel** construct is among the **constituent constructs** and the **list item** is not also specified in the **firstprivate** clause, then the effect of the **lastprivate** clause is as if the **shared** clause with the same **list item** is applied to the **parallel** construct. If the **teams** construct is among the **constituent constructs** and the **list item** is not also specified in the **firstprivate** clause, then the effect of the **lastprivate** clause is as if the **shared** clause with the same **list item** is applied to the **teams** construct. If the **target** construct is among the **constituent constructs** and the **list item** is not the **base variable** or **base pointer** of a **list item** that appears in a **map** clause, the effect of the **lastprivate** clause is as if the same **list item** appears in a **map** clause with a *map-type* of **tofrom**.

The effect of the **reduction** clause is as if it is applied to all **leaf constructs** that permit the **clause**, except for the following **constructs**:

- The **parallel** construct, when combined with the **sections**, **worksharing-loop**, **loop**, or **taskloop** construct; and
- The **teams** construct, when combined with the **loop** construct.

For the **parallel** and **teams** constructs above, the effect of the **reduction** clause instead is as if each **list item** or, for any **list item** that is an **array item**, its corresponding **base array** or corresponding **base pointer** appears in a **shared** clause for the **construct**. If the **task reduction-modifier** is specified, the effect is as if it only modifies the behavior of the **reduction** clause on the innermost **leaf construct** that accepts the **modifier** (see Section 8.10). If the **inscan reduction-modifier** is specified, the effect is as if it modifies the behavior of the **reduction** clause on all **constructs** of the **compound construct** to which the **clause** is applied and that accept the **modifier**. If a **list item** in a **reduction** clause on a **compound target construct** does not have the same **base variable** or **base pointer** as a **list item** in a **map** clause on the **construct**, then the effect is as if the **list item** in the **reduction** clause appears as a **list item** in a **map** clause with a *map-type* of **tofrom**.

The effect of the **linear** clause is as if it is applied to the innermost **leaf construct**. Additionally, if the **list item** is not the **loop-iteration variable** of a **construct** for which **simd** is a **constituent construct**, the effect on the outer **leaf constructs** is as if the **list item** was specified in **firstprivate** and **lastprivate** clauses on the **compound construct**, with the rules specified above applied. If a **list item** of the **linear** clause is the **loop-iteration variable** of a **construct** for which the **simd** construct is a **leaf construct** and the **variable** is not declared in the **construct**, the effect on the outer **leaf constructs** is as if the **list item** was specified in a **lastprivate** clause on the **compound construct** with the rules specified above applied.

If the **clauses** have expressions on them, such as for various **clauses** where the argument of the **clause** is an expression, or *lower-bound*, *length*, or *stride* expressions inside **array sections** (or *subscript* and *stride* expressions in *subscript-triplet* for Fortran), or *linear-step* or *alignment*

expressions, the expressions are evaluated immediately before the **construct** to which the **clause** has been split or duplicated per the above rules (therefore inside of the outer **leaf constructs**). However, the expressions inside the **num_teams** and **thread_limit** clauses are always evaluated before the outermost **leaf construct**.

The restriction that a **list item** may not appear in more than one **data-sharing attribute clause** with the exception of specifying a **variable** in both **firstprivate** and **lastprivate** clauses applies after the **clauses** are split or duplicated per the above rules.

Restrictions

Restrictions to **clauses** on **compound constructs** are as follows:

- A **clause** that appears on a **compound construct** must apply to at least one of the **leaf constructs** per the rules defined in this section.

Cross References

- **distribute** Construct, see [Section 19.7](#)
- **firstprivate** Clause, see [Section 7.3.4](#)
- **lastprivate** Clause, see [Section 7.3.5](#)
- **linear** Clause, see [Section 8.14](#)
- **loop** Construct, see [Section 19.8](#)
- **map** Clause, see [Section 11.3](#)
- **num_teams** Clause, see [Section 18.2.1](#)
- **parallel** Construct, see [Section 18.1](#)
- **reduction** Clause, see [Section 8.10](#)
- **sections** Construct, see [Section 19.3](#)
- **shared** Clause, see [Section 7.3.2](#)
- **simd** Construct, see [Section 18.4](#)
- **target** Construct, see [Section 21.8](#)
- **taskloop** Construct, see [Section 20.2](#)
- **teams** Construct, see [Section 18.2](#)
- **thread_limit** Clause, see [Section 21.3](#)

25.3 Compound Construct Semantics

The semantics of **combined constructs** are identical to that of explicitly specifying the first **construct** containing one instance of the second **construct** and no other statements.

Most **composite constructs** compose **constructs** that otherwise cannot be immediately nested to apply multiple **loop-nest-associated constructs** to the same **canonical loop nest**. The semantics of each of these **composite constructs** first apply the semantics of the enclosing **construct** as specified by *directive-name-A* and any **clauses** that apply to it. For each **task** 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. 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.

For all **compound constructs**, **tool callbacks** are invoked as if the **leaf constructs** were explicitly nested. All **compound constructs** for which a **loop-nest-associated construct** is a **leaf construct** are themselves **loop-nest-associated constructs**.

Restrictions

Restrictions to **compound construct** are as follows:

- The restrictions of all **constituent directives** apply.
- If **distribute** is a **constituent-directive name**, the **linear clause** may only be specified for **loop-iteration variables** of loops that are associated with the **construct** and the **ordered clause** must not be specified.

Cross References

- **distribute** Construct, see [Section 19.7](#)
- **in_reduction** Clause, see [Section 8.12](#)
- **linear** Clause, see [Section 8.14](#)
- **ordered** Clause, see [Section 6.4.6](#)
- **parallel** Construct, see [Section 18.1](#)
- **reduction** Clause, see [Section 8.10](#)
- **simd** Construct, see [Section 18.4](#)
- **taskloop** Construct, see [Section 20.2](#)

1

Part IV

2

Runtime Library Routines

26 Runtime Library Definitions

This chapter defines the naming convention for the OpenMP API routines. It also defines several OpenMP types. The names of OpenMP API routines have an `omp_` prefix. Names that begin with the `omp_x_` prefix are reserved for routines that are implementation defined extensions.

For each base language, a compliant implementation must supply a set of definitions for the OpenMP API routines and the OpenMP types that are used for their arguments and return values. The C/C++ header file (`omp.h`) and the Fortran module file (`omp_lib`) or the deprecated Fortran include file (`omp_lib.h`) provide these definitions and must contain a declaration for each routine and predefined identifier as well as a definition of each OpenMP type. In addition, each set of definitions may specify other implementation defined values.

C / C++

The routines are external functions with “C” linkage. C/C++ prototypes for the routines shall be provided in the `omp.h` header file.

C / C++

Fortran

The Fortran OpenMP API routines are external procedures. The return values of these routines are of default kind, unless otherwise specified. Interface declarations for the Fortran routines shall be provided in the form of a Fortran module named `omp_lib` or the deprecated Fortran include file named `omp_lib.h`. Whether the `omp_lib.h` file provides derived-type definitions or those routines that require an explicit interface is implementation defined. Whether the `include` file or the `module` file (or both) is provided is also implementation defined. Whether any of the routines that take an argument are extended with a generic interface so arguments of different `KIND` type can be accommodated is implementation defined.

Fortran

Restrictions

The following restrictions apply to all routines and OpenMP types:

C++

- Enumeration OpenMP types provided in the `omp.h` header file shall not be scoped enumeration types unless explicitly allowed.

C++

- [Routines](#) may not be called from **PURE** or **ELEMENTAL** procedures.
- [Routines](#) may not be called in **DO CONCURRENT** constructs.

26.1 Predefined Identifiers

Predefined Identifiers

Name	Value	Properties
<code>omp_curr_progress_width</code>	see below	<i>default</i>
<code>omp_initial_device</code>	<code>-1</code>	constant
<code>omp_default_device</code>	<code>< -1</code>	constant
<code>omp_invalid_device</code>	<code>< -1</code>	constant
<code>omp_unassigned_thread</code>	<code>< -1</code>	constant
<code>openmp_version</code>	see below	constant , Fortran-only

In addition to the [predefined identifiers](#) of [OpenMP types](#) that are defined with their corresponding [OpenMP type](#), the OpenMP API includes the [predefined identifiers](#) shown above. The [predefined identifiers](#) `omp_invalid_device` and `omp_default_device` have distinct [implementation defined](#) values less than -1. The [predefined identifier](#) `omp_unassigned_thread` has an [implementation defined](#) value less than -1. The [predefined identifier](#) `omp_curr_progress_width` is a context-specific value that represents 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](#).

The [predefined identifiers](#) are represented as default integer named constants. The [predefined identifier](#) `openmp_version` has a value `yyyymm` where `yyyy` and `mm` are the year and month designations of the version of the OpenMP API that the implementation supports. This value matches that of the C preprocessor macro `_OPENMP`, when a macro preprocessor is supported (see [Section 5.3](#)).

26.2 Routine Bindings

Unless otherwise specified, the [binding task set](#) of any [routine region](#) is its [encountering task](#) and the [binding thread set](#) of any [routine region](#) is the [encountering thread](#). That is, the default [binding properties](#) for [routines](#) are the [encountering-task binding property](#) and the [encountering-thread binding property](#). However, the [binding task set](#) for all [lock routine regions](#) is [all tasks](#) in the [contention group](#) so all of those [routines](#) have the [all-contention-group-tasks binding property](#).

Further, the **binding region** of any **routine** that has a **binding region** for any type of **region** that is relevant to that **routine region** is the innermost enclosing **region** of that type. The **binding thread set** of several **routines** is **all threads** or **all threads** on the **current device**. Those **routine** have the **all-threads binding property** or the **all-device-threads binding property**.

26.3 Routine Argument Properties

Similarly to **directive** and **clause** arguments, **routine** arguments have **properties** that often specify constraints on their values. For all **routines**, if an argument is specified that does not conform to the constraints implied by its **properties** then the behavior is **implementation defined**. **Routine properties** include the **properties** that apply to the arguments of **directives** and **clauses** with the same meanings. The default **property** for all **routine** arguments is the **required property**. **Routine** arguments that have the **optional property** may be omitted in **base languages** for which a default value is defined. In addition, **routine** argument **properties** include ones that correspond to aspects of their **base language** prototypes, as shown in **Table 26.1**.

TABLE 26.1: Routine Argument Properties

Property	Property Description
C/C++-only property	An argument that is absent in Fortran
C/C++ pointer property	A pointer type in C/C++, but not a pointer in Fortran
intent(in) property	An intent (in) argument in Fortran and, if type corresponds to a pointer type but not pointer to char , a const argument in C/C++
intent(out) property	An intent (out) argument in Fortran
ISO C property	Binds to an ISO C type in Fortran
pointer property	A pointer type in C/C++ and an assumed-size array in Fortran
pointer-to-pointer property	A pointer-to-pointer type in C/C++
procedure property	A function pointer type in C/C++ and a procedure type in Fortran
value property	A value argument in Fortran

26.4 General OpenMP Types

This section describes general OpenMP types.

26.4.1 OpenMP intptr Type

Name: <code>intptr</code> Properties: <code>omp</code>	Base Type: <code>c_intptr_t</code>
---	------------------------------------

Type Definition

C / C++	<code>typedef intptr_t omp_intptr_t;</code>
C / C++	
Fortran	<code>integer (kind=omp_c_intptr_t_kind)</code>
Fortran	

The `intptr` OpenMP type is a signed integer type that is capable of holding a pointer on any device, and is equivalent to `intptr_t` on platforms that provide it.

26.4.2 OpenMP uintptr Type

Name: <code>uintptr</code> Properties: <code>C/C++-only</code> , <code>omp</code>	Base Type: <code>c_uintptr_t</code>
--	-------------------------------------

Type Definition

C / C++	<code>typedef uintptr_t omp_uintptr_t;</code>
C / C++	

The `uintptr` OpenMP type is an unsigned integer type that is capable of holding a pointer on any device, and is equivalent to `uintptr_t` on platforms that provide it.

26.5 OpenMP Parallel Region Support Types

This section describes OpenMP types that support parallel regions.

26.5.1 OpenMP sched Type

Name: <code>sched</code> Properties: <code>omp</code>	Base Type: <code>enumeration</code>
--	-------------------------------------

Values

Name	Value	Properties
<code>omp_sched_static</code>	<code>0x1</code>	omp
<code>omp_sched_dynamic</code>	<code>0x2</code>	omp
<code>omp_sched_guided</code>	<code>0x3</code>	omp
<code>omp_sched_auto</code>	<code>0x4</code>	omp
<code>omp_sched_monotonic</code>	<code>0x80000000u</code>	omp

Type Definition

C / C++

```
typedef enum omp_sched_t {
    omp_sched_static      = 0x1,
    omp_sched_dynamic     = 0x2,
    omp_sched_guided      = 0x3,
    omp_sched_auto        = 0x4,
    omp_sched_monotonic   = 0x80000000u
} omp_sched_t;
```

C / C++

Fortran

```
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)
integer (kind=omp_sched_kind), &
    parameter :: omp_sched_monotonic = &
        int(Z'80000000', kind=omp_sched_kind)
```

Fortran

The [sched](#) type is used in [routines](#) that modify or retrieve the value of the [run-sched-var](#) ICV. Each of [omp_sched_static](#), [omp_sched_dynamic](#), [omp_sched_guided](#), and [omp_sched_auto](#) can be combined with [omp_sched_monotonic](#) by using the + or | operator in C/C++ or the + operator in Fortran. If the [schedule](#) type is combined with the [omp_sched_monotonic](#), the value corresponds to a schedule that is modified with the [monotonic ordering-modifier](#). Otherwise, the value corresponds to a schedule that is modified with the [nonmonotonic ordering-modifier](#).

Cross References

- *run-sched-var* ICV, see [Table 3.1](#)

26.6 OpenMP Tasking Support Types

This section describes [OpenMP types](#) that support tasking mechanisms.

26.6.1 OpenMP event_handle Type

Name: event_handle Properties: named-handle , omp , opaque	Base Type: implementation-defined-int
---	---

Type Definition

C / C++	<code>typedef <implementation-defined-integral> omp_event_handle_t;</code>
C / C++	
Fortran	<code>integer (kind=omp_event_handle_kind)</code>
Fortran	

The [event_handle](#) OpenMP type is an [opaque type](#) that represents [events](#) related to [detachable tasks](#).

26.7 OpenMP Interoperability Support Types

This section describes [OpenMP types](#) that support interoperability mechanisms.

26.7.1 OpenMP interop Type

Name: interop Properties: named-handle , omp , opaque	Base Type: implementation-defined-int
--	---

Predefined Identifiers

Name	Value	Properties
<code>omp_interop_none</code>	0	default

Type Definition

C / C++

```
typedef <implementation-defined-integer> omp_interop_t;
```

C / C++

Fortran

```
integer (kind=omp_interop_kind)
```

Fortran

The **interop** OpenMP type is an **opaque type** that represents OpenMP **interoperability objects**, which thus have the **opaque property**. **Interoperability objects** may be initialized, destroyed or otherwise used by an **interop** construct and may be initialized to **omp_interop_none**.

Cross References

- **interop** Construct, see [Section 22.1](#)

26.7.2 OpenMP interop_fr Type

Name: <code>interop_fr</code> Properties: <code>omp</code>	Base Type: <code>enumeration</code>
---	--

Values

Name	Value	Properties
<code>omp_ifr_last</code>	<code>N</code>	<code>omp</code>

Type Definition

C / C++

```
typedef enum omp_interop_fr_t {  
    omp_ifr_last = N  
} omp_interop_fr_t;
```

C / C++

Fortran

```
integer (kind=omp_interop_fr_kind), &  
    parameter :: omp_ifr_last = N
```

Fortran

The **interop_fr** OpenMP type represents supported **foreign runtime environments**. Each value of the **interop_fr** OpenMP type that an implementation provides 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 value of the supported **foreign runtime environments** that are listed in the aforementioned document or are **implementation defined**.

Cross References

- OpenMP Contexts, see [Section 15.1](#)
- `omp_get_num_devices` Routine, see [Section 30.3](#)

26.7.3 OpenMP `interop_property` Type

Name: <code>interop_property</code> Properties: <code>omp</code>	Base Type: enumeration
---	--

Values

Name	Value	Properties
<code>omp_ipr_fr_id</code>	-1	omp
<code>omp_ipr_fr_name</code>	-2	omp
<code>omp_ipr_vendor</code>	-3	omp
<code>omp_ipr_vendor_name</code>	-4	omp
<code>omp_ipr_device_num</code>	-5	omp
<code>omp_ipr_platform</code>	-6	omp
<code>omp_ipr_device</code>	-7	omp
<code>omp_ipr_device_context</code>	-8	omp
<code>omp_ipr_targetsync</code>	-9	omp
<code>omp_ipr_first</code>	-9	omp

Type Definition

C / C++

```
typedef enum omp_interop_property_t {  
    omp_ipr_fr_id           = -1,  
    omp_ipr_fr_name        = -2,  
    omp_ipr_vendor         = -3,  
    omp_ipr_vendor_name    = -4,  
    omp_ipr_device_num     = -5,  
    omp_ipr_platform       = -6,  
    omp_ipr_device         = -7,  
    omp_ipr_device_context = -8,  
    omp_ipr_targetsync     = -9,  
    omp_ipr_first          = -9  
} omp_interop_property_t;
```

C / C++

Fortran

```

1  integer (kind=omp_interop_property_kind), &
2      parameter :: omp_ipr_fr_id = -1
3  integer (kind=omp_interop_property_kind), &
4      parameter :: omp_ipr_fr_name = -2
5  integer (kind=omp_interop_property_kind), &
6      parameter :: omp_ipr_vendor = -3
7  integer (kind=omp_interop_property_kind), &
8      parameter :: omp_ipr_vendor_name = -4
9  integer (kind=omp_interop_property_kind), &
10     parameter :: omp_ipr_device_num = -5
11 integer (kind=omp_interop_property_kind), &
12     parameter :: omp_ipr_platform = -6
13 integer (kind=omp_interop_property_kind), &
14     parameter :: omp_ipr_device = -7
15 integer (kind=omp_interop_property_kind), &
16     parameter :: omp_ipr_device_context = -8
17 integer (kind=omp_interop_property_kind), &
18     parameter :: omp_ipr_targetsync = -9
19 integer (kind=omp_interop_property_kind), &
20     parameter :: omp_ipr_first = -9

```

Fortran

The **interop_property** OpenMP type is used in [interoperability routines](#) to represent [interoperability properties](#). OpenMP reserves all negative values for [interoperability properties](#), as listed in Table 26.2; [implementation defined interoperability properties](#) may use [non-negative](#) values. The special [interoperability property](#), **omp_ipr_first**, will always have the lowest **interop_property** value, which may change in future versions of this specification. Valid values and types for the [properties](#) that Table 26.2 lists are specified in the [OpenMP Additional Definitions document](#) or are [implementation defined](#) unless otherwise specified. The **Contexts** column of Table 26.2 lists the [OpenMP context](#) that is relevant to the value.

Cross References

- OpenMP Contexts, see [Section 15.1](#)
- **omp_get_num_devices** Routine, see [Section 30.3](#)

26.7.4 OpenMP interop_rc Type

Name: **interop_rc**
Properties: [omp](#)

Base Type: [enumeration](#)

TABLE 26.2: Required Values of the `interop_property` OpenMP Type

Enum Name	Contexts	Name	Property
<code>omp_ipr_fr_id</code>	all	<code>fr_id</code>	An <code>intptr_t</code> value that represents the foreign runtime environment ID of context
<code>omp_ipr_fr_name</code>	all	<code>fr_name</code>	C string value that represents the name of the foreign runtime environment of context
<code>omp_ipr_vendor</code>	all	<code>vendor</code>	An <code>intptr_t</code> that represents the vendor of context
<code>omp_ipr_vendor_name</code>	all	<code>vendor_name</code>	C string value that represents the vendor of context
<code>omp_ipr_device_num</code>	all	<code>device_num</code>	The OpenMP device number for the device in the range 0 to <code>omp_get_num_devices</code> inclusive
<code>omp_ipr_platform</code>	<i>target</i>	<code>platform</code>	A foreign platform handle usually spanning multiple devices
<code>omp_ipr_device</code>	<i>target</i>	<code>device</code>	A foreign device handle
<code>omp_ipr_device_context</code>	<i>target</i>	<code>device_context</code>	A handle to an instance of a foreign device context
<code>omp_ipr_targetsync</code>	<i>targetsync</i>	<code>targetsync</code>	A handle to a synchronization object of a foreign execution context

Values

Name	Value	Properties
<code>omp_irc_no_value</code>	1	omp
<code>omp_irc_success</code>	0	omp
<code>omp_irc_empty</code>	-1	omp
<code>omp_irc_out_of_range</code>	-2	omp
<code>omp_irc_type_int</code>	-3	omp
<code>omp_irc_type_ptr</code>	-4	omp
<code>omp_irc_type_str</code>	-5	omp
<code>omp_irc_other</code>	-6	omp

Type Definition

C / C++

```
typedef enum omp_interop_rc_t {
    omp_irc_no_value      = 1,
    omp_irc_success       = 0,
    omp_irc_empty         = -1,
    omp_irc_out_of_range  = -2,
    omp_irc_type_int      = -3,
```

TABLE 26.3: Required Values for the `interop_rc` OpenMP Type

Enum Name	Description
<code>omp_irc_no_value</code>	Valid but no meaningful value available
<code>omp_irc_success</code>	Successful, value is usable
<code>omp_irc_empty</code>	The provided interoperability object is equal to <code>omp_interop_none</code>
<code>omp_irc_out_of_range</code>	Property ID is out of range, see Table 26.2
<code>omp_irc_type_int</code>	Property type is <code>int</code> ; use <code>omp_get_interop_int</code>
<code>omp_irc_type_ptr</code>	Property type is pointer; use <code>omp_get_interop_ptr</code>
<code>omp_irc_type_str</code>	Property type is string; use <code>omp_get_interop_str</code>
<code>omp_irc_other</code>	Other error; use <code>omp_get_interop_rc_desc</code>

```
1      omp_irc_type_ptr      = -4,  
2      omp_irc_type_str     = -5,  
3      omp_irc_other        = -6  
4  } omp_interop_rc_t;
```

C / C++

Fortran

```
5  integer (kind=omp_interop_rc_kind), &  
6      parameter :: omp_irc_no_value = 1  
7  integer (kind=omp_interop_rc_kind), &  
8      parameter :: omp_irc_success = 0  
9  integer (kind=omp_interop_rc_kind), &  
10     parameter :: omp_irc_empty = -1  
11 integer (kind=omp_interop_rc_kind), &  
12     parameter :: omp_irc_out_of_range = -2  
13 integer (kind=omp_interop_rc_kind), &  
14     parameter :: omp_irc_type_int = -3  
15 integer (kind=omp_interop_rc_kind), &  
16     parameter :: omp_irc_type_ptr = -4  
17 integer (kind=omp_interop_rc_kind), &  
18     parameter :: omp_irc_type_str = -5  
19 integer (kind=omp_interop_rc_kind), &  
20     parameter :: omp_irc_other = -6
```

Fortran

The `interop_rc` OpenMP type is used in several [interoperability routines](#) to specify their results. Table 26.3 describes the values that this type must include.

Cross References

- OpenMP `interop` Type, see [Section 26.7.1](#)
- OpenMP `interop_property` Type, see [Section 26.7.3](#)
- `omp_get_interop_int` Routine, see [Section 32.2](#)
- `omp_get_interop_ptr` Routine, see [Section 32.3](#)
- `omp_get_interop_rc_desc` Routine, see [Section 32.7](#)
- `omp_get_interop_str` Routine, see [Section 32.4](#)

26.8 OpenMP Memory Management Types

This section describes OpenMP types that support memory management.

26.8.1 OpenMP access Type

Name: <code>access</code> Properties: <code>omp</code>	Base Type: <code>enumeration</code>
---	-------------------------------------

Values

Name	Value	Properties
<code>omp_access_cgroup</code>	0	<code>omp</code>
<code>omp_access_pteam</code>	1	<code>omp</code>

Type Definition

C / C++

```
typedef enum omp_access_t {  
    omp_access_cgroup = 0,  
    omp_access_pteam  = 1  
} omp_access_t;
```

C / C++

Fortran

```
integer (kind=omp_access_kind), &  
    parameter :: omp_access_cgroup = 0  
integer (kind=omp_access_kind), &  
    parameter :: omp_access_pteam = 1
```

Fortran

The `access` OpenMP type defines the values for the `access group` types. If the value is `omp_access_cgroup`, the `access group` type groups the `tasks` based on the `contention group` to which they belong. If the value is `omp_access_pteam`, the `access group` type groups the `tasks` based on the parallel `region` to which they bind.

26.8.2 OpenMP allocator_handle Type

Name: <code>allocator_handle</code> Properties: <code>omp</code>	Base Type: <code>enumeration</code>
---	-------------------------------------

Values

Name	Value	Properties
<code>omp_null_allocator</code>	0	<code>omp</code>
<code>omp_default_mem_alloc</code>	1	<code>omp</code>
<code>omp_large_cap_mem_alloc</code>	2	<code>omp</code>
<code>omp_const_mem_alloc</code>	3	<code>omp</code>
<code>omp_high_bw_mem_alloc</code>	4	<code>omp</code>
<code>omp_low_lat_mem_alloc</code>	5	<code>omp</code>
<code>omp_cgroup_mem_alloc</code>	6	<code>omp</code>
<code>omp_pteam_mem_alloc</code>	7	<code>omp</code>
<code>omp_thread_mem_alloc</code>	8	<code>omp</code>

Type Definition

C / C++

```
typedef enum omp_allocator_handle_t {  
    omp_null_allocator = 0,  
    omp_default_mem_alloc = 1,  
    omp_large_cap_mem_alloc = 2,  
    omp_const_mem_alloc = 3,  
    omp_high_bw_mem_alloc = 4,  
    omp_low_lat_mem_alloc = 5,  
    omp_cgroup_mem_alloc = 6,  
    omp_pteam_mem_alloc = 7,  
    omp_thread_mem_alloc = 8  
} omp_allocator_handle_t;
```

C / C++

Fortran

```
integer (kind=omp_allocator_handle_kind), &  
    parameter :: omp_null_allocator = 0  
integer (kind=omp_allocator_handle_kind), &  
    parameter :: omp_default_mem_alloc = 1  
integer (kind=omp_allocator_handle_kind), &  
    parameter :: omp_large_cap_mem_alloc = 2  
integer (kind=omp_allocator_handle_kind), &  
    parameter :: omp_const_mem_alloc = 3  
integer (kind=omp_allocator_handle_kind), &  
    parameter :: omp_high_bw_mem_alloc = 4  
integer (kind=omp_allocator_handle_kind), &  
    parameter :: omp_low_lat_mem_alloc = 5
```

```
1 integer (kind=omp_allocator_handle_kind), &
2   parameter :: omp_cgroup_mem_alloc = 6
3 integer (kind=omp_allocator_handle_kind), &
4   parameter :: omp_pteam_mem_alloc = 7
5 integer (kind=omp_allocator_handle_kind), &
6   parameter :: omp_thread_mem_alloc = 8
```

Fortran

7 The **allocator_handle** OpenMP type represents an **allocator** as described in Table 13.3. This
8 OpenMP type must be an **implementation defined** (for C++ possibly scoped) enum type and its
9 valid constants must include those shown above.

C++

10 The **omp.h** header file also defines a class template that models the **memory allocator** concept in
11 the **omp::allocator** namespace for each value of the **allocator_handle** OpenMP type.
12 The names in this class do not include either the **omp_** prefix or the **_alloc** suffix.

C++

26.8.3 OpenMP alloctrail Type

Name: alloctrail Properties: omp	Base Type: structure
---	-----------------------------

Fields

Name	Type	Properties
<i>key</i>	alloctrail_key	omp
<i>value</i>	alloctrail_val	omp

Type Definition

C / C++

```
18 typedef struct omp_alloctrail_t {
19     omp_alloctrail_key_t key;
20     omp_alloctrail_val_t value;
21 } omp_alloctrail_t;
```

C / C++

Fortran

```
22 ! omp_alloctrail might not be provided
23 ! in deprecated include file omp_lib.h
24 type omp_alloctrail
25     integer (kind=omp_alloctrail_key_kind) key
26     integer (kind=omp_alloctrail_val_kind) value
27 end type omp_alloctrail;
```

Fortran

TABLE 26.4: Allowed Key-Values for `alloctrait` OpenMP Type

Trait	Key	Allowed Values
<code>sync_hint</code>	<code>omp_atk_sync_hint</code>	<code>omp_atv_contended</code> , <code>omp_atv_uncontended</code> , <code>omp_atv_serialized</code> , <code>omp_atv_private</code>
<code>alignment</code>	<code>omp_atk_alignment</code>	Positive property integer powers of 2
<code>access</code>	<code>omp_atk_access</code>	<code>omp_atv_all</code> , <code>omp_atv_memspace</code> , <code>omp_atv_device</code> , <code>omp_atv_cgroup</code> , <code>omp_atv_pteam</code> , <code>omp_atv_thread</code>
<code>pool_size</code>	<code>omp_atk_pool_size</code>	Any positive property integer
<code>fallback</code>	<code>omp_atk_fallback</code>	<code>omp_atv_default_mem_fb</code> , <code>omp_atv_null_fb</code> , <code>omp_atv_abort_fb</code> , <code>omp_atv_allocator_fb</code>
<code>fb_data</code>	<code>omp_atk_fb_data</code>	An allocator handle
<code>pinned</code>	<code>omp_atk_pinned</code>	<code>omp_atv_true</code> , <code>omp_atv_false</code>
<code>partition</code>	<code>omp_atk_partition</code>	<code>omp_atv_environment</code> , <code>omp_atv_nearest</code> , <code>omp_atv_blocked</code> , <code>omp_atv_interleaved</code> , <code>omp_atv_partitioner</code>
<code>pin_device</code>	<code>omp_atk_pin_device</code>	Any conforming device number
<code>preferred_device</code>	<code>omp_atk_preferred_device</code>	Any conforming device number
<code>target_access</code>	<code>omp_atk_target_access</code>	<code>omp_atv_single</code> , <code>omp_atv_multiple</code>
<code>atomic_scope</code>	<code>omp_atk_atomic_scope</code>	<code>omp_atv_all</code> , <code>omp_atv_device</code>
<code>part_size</code>	<code>omp_atk_part_size</code>	Any positive property integer value

table continued on next page

table continued from previous page

Trait	Key	Allowed Values
partitioner	omp_atk_partitioner	A memory partitioner handle
partitioner_arg	omp_atk_partitioner_arg	Any integer value

The [alloctrait](#) OpenMP type is a key-value pair that represents the name of an [allocator trait](#), as the key, and its value (see [Table 26.4](#)).

Cross References

- Memory Allocators, see [Section 13.2](#)

26.8.4 OpenMP `alloctrait_key` Type

Name: <code>alloctrait_key</code> Properties: omp	Base Type: enumeration
--	--

Values

Name	Value	Properties
<code>omp_atk_sync_hint</code>	1	omp
<code>omp_atk_alignment</code>	2	omp
<code>omp_atk_access</code>	3	omp
<code>omp_atk_pool_size</code>	4	omp
<code>omp_atk_fallback</code>	5	omp
<code>omp_atk_fb_data</code>	6	omp
<code>omp_atk_pinned</code>	7	omp
<code>omp_atk_partition</code>	8	omp
<code>omp_atk_pin_device</code>	9	omp
<code>omp_atk_preferred_device</code>	10	omp
<code>omp_atk_device_access</code>	11	omp
<code>omp_atk_target_access</code>	12	omp
<code>omp_atk_atomic_scope</code>	13	omp
<code>omp_atk_part_size</code>	14	omp
<code>omp_atk_partitioner</code>	15	omp
<code>omp_atk_partitioner_arg</code>	16	omp

Type Definition

	C / C++
<pre>typedef enum omp_alloctrait_key_t { omp_atk_sync_hint = 1, omp_atk_alignment = 2, omp_atk_access = 3, omp_atk_pool_size = 4,</pre>	

```

1      omp_atk_fallback      = 5,
2      omp_atk_fb_data      = 6,
3      omp_atk_pinned       = 7,
4      omp_atk_partition    = 8,
5      omp_atk_pin_device   = 9,
6      omp_atk_preferred_device = 10,
7      omp_atk_device_access = 11,
8      omp_atk_target_access = 12,
9      omp_atk_atomic_scope = 13,
10     omp_atk_part_size    = 14,
11     omp_atk_partitioner  = 15,
12     omp_atk_partitioner_arg = 16
13 } omp_alloctrail_key_t;

```

C / C++

Fortran

```

14 integer (kind=omp_alloctrail_key_kind), &
15     parameter :: omp_atk_sync_hint = 1
16 integer (kind=omp_alloctrail_key_kind), &
17     parameter :: omp_atk_alignment = 2
18 integer (kind=omp_alloctrail_key_kind), &
19     parameter :: omp_atk_access = 3
20 integer (kind=omp_alloctrail_key_kind), &
21     parameter :: omp_atk_pool_size = 4
22 integer (kind=omp_alloctrail_key_kind), &
23     parameter :: omp_atk_fallback = 5
24 integer (kind=omp_alloctrail_key_kind), &
25     parameter :: omp_atk_fb_data = 6
26 integer (kind=omp_alloctrail_key_kind), &
27     parameter :: omp_atk_pinned = 7
28 integer (kind=omp_alloctrail_key_kind), &
29     parameter :: omp_atk_partition = 8
30 integer (kind=omp_alloctrail_key_kind), &
31     parameter :: omp_atk_pin_device = 9
32 integer (kind=omp_alloctrail_key_kind), &
33     parameter :: omp_atk_preferred_device = 10
34 integer (kind=omp_alloctrail_key_kind), &
35     parameter :: omp_atk_device_access = 11
36 integer (kind=omp_alloctrail_key_kind), &
37     parameter :: omp_atk_target_access = 12
38 integer (kind=omp_alloctrail_key_kind), &
39     parameter :: omp_atk_atomic_scope = 13
40 integer (kind=omp_alloctrail_key_kind), &
41     parameter :: omp_atk_part_size = 14

```

```
1 integer (kind=omp_alloctrail_key_kind), &
2     parameter :: omp_atk_partitioner = 15
3 integer (kind=omp_alloctrail_key_kind), &
4     parameter :: omp_atk_partitioner_arg = 16
```

▲ Fortran ▲

5 The [alloctrail_key](#) OpenMP type represents an [allocator trait](#) as described in Table 26.4.
6 The valid constants for this [OpenMP type](#) must include those shown above.

7 **Cross References**

- 8 • Memory Allocators, see [Section 13.2](#)

9 **26.8.5 OpenMP alloctrail_value Type**

Name: alloctrail_value Properties: omp	Base Type: enumeration
---	--

11 **Values**

Name	Value	Properties
omp_atv_default	-1	omp
omp_atv_false	0	omp
omp_atv_true	1	omp
omp_atv_contended	3	omp
omp_atv_uncontended	4	omp
omp_atv_serialized	5	omp
omp_atv_private	6	omp
omp_atv_device	7	omp
omp_atv_thread	8	omp
omp_atv_pteam	9	omp
omp_atv_cgroup	10	omp
omp_atv_default_mem_fb	11	omp
omp_atv_null_fb	12	omp
omp_atv_abort_fb	13	omp
omp_atv_allocator_fb	14	omp
omp_atv_environment	15	omp
omp_atv_nearest	16	omp
omp_atv_blocked	17	omp
omp_atv_interleaved	18	omp
omp_atv_all	19	omp
omp_atv_single	20	omp
omp_atv_multiple	21	omp
omp_atv_memspace	22	omp
omp_atv_partitioner	23	omp

Type Definition

C / C++

```
typedef enum omp_alloctrail_value_t {
    omp_atv_default      = -1,
    omp_atv_false        = 0,
    omp_atv_true         = 1,
    omp_atv_contended    = 3,
    omp_atv_uncontended  = 4,
    omp_atv_serialized   = 5,
    omp_atv_private      = 6,
    omp_atv_device       = 7,
    omp_atv_thread       = 8,
    omp_atv_pteam        = 9,
    omp_atv_cgroup       = 10,
    omp_atv_default_mem_fb = 11,
    omp_atv_null_fb      = 12,
    omp_atv_abort_fb     = 13,
    omp_atv_allocator_fb = 14,
    omp_atv_environment  = 15,
    omp_atv_nearest      = 16,
    omp_atv_blocked      = 17,
    omp_atv_interleaved  = 18,
    omp_atv_all          = 19,
    omp_atv_single       = 20,
    omp_atv_multiple     = 21,
    omp_atv_memspace     = 22,
    omp_atv_partitioner  = 23
} omp_alloctrail_value_t;
```

C / C++

Fortran

```
integer (kind=omp_alloctrail_value_kind), &
    parameter :: omp_atv_default = -1
integer (kind=omp_alloctrail_value_kind), &
    parameter :: omp_atv_false = 0
integer (kind=omp_alloctrail_value_kind), &
    parameter :: omp_atv_true = 1
integer (kind=omp_alloctrail_value_kind), &
    parameter :: omp_atv_contended = 3
integer (kind=omp_alloctrail_value_kind), &
    parameter :: omp_atv_uncontended = 4
integer (kind=omp_alloctrail_value_kind), &
    parameter :: omp_atv_serialized = 5
integer (kind=omp_alloctrail_value_kind), &
```

```

1      parameter :: omp_atv_private = 6
2      integer (kind=omp_alloctrail_value_kind), &
3      parameter :: omp_atv_device = 7
4      integer (kind=omp_alloctrail_value_kind), &
5      parameter :: omp_atv_thread = 8
6      integer (kind=omp_alloctrail_value_kind), &
7      parameter :: omp_atv_pteam = 9
8      integer (kind=omp_alloctrail_value_kind), &
9      parameter :: omp_atv_cgroup = 10
10     integer (kind=omp_alloctrail_value_kind), &
11     parameter :: omp_atv_default_mem_fb = 11
12     integer (kind=omp_alloctrail_value_kind), &
13     parameter :: omp_atv_null_fb = 12
14     integer (kind=omp_alloctrail_value_kind), &
15     parameter :: omp_atv_abort_fb = 13
16     integer (kind=omp_alloctrail_value_kind), &
17     parameter :: omp_atv_allocator_fb = 14
18     integer (kind=omp_alloctrail_value_kind), &
19     parameter :: omp_atv_environment = 15
20     integer (kind=omp_alloctrail_value_kind), &
21     parameter :: omp_atv_nearest = 16
22     integer (kind=omp_alloctrail_value_kind), &
23     parameter :: omp_atv_blocked = 17
24     integer (kind=omp_alloctrail_value_kind), &
25     parameter :: omp_atv_interleaved = 18
26     integer (kind=omp_alloctrail_value_kind), &
27     parameter :: omp_atv_all = 19
28     integer (kind=omp_alloctrail_value_kind), &
29     parameter :: omp_atv_single = 20
30     integer (kind=omp_alloctrail_value_kind), &
31     parameter :: omp_atv_multiple = 21
32     integer (kind=omp_alloctrail_value_kind), &
33     parameter :: omp_atv_memspace = 22
34     integer (kind=omp_alloctrail_value_kind), &
35     parameter :: omp_atv_partitioner = 23

```

Fortran

The [`alloctrail_value`](#) OpenMP type represents semantic values of [allocator traits](#) as described in Table 26.4. The valid constants for this OpenMP type must include those shown above.





Cross References

- Memory Allocators, see [Section 13.2](#)

26.8.6 OpenMP alloctrail_val Type

Name: <code>alloctrail_val</code> Properties: <code>omp</code>	Base Type: <code>intptr</code>
---	--------------------------------

Type Definition





	<code>typedef omp_intptr_t omp_alloctrail_val_t;</code>
	
	<code>integer (kind=c_intptr_t)</code>
	

The `alloctrail_val` OpenMP type represents the values that may be assigned to the `value` field of the `alloctrail_val` OpenMP type. Any of the semantic values of the `alloctrail_value` OpenMP type may be used for the `alloctrail_val` OpenMP type; in addition, other numeric values may be used for it as appropriate for the specified `key` of the `alloctrail` OpenMP type.

26.8.7 OpenMP mempartition Type

Name: <code>mempartition</code> Properties: <code>named-handle</code> , <code>omp</code> , <code>opaque</code>	Base Type: <code>opaque</code>
---	--------------------------------

Type Definition

	<code>typedef <implementation-defined> omp_mempartition_t;</code>
	
	<code>integer (kind=omp_mempartition_kind)</code>
	

The `mempartition` OpenMP type is an `opaque type` that represents `memory partitions`.

26.8.8 OpenMP mempartitioner Type

Name: <code>mempartitioner</code> Properties: <code>named-handle</code> , <code>omp</code> , <code>opaque</code>	Base Type: <code>opaque</code>
---	--------------------------------

1 **Type Definition**

C / C++

```
2       typedef <implementation-defined> omp_mempartitioner_t;
```

C / C++

Fortran

```
3       integer (kind=omp_mempartitioner_kind)
```

Fortran

4 The **mempartitioner** OpenMP type is an **opaque type** that represents **memory partitioners**.

5 **26.8.9 OpenMP mempartitioner_lifetime Type**

6 Name: mempartitioner_lifetime	Base Type: enumeration
6 Properties: omp	

7 **Values**

7 Name	Value	Properties
8 omp_static_mempartition	1	omp
8 omp_allocator_mempartition	2	omp
8 omp_dynamic_mempartition	3	omp

9 **Type Definition**

C / C++

```
10       typedef enum omp_mempartitioner_lifetime_t {
```

```
11           omp_static_mempartition     = 1,
```

```
12           omp_allocator_mempartition = 2,
```

```
13           omp_dynamic_mempartition    = 3
```

```
14       } omp_mempartitioner_lifetime_t;
```

C / C++

Fortran

```
15       integer (kind=omp_mempartitioner_lifetime_kind), &
```

```
16           parameter :: omp_static_mempartition = 1
```

```
17       integer (kind=omp_mempartitioner_lifetime_kind), &
```

```
18           parameter :: omp_allocator_mempartition = 2
```

```
19       integer (kind=omp_mempartitioner_lifetime_kind), &
```

```
20           parameter :: omp_dynamic_mempartition = 3
```

Fortran

21 The **mempartitioner_lifetime** OpenMP type represents the lifetime of a **memory**

22 **partitioner**. The valid constants for the **mempartitioner_lifetime** OpenMP type must

23 include those shown above.

26.8.10 OpenMP mempartitioner_compute_proc Type

Name: <code>mempartitioner_compute_proc</code>	Properties: <code>iso_c_binding</code> , <code>omp</code>
Category: <code>subroutine</code> pointer	

Arguments

Name	Type	Properties
<code>memspace</code>	<code>memspace_handle</code>	<code>omp</code>
<code>allocation_size</code>	<code>c_size_t</code>	<code>iso_c</code> , <code>value</code>
<code>partitioner_arg</code>	<code>alloctrait_val</code>	<code>omp</code> , <code>value</code>
<code>partition</code>	<code>mempartition</code>	<code>C/C++ pointer</code> , <code>omp</code>

Type Signature

C / C++

```
typedef void (*omp_mempartitioner_compute_proc_t) (  
    omp_memspace_handle_t memspace, size_t allocation_size,  
    omp_alloctrait_val_t partitioner_arg,  
    omp_mempartition_t *partition);
```

C / C++

Fortran

```
abstract interface  
    subroutine omp_mempartitioner_compute_proc_t(memspace, &  
        allocation_size, partitioner_arg, partition) bind(c)  
        use, intrinsic :: iso_c_binding, only : c_size_t  
        integer (kind=omp_memspace_handle_kind) memspace  
        integer (kind=c_size_t), value :: allocation_size  
        integer (kind=omp_alloctrait_val_kind), value :: &  
            partitioner_arg  
        integer (kind=omp_mempartition_kind) partition  
    end subroutine  
end interface
```

Fortran

The `mempartitioner_compute_proc` OpenMP type represents a partition computation procedure. When used through the `omp_init_mempartition` and `omp_mempartition_set_part` routines, the procedure will be passed the following arguments in the listed order:

- The memory space associated with the allocator to be used for the memory allocation;
- The size of the allocation in bytes;
- If the `omp_atk_partitioner_arg` trait was specified for the allocator, its specified value, otherwise, the value zero; and
- A memory partition object to be initialized

If the sum of the sizes of the parts specified in the [memory partition](#) object after executing the [procedure](#) is not equal to the *allocation_size* argument, the behavior is unspecified.

If the associated [memory partitioner](#) has been created with a call to [omp_init_mempartitioner](#) with the value of the *lifetime* argument set to [omp_static_mempartition](#) then the [memory partition](#) object computed by an invocation to the [procedure](#) might be used for the allocations of any [allocators](#) that have the *partitioner memory partitioner* object associated with them if the allocations have the same size and the same [memory space](#). The number of times that the *compute_proc procedure* is invoked is unspecified.

Cross References

- OpenMP `alloctrait_val` Type, see [Section 26.8.6](#)
- OpenMP `mempartition` Type, see [Section 26.8.7](#)
- OpenMP `memspace_handle` Type, see [Section 26.8.12](#)
- `omp_init_mempartition` Routine, see [Section 33.5.3](#)
- `omp_mempartition_set_part` Routine, see [Section 33.5.5](#)

26.8.11 OpenMP `mempartitioner_release_proc` Type

Name: <code>mempartitioner_release_proc</code>	Properties: iso_c_binding , omp
Category: subroutine pointer	

Arguments

Name	Type	Properties
<i>partition</i>	<code>mempartition</code>	C/C++ pointer , omp

Type Signature

C / C++

```
typedef void (*omp_mempartitioner_release_proc_t) (  
    omp_mempartition_t *partition);
```

C / C++

Fortran

```
abstract interface  
    subroutine omp_mempartitioner_release_proc_t(partition) &  
        bind(c)  
        integer (kind=omp_mempartition_kind) partition  
    end subroutine  
end interface
```

Fortran

The `mempartitioner_release_proc` OpenMP type represents a partition release procedure. When an implementation finishes using a `memory partition` object that was created with the procedure used as the `compute_proc` argument for a call to the `omp_init_mempartitioner` routine to which the represented release procedure was the `release_proc` argument, that release procedure will be called with the `memory partition` object as its argument. The procedure can then release the object and its resources using the `omp_destroy_mempartition` routine. The implementation will invoke the `release_proc` at most once for each `memory partition` object.

Cross References

- OpenMP `mempartition` Type, see [Section 26.8.7](#)
- `omp_init_mempartitioner` Routine, see [Section 33.5.1](#)

26.8.12 OpenMP `memspace_handle` Type

Name: <code>memspace_handle</code> Properties: <code>omp</code>	Base Type: <code>enumeration</code>
--	-------------------------------------

Values

Name	Value	Properties
<code>omp_default_mem_space</code>	-1	<code>omp</code>
<code>omp_null_mem_space</code>	0	<code>omp</code>
<code>omp_large_cap_mem_space</code>	1	<code>omp</code>
<code>omp_const_mem_space</code>	2	<code>omp</code>
<code>omp_high_bw_mem_space</code>	3	<code>omp</code>
<code>omp_low_lat_mem_space</code>	4	<code>omp</code>

Type Definition

C / C++

```
typedef enum omp_memspace_handle_t {
    omp_default_mem_space = -1,
    omp_null_mem_space    = 0,
    omp_large_cap_mem_space = 1,
    omp_const_mem_space    = 2,
    omp_high_bw_mem_space  = 3,
    omp_low_lat_mem_space  = 4
} omp_memspace_handle_t;
```

C / C++

Fortran

```
integer (kind=omp_memspace_handle_kind), &  
    parameter :: omp_default_mem_space = -1  
integer (kind=omp_memspace_handle_kind), &  
    parameter :: omp_null_mem_space = 0  
integer (kind=omp_memspace_handle_kind), &  
    parameter :: omp_large_cap_mem_space = 1  
integer (kind=omp_memspace_handle_kind), &  
    parameter :: omp_const_mem_space = 2  
integer (kind=omp_memspace_handle_kind), &  
    parameter :: omp_high_bw_mem_space = 3  
integer (kind=omp_memspace_handle_kind), &  
    parameter :: omp_low_lat_mem_space = 4
```

Fortran

The **memspace_handle** OpenMP type represents an **allocator** as described in Table 13.1. This OpenMP type must be an **implementation defined** (for C++ possibly scoped) enum type and its valid constants must include those shown above.

26.9 OpenMP Synchronization Types

This section describes OpenMP types related to synchronization, including **locks**.

26.9.1 OpenMP depend Type

Name: depend	Base Type:
Properties: named-handle , omp , opaque	implementation-defined-int

Type Definition

C / C++

```
typedef <implementation-defined-integral> omp_depend_t;
```

C / C++

Fortran

```
integer (kind=omp_depend_kind)
```

Fortran

The **depend** OpenMP type is an **opaque type** that represents **depend objects**.

26.9.2 OpenMP impex Type

Name: impex	Base Type: enumeration
Properties: omp	

1 **Values**

Name	Value	Properties
<code>omp_not_impex</code>	0	omp
<code>omp_import</code>	1	omp
<code>omp_export</code>	2	omp
<code>omp_impex</code>	3	omp

3 **Type Definition**

C / C++

```
typedef enum omp_impex_t {
    omp_not_impex = 0,
    omp_import     = 1,
    omp_export     = 2,
    omp_impex      = 3
} omp_impex_t;
```

C / C++

Fortran

```
integer (kind=omp_impex_kind), &
    parameter :: omp_not_impex = 0
integer (kind=omp_impex_kind), &
    parameter :: omp_import = 1
integer (kind=omp_impex_kind), &
    parameter :: omp_export = 2
integer (kind=omp_impex_kind), &
    parameter :: omp_impex = 3
```

Fortran

18 The [impex](#) OpenMP type is an [enumeration](#) type that is used to specify whether the [child tasks](#) of
19 a [task](#) may form a [task dependence](#) with respect to its [dependence-compatible tasks](#). In particular, it
20 is used to identify whether a [task](#) is an [importing task](#) and/or an [exporting task](#). The valid constants
21 must include those shown above.

22 **Cross References**

- `transparent` Clause, see [Section 23.9.6](#)

24 **26.9.3 OpenMP lock Type**

Name: <code>lock</code> Properties : <code>named-handle</code> , <code>opaque</code>	Base Type: <code>opaque</code>
---	--------------------------------

Type Definition

C / C++

```
typedef <implementation-defined> omp_lock_t;
```

C / C++

Fortran

```
integer (kind=omp_lock_kind)
```

Fortran

The **lock** OpenMP type is an **opaque type** that represents **simple locks** used in **simple lock routines**.

26.9.4 OpenMP nest_lock Type

Name: nest_lock Properties: named-handle, opaque	Base Type: opaque
---	--------------------------

Type Definition

C / C++

```
typedef <implementation-defined> omp_nest_lock_t;
```

C / C++

Fortran

```
integer (kind=omp_nest_lock_kind)
```

Fortran

The **nest_lock** OpenMP type is an **opaque type** that represents **nestable locks** used in **nestable lock routines**.

26.9.5 OpenMP sync_hint Type

Name: sync_hint Properties: omp	Base Type: enumeration
--	-------------------------------

Values

Name	Value	Properties
omp_sync_hint_none	0x0	omp
omp_sync_hint_uncontended	0x1	omp
omp_sync_hint_contended	0x2	omp
omp_sync_hint_nonspeculative	0x4	omp
omp_sync_hint_speculative	0x8	omp

Type Definition

C / C++

```
typedef enum omp_sync_hint_t {  
    omp_sync_hint_none          = 0x0,  
    omp_sync_hint_uncontended   = 0x1,  
    omp_sync_hint_contended     = 0x2,  
    omp_sync_hint_nonspeculative = 0x4,  
    omp_sync_hint_speculative    = 0x8  
} omp_sync_hint_t;
```

C / C++

Fortran

```
integer (kind=omp_sync_hint_kind), &  
    parameter :: omp_sync_hint_none = &  
        int(Z'0', kind=omp_sync_hint_kind)  
integer (kind=omp_sync_hint_kind), &  
    parameter :: omp_sync_hint_uncontended = &  
        int(Z'1', kind=omp_sync_hint_kind)  
integer (kind=omp_sync_hint_kind), &  
    parameter :: omp_sync_hint_contended = &  
        int(Z'2', kind=omp_sync_hint_kind)  
integer (kind=omp_sync_hint_kind), &  
    parameter :: omp_sync_hint_nonspeculative = &  
        int(Z'4', kind=omp_sync_hint_kind)  
integer (kind=omp_sync_hint_kind), &  
    parameter :: omp_sync_hint_speculative = &  
        int(Z'8', kind=omp_sync_hint_kind)
```

Fortran

The **sync_hint** OpenMP type is used to specify **synchronization hints**. The **omp_init_lock_with_hint** and **omp_init_nest_lock_with_hint** routines provide hints about the expected dynamic behavior or suggested implementation of a **lock**. **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** or **routine**; if ignoring the hint changes the program semantics, the result is **unspecified**.

Synchronization hints can be combined by using the **+** or **|** operators in C/C++ or the **+** operator in Fortran. Combining **omp_sync_hint_none** with any other **synchronization hint** is equivalent to specifying the other **synchronization hint**.

The intended meaning of each **synchronization hint** is:

- **omp_sync_hint_uncontended**: low contention is expected in this operation, that is, few **threads** are expected to perform the operation simultaneously in a manner that requires synchronization;

- **omp_sync_hint_contended**: high contention is expected in this operation, that is, many **threads** are expected to perform the operation simultaneously in a manner that requires synchronization;
- **omp_sync_hint_speculative**: the programmer suggests that the operation should be 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** OpenMP 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 **omp_sync_hint_uncontended** and **omp_sync_hint_contended** values may not be combined.
- The **omp_sync_hint_nonspeculative** and **omp_sync_hint_speculative** values may not be combined.

Cross References

- **atomic** Construct, see [Section 23.8.5](#)
- **critical** Construct, see [Section 23.2](#)
- **hint** Clause, see [Section 23.1](#)
- **omp_init_lock_with_hint** Routine, see [Section 34.1.3](#)
- **omp_init_nest_lock_with_hint** Routine, see [Section 34.1.4](#)

26.10 OpenMP Affinity Support Types

This section describes **OpenMP types** that support affinity mechanisms.

26.10.1 OpenMP **proc_bind** Type

Name: proc_bind	Base Type: enumeration
Properties: omp	

1 **Values**

Name	Value	Properties
<code>omp_proc_bind_false</code>	0	omp
<code>omp_proc_bind_true</code>	1	omp
<code>omp_proc_bind_primary</code>	2	omp
<code>omp_proc_bind_close</code>	3	omp
<code>omp_proc_bind_spread</code>	4	omp

3 **Type Definition**

C / C++

```
4 typedef enum omp_proc_bind_t {  
5     omp_proc_bind_false    = 0,  
6     omp_proc_bind_true     = 1,  
7     omp_proc_bind_primary  = 2,  
8     omp_proc_bind_close    = 3,  
9     omp_proc_bind_spread   = 4  
10 } omp_proc_bind_t;
```

C / C++

Fortran

```
11 integer (kind=omp_proc_bind_kind), &  
12     parameter :: omp_proc_bind_false = 0  
13 integer (kind=omp_proc_bind_kind), &  
14     parameter :: omp_proc_bind_true = 1  
15 integer (kind=omp_proc_bind_kind), &  
16     parameter :: omp_proc_bind_primary = 2  
17 integer (kind=omp_proc_bind_kind), &  
18     parameter :: omp_proc_bind_close = 3  
19 integer (kind=omp_proc_bind_kind), &  
20     parameter :: omp_proc_bind_spread = 4
```

Fortran

21 The [proc_bind](#) OpenMP type is used in [routines](#) that modify or retrieve the value of the *bind-var*
22 [ICV](#). The valid constants for the [proc_bind](#) type must include those shown above.

23 **Cross References**

- *bind-var* ICV, see [Table 3.1](#)

26.11 OpenMP Resource Relinquishing Types

This section describes OpenMP types related to resource-relinquishing routines.

26.11.1 OpenMP pause_resource Type

Name: <code>pause_resource</code> Properties: <code>omp</code>	Base Type: <code>enumeration</code>
---	-------------------------------------

Values

Name	Value	Properties
<code>omp_pause_soft</code>	1	<code>omp</code>
<code>omp_pause_hard</code>	2	<code>omp</code>
<code>omp_pause_stop_tool</code>	3	<code>omp</code>

Type Definition

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 `pause_resource` OpenMP type is used in resource-relinquishing routines to specify the resources that the instance of the routine relinquishes. The valid constants for the `pause_resource` OpenMP type must include those shown above.

When specified and successful, the `omp_pause_hard` value results in a **hard pause**, which implies that the OpenMP state is not guaranteed to persist across the resource-relinquishing routine call. A **hard pause** may relinquish any data allocated by OpenMP on specified devices, including data allocated by device memory routines as well as data present on the devices as a result of a `declare target` directive or map-entering constructs. A **hard pause** may also relinquish any data associated with a `threadprivate` directive. When relinquished and when applicable, base language appropriate deallocation/finalization is performed. When relinquished and when applicable, mapped variables on a device will not be copied back from the device to the host device.

When specified and successful, the `omp_pause_soft` value results in a `soft pause` for which the OpenMP state is guaranteed to persist across the `resource-relinquishing routine` call, with the exception of any data associated with a `threadprivate` directive, which may be relinquished across the call. When relinquished and when applicable, `base language` appropriate deallocation/finalization is performed.

Note – A `hard pause` may relinquish more resources, but may resume processing `regions` more slowly. A `soft pause` allows `regions` to restart more quickly, but may relinquish fewer resources. An OpenMP implementation will reclaim resources as needed for `regions` encountered after the `resource-relinquishing routine region`. Since a `hard pause` may unmap data on the specified `devices`, appropriate `mapping operations` are required before using data on the specified `devices` after the `resource-relinquishing routine region`.

When specified and successful, the `omp_pause_stop_tool` value implies the effects described above for the `omp_pause_hard` value. Additionally, unless otherwise specified, the value implies that the implementation will shutdown the `OMPT` interface as if program execution is ending.

26.12 OpenMP Tool Types

This section describes `OpenMP types` that support the use of `tools`.

26.12.1 OpenMP `control_tool` Type

Name: <code>control_tool</code> Properties: <code>omp</code>	Base Type: <code>enumeration</code>
---	--

Values

Name	Value	Properties
<code>omp_control_tool_start</code>	1	<code>omp</code>
<code>omp_control_tool_pause</code>	2	<code>omp</code>
<code>omp_control_tool_flush</code>	3	<code>omp</code>
<code>omp_control_tool_end</code>	4	<code>omp</code>
<code>omp_control_tool_max</code>	<code>INT32_MAX</code>	<code>omp</code>

Type Definition

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,  
}
```

1 omp_control_tool_max = INT32_MAX

2 } omp_control_tool_t;

C / C++

Fortran

3 integer (kind=omp_control_tool_kind), &

4 parameter :: omp_control_tool_start = 1

5 integer (kind=omp_control_tool_kind), &

6 parameter :: omp_control_tool_pause = 2

7 integer (kind=omp_control_tool_kind), &

8 parameter :: omp_control_tool_flush = 3

9 integer (kind=omp_control_tool_kind), &

10 parameter :: omp_control_tool_end = 4

11 integer (kind=omp_control_tool_kind), &

12 parameter :: omp_control_tool_max = INT32_MAX

Fortran

The `control_tool` OpenMP type is used in `tool` support routines to specify `tool` commands. Table 26.5 describes the actions that standard commands request from a `tool`. The valid constants for the `control_tool` OpenMP type must include those shown above.

Tool-defined values for the `control_tool` OpenMP type must be greater than or equal to 64 and less than or equal to 2147483647 (`INT32_MAX`). Tools must ignore `control_tool` values that they are not explicitly designed to handle. Other values accepted by a `tool` for the `control_tool` OpenMP type are `tool` defined.

TABLE 26.5: 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 <code>tool</code> . This command may be applied whether monitoring is on or off.
<code>omp_control_tool_end</code>	Turn monitoring off permanently; the <code>tool</code> finalizes itself and flushes all output.

26.12.2 OpenMP control_tool_result Type

Name: control_tool_result Properties: omp	Base Type: enumeration
--	------------------------

Values

Name	Value	Properties
omp_control_tool_notool	-2	omp
omp_control_tool_nocallback	-1	omp
omp_control_tool_success	0	omp
omp_control_tool_ignored	1	omp

Type Definition

C / C++

```
typedef enum omp_control_tool_result_t {  
    omp_control_tool_notool      = -2,  
    omp_control_tool_nocallback = -1,  
    omp_control_tool_success     = 0,  
    omp_control_tool_ignored     = 1  
} omp_control_tool_result_t;
```

C / C++

Fortran

```
integer (kind=omp_control_tool_result_kind), &  
    parameter :: omp_control_tool_notool = -2  
integer (kind=omp_control_tool_result_kind), &  
    parameter :: omp_control_tool_nocallback = -1  
integer (kind=omp_control_tool_result_kind), &  
    parameter :: omp_control_tool_success = 0  
integer (kind=omp_control_tool_result_kind), &  
    parameter :: omp_control_tool_ignored = 1
```

Fortran

The **control_tool_result** OpenMP type is used in **tool** support routines to specify the results of **tool** commands. The valid constants for the **control_tool_result** OpenMP type must include those shown above.

27 Parallel Region Support Routines

This chapter describes routines that support execution of parallel regions, including routines to determine the number of OpenMP threads for parallel regions and that query the nesting of parallel regions at runtime.

27.1 omp_set_num_threads Routine

Name: omp_set_num_threads Category: subroutine	Properties: ICV-modifying
--	---

Arguments

Name	Type	Properties
<i>num_threads</i>	integer	positive

Prototypes

C / C++	<pre>void omp_set_num_threads(int num_threads);</pre>
C / C++	
Fortran	<pre>subroutine omp_set_num_threads(num_threads) integer num_threads</pre>
Fortran	

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. Thus, the routine has the ICV modifying property, through which it affects the number of threads to be used for subsequent parallel regions that do not specify a **num_threads** clause.

Cross References

- *nthreads-var* ICV, see [Table 3.1](#)
- **num_threads** Clause, see [Section 18.1.2](#)
- **parallel** Construct, see [Section 18.1](#)
- Determining the Number of Threads for a **parallel** Region, see [Section 18.1.1](#)



27.2 omp_get_num_threads Routine

Name: <code>omp_get_num_threads</code> Category: <code>function</code>	Properties: <i>default</i>
---	----------------------------

Return Type

Name	Type	Properties
<code><return type></code>	integer	<i>default</i>

Prototypes

	<code>int omp_get_num_threads(void);</code>
	<code>integer function omp_get_num_threads()</code>

Effect

The `omp_get_num_threads` routine returns the number of threads in the team that is executing the parallel region to which the routine region binds.



27.3 omp_get_thread_num Routine

Name: <code>omp_get_thread_num</code> Category: <code>function</code>	Properties: <i>default</i>
--	----------------------------

Return Type

Name	Type	Properties
<code><return type></code>	integer	<i>default</i>

Prototypes

	<code>int omp_get_thread_num(void);</code>
	<code>integer function omp_get_thread_num()</code>

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

- Predefined Identifiers, see [Section 26.1](#)
- `omp_get_num_threads` Routine, see [Section 27.2](#)









27.4 `omp_get_max_threads` Routine

Name: <code>omp_get_max_threads</code>	Properties: <code>ICV-retrieving</code>
Category: <code>function</code>	

Return Type

Name	Type	Properties
<code><return type></code>	integer	<i>default</i>

Prototypes

	
<code>int omp_get_max_threads(void);</code>	
	
	
<code>integer function omp_get_max_threads()</code>	
	

Effect

The value returned by `omp_get_max_threads` is the value of the first element of the `nthreads-var ICV` of the `current task`; thus, the `routine` has the `ICV retrieving property`. Its return value is 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`.

Cross References

- `nthreads-var ICV`, see [Table 3.1](#)
- `num_threads` Clause, see [Section 18.1.2](#)
- `parallel` Construct, see [Section 18.1](#)
- Determining the Number of Threads for a `parallel` Region, see [Section 18.1.1](#)

27.5 omp_get_thread_limit Routine

Name: <code>omp_get_thread_limit</code> Category: function	Properties: ICV-retrieving
---	--

Return Type

Name	Type	Properties
<code><return type></code>	integer	default

Prototypes

C / C++	<code>int omp_get_thread_limit(void);</code>
C / C++	
Fortran	<code>integer function omp_get_thread_limit()</code>
Fortran	

Effect

The `omp_get_thread_limit` routine returns the value of the *thread-limit-var* ICV. Thus, it returns the maximum number of threads available to execute tasks in the current contention group.

Cross References

- thread-limit-var* ICV, see [Table 3.1](#)

27.6 omp_in_parallel Routine

Name: <code>omp_in_parallel</code> Category: function	Properties: default
--	-------------------------------------

Return Type

Name	Type	Properties
<code><return type></code>	logical	default

Prototypes

C / C++	<code>int omp_in_parallel(void);</code>
C / C++	
Fortran	<code>logical function omp_in_parallel()</code>
Fortran	

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*. That is, it returns *true* if the *active-levels-var* ICV is greater than zero. Otherwise, it returns *false*.

Cross References

- *active-levels-var* ICV, see [Table 3.1](#)
- `parallel` Construct, see [Section 18.1](#)

27.7 omp_set_dynamic Routine

Name: <code>omp_set_dynamic</code>	Properties: ICV-modifying
Category: <code>subroutine</code>	

Arguments

Name	Type	Properties
<i>dynamic_threads</i>	logical	<i>default</i>

Prototypes

C / C++
▼
void omp_set_dynamic(int dynamic_threads);
▲

C / C++
▼
Fortran
▼
subroutine omp_set_dynamic(dynamic_threads)
 logical dynamic_threads
▲
Fortran
▲

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* by setting the value of the *dyn-var* ICV to *true*; otherwise, dynamic adjustment is disabled for the *current task* by setting the value of the *dyn-var* ICV to *false*. 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 3.1](#)

27.8 omp_get_dynamic Routine

Name: <code>omp_get_dynamic</code>	Properties: ICV-retrieving
Category: function	

Return Type

Name	Type	Properties
<code><return type></code>	logical	default

Prototypes

C / C++	<code>int omp_get_dynamic(void);</code>
C / C++	
Fortran	<code>logical function omp_get_dynamic()</code>
Fortran	

Effect

The `omp_get_dynamic` routine returns the value of the *dyn-var* ICV. Thus, 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 3.1](#)

27.9 omp_set_schedule Routine

Name: <code>omp_set_schedule</code>	Properties: ICV-modifying
Category: subroutine	

Arguments

Name	Type	Properties
<i>kind</i>	sched	omp
<i>chunk_size</i>	integer	default

1

2

3

4

5

6

7

8

0

1

2

3

4

5

6

- 7

9

20

21

22

1 **Prototypes**

C / C++

2 **void omp_get_schedule(omp_sched_t *kind, int *chunk_size);**

C / C++

Fortran

3 **subroutine omp_get_schedule(kind, chunk_size)**
4 integer (kind=omp_sched_kind) kind
5 integer chunk_size

Fortran

6 **Effect**

7 The **omp_get_schedule** routine returns the *run-sched-var* ICV in the task to which the routine
8 binds. Thus, the routine returns the schedule that is applied when the **runtime** schedule type is
9 used. The first argument *kind* returns the *schedule type* to be used. If the returned *schedule type* is
10 **omp_sched_static**, **omp_sched_dynamic**, or **omp_sched_guided**, the second
11 argument, *chunk_size*, returns the *chunk* size to be used, or a value less than 1 if the default *chunk*
12 size is to be used. The value returned by the second argument is *implementation defined* for any
13 other *schedule types*.

14 **Cross References**

- *run-sched-var* ICV, see [Table 3.1](#)
- OpenMP **sched** Type, see [Section 26.5.1](#)

17 **27.11 omp_get_supported_active_levels**
18 **Routine**

Name: omp_get_supported_active_levels Category: function	Properties: <i>default</i>
---	-----------------------------------

20 **Return Type**

Name	Type	Properties
<return type>	integer	<i>default</i>

22 **Prototypes**

C / C++

23 **int omp_get_supported_active_levels(void);**

C / C++

Fortran

24 **integer function omp_get_supported_active_levels()**

Fortran

Effect

The `omp_get_supported_active_levels` routine returns the number of supported active levels. 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 3.1

27.12 omp_set_max_active_levels Routine

Name: <code>omp_set_max_active_levels</code>	Properties: ICV-modifying
Category: subroutine	

Arguments

Name	Type	Properties
<i>max_levels</i>	integer	non-negative

Prototypes

C / C++	
<code>void omp_set_max_active_levels(int max_levels);</code>	
C / C++	
Fortran	
<code>subroutine omp_set_max_active_levels(max_levels)</code>	
<code>integer max_levels</code>	
Fortran	

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. Thus, the routine limits the number of nested active parallel regions when a new nested parallel region is generated by the current task.

If the number of active levels requested exceeds the number of supported active levels, the value of the *max-active-levels-var* ICV will be set to the number of supported active levels. 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

- *active-levels-var* ICV, see Table 3.1
- *max-active-levels-var* ICV, see Table 3.1
- `parallel` Construct, see Section 18.1

27.13 omp_get_max_active_levels Routine

Name: <code>omp_get_max_active_levels</code> Category: function	Properties: ICV-retrieving
--	--

Return Type

Name	Type	Properties
<code><return type></code>	integer	default

Prototypes

C / C++	<code>int omp_get_max_active_levels(void);</code>
C / C++	
Fortran	<code>integer function omp_get_max_active_levels()</code>
Fortran	

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 3.1](#)

27.14 omp_get_level Routine

Name: <code>omp_get_level</code> Category: function	Properties: ICV-retrieving
--	--

Return Type

Name	Type	Properties
<code><return type></code>	integer	default

Prototypes

C / C++	<code>int omp_get_level(void);</code>
C / C++	
Fortran	<code>integer function omp_get_level()</code>
Fortran	

Effect

The `omp_get_level` routine returns the value of the *levels-var* ICV. Thus, its effect is to return the number of nested `parallel` regions (whether *active* or *inactive*) that enclose the *current task* such that all of the `parallel` regions are enclosed by the outermost *initial task region* on the *current device*.

Cross References

- *levels-var* ICV, see [Table 3.1](#)
- `parallel` Construct, see [Section 18.1](#)









27.15 `omp_get_ancestor_thread_num` Routine

Name: <code>omp_get_ancestor_thread_num</code>	Properties: <i>default</i>
Category: <i>function</i>	

Return Type and Arguments

Name	Type	Properties
<return type>	integer	<i>default</i>
<i>level</i>	integer	<i>default</i>

Prototypes

	
<code>int omp_get_ancestor_thread_num(int <i>level</i>) ;</code>	
	
	
<code>integer function omp_get_ancestor_thread_num(<i>level</i>)</code>	
<code>integer <i>level</i></code>	
	

Effect

The `omp_get_ancestor_thread_num` routine returns the *thread number* of the *ancestor thread* at a given nest level of the *encountering thread* or the *thread number* of the *encountering thread*. If the requested nest level is outside the range of 0 and the nest level of the *encountering thread*, as returned by the `omp_get_level` routine, the *routine* returns -1.

Note – When the `omp_get_ancestor_thread_num` routine is called with value of *level* =0, the *routine* always returns 0. If *level* =`omp_get_level()`, the *routine* has the same effect as the `omp_get_thread_num` routine.

Cross References

- `omp_get_level` Routine, see [Section 27.14](#)
- `omp_get_thread_num` Routine, see [Section 27.3](#)

27.16 `omp_get_team_size` Routine

Name: <code>omp_get_team_size</code> Category: function	Properties: <i>default</i>
--	----------------------------

Return Type and Arguments

Name	Type	Properties
<return type>	integer	<i>default</i>
<i>level</i>	integer	<i>default</i>

Prototypes

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

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

- `omp_get_level` Routine, see [Section 27.14](#)
- `omp_get_num_threads` Routine, see [Section 27.2](#)









27.17 omp_get_active_level Routine

Name: <code>omp_get_active_level</code> Category: function	Properties: ICV-retrieving
---	--

Return Type

Name	Type	Properties
<code><return type></code>	integer	default

Prototypes

	
<code>int omp_get_active_level(void);</code>	
	
	
<code>integer function omp_get_active_level()</code>	
	

Effect

The effect of the `omp_get_active_level` routine is to return the number of nested [active parallel regions](#) that enclose the [current task](#) such that all [parallel regions](#) are enclosed by the outermost [initial task region](#) on the [current device](#). Thus, the routine returns the value of the [active-levels-var](#) ICV.

Cross References

- [active-levels-var](#) ICV, see [Table 3.1](#)
- `parallel` Construct, see [Section 18.1](#)

28 Teams Region Routines

This chapter describes routines that affect and monitor the league of teams that may execute a teams region.

28.1 omp_get_num_teams Routine

Name: <code>omp_get_num_teams</code> Category: function	Properties: ICV-retrieving , teams-nestable
--	---

Return Type

Name	Type	Properties
<code><return type></code>	integer	default

Prototypes

<code>int omp_get_num_teams(void);</code>	C / C++
<code>integer function omp_get_num_teams()</code>	Fortran

Effect

The `omp_get_num_teams` routine returns the value of the *league-size-var* ICV, which is 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

- *league-size-var* ICV, see [Table 3.1](#)
- **teams** Construct, see [Section 18.2](#)

28.2 omp_set_num_teams Routine

Name: <code>omp_set_num_teams</code> Category: subroutine	Properties: ICV-modifying
--	---

1 **Arguments**

2

Name	Type	Properties
<i>num_teams</i>	integer	non-negative

3 **Prototypes**

4

		C / C++		
void omp_set_num_teams (int <i>num_teams</i>) ;				
		C / C++		
		Fortran		
subroutine omp_set_num_teams (<i>num_teams</i>)				
integer <i>num_teams</i>				
		Fortran		

7 **Effect**

8 The effect of the **omp_set_num_teams** routine is to set the value of the *ntteams-var* ICV of the

9 **host device** to the value specified in the *num_teams* argument.

10 **Restrictions**

11 Restrictions to the **omp_set_num_teams** routine are as follows:

- 12 • An **omp_set_num_teams** region must be a strictly nested region of the implicit parallel
- 13 region that surrounds the whole OpenMP program.

14 **Cross References**

- 15 • *ntteams-var* ICV, see [Table 3.1](#)
- 16 • **num_teams** Clause, see [Section 18.2.1](#)
- 17 • **teams** Construct, see [Section 18.2](#)

18 **28.3 omp_get_team_num Routine**

19

Name: omp_get_team_num	Properties: ICV-retrieving, teams-nestable
Category: function	

20 **Return Type**

21

Name	Type	Properties
<return type>	integer	default

1 **Prototypes**

C / C++

int omp_get_team_num(void);

C / C++

Fortran

integer function omp_get_team_num()

Fortran

4 **Effect**

5 The `omp_get_team_num` routine returns the value of the *team-num-var* ICV, which is the *team*

6 *number* of the current *team* and 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* region.

8 **Cross References**

- 9
 - *team-num-var* ICV, see [Table 3.1](#)
 - `omp_get_num_teams` Routine, see [Section 28.1](#)
 - *teams* Construct, see [Section 18.2](#)

12 **28.4 omp_get_max_teams Routine**

Name: <code>omp_get_max_teams</code> Category: function	Properties: ICV-retrieving
--	---

14 **Return Type**

Name	Type	Properties
<return type>	integer	<i>default</i>

16 **Prototypes**

C / C++

int omp_get_max_teams(void);

C / C++

Fortran

integer function omp_get_max_teams()

Fortran

19 **Effect**

20 The `omp_get_max_teams` routine returns the value of the *nteam*s-var ICV of the *current*

21 *device*. If *positive*, this value is also an upper bound on the number of *teams* that can be created by

22 a *teams* construct without a *num_teams* clause that is encountered after execution returns from

23 this routine.

Cross References

- *ntteams-var* ICV, see [Table 3.1](#)
- **num_teams** Clause, see [Section 18.2.1](#)
- **teams** Construct, see [Section 18.2](#)

28.5 omp_get_teams_thread_limit Routine

Name: <code>omp_get_teams_thread_limit</code> Category: function	Properties: ICV-retrieving
---	--

Return Type

Name	Type	Properties
<i><return type></i>	integer	<i>default</i>

Prototypes

C / C++	<code>int omp_get_teams_thread_limit(void);</code>
C / C++	
Fortran	<code>integer function omp_get_teams_thread_limit()</code>
Fortran	

Effect

The `omp_get_teams_thread_limit` routine returns the value of the *teams-thread-limit-var* ICV, which is the maximum number of threads available to execute tasks in each contention group that a **teams** construct creates.

Cross References

- *teams-thread-limit-var* ICV, see [Table 3.1](#)
- **teams** Construct, see [Section 18.2](#)

28.6 omp_set_teams_thread_limit Routine

Name: <code>omp_set_teams_thread_limit</code> Category: subroutine	Properties: ICV-modifying
---	---

Arguments

Name	Type	Properties
<i>thread_limit</i>	integer	positive

Prototypes

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

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 and thus defines the maximum number of threads that can execute tasks in each contention group that a **teams** construct creates on the host device. If the value of *thread_limit* exceeds the number of 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:

- An **omp_set_num_teams** region must be a strictly nested region of the implicit parallel region that surrounds the whole OpenMP program.

Cross References

- *teams-thread-limit-var* ICV, see Table 3.1
- **teams** Construct, see Section 18.2
- **thread_limit** Clause, see Section 21.3

29 Tasking Support Routines

This chapter specifies OpenMP API routines that support task execution:

- Tasking routines that query general task execution properties; and
- The event routine to fulfill task dependences.

29.1 Tasking Routines

This section describes routines that pertain to OpenMP explicit tasks.

29.1.1 omp_get_max_task_priority Routine

Name: <code>omp_get_max_task_priority</code> Category: <code>function</code>	Properties: <code>all-device-threads-binding</code> , <code>ICV-retrieving</code>
---	--

Return Type

Name	Type	Properties
<code><return type></code>	integer	<code>default</code>

Prototypes

<code>int omp_get_max_task_priority(void);</code>	C / C++
<code>integer function omp_get_max_task_priority()</code>	Fortran

Effect

The `omp_get_max_task_priority` routine returns the value of the `max-task-priority-var` ICV, which determines the maximum value that can be specified in the `priority` clause.

Cross References

- `max-task-priority-var` ICV, see Table 3.1
- `priority` Clause, see Section 20.9

29.1.2 omp_in_explicit_task Routine

Name: <code>omp_in_explicit_task</code> Category: <code>function</code>	Properties: <code>ICV-retrieving</code>
--	---

Return Type

Name	Type	Properties
<code><return type></code>	logical	<code>default</code>

Prototypes

<code>int omp_in_explicit_task(void);</code>	C / C++
<code>logical function omp_in_explicit_task()</code>	Fortran

Effect

The `omp_in_explicit_task` routine returns the value of the `explicit-task-var` ICV, which indicates whether the encountering task is an `explicit task region`.

Cross References

- `explicit-task-var` ICV, see [Table 3.1](#)
- `task` Construct, see [Section 20.1](#)

29.1.3 omp_in_final Routine

Name: <code>omp_in_final</code> Category: <code>function</code>	Properties: <code>ICV-retrieving</code>
--	---

Return Type

Name	Type	Properties
<code><return type></code>	logical	<code>default</code>

Prototypes

<code>int omp_in_final(void);</code>	C / C++
<code>logical function omp_in_final()</code>	Fortran

Effect

The `omp_in_final` routine returns the value of the *final-task-var* ICV, which indicates whether the encountering task is a final task region.

Cross References

- `final` Clause, see [Section 20.7](#)
- *final-task-var* ICV, see [Table 3.1](#)
- `task` Construct, see [Section 20.1](#)









29.1.4 `omp_is_free_agent` Routine

Name: <code>omp_is_free_agent</code> Category: function	Properties: ICV-retrieving
--	---

Return Type

Name	Type	Properties
<i><return type></i>	logical	<i>default</i>

Prototypes

	
<code>int omp_is_free_agent(void);</code>	
	
	
<code>logical function omp_is_free_agent()</code>	
	

Effect

The `omp_is_free_agent` routine returns the value of the *free-agent-var* ICV, which indicates whether a *free-agent thread* is executing the enclosing *task region* at the time the routine is called.

Cross References

- *free-agent-var* ICV, see [Table 3.1](#)
- `task` Construct, see [Section 20.1](#)
- `threadset` Clause, see [Section 20.8](#)

29.1.5 `omp_ancestor_is_free_agent` Routine

Name: <code>omp_ancestor_is_free_agent</code> Category: function	Properties: <i>default</i>
---	-----------------------------------

1 **Return Type and Arguments**

Name	Type	Properties
<return type>	logical	<i>default</i>
<i>level</i>	integer	<i>default</i>

3 **Prototypes**

		C / C++	
int omp_ancestor_is_free_agent (int <i>level</i>);			
		C / C++	
		Fortran	
logical function omp_ancestor_is_free_agent (<i>level</i>)			
integer <i>level</i>			
		Fortran	

7 **Effect**

8 The **omp_ancestor_is_free_agent** routine returns *true* if the **ancestor thread** of the
9 **encountering thread** is a **free-agent thread**, for a given nested level of the **encountering thread**;
10 otherwise, it returns *false*. If the requested nesting level is outside the range of 0 and the nesting
11 level of the **current task**, as returned by the **omp_get_level** routine, the routine returns *false*.

13 Note – When the **omp_ancestor_is_free_agent** routine is called with a value of *level*
14 =**omp_get_level**, the routine has the same effect as the **omp_is_free_agent** routine.

16 **Cross References**

- **omp_get_level** Routine, see [Section 27.14](#)
- **omp_is_free_agent** Routine, see [Section 29.1.4](#)
- **task** Construct, see [Section 20.1](#)
- **threadset** Clause, see [Section 20.8](#)

21 **29.2 Event Routine**

22 This section describes **routines** that support OpenMP **event** objects.

23 **29.2.1 omp_fulfill_event Routine**

Name: omp_fulfill_event	Properties: <i>default</i>
Category: <i>subroutine</i>	

1 **Arguments**

2

Name	Type	Properties
<i>event</i>	event_handle	<i>default</i>

3 **Prototypes**

4

C / C++

void omp_fulfill_event(omp_event_handle_t event);

5

C / C++

6

Fortran

subroutine omp_fulfill_event(event)
 integer (kind=omp_event_handle_kind) event

Fortran

7 **Effect**

8 The effect of this routine is to fulfill the event associated with the event argument. The effect of
9 fulfilling the event will depend on how the event object was created. The event object is destroyed
10 and cannot be accessed after calling this routine, and the event handle becomes unassociated with
11 any event object. This routine has no effect if the event argument corresponds to a completed task.

12 **Execution Model Events**

13 The task-fulfill event occurs in a thread that executes an omp_fulfill_event region before the
14 event is fulfilled if the OpenMP event object was created by a detach clause on a task.

15 **Tool Callbacks**

16 A thread dispatches a registered task_schedule callback with NULL as its next_task_data
17 argument while the argument prior_task_data binds to the detachable task for each occurrence of a
18 task-fulfill event. If the task-fulfill event occurs before the detachable task finished execution of the
19 associated structured block, the callback has ompt_task_early_fulfill as its
20 prior_task_status argument; otherwise the callback has ompt_task_late_fulfill as its
21 prior_task_status argument.

22 **Restrictions**

23 Restrictions to the omp_fulfill_event routine are as follows:

- 24
 - The event that corresponds to the event argument must not have already been fulfilled.
 - The event handle that the event argument identifies must have been created by the effect of a detach clause.
 - The event handle passed to the routine must refer to an event object that was created by a thread in the same device as the thread that invoked the routine.
 - An event handle must be fulfilled before execution continues beyond the next barrier of the current team after a detach clause creates the event that the event argument represents.

Cross References

- **detach** Clause, see [Section 20.11](#)
- OpenMP **event_handle** Type, see [Section 26.6.1](#)
- **task_schedule** Callback, see [Section 40.5.2](#)
- OMPT **task_status** Type, see [Section 39.38](#)

30 Device Information Routines

This chapter describes [device-information routines](#), which are [routines](#) that have the [device-information](#) property. These [routines](#) modify or retrieve information that supports the use of the set of [devices](#) that are available to an [OpenMP program](#).

Restrictions

Restrictions to [device-information routines](#) are as follows.

- Any *device_num* argument must be a [conforming device number](#) unless otherwise specified.









30.1 omp_set_default_device Routine

Name: <code>omp_set_default_device</code>	Properties: device-information , ICV-modifying
Category: subroutine	

Arguments

Name	Type	Properties
<i>device_num</i>	integer	default

Prototypes

	C / C++	
<code>void omp_set_default_device(int device_num);</code>		
	C / C++	
	Fortran	
<code>subroutine omp_set_default_device(device_num) integer device_num</code>		
	Fortran	

Effect

The effect of the [omp_set_default_device](#) routine is to set the value of the [default-device-var ICV](#) of the [current task](#) to the value specified in the *device-num* argument, thus determining the default [target device](#). If the value specified in the *device-num* argument is the value of the [predefined identifier](#) `omp_default_device`, the [default-device-var ICV](#) is not modified. When called from within a [target region](#), the effect of this [routine](#) is [unspecified](#).

Cross References

- *default-device-var* ICV, see [Table 3.1](#)
- **target** Construct, see [Section 21.8](#)

30.2 omp_get_default_device Routine

Name: <code>omp_get_default_device</code>	Properties: device-information , ICV-retrieving
Category: function	

Return Type

Name	Type	Properties
<code><return type></code>	integer	default

Prototypes

C / C++		
<code>int omp_get_default_device(void);</code>		
C / C++		
Fortran		
<code>integer function omp_get_default_device()</code>		
Fortran		

Effect

The `omp_get_default_device` routine returns the value of the *default-device-var* ICV of the current task, which is the device number of the default target device. When called from within a target region the effect of this routine is unspecified.

Cross References

- *default-device-var* ICV, see [Table 3.1](#)
- **target** Construct, see [Section 21.8](#)

30.3 omp_get_num_devices Routine

Name: <code>omp_get_num_devices</code>	Properties: device-information , ICV-retrieving
Category: function	

Return Type

Name	Type	Properties
<code><return type></code>	integer	default

1 **Prototypes**

C / C++

int omp_get_num_devices(void);

C / C++

Fortran

integer function omp_get_num_devices()

Fortran

4 **Effect**

5 The `omp_get_num_devices` routine returns the value of the *num-devices-var* ICV, which is
6 the number of available *non-host devices* onto which code or data may be offloaded. When called
7 from within a *target region* the effect of this routine is *unspecified*.

8 **Cross References**

- 9 • *num-devices-var* ICV, see [Table 3.1](#)
10 • *target* Construct, see [Section 21.8](#)

11 **30.4 omp_get_device_num Routine**

Name: <code>omp_get_device_num</code> Category: function	Properties: device-information
---	---

13 **Return Type**

Name	Type	Properties
<i><return type></i>	integer	<i>default</i>

15 **Prototypes**

C / C++

int omp_get_device_num(void);

C / C++

Fortran

integer function omp_get_device_num()

Fortran

18 **Effect**

19 The `omp_get_device_num` routine returns the value of the *device-num-var* ICV, which is the
20 *device number* of the *device* on which the *encountering thread* is executing. When called on the
21 *host device*, it will return the same value as the `omp_get_initial_device` routine.

Cross References

- *device-num-var* ICV, see [Table 3.1](#)
- **target** Construct, see [Section 21.8](#)

30.5 omp_get_num_procs Routine

Name: omp_get_num_procs	Properties: all-device-threads-binding , device-information , ICV-retrieving
Category: function	

Return Type

Name	Type	Properties
<i><return type></i>	integer	default

Prototypes

	C / C++	
int omp_get_num_procs(void);		
	C / C++	
	Fortran	
integer function omp_get_num_procs()		
	Fortran	

Effect

The **omp_get_num_procs** routine returns the value of the *num-procs-var* ICV. Thus, this routine returns the number of [processors](#) that are available to the [device](#) at the time the routine is called. This value may change between the time that it is determined by the **omp_get_num_procs** routine and the time that it is read in the calling context due to system actions outside the control of the OpenMP implementation.

Cross References

- *num-procs-var* ICV, see [Table 3.1](#)

30.6 omp_get_max_progress_width Routine

Name: omp_get_max_progress_width	Properties: device-information
Category: function	

Return Type and Arguments

Name	Type	Properties
<i><return type></i>	integer	default
<i>device_num</i>	integer	default

1 **Prototypes**

C / C++

2 **int omp_get_max_progress_width(int device_num);**

C / C++

Fortran

3 **integer function omp_get_max_progress_width(device_num)**
4 **integer device_num**

Fortran

5 **Effect**

6 The `omp_get_max_progress_width` routine returns the maximum size, in terms of
7 hardware threads, of progress units on the device specified by `device_num`. When called from
8 within a `target` region the effect of this routine is unspecified.

9 **30.7 omp_get_device_from_uid Routine**

Name: <code>omp_get_device_from_uid</code> Category: <code>function</code>	Properties: <code>device-information</code>
---	---

11 **Return Type and Arguments**

Name	Type	Properties
<return type>	integer	default
<code>uid</code>	char	pointer, intent(in)

13 **Prototypes**

C / C++

14 **int omp_get_device_from_uid(const char *uid);**

C / C++

Fortran

15 **integer function omp_get_device_from_uid(uid)**
16 **character(len=*) , intent(in) :: uid**

Fortran

17 **Effect**

18 The `omp_get_device_from_uid` routine returns the device number associated with the device
19 specified by the `uid`; if no device with that `uid` is available, the value of `omp_invalid_device`
20 is returned. When called from within a `target` region, the effect is unspecified.

Cross References

- *available-devices-var* ICV, see [Table 3.1](#)
- *default-device-var* ICV, see [Table 3.1](#)
- `omp_get_uid_from_device` Routine, see [Section 30.8](#)

30.8 `omp_get_uid_from_device` Routine

Name: <code>omp_get_uid_from_device</code> Category: function	Properties: device-information
--	--

Return Type and Arguments

Name	Type	Properties
<return type>	const char	pointer
<i>device_num</i>	integer	intent(in)

Prototypes

C / C++

const char *omp_get_uid_from_device(int device_num);

C / C++

Fortran

character(:) function omp_get_uid_from_device(device_num)
 pointer :: omp_get_uid_from_device
 integer, intent(in) :: device_num

Fortran

Effect

The `omp_get_uid_from_device` routine returns the [implementation defined](#) unique identifier string that identifies the [device](#) specified by *device_num*. If the *device_num* argument has a value of [omp_invalid_device](#), the routine returns `NULL`. When called from within a [target region](#), the effect is [unspecified](#).

Cross References

- *available-devices-var* ICV, see [Table 3.1](#)
- *default-device-var* ICV, see [Table 3.1](#)
- `omp_get_device_from_uid` Routine, see [Section 30.7](#)

30.9 omp_is_initial_device Routine

Name: <code>omp_is_initial_device</code> Category: <code>function</code>	Properties: <code>device-information</code>
---	---

Return Type

Name	Type	Properties
<code><return type></code>	logical	<code>default</code>

Prototypes

<code>int omp_is_initial_device(void);</code>	C / C++
<code>logical function omp_is_initial_device()</code>	Fortran

Effect

The `omp_is_initial_device` routine returns `true` if the `current task` is executing on the `host device`; otherwise, it returns `false`.

30.10 omp_get_initial_device Routine

Name: <code>omp_get_initial_device</code> Category: <code>function</code>	Properties: <code>device-information</code>
--	---

Return Type

Name	Type	Properties
<code><return type></code>	integer	<code>default</code>

Prototypes

<code>int omp_get_initial_device(void);</code>	C / C++
<code>integer function omp_get_initial_device()</code>	Fortran

Effect

The effect of the `omp_get_initial_device` routine is to return the `device number` of the `host device`. The value of the `device number` is the value of `omp_initial_device` or 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` Construct, see [Section 21.8](#)

30.11 `omp_get_device_num_teams` Routine

Name: <code>omp_get_device_num_teams</code> Category: function	Properties: device-information , ICV-retrieving
---	---

Return Type and Arguments

Name	Type	Properties
<return type>	integer	default
<code>device_num</code>	integer	default

Prototypes

C / C++

`int omp_get_device_num_teams(int device_num);`

C / C++

Fortran

`integer function omp_get_device_num_teams(device_num)`
`integer device_num`

Fortran

Effect

The `omp_get_device_num_teams` routine returns the value of the `nteam`s-var ICV in the device data environment of device `device_num`. Thus, the routine returns the number of teams that will be requested for a teams region on device `device_num` if the `num_teams` clause is not specified. If `device_num` is the device number of the host device, `omp_get_device_num_teams` is equivalent to `omp_get_num_teams`. If the `device_num` argument has the value of `omp_invalid_device` or is not a conforming device number, the routine returns zero. When called from within a target region, the effect of this routine is unspecified.

Cross References

- `nteam`s-var ICV, see [Table 3.1](#)
- `num_teams` Clause, see [Section 18.2.1](#)
- `teams` Construct, see [Section 18.2](#)









30.12 omp_set_device_num_teams Routine

Name: <code>omp_set_device_num_teams</code> Category: subroutine	Properties: device-information , ICV-modifying
---	--

Arguments

Name	Type	Properties
<code>num_teams</code>	integer	non-negative
<code>device_num</code>	integer	default

Prototypes

	C / C++	
	C / C++	
	Fortran	
	Fortran	

```
void omp_set_device_num_teams(int num_teams, int device_num);  
  
subroutine omp_set_device_num_teams(num_teams, device_num)  
  integer num_teams, device_num
```

Effect

The effect of the `omp_set_device_num_teams` routine is to set the value of the *ntteams-var* ICV of device `device_num` to the value specified in the `num_teams` argument. Thus, the routine determines the number of *teams* that will be requested for a *teams* region on device `device_num` if the *num_teams* clause is not specified. If `device_num` is the device number of the host device, `omp_set_device_num_teams` is equivalent to `omp_set_num_teams`. If the `device_num` argument has the value of `omp_invalid_device` or is not a conforming device number, runtime error termination occurs. When called from within a *target* region, the effect of this routine is unspecified.

Restrictions

Restrictions to the `omp_set_device_num_teams` routine are as follows:

- The routine must not execute concurrently with any device-affecting construct on device `device_num`.
- If device `device_num` is the host device, an `omp_set_device_num_teams` region must be a strictly nested region of the implicit parallel region that surrounds the whole OpenMP program.

Cross References

- *ntteams-var* ICV, see [Table 3.1](#)
- `num_teams` Clause, see [Section 18.2.1](#)
- *teams* Construct, see [Section 18.2](#)

1

2

30.13 omp_get_device_teams_thread_limit Routine

3

Name: omp_get_device_teams_thread_limit Category: function	Properties: device-information, ICV-retrieving
--	---

4

Return Type and Arguments

5

Name	Type	Properties
<return type>	integer	default
device_num	integer	default

6

Prototypes

7

C / C++	
int omp_get_device_teams_thread_limit(int device_num);	
C / C++	
Fortran	
integer function omp_get_device_teams_thread_limit(device_num) integer device_num	
Fortran	

10

Effect

11

12

13

14

15

16

17

18

The `omp_get_device_teams_thread_limit` routine returns the value of the `teams-thread-limit-var` ICV in the device data environment of device `device_num`, which is the maximum number of threads available to execute tasks in each contention group that a `teams` construct creates on that device. If `device_num` is the device number of the host device, `omp_get_device_teams_thread_limit` is equivalent to `omp_get_teams_thread_limit`. If the `device_num` argument has the value of `omp_invalid_device` or is not a conforming device number, the routine returns zero. When called from within a `target` region, the effect of this routine is unspecified.

19

Cross References

- 20
- 21
- `teams-thread-limit-var` ICV, see Table 3.1
 - `teams` Construct, see Section 18.2

22

23

30.14 omp_set_device_teams_thread_limit Routine

24

Name: omp_set_device_teams_thread_limit Category: subroutine	Properties: device-information, ICV-modifying
--	--

Arguments

Name	Type	Properties
<i>thread_limit</i>	integer	positive
<i>device_num</i>	integer	default

Prototypes

C / C++

void omp_set_device_teams_thread_limit(int thread_limit,
int device_num);

C / C++
Fortran

subroutine omp_set_device_teams_thread_limit(thread_limit, &
device_num)
integer thread_limit, device_num

Fortran

Effect

The `omp_set_device_teams_thread_limit` routine sets the value of the `teams-thread-limit-var` ICV in the device data environment of device `device_num` to the value of the `thread_limit` argument and thus defines the maximum number of threads that can execute tasks in each contention group that a `teams` construct creates on that device. If the value of `thread_limit` exceeds the number of threads that an implementation supports for each contention group created by a `teams` construct on device `device_num`, the value of the `teams-thread-limit-var` ICV will be set to the number that is supported by the implementation. If `device_num` is the device number of the host device, `omp_set_device_teams_thread_limit` is equivalent to `omp_set_teams_thread_limit`. If the `device_num` argument has the value of `omp_invalid_device` or is not a conforming device number, runtime error termination occurs. When called from within a `target` region, the effect of this routine is unspecified.

Restrictions

- Restrictions to the `omp_set_device_teams_thread_limit` routine are as follows:
- The routine must not execute concurrently with any device-affecting construct on device `device_num`.
 - If device `device_num` is the host device, an `omp_set_device_teams_thread_limit` region must be a strictly nested region of the implicit parallel region that surrounds the whole OpenMP program.

Cross References

- `teams-thread-limit-var` ICV, see Table 3.1
- `teams` Construct, see Section 18.2
- `thread_limit` Clause, see Section 21.3

31 Device Memory Routines

This chapter describes [device memory routines](#) that support allocation of [memory](#) and management of pointers in the [data environments](#) of [target devices](#), and therefore the [routines](#) have the [device memory routine property](#).

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.

[Device memory routines](#) that are not [device-memory-information routines](#) execute as if part of a [target task](#) that is generated by the call to the [routine](#). This [target task](#), which is an [included task](#) if the [routine](#) is not an [asynchronous device routine](#), is the [generating task](#) of the [region](#) associated with the [routine](#). Since the [target task](#) provides the execution context for any execution that occurs on the [device](#), it is the [binding task set](#) for the [routine](#). Thus, all of these [routines](#) have the [generating-task binding property](#).

Fortran

The Fortran version of all [device memory routines](#) have ISO C bindings so the [routines](#) have the [ISO C binding property](#). Thus, each [device memory routine](#) requires an explicit interface and so might not be provided in the [deprecated](#) include file `omp_lib.h`.

Fortran

Execution Model Events

[Events](#) associated with a [target task](#) are the same as for the [task construct](#) defined in [Section 20.1](#).

Tool Callbacks

[Callbacks](#) associated with events for [target tasks](#) are the same as for the [task construct](#) defined in [Section 20.1](#); (`flags & omp_t_task_target`) always evaluates to `true` in the dispatched [callback](#).

Restrictions

Restrictions to [device memory routines](#) are as follows:

- Any `device_num`, `src_device_num`, and `dst_device_num` arguments must be [conforming device numbers](#).
- When called from within a [target region](#), the effect is [unspecified](#).

1 **Cross References**

- 2 • **target** Construct, see [Section 21.8](#)
- 3 • **task** Construct, see [Section 20.1](#)
- 4 • OMPT **task_flag** Type, see [Section 39.37](#)

5 **31.1 Asynchronous Device Memory Routines**

6 Some [device memory routines](#) have the [asynchronous-device routine property](#). The execution of the

7 [target task](#) that is generated by the call to an [asynchronous device routines](#) may be deferred. [Task](#)

8 [dependences](#) are expressed with zero or more OpenMP [depend objects](#). The [dependences](#) are

9 specified by passing the number of [depend objects](#) followed by an array of the objects. The

10 generated [target task](#) is not a [dependent task](#) if the program passes in a count of zero for

11 *depobj_count*. The *depobj_list* argument is ignored if the value of *depobj_count* is zero.

12 **Execution Model Events**

13 [Events](#) associated with [task dependences](#) that result from *depobj_list* are the same as for a [depend](#)

14 [clause](#) with the [depobj task-dependence-type](#) defined in [Section 23.9.5](#).

15 **Tool Callbacks**

16 [Callbacks](#) associated with events for [task dependences](#) are the same as for the [depend clause](#)

17 defined in [Section 23.9.5](#).

18 **Cross References**

- 19 • **depend** Clause, see [Section 23.9.5](#)
- 20 • **depobj** Construct, see [Section 23.9.3](#)

21 **31.2 Device Memory Information Routines**

22 This section describes [routines](#) that have the [device-memory-information routine property](#). These

23 [device-memory-information routines](#) provide information about [device pointers](#), which can be

24 determined without directly accessing the [target device](#); thus, they do not create a [target task](#).

25 **31.2.1 omp_target_is_present Routine**

26	Name: <code>omp_target_is_present</code> Category: <code>function</code>	Properties: <code>device-memory-</code> <code>information-routine, device-memory-</code> <code>routine, iso_c_binding</code>
----	---	---

1 **Return Type and Arguments**

Name	Type	Properties
<return type>	c_int	default
ptr	c_ptr	intent(in), iso_c, value
device_num	c_int	iso_c, value

3 **Prototypes**

C / C++	
int omp_target_is_present(const void *ptr, int device_num);	
C / C++	
Fortran	
integer (kind=c_int) function omp_target_is_present (ptr, & device_num) bind(c) use, intrinsic :: iso_c_binding, only : c_int, c_ptr type (c_ptr), value, intent(in) :: ptr integer (kind=c_int), value :: device_num	
Fortran	

10 **Effect**

11 The **omp_target_is_present** routine returns a non-zero value if *device_num* refers to the
12 **host device** or if *ptr* refers to storage that has **corresponding storage** in the **device data environment**
13 of **device** *device_num*. Otherwise, the **routine** returns zero. If *ptr* is **NULL**, the **routine** returns zero.
14 Thus, the **omp_target_is_present** routine tests whether a **host pointer** refers to storage that
15 is mapped to a given **device**.

16 **Restrictions**

17 Restrictions to the **omp_target_is_present** routine are as follows:

- 18 • The value of *ptr* must be a valid **host pointer** or **NULL**.

19 **31.2.2 omp_target_is_accessible Routine**

Name: omp_target_is_accessible Category: function	Properties: device-memory- information-routine, device-memory- routine, iso_c_binding
--	--

21 **Return Type and Arguments**

Name	Type	Properties
<return type>	c_int	default
ptr	c_ptr	intent(in), iso_c, value
size	c_size_t	iso_c, positive, value
device_num	c_int	iso_c, value

1 Prototypes

C / C++

2 `int omp_target_is_accessible(const void *ptr, size_t size,`
3 `int device_num);`

C / C++

Fortran

4 `integer (kind=c_int) function omp_target_is_accessible(ptr, &`
5 `size, device_num) bind(c)`
6 `use, intrinsic :: iso_c_binding, only : c_int, c_ptr, &`
7 `c_size_t`
8 `type (c_ptr), value, intent(in) :: ptr`
9 `integer (kind=c_size_t), value :: size`
10 `integer (kind=c_int), value :: device_num`

Fortran

11 Effect

12 The `omp_target_is_accessible` routine returns a non-zero value if the storage of *size* bytes
13 that corresponds to the `address range` starting at the address given by *ptr* is `accessible` from `device`
14 `device_num`. Otherwise, it returns zero. If *ptr* is `NULL` or the implementation cannot guarantee
15 accessibility, the `routine` returns zero. The value of *ptr* is interpreted as an address in the `address`
16 `space` of the specified `device`.

17 31.2.3 omp_get_mapped_ptr Routine

Name: <code>omp_get_mapped_ptr</code>	Properties: <code>device-memory-</code>
Category: <code>function</code>	<code>information-routine</code> , <code>device-memory-</code>
	<code>routine</code> , <code>iso_c_binding</code>

19 Return Type and Arguments

Name	Type	Properties
<code><return type></code>	<code>c_ptr</code>	<code>default</code>
<code>ptr</code>	<code>c_ptr</code>	<code>intent(in)</code> , <code>iso_c</code> , <code>value</code>
<code>device_num</code>	<code>c_int</code>	<code>iso_c</code> , <code>value</code>

1 **Prototypes**

C / C++

2 **void *omp_get_mapped_ptr(const void *ptr, int device_num);**

C / C++

Fortran

3 **type (c_ptr) function omp_get_mapped_ptr(ptr, device_num) &**
4 **bind(c)**
5 **use, intrinsic :: iso_c_binding, only : c_ptr, c_int**
6 **type (c_ptr), value, intent(in) :: ptr**
7 **integer (kind=c_int), value :: device_num**

Fortran

8 **Effect**

9 The **omp_get_mapped_ptr** routine returns the associated **device pointer** for **host pointer** *ptr* on
10 **device** *device_num*. A call to this **routine** for a pointer that is not **NULL** and does not have an
11 associated pointer on the given **device** will return **NULL**. The **routine** returns **NULL** if unsuccessful.
12 Otherwise it returns the **device pointer**, which is *ptr* if *device_num* specifies the **host device**.

13 **Cross References**

- 14 • **omp_get_initial_device** Routine, see [Section 30.10](#)

15 **31.3 omp_target_alloc Routine**

16 **Name:** **omp_target_alloc**
 Category: [function](#)

Properties: [device-memory-routine](#),
[generating-task-binding](#), [iso_c_binding](#)

17 **Return Type and Arguments**

Name	Type	Properties
<i><return type></i>	c_ptr	default
<i>size</i>	c_size_t	iso_c , value
<i>device_num</i>	c_int	iso_c , value

Prototypes

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 (kind=c_size_t), value :: size  
    integer (kind=c_int), value :: device_num
```

Fortran

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 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 7.3.6).

Execution Model Events

The *target-data-allocation-begin* **event** occurs before a **thread** initiates a data allocation on a **target device**. The *target-data-allocation-end* **event** occurs after a **thread** initiates a data allocation on a **target device**.

Tool Callbacks

A **thread** dispatches a registered **target_data_op_emi** callback with **ompt_scope_begin** as its *endpoint* argument for each occurrence of a *target-data-allocation-begin* **event** in that **thread**. Similarly, a **thread** dispatches a registered **target_data_op_emi** callback with **ompt_scope_end** as its *endpoint* argument for each occurrence of a *target-data-allocation-end* **event** in that **thread**.

Restrictions

Restrictions to the **omp_target_alloc** routine are as follows:

- Freeing the storage returned by **omp_target_alloc** with any **routine** other than **omp_target_free** results in **unspecified behavior**.

C / C++

- Unless the **unified_address** clause appears on a **requires** directive in the **compilation unit**, pointer arithmetic is not supported on the **device pointer** returned by **omp_target_alloc**.

C / C++

Cross References

- `is_device_ptr` Clause, see [Section 7.3.6](#)
- `omp_target_free` Routine, see [Section 31.4](#)
- OMPT `scope_endpoint` Type, see [Section 39.27](#)
- `target_data_op_emi` Callback, see [Section 41.7](#)

31.4 `omp_target_free` Routine

Name: <code>omp_target_free</code>	Properties: device-memory-routine , generating-task-binding , iso_c_binding
Category: subroutine	

Arguments

Name	Type	Properties
<i>device_ptr</i>	<code>c_ptr</code>	iso_c , value
<i>device_num</i>	<code>c_int</code>	iso_c , value

Prototypes

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 (kind=c_int), value :: device_num`

Fortran

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.

Execution Model Events

The *target-data-free-begin* [event](#) occurs before a [thread](#) initiates a data free on a [target device](#). The *target-data-free-end* [event](#) occurs after a [thread](#) initiates a data free on a [target device](#).

Tool Callbacks

A [thread](#) dispatches a registered `target_data_op_emi` callback with `ompt_scope_begin` as its *endpoint* argument for each occurrence of a *target-data-free-begin* [event](#) in that [thread](#). Similarly, a [thread](#) dispatches a registered `target_data_op_emi` callback with `ompt_scope_end` as its *endpoint* argument for each occurrence of a *target-data-free-end* [event](#) in that [thread](#).

1 **Restrictions**

2 Restrictions to the `omp_target_free` routine are as follows:

- 3 • The value of `device_ptr` must be `NULL` or have been returned by `omp_target_alloc`.

4 **Cross References**

- 5 • `omp_target_alloc` Routine, see [Section 31.3](#)
6 • OMPT `scope_endpoint` Type, see [Section 39.27](#)
7 • `target_data_op_emi` Callback, see [Section 41.7](#)

8 **31.5 omp_target_associate_ptr Routine**

9

Name: <code>omp_target_associate_ptr</code>	Properties: device-memory-routine , generating-task-binding , iso_c_binding
Category: function	

10 **Return Type and Arguments**

11

Name	Type	Properties
<i><return type></i>	<code>c_int</code>	default
<i>host_ptr</i>	<code>c_ptr</code>	intent(in) , iso_c , value
<i>device_ptr</i>	<code>c_ptr</code>	intent(in) , iso_c , value
<i>size</i>	<code>c_size_t</code>	iso_c , value
<i>device_offset</i>	<code>c_size_t</code>	iso_c , value
<i>device_num</i>	<code>c_int</code>	iso_c , value

12 **Prototypes**

13 **C / C++**

```
14 int omp_target_associate_ptr(const void *host_ptr,  
15   const void *device_ptr, size_t size, size_t device_offset,  
   int device_num);
```

16 **C / C++**

17 **Fortran**

```
18 integer (kind=c_int) function omp_target_associate_ptr(host_ptr, &  
19   device_ptr, size, device_offset, device_num) bind(c)  
20   use, intrinsic :: iso_c_binding, only : c_int, c_ptr, &  
21     c_size_t  
22   type (c_ptr), value, intent(in) :: host_ptr, device_ptr  
   integer (kind=c_size_t), value :: size, device_offset  
   integer (kind=c_int), value :: device_num
```

Fortran

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 device pointer appears in a subsequent `map` clause, the associated device pointer is used as the target for data motion associated with that host pointer. Thus, the `omp_target_associate_ptr` routine maps a device pointer, which may be returned from `omp_target_alloc` or implementation defined routine, to a host pointer. The `device_offset` argument 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. The association between the host pointer and the device pointer can be removed by using the `omp_target_disassociate_ptr` routine. The routine returns zero if successful. Otherwise it returns a non-zero value.

Only one device buffer can be associated with a given host pointer value and device number pair. Attempting to associate a second buffer will return non-zero. Associating the same pair of pointers on the same device with the same offset has no effect and returns zero. Associating pointers that share underlying storage will result in unspecified behavior. The `omp_target_is_present` routine can be used to test whether a given host pointer has a corresponding list item in the device data environment.

Execution Model Events

The `target-data-associate` event occurs before a thread initiates a device pointer association on a target device.

Tool Callbacks

A thread dispatches a registered `target_data_op_emi` callback with `ompt_scope_beginend` as its `endpoint` argument for each occurrence of a `target-data-associate` event in that thread.

Cross References

- `omp_target_alloc` Routine, see Section 31.3
- `omp_target_disassociate_ptr` Routine, see Section 31.6
- `omp_target_is_present` Routine, see Section 31.2.1
- OMPT `scope_endpoint` Type, see Section 39.27
- `target_data_op_emi` Callback, see Section 41.7

31.6 omp_target_disassociate_ptr Routine

Name: <code>omp_target_disassociate_ptr</code> Category: <code>function</code>	Properties: <code>device-memory-routine</code> , <code>generating-task-binding</code> , <code>iso_c_binding</code>
---	---

1 **Return Type and Arguments**

Name	Type	Properties
<return type>	c_int	default
ptr	c_ptr	intent(in), iso_c, value
device_num	c_int	iso_c, value

3 **Prototypes**

4 **C / C++**

```
int omp_target_disassociate_ptr(const void *ptr, int device_num);
```

5 **C / C++**

6 **Fortran**

```
integer (kind=c_int) function omp_target_disassociate_ptr(ptr, &  
device_num) bind(c)  
use, intrinsic :: iso_c_binding, only : c_int, c_ptr  
type (c_ptr), value, intent(in) :: ptr  
integer (kind=c_int), value :: device_num
```

7 **Fortran**

10 **Effect**

11 The **omp_target_disassociate_ptr** removes the associated **device** data on **device**
12 **device_num** from the presence table for **host pointer** **ptr**. A call to this **routine** on a pointer that is
13 not **NULL** and does not have associated data on the given **device** results in **unspecified behavior**.
14 The reference count of the mapping is reduced to zero, regardless of its current value. The **routine**
15 returns zero if successful. Otherwise it returns a non-zero value.

16 **Execution Model Events**

17 The **target-data-disassociate** **event** occurs before a **thread** initiates a **device pointer** disassociation
18 on a **target device**.

19 **Tool Callbacks**

20 A **thread** dispatches a registered **target_data_op_emi** callback with
21 **ompt_scope_beginend** as its **endpoint** argument for each occurrence of a
22 **target-data-disassociate** **event** in that **thread**.

23 **Cross References**

- 24 • OMPT **scope_endpoint** Type, see [Section 39.27](#)
- 25 • **target_data_op_emi** Callback, see [Section 41.7](#)

31.7 Memory Copying Routines

This section describes [memory-copying routines](#), which are [routines](#) that have the [memory-copying property](#). These [routines](#) copy memory from the [device data environment](#) of a [src_device_num](#) [device](#) to the [device data environment](#) of a [dst_device_num](#) [device](#). OpenMP provides two varieties of [memory-copying routines](#): [flat-memory-copying routines](#), which have the [flat-memory-copying property](#); and [rectangular-memory-copying routines](#), which have the [rectangular-memory-copying property](#).

Each [flat-memory-copying 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](#).

Each [rectangular-memory-copying routine](#) performs a copy between any combination of [host pointers](#) and [device pointers](#). Specifically, the [routine](#) copies a rectangular subvolume from a multi-dimensional array *src*, in the [device data environment](#) of [device src_device_num](#), to another multi-dimensional array *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*) argument specifies the number of elements from the origin of *dst* (*src*) in elements. The *dst_dimensions* (*src_dimensions*) argument specifies the length of each dimension of *dst* (*src*).

An [OpenMP program](#) can determine the inclusive number of dimensions that an implementation supports for a [rectangular-memory-copying routine](#) 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

Because the interface of each [rectangular-memory-copying routine](#) binds directly to a C language [routine](#), each of these [routines](#) assumes C [memory](#) ordering.

Fortran

Each [memory-copying routine](#) contains a [task scheduling point](#). These [routines](#) return zero on success and non-zero on failure.

Execution Model Events

The *target-data-op-begin* [event](#) occurs before a [thread](#) initiates a data transfer in a [memory-copying routine](#) region. The *target-data-op-end* [event](#) occurs after a [thread](#) initiates a data transfer in a [memory-copying routine](#) region.

Tool Callbacks

A `thread` dispatches a registered `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 `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` occur in the context of the `target` task.

Restrictions

Restrictions to the `memory-copying` routines are as follows:

- The value of *src* must be a valid `device pointer` for the `device` *src_device_num*.
- The value of *dst* must be a valid `device pointer` for the `device` *dst_device_num*.
- The value of *num_dims* must be between 1 and the `implementation defined` limit, which must be at least three.
- The length of the offset (*src_offset* and *dst_offset*) and dimension (*src_dimensions* and *dst_dimensions*) arrays must be at least the value of *num_dims*.

Cross References

- OMPT `scope_endpoint` Type, see [Section 39.27](#)
- `target_data_op_emi` Callback, see [Section 41.7](#)

31.7.1 omp_target_memcpy Routine

Name: <code>omp_target_memcpy</code> Category: <code>function</code>	Properties: <code>device-memory-routine</code> , <code>flat-memory-copying</code> , <code>generating-task-binding</code> , <code>iso_c_binding</code> , <code>memory-copying</code>
---	--

Return Type and Arguments

Name	Type	Properties
<i><return type></i>	<code>c_int</code>	<i>default</i>
<i>dst</i>	<code>c_ptr</code>	<code>iso_c</code> , <code>value</code>
<i>src</i>	<code>c_ptr</code>	<code>intent(in)</code> , <code>iso_c</code> , <code>value</code>
<i>length</i>	<code>c_size_t</code>	<code>iso_c</code> , <code>value</code>
<i>dst_offset</i>	<code>c_size_t</code>	<code>iso_c</code> , <code>value</code>
<i>src_offset</i>	<code>c_size_t</code>	<code>iso_c</code> , <code>value</code>
<i>dst_device_num</i>	<code>c_int</code>	<code>iso_c</code> , <code>value</code>
<i>src_device_num</i>	<code>c_int</code>	<code>iso_c</code> , <code>value</code>

1 Prototypes

C / C++

2 `int omp_target_memcpy(void *dst, const void *src, size_t length,`
3 `size_t dst_offset, size_t src_offset, int dst_device_num,`
4 `int src_device_num);`

C / C++

Fortran

5 `integer (kind=c_int) function omp_target_memcpy(dst, src, &`
6 `length, dst_offset, src_offset, dst_device_num, &`
7 `src_device_num) bind(c)`
8 `use, intrinsic :: iso_c_binding, only : c_int, c_ptr, &`
9 `c_size_t`
10 `type (c_ptr), value :: dst`
11 `type (c_ptr), value, intent(in) :: src`
12 `integer (kind=c_size_t), value :: length, dst_offset, &`
13 `src_offset`
14 `integer (kind=c_int), value :: dst_device_num, src_device_num`

Fortran

15 Effect

16 As a [flat-memory-copying routine](#), the effect of the [omp_target_memcpy routine](#) is as described
17 in [Section 31.7](#). This effect includes the associated [tool events](#) and [callbacks](#) defined in that section.

18 Cross References

- 19
 - Memory Copying Routines, see [Section 31.7](#)

20 31.7.2 omp_target_memcpy_rect Routine

21 Name: `omp_target_memcpy_rect`
Category: [function](#)

Properties: [device-memory-routine](#),
[generating-task-binding](#), [iso_c_bind-](#)
[ing](#), [memory-copying](#), [rectangular-](#)
[memory-copying](#)

1 **Return Type and Arguments**

Name	Type	Properties
<i><return type></i>	c_int	<i>default</i>
<i>dst</i>	c_ptr	iso_c, value
<i>src</i>	c_ptr	intent(in), iso_c, value
<i>element_size</i>	c_size_t	iso_c, value
<i>num_dims</i>	c_int	iso_c, positive, value
<i>volume</i>	c_size_t	intent(in), iso_c, pointer
<i>dst_offsets</i>	c_size_t	intent(in), iso_c, pointer
<i>src_offsets</i>	c_size_t	intent(in), iso_c, pointer
<i>dst_dimensions</i>	c_size_t	intent(in), iso_c, pointer
<i>src_dimensions</i>	c_size_t	intent(in), iso_c, pointer
<i>dst_device_num</i>	c_int	iso_c, value
<i>src_device_num</i>	c_int	iso_c, value

3 **Prototypes**

4 **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);
```

5 **C / C++**

6 **Fortran**

```
integer (kind=c_int) function omp_target_memcpy_rect (dst, src, &
    element_size, num_dims, volume, dst_offsets, src_offsets, &
    dst_dimensions, src_dimensions, dst_device_num, &
    src_device_num) bind(c)
    use, intrinsic :: iso_c_binding, only : c_int, c_ptr, &
        c_size_t
    type (c_ptr), value :: dst
    type (c_ptr), value, intent(in) :: src
    integer (kind=c_size_t), value :: element_size
    integer (kind=c_int), value :: num_dims, dst_device_num, &
        src_device_num
    integer (kind=c_size_t), intent(in) :: volume(*), dst_offsets&
        (*), src_offsets(*), dst_dimensions(*), src_dimensions(*)
```

7 **Fortran**

22 **Effect**

23 As a [rectangular-memory-copying routine](#), the effect of the [omp_target_memcpy_rect](#)
24 [routine](#) is as described in [Section 31.7](#). This effect includes the associated [tool events](#) and [callbacks](#)
25 defined in that section.

Cross References

- Memory Copying Routines, see [Section 31.7](#)

31.7.3 omp_target_memcpy_async Routine

Name: <code>omp_target_memcpy_async</code> Category: function	Properties: asynchronous-device-routine , device-memory-routine , flat-memory-copying , generating-task-binding , iso_c_binding , memory-copying
--	---

Return Type and Arguments

Name	Type	Properties
<return type>	<code>c_int</code>	default
<i>dst</i>	<code>c_ptr</code>	iso_c , value
<i>src</i>	<code>c_ptr</code>	intent(in) , iso_c , value
<i>length</i>	<code>c_size_t</code>	iso_c , value
<i>dst_offset</i>	<code>c_size_t</code>	iso_c , value
<i>src_offset</i>	<code>c_size_t</code>	iso_c , value
<i>dst_device_num</i>	<code>c_int</code>	iso_c , value
<i>src_device_num</i>	<code>c_int</code>	iso_c , value
<i>depobj_count</i>	<code>c_int</code>	iso_c , value
<i>depobj_list</i>	<code>depend</code>	optional , pointer

Prototypes

C / C++

```
int omp_target_memcpy_async(void *dst, const void *src,
    size_t length, size_t dst_offset, size_t src_offset,
    int dst_device_num, int src_device_num, int depobj_count,
    omp_depend_t *depobj_list);
```

C / C++

Fortran

```
integer (kind=c_int) function omp_target_memcpy_async(dst, src, &
    length, dst_offset, src_offset, dst_device_num, &
    src_device_num, depobj_count, depobj_list) bind(c)
    use, intrinsic :: iso_c_binding, only : c_int, c_ptr, &
    c_size_t
    type (c_ptr), value :: dst
    type (c_ptr), value, intent(in) :: src
    integer (kind=c_size_t), value :: length, dst_offset, &
    src_offset
```

```
1 integer (kind=c_int), value :: dst_device_num, src_device_num, &
2 depobj_count
3 integer (kind=omp_depend_kind), optional :: depobj_list(*)
```

Fortran

Effect

As a flat-memory-copying routine, the effect of the `omp_target_memcpy_async` routine is as described in Section 31.7. This effect includes the tool events and callbacks defined in that section. As it is also an asynchronous device routine, the routine also includes the tool events and callbacks defined in Section 31.1.

Cross References

- Asynchronous Device Memory Routines, see Section 31.1
- Memory Copying Routines, see Section 31.7

31.7.4 omp_target_memcpy_rect_async Routine

Name: <code>omp_target_memcpy_rect_async</code> Category: <code>function</code>	Properties: asynchronous-device-routine, device-memory-routine, generating-task-binding, iso_c_binding, memory-copying, rectangular-memory-copying
--	--

Return Type and Arguments

Name	Type	Properties
<return type>	c_int	default
dst	c_ptr	iso_c, value
src	c_ptr	intent(in), iso_c, value
element_size	c_size_t	iso_c, value
num_dims	c_int	iso_c, positive, value
volume	c_size_t	intent(in), iso_c, pointer
dst_offsets	c_size_t	intent(in), iso_c, pointer
src_offsets	c_size_t	intent(in), iso_c, pointer
dst_dimensions	c_size_t	intent(in), iso_c, pointer
src_dimensions	c_size_t	intent(in), iso_c, pointer
dst_device_num	c_int	iso_c, value
src_device_num	c_int	iso_c, value
depobj_count	c_int	iso_c, value
depobj_list	depend	optional, pointer

Prototypes

C / C++

```
int omp_target_memcpy_rect_async(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, int depobj_count,
    omp_depend_t *depobj_list);
```

C / C++

Fortran

```
integer (kind=c_int) function omp_target_memcpy_rect_async(dst, &
    src, element_size, num_dims, volume, dst_offsets, src_offsets, &
    dst_dimensions, src_dimensions, dst_device_num, &
    src_device_num, depobj_count, depobj_list) bind(c)
    use, intrinsic :: iso_c_binding, only : c_int, c_ptr, &
    c_size_t
    type (c_ptr), value :: dst
    type (c_ptr), value, intent(in) :: src
    integer (kind=c_size_t), value :: element_size
    integer (kind=c_int), value :: num_dims, dst_device_num, &
    src_device_num, depobj_count
    integer (kind=c_size_t), intent(in) :: volume(*), dst_offsets&
    (*), src_offsets(*), dst_dimensions(*), src_dimensions(*)
    integer (kind=omp_depend_kind), optional :: depobj_list(*)
```

Fortran

Effect

As a [rectangular-memory-copying routine](#), the effect of the [omp_target_memcpy_rect_async](#) routine is as described in [Section 31.7](#). This effect includes the [tool events](#) and [callbacks](#) defined in that section. As it is also an [asynchronous device routine](#), the routine also includes the [tool events](#) and [callbacks](#) defined in [Section 31.1](#).

Cross References

- Asynchronous Device Memory Routines, see [Section 31.1](#)
- Memory Copying Routines, see [Section 31.7](#)

31.8 Memory Setting Routines

This section describes the [memory-setting routines](#), which are [routines](#) that have the [memory-setting property](#). These [routines](#) fill [memory](#) in a [device data environment](#) with a given value. The effect of a [memory-setting routine](#) is to fill 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 [memory-setting routines](#) return *ptr*. Each [memory-setting routine](#) contains a [task scheduling point](#).

Execution Model Events

The *target-data-op-begin* [event](#) occurs before a [thread](#) initiates filling the [memory](#) in a [memory-setting routine region](#). The *target-data-op-end* [event](#) occurs after a [thread](#) initiates filling the [memory](#) in a [memory-setting routine region](#).

Tool Callbacks

A [thread](#) dispatches a registered [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 [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](#) occur in the context of the [target task](#).

Restrictions

The restrictions to the [memory-setting routines](#) are as follows:

- The value of the *ptr* argument must be a valid pointer to [device memory](#) for the [device](#) denoted by the value of the *device_num* argument.

Cross References

- OMPT [scope_endpoint](#) Type, see [Section 39.27](#)
- [target_data_op_emi](#) Callback, see [Section 41.7](#)

31.8.1 omp_target_memset Routine

Name: omp_target_memset	Properties: device-memory-routine , generating-task-binding , iso_c_binding , memory-setting
Category: function	

Return Type and Arguments

Name	Type	Properties
<i><return type></i>	c_ptr	default
<i>ptr</i>	c_ptr	iso_c , value
<i>val</i>	c_int	iso_c , value
<i>count</i>	c_size_t	iso_c , value
<i>device_num</i>	c_int	iso_c , value

1 **Prototypes**

C / C++

2 **void** *omp_target_memset(void *ptr, int val, size_t count,
3 int device_num);

C / C++

Fortran

4 **type** (c_ptr) **function** omp_target_memset(ptr, val, count, &
5 device_num) **bind**(c)
6 **use**, **intrinsic** :: iso_c_binding, **only** : c_ptr, c_int, &
7 c_size_t
8 **type** (c_ptr), **value** :: ptr
9 **integer** (kind=c_int), **value** :: val, device_num
10 **integer** (kind=c_size_t), **value** :: count

Fortran

11 **Effect**

12 As a [memory-setting routine](#), the effect of the **omp_target_memset** routine is as described in
13 [Section 31.8](#). This effect includes the [tool events](#) and [callbacks](#) defined in that section.

14 **Cross References**

- 15
 - Memory Setting Routines, see [Section 31.8](#)

16 **31.8.2 omp_target_memset_async Routine**

17 **Name:** omp_target_memset_async
 Category: [function](#)

Properties: [asynchronous-device-](#)
[routine](#), [device-memory-routine](#),
[generating-task-binding](#), [iso_c_bind-](#)
[ing](#), [memory-setting](#)

18 **Return Type and Arguments**

Name	Type	Properties
<return type>	c_ptr	default
ptr	c_ptr	iso_c , value
val	c_int	iso_c , value
count	c_size_t	iso_c , value
device_num	c_int	iso_c , value
depobj_count	c_int	iso_c , value
depobj_list	depend	optional , pointer

Prototypes

C / C++

```
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

```
type (c_ptr) function omp_target_memset_async(ptr, val, count, &  
device_num, depobj_count, depobj_list) bind(c)  
use, intrinsic :: iso_c_binding, only : c_ptr, c_int, &  
c_size_t  
type (c_ptr), value :: ptr  
integer (kind=c_int), value :: val, device_num, depobj_count  
integer (kind=c_size_t), value :: count  
integer (kind=omp_depend_kind), optional :: depobj_list(*)
```

Fortran

Effect

As a [memory-setting routine](#), the effect of the `omp_target_memset_async` routine is as described in [Section 31.8](#). This effect includes the [tool events](#) and [callbacks](#) defined in that section. As it is also an [asynchronous device routine](#), the routine also includes the [tool events](#) and [callbacks](#) defined in [Section 31.1](#).

Cross References

- Asynchronous Device Memory Routines, see [Section 31.1](#)
- Memory Setting Routines, see [Section 31.8](#)

32 Interoperability Routines

This section describes [interoperability routines](#), which have the [interoperability-routine](#) property. These [routines](#) provide mechanisms to inspect the [properties](#) associated with an [interoperability object](#). Each [interoperability routine](#) takes an *interop* argument of the [interop](#) OpenMP type. Most [interoperability routines](#) also take a *property_id* argument of the [interop_property](#) OpenMP type and a *ret_code* argument of (pointer to) [interop_rc](#) OpenMP type.

[Interoperability-property-retrieving routines](#), which have the [interoperability-property-retrieving](#) property, retrieve an [interoperability property](#) from an [interoperability object](#). For these [routines](#), if a [non-null pointer](#) is passed to the *ret_code* argument, an [interop_rc](#) OpenMP type value that indicates the return code is stored in the object to which *ret_code* points. If an error occurred, the stored value is negative and matches the error as defined in Table 26.3. On success, [omp_irc_success](#) is stored. If no error occurred but no meaningful value can be returned, [omp_irc_no_value](#) is stored.

[Interoperability-property-retrieving routines](#) return the requested [interoperability property](#), if available, and zero if an error occurs or no value is available. If the *interop* argument is [omp_interop_none](#), an empty error occurs. If the *property_id* argument is greater than or equal to [omp_get_num_interop_properties](#) (*interop*) or less than [omp_ipr_first](#), an out-of-range error occurs. If the requested property value is not convertible into a value of the type that the specific [interoperability-property-retrieving routine](#) retrieves, a type error occurs.

Restrictions

Restrictions to [interoperability routines](#) are as follows:

- Providing an invalid [interoperability object](#) for the *interop* argument results in [unspecified behavior](#).
- For any [interoperability routine](#) that returns a pointer, memory referenced by the pointer is managed by the OpenMP implementation and should not be freed or modified and memory referenced by that pointer cannot be accessed after the [interoperability object](#) that was used to obtain the pointer is destroyed.

Cross References

- OpenMP Interoperability Support Types, see [Section 26.7](#)

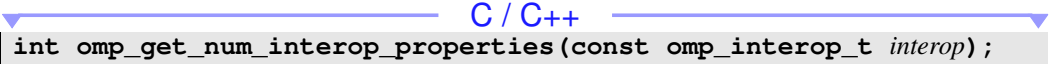
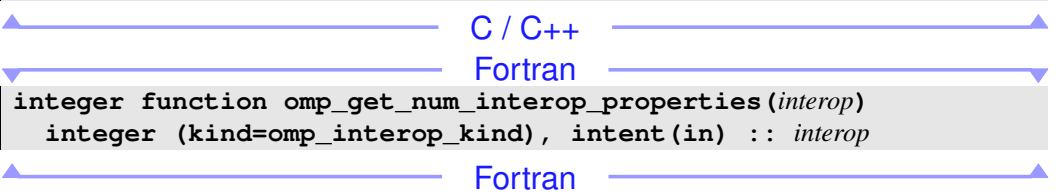
32.1 omp_get_num_interop_properties Routine

Name: <code>omp_get_num_interop_properties</code> Category: function	Properties: interoperability-routine
---	---

Return Type and Arguments

Name	Type	Properties
<return type>	integer	default
<i>interop</i>	interop	intent(in)

Prototypes

 <code>int omp_get_num_interop_properties(const omp_interop_t interop);</code>
 <code>integer function omp_get_num_interop_properties(interop) integer (kind=omp_interop_kind), intent(in) :: interop</code>

Effect

The `omp_get_num_interop_properties` routine returns the number of [implementation defined interoperability properties](#) available for *interop*. The total number of [properties](#) available for *interop* is the returned value minus `omp_ipr_first`.

Cross References

- OpenMP `interop` Type, see [Section 26.7.1](#)

32.2 omp_get_interop_int Routine

Name: <code>omp_get_interop_int</code> Category: function	Properties: interoperability-property-retrieving , interoperability-routine
--	--

Return Type and Arguments

Name	Type	Properties
<return type>	intptr	default
<i>interop</i>	interop	omp , opaque , intent(in)
<i>property_id</i>	interop_property	omp
<i>ret_code</i>	interop_rc	omp , intent(out) , optional

Prototypes

C / C++

```
omp_intptr_t omp_get_interop_int(const omp_interop_t interop,
    omp_interop_property_t property_id, omp_interop_rc_t *ret_code);
```

C / C++

Fortran

```
integer (kind=c_intptr_t) function omp_get_interop_int(interop, &
  property_id, ret_code)
  use, intrinsic :: iso_c_binding, only : c_intptr_t
  integer (kind=omp_interop_kind), intent(in) :: interop
  integer (kind=omp_interop_property_kind) property_id
  integer (kind=omp_interop_rc_kind), intent(out), optional :: &
    ret_code
```

Fortran

Effect

The `omp_get_interop_int` routine is an interoperability-property-retrieving routine that retrieves an interoperability property of integer type, if available.

Cross References

- OpenMP **interop** Type, see [Section 26.7.1](#)
- OpenMP **interop_property** Type, see [Section 26.7.3](#)
- OpenMP **interop_rc** Type, see [Section 26.7.4](#)
- **omp_get_num_interop_properties** Routine, see [Section 32.1](#)

32.3 omp_get_interop_ptr Routine

Name: <code>omp_get_interop_ptr</code> Category: function	Properties: interoperability-property-retrieving , interoperability-routine
--	--

Return Type and Arguments

Name	Type	Properties
<i><return type></i>	c_ptr	<i>default</i>
<i>interop</i>	interop	omp, opaque, intent(in)
<i>property_id</i>	interop_property	omp
<i>ret_code</i>	interop_rc	omp, intent(out), optional

1

2

3

4

5

1

2

3

4

- 5

9

20

21

22

1 Prototypes

C / C++

2 `const char *omp_get_interop_str(const omp_interop_t interop,`
3 `omp_interop_property_t property_id, omp_interop_rc_t *ret_code);`

C / C++

Fortran

4 `character(:) function omp_get_interop_str(interop, property_id, &`
5 `ret_code)`
6 `pointer :: omp_get_interop_str`
7 `integer (kind=omp_interop_kind), intent(in) :: interop`
8 `integer (kind=omp_interop_property_kind) property_id`
9 `integer (kind=omp_interop_rc_kind), intent(out), optional :: &`
10 `ret_code`

Fortran

11 Effect

12 The `omp_get_interop_str` routine is an interoperability-property-retrieving routine that
13 retrieves an interoperability string `property` type as a string, if available.

14 Cross References

- 15 • OpenMP `interop` Type, see [Section 26.7.1](#)
- 16 • OpenMP `interop_property` Type, see [Section 26.7.3](#)
- 17 • OpenMP `interop_rc` Type, see [Section 26.7.4](#)
- 18 • `omp_get_num_interop_properties` Routine, see [Section 32.1](#)

19 32.5 `omp_get_interop_name` Routine

20 **Name:** `omp_get_interop_name`
Category: [function](#)

Properties: [interoperability-routine](#)

21 Return Type and Arguments

22

Name	Type	Properties
<code><return type></code>	<code>const char</code>	pointer
<code>interop</code>	<code>interop</code>	omp , opaque , intent(in)
<code>property_id</code>	<code>interop_property</code>	omp

1 Prototypes

C / C++

2 `const char *omp_get_interop_name(const omp_interop_t interop,`
3 `omp_interop_property_t property_id);`

C / C++

Fortran

4 `character(:) function omp_get_interop_name(interop, property_id)`
5 `pointer :: omp_get_interop_name`
6 `integer (kind=omp_interop_kind), intent(in) :: interop`
7 `integer (kind=omp_interop_property_kind) property_id`

Fortran

8 Effect

9 The `omp_get_interop_name` routine returns, as a string, the name of the interoperability
10 property identified by `property_id`. Property names for non-implementation defined interoperability
11 properties are listed in Table 26.2. If the `property_id` is less than `omp_ipr_first` or greater than
12 or equal to `omp_get_num_interop_properties` (`interop`), NULL is returned.

13 Cross References

- 14 • OpenMP `interop` Type, see Section 26.7.1
- 15 • OpenMP `interop_property` Type, see Section 26.7.3
- 16 • `omp_get_num_interop_properties` Routine, see Section 32.1

17 32.6 omp_get_interop_type_desc Routine

Name: <code>omp_get_interop_type_desc</code>	Properties: interoperability-routine
Category: <code>function</code>	

19 Return Type and Arguments

Name	Type	Properties
<return type>	const char	pointer
<code>interop</code>	interop	omp, opaque, intent(in)
<code>property_id</code>	interop_property	omp

1 **Prototypes**

C / C++

2 **const char *omp_get_interop_type_desc**(
3 **const omp_interop_t** *interop*, **omp_interop_property_t** *property_id*);

C / C++

Fortran

4 **character(:) function omp_get_interop_type_desc**(*interop*, &
5 *property_id*)
6 **pointer :: omp_get_interop_type_desc**
7 **integer (kind=omp_interop_kind), intent(in) ::** *interop*
8 **integer (kind=omp_interop_property_kind)** *property_id*

Fortran

9 **Effect**

10 The **omp_get_interop_type_desc** routine returns a string that describes the type of the
11 **interoperability property** identified by *property_id* in human-readable form. The description may
12 contain a valid type declaration, possibly followed by a description or name of the type. If *interop*
13 has the value **omp_interop_none**, or if the *property_id* is less than **omp_ipr_first** or
14 greater than or equal to **omp_get_num_interop_properties** (*interop*), **NULL** is returned.

15 **Cross References**

- 16 • OpenMP **interop** Type, see [Section 26.7.1](#)
- 17 • OpenMP **interop_property** Type, see [Section 26.7.3](#)
- 18 • **omp_get_num_interop_properties** Routine, see [Section 32.1](#)

19 **32.7 omp_get_interop_rc_desc Routine**

20 **Name:** **omp_get_interop_rc_desc**
 Category: [function](#)

Properties: [interoperability-routine](#)

21 **Return Type and Arguments**

Name	Type	Properties
<return type>	const char	pointer
<i>interop</i>	interop	omp , opaque , intent(in)
<i>ret_code</i>	interop_rc	omp

Prototypes

C / C++

```
const char *omp_get_interop_rc_desc(const omp_interop_t interop,  
omp_interop_rc_t ret_code);
```

C / C++

Fortran

```
character(:) function omp_get_interop_rc_desc(interop, ret_code)  
  pointer :: omp_get_interop_rc_desc  
  integer (kind=omp_interop_kind), intent(in) :: interop  
  integer (kind=omp_interop_rc_kind) ret_code
```

Fortran

Effect

The `omp_get_interop_rc_desc` routine returns a string that describes the return code `ret_code` associated with an `interoperability object` 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 `ret_code` was not last written by an `interoperability routine` invoked with the `interoperability object` `interop`.

Cross References

- OpenMP `interop` Type, see [Section 26.7.1](#)
- OpenMP `interop_property` Type, see [Section 26.7.3](#)
- OpenMP `interop_rc` Type, see [Section 26.7.4](#)
- `omp_get_num_interop_properties` Routine, see [Section 32.1](#)

33 Memory Management Routines

This chapter describes OpenMP [memory-management routines](#), which are [OpenMP API routines](#) that have the [memory-management-routine](#) property. These [routines](#) support [memory](#) management on the [current device](#).

Fortran

The Fortran versions of the [memory-management routines](#) require an explicit interface and thus might not be provided in the deprecated include file `omp_lib.h`.

Fortran

33.1 Memory Space Retrieving Routines

This section describes the [memory-space-retrieving routines](#), which are [routines](#) that have the [memory-space-retrieving](#) property. Each of these [routines](#) returns a [handle](#) to a [memory space](#) that represents a set of storage resources accessible by one or more [devices](#). 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 the [memory space](#) represented by the *memspace* argument 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. These [routines](#) have the [all-device-threads binding](#) property for each [device](#) selected by the [routine](#). Thus, the [binding thread set](#) for a [region](#) that corresponds to a [memory-space-retrieving routine](#) is [all threads](#) on the [devices](#) selected by the [routine](#).

The [memory spaces](#) returned by these [routines](#) are [target memory spaces](#) if any of the selected [devices](#) is not the [current device](#).

For any [memory-space-retrieving routine](#) that takes a *devs* argument, if the array to which the argument points has more than *ndevs* values, the additional values are ignored.

1 **Restrictions**

2 The restrictions to [memory-space-retrieving routines](#) are as follows:

- 3 • These [routines](#) must only be invoked on the [host device](#).
- 4 • The *memspace* argument must be one of the predefined [memory spaces](#).
- 5 • For any [memory-space-retrieving routine](#) that has a *devs* argument, the argument must point
- 6 to an array that contains at least *ndevs* values.
- 7 • For any [memory-space-retrieving routine](#) that has a *dev* or *devs* argument, the value of the
- 8 *dev* argument the *ndevs* values of the array to which *devs* points must be [conforming device](#)
- 9 [numbers](#).

10 **Cross References**

- 11 • Memory Spaces, see [Section 13.1](#)
- 12 • **requires** Directive, see [Section 16.5](#)
- 13 • **target** Construct, see [Section 21.8](#)

14 **33.1.1 omp_get_devices_memspace Routine**

Name: <code>omp_get_devices_memspace</code> Category: function	Properties: all-device-threads-binding , memory-management-routine , memory-space-retrieving
---	---

16 **Return Type and Arguments**

Name	Type	Properties
<return type>	<code>memspace_handle</code>	default
<i>ndevs</i>	integer	intent(in) , positive
<i>devs</i>	integer	intent(in) , pointer
<i>memspace</i>	<code>memspace_handle</code>	intent(in) , omp

18 **Prototypes**

C / C++

19 `omp_memspace_handle_t omp_get_devices_memspace(int ndevs,`

20 `const int *devs, omp_memspace_handle_t memspace);`

C / C++

Fortran

21 `integer (kind=omp_memspace_handle_kind) function &`

22 `omp_get_devices_memspace(ndevs, devs, memspace)`

23 `integer, intent(in) :: ndevs, devs(*)`

24 `integer (kind=omp_memspace_handle_kind), intent(in) :: memspace`

Fortran

Effect

The `omp_get_devices_memspace` routine is a [memory-space-retrieving routine](#). The `devices` selected by the [routine](#) are those specified in the `devs` argument.

Cross References

- [Memory Space Retrieving Routines](#), see [Section 33.1](#)
- [OpenMP `memspace_handle` Type](#), see [Section 26.8.12](#)

33.1.2 `omp_get_device_memspace` Routine

Name: <code>omp_get_device_memspace</code> Category: function	Properties: all-device-threads-binding , memory-management-routine , memory-space-retrieving
--	---

Return Type and Arguments

Name	Type	Properties
<i><return type></i>	<code>memspace_handle</code>	default
<i>dev</i>	integer	intent(in)
<i>memspace</i>	<code>memspace_handle</code>	intent(in) , omp

Prototypes

C / C++

omp_memspace_handle_t omp_get_device_memspace(int dev,
omp_memspace_handle_t memspace);

C / C++

Fortran

integer (kind=omp_memspace_handle_kind) function &
omp_get_device_memspace(dev, memspace)
integer, intent(in) :: dev
integer (kind=omp_memspace_handle_kind), intent(in) :: memspace

Fortran

Effect

The `omp_get_device_memspace` routine is a [memory-space-retrieving routine](#). The `device` selected by the [routine](#) is the `device` specified in the `dev` argument.

Cross References

- [Memory Space Retrieving Routines](#), see [Section 33.1](#)
- [OpenMP `memspace_handle` Type](#), see [Section 26.8.12](#)

33.1.3 omp_get_devices_and_host_memspace Routine

Name: <code>omp_get_devices_and_host_memspace</code> Category: function	Properties: all-device-threads-binding , memory-management-routine , memory-space-retrieving
---	---

Return Type and Arguments

Name	Type	Properties
<i><return type></i>	<code>memspace_handle</code>	default
<i>ndevs</i>	<code>integer</code>	intent(in) , positive
<i>devs</i>	<code>integer</code>	intent(in) , pointer
<i>memspace</i>	<code>memspace_handle</code>	intent(in) , omp

Prototypes

C / C++

```
omp_memspace_handle_t omp_get_devices_and_host_memspace(  
    int ndevs, const int *devs, omp_memspace_handle_t memspace);
```

C / C++

Fortran

```
integer (kind=omp_memspace_handle_kind) function &  
    omp_get_devices_and_host_memspace(ndevs, devs, memspace)  
integer, intent(in) :: ndevs, devs(*)  
integer (kind=omp_memspace_handle_kind), intent(in) :: memspace
```

Fortran

Effect

The `omp_get_devices_and_host_memspace` routine is a [memory-space-retrieving routine](#). The [devices](#) selected by the [routine](#) are the [host device](#) and those specified in the *devs* argument.

Cross References

- [Memory Space Retrieving Routines](#), see [Section 33.1](#)
- OpenMP `memspace_handle` Type, see [Section 26.8.12](#)

33.1.4 omp_get_device_and_host_memspace Routine

Name: <code>omp_get_device_and_host_memspace</code> Category: function	Properties: all-device-threads-binding , memory-management-routine , memory-space-retrieving
--	---

1 **Return Type and Arguments**

Name	Type	Properties
<return type>	memspace_handle	default
<i>dev</i>	integer	intent(in)
<i>memspace</i>	memspace_handle	intent(in), omp

3 **Prototypes**

C / C++

4 **omp_memspace_handle_t** **omp_get_device_and_host_memspace**(int *dev*,
5 **omp_memspace_handle_t** *memspace*);

C / C++

Fortran

6 **integer** (kind=omp_memspace_handle_kind) **function** &
7 **omp_get_device_and_host_memspace**(*dev*, *memspace*)
8 **integer**, **intent(in)** :: *dev*
9 **integer** (kind=omp_memspace_handle_kind), **intent(in)** :: *memspace*

Fortran

10 **Effect**

11 The **omp_get_device_and_host_memspace** routine is a [memory-space-retrieving routine](#).
12 The [devices](#) selected by the [routine](#) are the [host device](#) and the [device](#) specified in the *dev* argument.

13 **Cross References**

- 14 • Memory Space Retrieving Routines, see [Section 33.1](#)
- 15 • OpenMP **memspace_handle** Type, see [Section 26.8.12](#)

16 **33.1.5 omp_get_devices_all_memspace Routine**

Name: omp_get_devices_all_memspace	Properties: all-device-threads-binding , memory-management-routine , memory-space-retrieving
Category: function	

18 **Return Type and Arguments**

Name	Type	Properties
<return type>	memspace_handle	default
<i>memspace</i>	memspace_handle	intent(in), omp

1

2

3

4

5

6

7

8

0

- 1

3

4

5

6

7

8

9

20

21

656

Effect

The `omp_get_memspace_num_resources` routine is a memory-management routine that returns the number of distinct storage resources that are associated with the memory space represented by the *memspace handle*.

Restrictions

The restrictions to the `omp_get_memspace_num_resources` routine are as follows:

- The *memspace* argument must be a valid memory space.

Cross References

- Memory Spaces, see [Section 13.1](#)
- OpenMP `memspace_handle` Type, see [Section 26.8.12](#)

33.3 omp_get_memspace_pagesize Routine

Name: <code>omp_get_memspace_pagesize</code>	Properties: <code>all-device-threads-binding</code> , <code>iso_c_binding</code> , <code>memory-management-routine</code>
Category: <code>function</code>	

Return Type and Arguments

Name	Type	Properties
<i><return type></i>	<code>c_size_t</code>	<code>default</code>
<i>memspace</i>	<code>memspace_handle</code>	<code>intent(in)</code> , <code>omp</code>

Prototypes

C / C++

```
size_t omp_get_memspace_pagesize(omp_memspace_handle_t memspace);
```

C / C++

Fortran

```
integer (kind=c_size_t) function omp_get_memspace_pagesize(&  
    memspace) bind(c)  
    use, intrinsic :: iso_c_binding, only : c_size_t  
    integer (kind=omp_memspace_handle_kind), intent(in) :: memspace
```

Fortran

Effect

The `omp_get_memspace_pagesize` routine is a memory-management routine that returns the page size that the memory space represented by the *memspace handle* supports.

1 **Restrictions**

2 The restrictions to the `omp_get_memspace_pagesize` routine are as follows:

- 3 • The *memspace* argument must be a valid [memory space](#).

4 **Cross References**

- 5 • Memory Spaces, see [Section 13.1](#)
- 6 • OpenMP `memspace_handle` Type, see [Section 26.8.12](#)

7

33.4 omp_get_submemspace Routine

8

Name: <code>omp_get_submemspace</code>	Properties: all-device-threads-binding , memory-management-routine
Category: function	

9 **Return Type and Arguments**

10

Name	Type	Properties
<i><return type></i>	<code>memspace_handle</code>	default
<i>memspace</i>	<code>memspace_handle</code>	intent(in) , omp
<i>num_resources</i>	<code>integer</code>	intent(in) , non-negative
<i>resources</i>	<code>integer</code>	intent(in) , pointer

11 **Prototypes**

12

13

14

<div><div>C / C++</div><div><pre>omp_memspace_handle_t omp_get_submemspace(omp_memspace_handle_t memspace, int num_resources, const int *resources);</pre></div></div>	<div>C / C++</div>
<div><div>Fortran</div><div><pre>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, resources(*)</pre></div></div>	<div>Fortran</div>

19 **Effect**

20 The `omp_get_submemspace` routine is a [memory-management routine](#) that returns a new
21 [memory space](#) that contains a subset of the resources of the original [memory space](#). The new
22 [memory space](#) represents only the resources of the [memory space](#) represented by the *memspace*
23 [handle](#) that are specified by the *resources* argument. If *num_resources* is zero or a [memory space](#)
24 cannot be created for the requested resources, the special value `omp_null_mem_space` is
25 returned.

Restrictions

The restrictions to the `omp_get_submemspace` routine are as follows:

- The *memspace* argument must be a valid [memory space](#).
- The *resources* array must contain at least as many entries as specified by the *num_resources* argument.
- The value of each entry of the *resources* array must be between 0 and one less than the number of resources associated with the [memory space](#) represented by the *memspace* argument.

Cross References

- Memory Spaces, see [Section 13.1](#)
- OpenMP `memspace_handle` Type, see [Section 26.8.12](#)

33.5 OpenMP Memory Partitioning Routines

This section describes the [memory-partitioning routines](#), which are [routines](#) that have the [memory-partitioning property](#). These [routines](#) provide mechanisms to create and to use [memory partitioners](#).

33.5.1 `omp_init_mempartitioner` Routine

Name: <code>omp_init_mempartitioner</code>	Properties: all-device-threads-binding , memory-management-routine , memory-partitioning
Category: subroutine	

Arguments

Name	Type	Properties
<i>partitioner</i>	<code>mempartitioner</code>	C/C++ pointer , omp , intent(out)
<i>lifetime</i>	<code>mempartitioner_lifetime</code>	omp , intent(in)
<i>compute_proc</i>	<code>mempartitioner_compute_proc</code>	omp , procedure
<i>release_proc</i>	<code>mempartitioner_release_proc</code>	omp , procedure

Prototypes

C / C++

```
void omp_init_mempartitioner(omp_mempartitioner_t *partitioner,  
    omp_mempartitioner_lifetime_t lifetime,  
    omp_mempartitioner_compute_proc_t compute_proc,  
    omp_mempartitioner_release_proc_t release_proc);
```

C / C++

Fortran

```
subroutine omp_init_mempartitioner(partitioner, lifetime, &  
    compute_proc, release_proc)  
    integer (kind=omp_mempartitioner_kind), intent(out) :: &  
        partitioner  
    integer (kind=omp_mempartitioner_lifetime_kind), intent(in) ::  
        &  
        lifetime  
    procedure (omp_mempartitioner_compute_proc_t) compute_proc  
    procedure (omp_mempartitioner_release_proc_t) release_proc
```

Fortran

Effect

The `omp_init_mempartitioner` routine initializes the `memory partitioner` that the `partitioner` object represents with the lifetime specified by the `lifetime` argument, and the `compute_proc` partition computation `procedure` and the `release_proc` partition release `procedure`.

Once initialized the `partitioner` object can be associated with an `allocator` when the `allocator` is initialized with `omp_init_allocator` by using the `omp_atk_partitioner` trait. If the `omp_atk_partition` allocator trait is set to `omp_atv_partitioner`, then, for allocations that use the `allocator`, the number of `memory` parts of an allocation and how they are distributed across the storage resources are defined by a `memory partition` object that must be initialized in the `compute_proc` provided in this routine through calls to the `omp_init_mempartition` and `omp_mempartition_set_part` routines.

If the value of the `lifetime` argument is `omp_allocator_mempartition` then the `memory partition` object that is created through the `compute_proc` `procedure` might be used for all allocations of an `allocator` that has the same allocation size. If the value of the `lifetime` argument is `omp_dynamic_mempartition` then a `memory partition` object will be initialized for every allocation.

Restrictions

The restrictions to the `omp_init_mempartitioner` routine are as follows:

- The `memory partitioner` represented by the `partitioner` argument must be in the `uninitialized` state.

Cross References

- Memory Allocators, see [Section 13.2](#)
- Memory Spaces, see [Section 13.1](#)
- OpenMP `mempartitioner` Type, see [Section 26.8.8](#)
- OpenMP `mempartitioner_compute_proc` Type, see [Section 26.8.10](#)
- OpenMP `mempartitioner_lifetime` Type, see [Section 26.8.9](#)
- OpenMP `mempartitioner_release_proc` Type, see [Section 26.8.11](#)

33.5.2 `omp_destroy_mempartitioner` Routine

Name: <code>omp_destroy_mempartitioner</code> Category: subroutine	Properties: all-device-threads-binding , memory-management-routine , memory-partitioning
---	--

Arguments

Name	Type	Properties
<i>partitioner</i>	<code>mempartitioner</code>	C/C++ pointer , omp , intent(in)

Prototypes

C / C++	<pre>void omp_destroy_mempartitioner(const omp_mempartitioner_t *partitioner);</pre>
C / C++	
Fortran	<pre>subroutine omp_destroy_mempartitioner(partitioner) integer (kind=omp_mempartitioner_kind), intent(in) :: & partitioner</pre>
Fortran	

Effect

The effect of the `omp_destroy_mempartitioner` routine is to uninitialized a [memory partitioner](#). Thus, the routine changes the state of the [memory partitioner](#) object represented by the *partitioner* argument to uninitialized and releases all resources associated with it.

Restrictions

The restrictions to the `omp_destroy_mempartitioner` routine are as follows:

- The [memory partitioner](#) represented by the *partitioner* argument must be in the initialized state.
- Any [allocator](#) that references the [memory partitioner](#) object represented by the *partitioner* argument must be destroyed before this routine is called.

Cross References

- Memory Allocators, see [Section 13.2](#)
- OpenMP `mempartitioner` Type, see [Section 26.8.8](#)

33.5.3 `omp_init_mempartition` Routine

Name: <code>omp_init_mempartition</code> Category: subroutine	Properties: all-device-threads-binding , iso_c_binding , memory-management-routine , memory-partitioning
--	---

Arguments

Name	Type	Properties
<i>partition</i>	<code>mempartition</code>	C/C++ pointer , omp , intent(out)
<i>nparts</i>	<code>c_size_t</code>	intent(in) , iso_c , intent(in)
<i>user_data</i>	<code>c_ptr</code>	intent(in) , iso_c , intent(in)

Prototypes

C / C++

```
void omp_init_mempartition(omp_mempartition_t *partition,  
    size_t nparts, const void *user_data);
```

C / C++

Fortran

```
subroutine omp_init_mempartition(partition, nparts, user_data) &  
    bind(c)  
    use, intrinsic :: iso_c_binding, only : c_size_t, c_ptr  
    integer (kind=omp_mempartition_kind), intent(out) :: partition  
    integer (kind=c_size_t), intent(in) :: nparts  
    type (c_ptr), intent(in) :: user_data
```

Fortran

Effect

The effect of the `omp_init_mempartition` routine is to initialize a [memory partition](#) object. Thus, the routine sets the [memory partition](#) object indicated by the *partition* argument to represent a [memory partition](#) of *nparts* parts and associates the user data indicated by the *user_data* argument with it.

Restrictions

The restrictions to the `omp_init_mempartition` routine are as follows:

- The `memory partition` represented by the `partition` argument must be in the `uninitialized state`.
- This `routine` must only be called by a `procedure` that is associated with the `memory partitioner` object that allocated the `memory partition` indicated by the `partition` argument.

Cross References

- OpenMP Memory Management Types, see [Section 26.8](#)
- OpenMP `mempartitioner` Type, see [Section 26.8.8](#)

33.5.4 omp_destroy_mempartition Routine

Name: <code>omp_destroy_mempartition</code>	Properties: <code>all-device-threads-binding</code> , <code>memory-management-routine</code> , <code>memory-partitioning</code>
Category: <code>subroutine</code>	

Arguments

Name	Type	Properties
<code>partition</code>	<code>mempartition</code>	<code>C/C++ pointer</code> , <code>omp</code> , <code>intent(in)</code>

Prototypes

	C / C++	
<code>void omp_destroy_mempartition(</code>		
<code> const omp_mempartition_t *partition);</code>		
	C / C++	
	Fortran	
<code>subroutine omp_destroy_mempartition(partition)</code>		
<code> integer (kind=omp_mempartition_kind), intent(in) :: partition</code>		
	Fortran	

Effect

The effect of the `omp_destroy_mempartition` routine is to uninitialized a `memory partition` object. Thus, the `routine` releases the `memory partition` indicated by the `partition` argument and all resources associated with it.

Restrictions

The restrictions to the `omp_destroy_mempartition` routine are as follows:

- The `memory partition` represented by the `partition` argument must be in the `initialized state`.
- This `routine` must only be called by a `procedure` that is associated with the `memory partitioner` object that allocated the `memory partition` indicated by the `partition` argument.

Cross References

- OpenMP Memory Management Types, see [Section 26.8](#)
- OpenMP `mempartitioner` Type, see [Section 26.8.8](#)

33.5.5 `omp_mempartition_set_part` Routine

Name: <code>omp_mempartition_set_part</code> Category: function	Properties: all-device-threads-binding , iso_c_binding , memory-management-routine , memory-partitioning
--	---

Return Type and Arguments

Name	Type	Properties
<i><return type></i>	integer	default
<i>partition</i>	mempartition	C/C++ pointer , omp , intent(out)
<i>part</i>	<code>c_size_t</code>	intent(in) , iso_c
<i>resource</i>	integer	intent(in) , iso_c
<i>size</i>	<code>c_size_t</code>	intent(in) , iso_c

Prototypes

C / C++

```
int omp_mempartition_set_part(omp_mempartition_t *partition,  
    size_t part, int resource, size_t size);
```

C / C++

Fortran

```
integer function omp_mempartition_set_part(partition, part, &  
    resource, size) bind(c)  
    use, intrinsic :: iso_c_binding, only : c_size_t  
    integer (kind=omp_mempartition_kind), intent(out) :: partition  
    integer (kind=c_size_t), intent(in) :: part, size  
    integer, intent(in) :: resource
```

Fortran

Effect

The effect of the `omp_mempartition_set_part` routine is to define the size and resource of a given part of a [memory partition](#). Thus the routine defines the part number indicated by the *part* argument of the [memory partition](#) object indicated by the *partition* argument to be associated to the resource indicated by the *resource* argument and to be of size indicated by the *size* argument.

The size of all parts of a [memory partition](#), except the last one, need to be a multiple of the page size that the [memory space](#) where the [memory](#) is being allocated supports. If the specified *size* cannot be supported by the specified *resource*, this routine returns negative one. Otherwise, it returns zero.

Restrictions

The restrictions to the `omp_mempartition_set_part` routine are as follows:

- The `memory partition` represented by the `partition` argument must be in the initialized state.
- This routine must only be called by a `procedure` that is associated with the `memory partitioner` object that allocated the `memory partition` indicated by the `partition` argument.

Cross References

- Memory Spaces, see [Section 13.1](#)
- OpenMP Memory Management Types, see [Section 26.8](#)
- OpenMP `mempartitioner` Type, see [Section 26.8.8](#)

33.5.6 omp_mempartition_get_user_data Routine

Name: <code>omp_mempartition_get_user_data</code> Category: function	Properties: all-device-threads-binding , iso_c_binding , memory-management-routine , memory-partitioning
---	---

Return Type and Arguments

Name	Type	Properties
<code><return type></code>	<code>c_ptr</code>	default
<code>partition</code>	<code>mempartition</code>	intent(in) , C/C++ pointer , omp

Prototypes

C / C++

```
void *omp_mempartition_get_user_data(  
    const omp_mempartition_t *partition);
```

C / C++

Fortran

```
type (c_ptr) function omp_mempartition_get_user_data(partition) &  
    bind(c)  
    use, intrinsic :: iso_c_binding, only : c_ptr  
    integer (kind=omp_mempartition_kind), intent(in) :: partition
```

Fortran

Effect

The effect of the `omp_mempartition_get_user_data` routine is to retrieve the user data that was associated with the `memory partition` when it was created. Thus, the routine returns the data associated with the `memory partition` object indicated by the `partition` argument.

1 **Restrictions**

2 The restrictions to the `omp_mempartition_get_user_data` routine are as follows:

- 3
 - The `memory partition` represented by the `partition` argument must be in the initialized state.
 - This routine must only be called by a `procedure` that is associated with the `memory partitioner` object that allocated the `memory partition` indicated by the `partition` argument.

6 **Cross References**

- 7
 - OpenMP Memory Management Types, see [Section 26.8](#)
 - OpenMP `mempartitioner` Type, see [Section 26.8.8](#)

9 **33.6 omp_init_allocator Routine**

Name: <code>omp_init_allocator</code> Category: <code>function</code>	Properties: <code>all-device-threads-binding</code> , <code>memory-management-routine</code>
--	--

11 **Return Type and Arguments**

Name	Type	Properties
<code><return type></code>	<code>allocator_handle</code>	<code>default</code>
<code>memspace</code>	<code>memspace_handle</code>	<code>intent(in)</code> , <code>omp</code>
<code>ntraits</code>	<code>integer</code>	<code>intent(in)</code>
<code>traits</code>	<code>alloctrait</code>	<code>intent(in)</code> , <code>pointer</code> , <code>omp</code>

13 **Prototypes**

C / C++

```
omp_allocator_handle_t omp_init_allocator(  
    omp_memspace_handle_t memspace, int ntraits,  
    const omp_alloctrait_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  
    integer (kind=omp_alloctrait_kind), intent(in) :: traits(*)
```

Fortran

22 **Effect**

23 The `omp_init_allocator` routine creates a new `allocator` that is associated with the
24 `memspace` memory space and returns a `handle` to it. All allocations through the created `allocator`
25 will behave according to the `allocator traits` specified in the `traits` argument. The number of `traits` in
26 the `traits` argument is specified by the `ntraits` argument. If the special `omp_atv_default` value
27 is used for a given `trait`, then its value will be the default value specified in Table 13.2 for that `trait`.

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`. 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.

Restrictions

The restrictions to the `omp_init_allocator` routine are as follows:

- Each *allocator trait* must be specified at most once.
- The *memspace* argument must be a valid *memory space handle* or the value `omp_null_mem_space`.
- If the *ntraits* argument is *positive* then the *traits* argument must specify at least *ntraits traits*.
- The use of an *allocator* returned by this routine on *devices* 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 the *memspace handle* represents a *target memory space*, the values `omp_atv_device`, `omp_atv_cgroup`, `omp_atv_pteam` or `omp_atv_thread` must not be specified for the `omp_atk_access` allocator trait.

Cross References

- OpenMP `allocator_handle` Type, see [Section 26.8.2](#)
- Memory Allocators, see [Section 13.2](#)
- Memory Spaces, see [Section 13.1](#)
- OpenMP `memspace_handle` Type, see [Section 26.8.12](#)
- *requires* Directive, see [Section 16.5](#)
- *target* Construct, see [Section 21.8](#)

33.7 omp_destroy_allocator Routine

Name: <code>omp_destroy_allocator</code> Category: subroutine	Properties: all-device-threads-binding , memory-management-routine
--	--

Arguments

Name	Type	Properties
<i>allocator</i>	<code>allocator_handle</code>	intent(in) , omp

Prototypes

C / C++
void omp_destroy_allocator(omp_allocator_handle_t allocator);

C / C++
Fortran
subroutine omp_destroy_allocator(allocator)
 integer (kind=omp_allocator_handle_kind), intent(in) :: &
 allocator

Fortran

Effect

The **omp_destroy_allocator** routine releases all resources used to implement the *allocator handle*. If *allocator* is **omp_null_allocator** then this routine has no effect.

Restrictions

The restrictions to the **omp_destroy_allocator** routine are as follows:

- The *allocator* argument must not represent a predefined **memory allocator**.
- 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

- OpenMP **allocator_handle** Type, see [Section 26.8.2](#)
- Memory Allocators, see [Section 13.2](#)
- **requires** Directive, see [Section 16.5](#)
- **target** Construct, see [Section 21.8](#)

33.8 Memory Allocator Retrieving Routines

This section describes the **memory-allocator-retrieving routines**, which are routines that have the **memory-allocator-retrieving property**. Each of these routines returns a **handle** to a predefined **memory allocator** that represents the default **memory allocator** for a given **device** for a certain kind of **memory**. If the implementation does not have a predefined **allocator** that satisfies the request, then the special value **omp_null_allocator** is returned. For any **memory-allocator-retrieving routine** that takes a *devs* argument, if the array to which the argument points has more than *ndevs* values, the additional values are ignored. Each of these routines returns an **allocator** that may be used anywhere that requires a predefined **allocator** specified in Table 13.3. The **allocator** is associated with a **target memory space** if any of the selected **devices** is not the **current device**.

Restrictions

The restrictions to [memory-allocator-retrieving routines](#) are as follows:

- These [routines](#) must only be invoked on the [host device](#).
- The *memspace* argument must not be one of the predefined [memory spaces](#).
- For any [memory-allocator-retrieving routine](#) that has a *devs* argument, the argument must point to an array that contains at least *ndevs* values.
- For any [memory-allocator-retrieving routine](#) that has a *dev* or *devs* argument, the value of the *dev* argument the *ndevs* values of the array to which *devs* points must be [conforming device numbers](#).

Cross References

- Memory Allocators, see [Section 13.2](#)
- Memory Spaces, see [Section 13.1](#)

33.8.1 omp_get_devices_allocator Routine

Name: <code>omp_get_devices_allocator</code>	Properties: all-device-threads-binding , memory-management-routine , memory-allocator-retrieving
Category: function	

Return Type and Arguments

Name	Type	Properties
<i><return type></i>	<code>allocator_handle</code>	default
<i>ndevs</i>	integer	intent(in) , positive
<i>devs</i>	integer	intent(in) , pointer
<i>memspace</i>	<code>memspace_handle</code>	intent(in) , omp

Prototypes

C / C++

`omp_allocator_handle_t omp_get_devices_allocator(int ndevs,
const int *devs, omp_memspace_handle_t memspace);`

C / C++

Fortran

`integer (kind=omp_allocator_handle_kind) function &
omp_get_devices_allocator(ndevs, devs, memspace)
integer, intent(in) :: ndevs, devs(*)
integer (kind=omp_memspace_handle_kind), intent(in) :: memspace`

Fortran

Effect

The `omp_get_devices_allocator` routine is a [memory-allocator-retrieving routine](#). The [devices](#) selected by the [routine](#) are those specified in the *devs* argument.

Cross References

- OpenMP `allocator_handle` Type, see [Section 26.8.2](#)
- Memory Allocator Retrieving Routines, see [Section 33.8](#)
- OpenMP `memspace_handle` Type, see [Section 26.8.12](#)

33.8.2 omp_get_device_allocator Routine

Name: <code>omp_get_device_allocator</code> Category: function	Properties: all-device-threads-binding , memory-management-routine , memory-allocator-retrieving
---	---

Return Type and Arguments

Name	Type	Properties
<i><return type></i>	<code>allocator_handle</code>	default
<i>dev</i>	<code>integer</code>	intent(in)
<i>memspace</i>	<code>memspace_handle</code>	intent(in) , omp

Prototypes

C / C++

`omp_allocator_handle_t omp_get_device_allocator(int dev,
omp_memspace_handle_t memspace);`

C / C++

Fortran

`integer (kind=omp_allocator_handle_kind) function &
omp_get_device_allocator(dev, memspace)
integer, intent(in) :: dev
integer (kind=omp_memspace_handle_kind), intent(in) :: memspace`

Fortran

Effect

The `omp_get_device_allocator` routine is a [memory-allocator-retrieving routine](#). The [device](#) selected by the [routine](#) is the [device](#) specified in the *dev* argument.

Cross References

- OpenMP `allocator_handle` Type, see [Section 26.8.2](#)
- Memory Allocator Retrieving Routines, see [Section 33.8](#)
- OpenMP `memspace_handle` Type, see [Section 26.8.12](#)

1 **33.8.3 omp_get_devices_and_host_allocator Routine**

2 Name: omp_get_devices_and_host_allocator Category: function	Properties: all-device-threads-binding , memory-management-routine , memory-allocator-retrieving
--	---

3 **Return Type and Arguments**

Name	Type	Properties
<return type>	allocator_handle	default
<i>ndevs</i>	integer	intent(in) , positive
<i>devs</i>	integer	intent(in) , pointer
<i>memspace</i>	memspace_handle	intent(in) , omp

5 **Prototypes**

C / C++

```
omp_allocator_handle_t omp_get_devices_and_host_allocator(  
    int ndevs, const int *devs, omp_memspace_handle_t memspace);
```

C / C++

Fortran

```
integer (kind=omp_allocator_handle_kind) function &  
    omp_get_devices_and_host_allocator(ndevs, devs, memspace)  
    integer, intent(in) :: ndevs, devs(*)  
    integer (kind=omp_memspace_handle_kind), intent(in) :: memspace
```

Fortran

12 **Effect**

13 The [omp_get_devices_and_host_allocator](#) routine is a [memory-allocator-retrieving](#)
14 [routine](#). The [devices](#) selected by the [routine](#) are the [host device](#) and those specified in the *devs*
15 argument.

16 **Cross References**

- 17 • OpenMP `allocator_handle` Type, see [Section 26.8.2](#)
- 18 • Memory Allocator Retrieving Routines, see [Section 33.8](#)
- 19 • OpenMP `memspace_handle` Type, see [Section 26.8.12](#)

20 **33.8.4 omp_get_device_and_host_allocator Routine**

21 Name: omp_get_device_and_host_allocator Category: function	Properties: all-device-threads-binding , memory-management-routine , memory-allocator-retrieving
--	---

1 **Return Type and Arguments**

Name	Type	Properties
<return type>	allocator_handle	default
dev	integer	intent(in)
memspace	memspace_handle	intent(in), omp

3 **Prototypes**

4

C / C++

omp_allocator_handle_t omp_get_device_and_host_allocator(int dev,
omp_memspace_handle_t memspace);

6

C / C++

Fortran

integer (kind=omp_allocator_handle_kind) function &
omp_get_device_and_host_allocator(dev, memspace)
integer, intent(in) :: dev
integer (kind=omp_memspace_handle_kind), intent(in) :: memspace

10 **Effect**

11 The **omp_get_device_and_host_allocator** routine is a [memory-allocator-retrieving](#)
12 [routine](#). The [devices](#) selected by the [routine](#) are the [host device](#) and the [device](#) specified in the *dev*
13 argument.

14 **Cross References**

- 15 • OpenMP **allocator_handle** Type, see [Section 26.8.2](#)
- 16 • Memory Allocator Retrieving Routines, see [Section 33.8](#)
- 17 • OpenMP **memspace_handle** Type, see [Section 26.8.12](#)

18 **33.8.5 omp_get_devices_all_allocator Routine**

Name: <code>omp_get_devices_all_allocator</code> Category: function	Properties: all-device-threads-binding , memory-management-routine , memory-allocator-retrieving
--	---

20 **Return Type and Arguments**

Name	Type	Properties
<return type>	allocator_handle	default
memspace	memspace_handle	intent(in), omp

1

2

3

4

5

6

7

8

0

- 1

- 2

- 3

4

5

6

7

8

9

9

20

21

22

673

Effect

The effect of the `omp_set_default_allocator` is to set the value of the *def-allocator-var* ICV of the binding implicit task to the value specified in the *allocator* argument. Thus, it sets the default memory allocator to be used by allocation calls, `allocate` clauses and `allocate` and `allocators` directives that do not specify an *allocator*. This routine has the binding-implicit-task binding property so the binding task set for an `omp_set_default_allocator` region is the binding implicit task.

Restrictions

The restrictions to the `omp_set_default_allocator` routine are as follows:

- The *allocator* argument must be a valid memory allocator handle.

Cross References

- `allocate` Clause, see Section 13.6
- `allocate` Directive, see Section 13.5
- OpenMP `allocator_handle` Type, see Section 26.8.2
- `allocators` Construct, see Section 13.7
- Memory Allocators, see Section 13.2
- *def-allocator-var* ICV, see Table 3.1

33.10 omp_get_default_allocator Routine

Name: <code>omp_get_default_allocator</code>	Properties: binding-implicit-task-binding, memory-management-routine
Category: <code>function</code>	

Return Type

Name	Type	Properties
<return type>	<code>allocator_handle</code>	<i>default</i>

Prototypes

C / C++	
<code>omp_allocator_handle_t</code>	<code>omp_get_default_allocator(void);</code>
C / C++	
Fortran	
<code>integer (kind=omp_allocator_handle_kind)</code>	<code>function &</code>
<code>omp_get_default_allocator()</code>	
Fortran	

Effect

The `omp_get_default_allocator` routine returns the value of the *def-allocator-var* ICV of the binding implicit task, which is 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. This routine has the binding-implicit-task binding property, so the binding task set for an `omp_get_default_allocator` region is the binding implicit task.

Cross References

- `allocate` Clause, see [Section 13.6](#)
- `allocate` Directive, see [Section 13.5](#)
- OpenMP `allocator_handle` Type, see [Section 26.8.2](#)
- `allocators` Construct, see [Section 13.7](#)
- Memory Allocators, see [Section 13.2](#)
- *def-allocator-var* ICV, see [Table 3.1](#)

33.11 Memory Allocating Routines

This section describes the *memory-allocating routines*, which are routines that have the *memory-allocating-routine* property. Each of these routines requests a memory allocation from the memory allocator that its *allocator* argument specifies. If the *allocator* argument is `omp_null_allocator`, the routine uses the memory allocator specified by the *def-allocator-var* ICV of the binding implicit task. Upon success, these routines return a pointer to the allocated memory. Otherwise, the behavior that the `omp_atk_fallback` trait of the allocator specifies is followed. Pointers returned by these routines are considered *device pointers* if at least one of the *devices* associated with the allocator that the *allocator* argument represents is not the *current device*.

OpenMP provides several kinds of *memory-allocating routines*. The memory allocated by *raw-memory-allocating routines*, which have the *raw-memory-allocating-routine* property, is uninitialized. The memory allocated by *zeroed-memory-allocating routines*, which have the *zeroed-memory-allocating-routine* property, is set to zero before the routine returns.

The memory allocated by *aligned-memory-allocating routines*, which have the *aligned-memory-allocating-routine* property, is byte-aligned to at least the maximum of the alignment required by `malloc`, the `omp_atk_alignment` trait of the allocator and the value of their *alignment* argument. The memory allocated by all other *memory-allocating routines* is byte-aligned to at least the maximum of the alignment required by `malloc` and the `omp_atk_alignment` trait of the allocator.

Raw-memory-allocating routines request a memory allocation of *size* bytes from the specified memory allocator. *Zeroed-memory-allocating routines* request a memory allocation for an array of

nmemb elements, each of which has a size of *size* bytes. If any of the *size* or *nmemb* arguments are zero, these [routines](#) return `NULL`.

[Memory-reallocating routines](#) deallocate the [memory](#) to which the *ptr* argument points and request a new [memory](#) allocation of *size* bytes from the [memory allocator](#) that is specified by the *allocator* argument. 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, each of these [routines](#) returns a (possibly moved) pointer to the allocated [memory](#) and the contents of the new object will be the same as that of the old object prior to deallocation, up to the minimum size of the 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 `omp_atk_fallback` trait of the *allocator* specifies will be followed. If *ptr* is `NULL`, a [memory-reallocating routine](#) behaves the same as a [raw-memory-allocating routine](#) with the same *size* and *allocator* arguments. If *size* is zero, a [memory-reallocating routine](#) returns `NULL` and the old allocation is deallocated. If *size* is not zero, the old allocation will be deallocated if and only if the [routine](#) returns a [non-null value](#).

C++

The C++ version of all [memory-allocating routines](#) have the [overloaded property](#) since they are [overloaded routines](#) for which the *allocator* argument may be omitted, in which case the effect is as if `omp_null_allocator` is specified.

C++

Restrictions

The restrictions to [memory-allocating routines](#) are as follows:

- Unless the [unified_address clause](#) is specified or the [current device](#) is an associated [device](#) of the *allocator*, pointer arithmetic is not supported on the pointer that a [memory-allocating routine](#) returns.
- Each *allocator* and *free_allocator* argument must be an expression that evaluates to a [handle](#) that represents a [memory allocator](#).
- The value of the *alignment* argument to an [aligned-memory-allocating routine](#) must be a power of two.
- The value of a *size* argument to an [aligned-memory-allocating routine](#) must be a multiple of the *alignment* argument.
- The value of the *ptr* argument to a [memory-reallocating routine](#) must have been returned by a [memory-allocating routine](#).
- If the *free_allocator* argument is specified for a [memory-reallocating routine](#), it must be the [memory allocator](#) to which the previous allocation request was made.

- Using a [memory-reallocating routine](#) 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](#).
- Unless a `requires` directive with the `dynamic_allocators` clause is present in the same [compilation unit](#), [memory-allocating routines](#) that appear in [target regions](#) must not pass `omp_null_allocator` as the *allocator* or *free_allocator* argument.

Cross References

- [Memory Allocators](#), see [Section 13.2](#)
- `def-allocator-var` ICV, see [Table 3.1](#)
- `omp_destroy_allocator` Routine, see [Section 33.7](#)
- `requires` Directive, see [Section 16.5](#)
- `target` Construct, see [Section 21.8](#)

33.11.1 omp_alloc Routine

Name: <code>omp_alloc</code> Category: function	Properties: iso_c_binding , memory-allocating-routine , memory-management-routine , overloaded , raw-memory-allocating-routine
--	---

Return Type and Arguments

Name	Type	Properties
<i><return type></i>	<code>c_ptr</code>	default
<i>size</i>	<code>c_size_t</code>	iso_c , value
<i>allocator</i>	<code>allocator_handle</code>	value , omp

Prototypes

C
<code>void *omp_alloc(size_t size, omp_allocator_handle_t allocator);</code>
C
C++
<code>void *omp_alloc(size_t size, omp_allocator_handle_t allocator = omp_null_allocator);</code>
C++
Fortran
<code>type (c_ptr) function omp_alloc(size, allocator) bind(c)</code>
<code>use, intrinsic :: iso_c_binding, only : c_ptr, c_size_t</code>
<code>integer (kind=c_size_t), value :: size</code>
<code>integer (kind=omp_allocator_handle_kind), value :: allocator</code>
Fortran

Effect

The `omp_alloc` routine is a raw-memory-allocating routine.

Cross References

- OpenMP `allocator_handle` Type, see [Section 26.8.2](#)
- Memory Allocating Routines, see [Section 33.11](#)

33.11.2 omp_aligned_alloc Routine

Name: <code>omp_aligned_alloc</code> Category: function	Properties: aligned-memory-allocating-routine , iso_c_binding , memory-allocating-routine , memory-management-routine , overloaded , raw-memory-allocating-routine
--	---

Return Type and Arguments

Name	Type	Properties
<i><return type></i>	<code>c_ptr</code>	default
<i>alignment</i>	<code>c_size_t</code>	iso_c , value
<i>size</i>	<code>c_size_t</code>	iso_c , value
<i>allocator</i>	<code>allocator_handle</code>	value , omp

Prototypes

C

```
void *omp_aligned_alloc(size_t alignment, size_t size,
    omp_allocator_handle_t allocator);
```

C

C++

```
void *omp_aligned_alloc(size_t alignment, size_t size,
    omp_allocator_handle_t allocator = omp_null_allocator);
```

C++

Fortran

```
type (c_ptr) function omp_aligned_alloc(alignment, size, &
    allocator) bind(c)
    use, intrinsic :: iso_c_binding, only : c_ptr, c_size_t
    integer (kind=c_size_t), value :: alignment, size
    integer (kind=omp_allocator_handle_kind), value :: allocator
```

Fortran

Effect

The `omp_aligned_alloc` routine is a raw-memory-allocating routine and an aligned-memory-allocating routine.

Cross References

- OpenMP `allocator_handle` Type, see [Section 26.8.2](#)
- Memory Allocating Routines, see [Section 33.11](#)

33.11.3 `omp_calloc` Routine

Name: <code>omp_calloc</code> Category: function	Properties: iso_c_binding , memory-allocating-routine , memory-management-routine , overloaded , zeroed-memory-allocating-routine
---	--

Return Type and Arguments

Name	Type	Properties
<i><return type></i>	<code>c_ptr</code>	default
<i>nmemb</i>	<code>c_size_t</code>	iso_c , value
<i>size</i>	<code>c_size_t</code>	iso_c , value
<i>allocator</i>	<code>allocator_handle</code>	value , omp

Prototypes

C

```
void *omp_calloc(size_t nmemb, size_t size,
                 omp_allocator_handle_t allocator);
```

C

C++

```
void *omp_calloc(size_t nmemb, size_t size,
                 omp_allocator_handle_t allocator = omp_null_allocator);
```

C++

Fortran

```
type (c_ptr) function omp_calloc(nmemb, size, allocator) &
  bind(c)
  use, intrinsic :: iso_c_binding, only : c_ptr, c_size_t
  integer (kind=c_size_t), value :: nmemb, size
  integer (kind=omp_allocator_handle_kind), value :: allocator
```

Fortran

Effect

The `omp_calloc` routine is a [zeroed-memory-allocating routines](#).

Cross References

- OpenMP `allocator_handle` Type, see [Section 26.8.2](#)
- Memory Allocating Routines, see [Section 33.11](#)

33.11.4 omp_aligned_malloc Routine

Name: <code>omp_aligned_malloc</code> Category: function	Properties: aligned-memory-allocating-routine , iso_c_binding , memory-allocating-routine , memory-management-routine , overloaded , zeroed-memory-allocating-routine
---	--

Return Type and Arguments

Name	Type	Properties
<i><return type></i>	<code>c_ptr</code>	<i>default</i>
<i>alignment</i>	<code>c_size_t</code>	iso_c , value
<i>nmemb</i>	<code>c_size_t</code>	iso_c , value
<i>size</i>	<code>c_size_t</code>	iso_c , value
<i>allocator</i>	<code>allocator_handle</code>	value , omp

Prototypes

C

```
void *omp_aligned_malloc(size_t alignment, size_t nmemb,
    size_t size, omp_allocator_handle_t allocator);
```

C

C++

```
void *omp_aligned_malloc(size_t alignment, size_t nmemb,
    size_t size,
    omp_allocator_handle_t allocator = omp_null_allocator);
```

C++

Fortran

```
type (c_ptr) function omp_aligned_malloc(alignment, nmemb, size, &
    allocator) bind(c)
    use, intrinsic :: iso_c_binding, only : c_ptr, c_size_t
    integer (kind=c_size_t), value :: alignment, nmemb, size
    integer (kind=omp_allocator_handle_kind), value :: allocator
```

Fortran

Effect

The [omp_aligned_malloc](#) routine is a [zeroed-memory-allocating routine](#) and an [aligned-memory-allocating routine](#).

Cross References

- OpenMP `allocator_handle` Type, see [Section 26.8.2](#)
- Memory Allocating Routines, see [Section 33.11](#)

33.11.5 omp_realloc Routine

Name: <code>omp_realloc</code> Category: function	Properties: iso_c_binding , memory-allocating-routine , memory-management-routine , memory-reallocating-routine , overloaded
--	---

Return Type and Arguments

Name	Type	Properties
<return type>	<code>c_ptr</code>	default
<i>ptr</i>	<code>c_ptr</code>	iso_c , value
<i>size</i>	<code>c_size_t</code>	iso_c , value
<i>allocator</i>	<code>allocator_handle</code>	value , omp
<i>free_allocator</i>	<code>allocator_handle</code>	value , omp

Prototypes

C

```
void *omp_realloc(void *ptr, size_t size,  
    omp_allocator_handle_t allocator,  
    omp_allocator_handle_t free_allocator);
```

C

C++

```
void *omp_realloc(void *ptr, size_t size,  
    omp_allocator_handle_t allocator = omp_null_allocator,  
    omp_allocator_handle_t free_allocator = omp_null_allocator);
```

C++

Fortran

```
type (c_ptr) function omp_realloc(ptr, size, allocator, &  
    free_allocator) bind(c)  
    use, intrinsic :: iso_c_binding, only : c_ptr, c_size_t  
    type (c_ptr), value :: ptr  
    integer (kind=c_size_t), value :: size  
    integer (kind=omp_allocator_handle_kind), value :: allocator, &  
    free_allocator
```

Fortran

Effect

The `omp_realloc` routine is a [memory-reallocating routine](#).

Cross References

- OpenMP `allocator_handle` Type, see [Section 26.8.2](#)
- Memory Allocating Routines, see [Section 33.11](#)

33.12 omp_free Routine

Name: <code>omp_free</code> Category: <code>subroutine</code>	Properties: <code>iso_c_binding</code> , <code>memory-management-routine</code> , <code>overloaded</code>
--	--

Arguments

Name	Type	Properties
<i>ptr</i>	<code>c_ptr</code>	<code>iso_c</code> , <code>value</code>
<i>allocator</i>	<code>allocator_handle</code>	<code>value</code> , <code>omp</code>

Prototypes

C
<code>void omp_free(void *ptr, omp_allocator_handle_t allocator);</code>
C++
<code>void omp_free(void *ptr, omp_allocator_handle_t allocator = omp_null_allocator);</code>
C++
Fortran
<code>subroutine omp_free(ptr, allocator) bind(c) use, intrinsic :: iso_c_binding, only : c_ptr type (c_ptr), value :: ptr integer (kind=omp_allocator_handle_kind), value :: allocator</code>
Fortran

Effect

The `omp_free` routine deallocates the memory to which the *ptr* argument points. If the *allocator* argument is `omp_null_allocator`, the implementation will determine that value automatically. If *ptr* is `NULL`, no operation is performed.

C++
The C++ version of the <code>omp_free</code> routine has the <code>overloaded</code> property since it is an <code>overloaded routine</code> for which the <i>allocator</i> argument may be omitted, in which case the effect is as if <code>omp_null_allocator</code> is specified.
C++

Restrictions

The restrictions to the `omp_free` routine are as follows:

- The *ptr* argument must have been returned by a `memory-allocating routine`.
- If the *allocator* argument is specified it must be the `memory allocator` to which the allocation request was made.
- 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

- OpenMP `allocator_handle` Type, see [Section 26.8.2](#)
- Memory Allocating Routines, see [Section 33.11](#)
- Memory Allocators, see [Section 13.2](#)
- `omp_destroy_allocator` Routine, see [Section 33.7](#)

33.13 Groupprivate Memory Routines

This section describes the [groupprivate-information routines](#), which are [routines](#) that have the [groupprivate-information-routine](#) property. Some of these [routines](#) are [dynamic-groupprivate-information routines](#), which are [routines](#) that have the [dynamic-groupprivate-information-routine](#) property. The [dynamic-groupprivate-information routines](#) include [routines](#) to retrieve a pointer to and the size in bytes of the [dynamic groupprivate block](#).

All [dynamic-groupprivate-information routines](#) take a `access_group` argument, which specifies the [access group type](#) of the [dynamic groupprivate block](#) for which information is requested. Any [task](#) that belongs to a specific [access group](#) may call these [routines](#) by passing the corresponding [access group type](#) as the argument to obtain information about the [dynamic groupprivate block](#) associated with that [access group](#).

If the value of the `access_group` argument is `omp_access_cgroup`, the [dynamic-groupprivate-information routine](#) returns information about the [dynamic groupprivate block](#) associated with the [contention group](#) of the calling [task](#).

C++

The C++ versions of the [groupprivate-information routines](#) have the [overloaded](#) property since they are [overloaded routines](#) for which the `access_group` argument may be omitted, in which case the effect is as if `omp_access_cgroup` is specified. For the C++ version of the `omp_get_dyn_groupprivate_ptr` routine, the `offset` and `is_fallback` arguments may also be omitted, in which case the effect is as if zero and `NULL` are specified, respectively.

C++

Fortran

For the Fortran versions of the [groupprivate-information routines](#), the `access_group` argument is [optional](#) and may be omitted, in which case the effect is as if `omp_access_cgroup` is specified. For the Fortran version of the `omp_get_dyn_groupprivate_ptr` routine, the `offset` and `is_fallback` arguments are also [optional](#) and may be omitted. If the `offset` argument is omitted, the effect is as if zero is specified as the argument.

Fortran

1 **Restrictions**

2 Restrictions to the [dynamic-groupprivate-information routines](#) are as follows:

- 3 • A [task](#) may call a [dynamic-groupprivate-information routine](#) with an [access group type](#) as the
- 4 *access_group* argument only if it belongs to an [access group](#) of that [access group type](#).

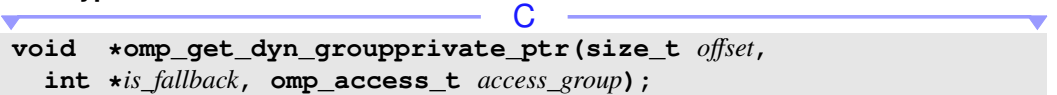
5 **33.13.1 omp_get_dyn_groupprivate_ptr Routine**

Name: <code>omp_get_dyn_groupprivate_ptr</code> Category: function	Properties: dynamic-groupprivate-information-routine , groupprivate-information-routine , memory-management-routine , overloaded
---	---

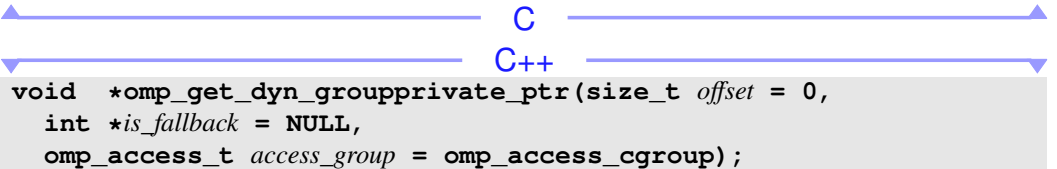
7 **Return Type and Arguments**

Name	Type	Properties
<i><return type></i>	<code>c_ptr</code>	default
<i>offset</i>	<code>c_size_t</code>	iso_c , value , optional
<i>is_fallback</i>	logical	C/C++ pointer , intent(out) , optional
<i>access_group</i>	<code>access</code>	value , omp , optional

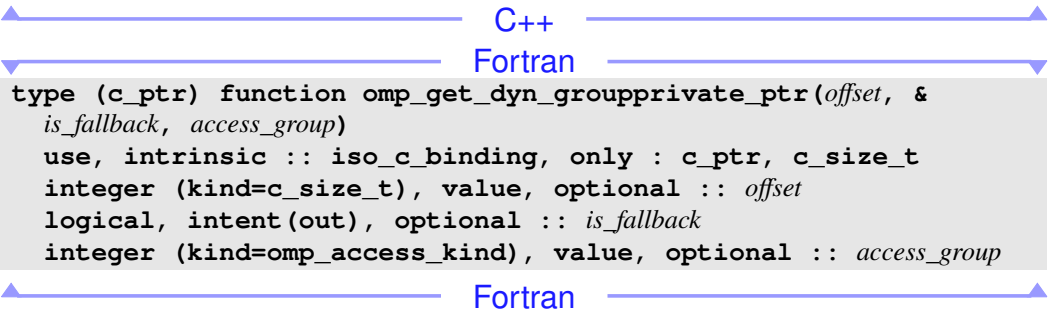
9 **Prototypes**

10  C

```
11 void *omp_get_dyn_groupprivate_ptr(size_t offset,  
    int *is_fallback, omp_access_t access_group);
```

12  C++

```
13 void *omp_get_dyn_groupprivate_ptr(size_t offset = 0,  
    int *is_fallback = NULL,  
    omp_access_t access_group = omp_access_cgroup);
```

15  Fortran

```
16 type (c_ptr) function omp_get_dyn_groupprivate_ptr(offset, &  
17     is_fallback, access_group)  
18     use, intrinsic :: iso_c_binding, only : c_ptr, c_size_t  
19     integer (kind=c_size_t), value, optional :: offset  
20     logical, intent(out), optional :: is_fallback  
    integer (kind=omp_access_kind), value, optional :: access_group
```

Effect

The `omp_get_dyn_groupprivate_ptr` routine returns a pointer that contains the resulting address of applying an offset of *offset* bytes to the *base address* of the *dynamic groupprivate block*.

C / C++

If a *non-null pointer* is passed to the *is_fallback* argument, a *true* value is stored in the object to which *is_fallback* points if a *successful groupprivate fallback situation* occurred for the *dynamic groupprivate block*, and otherwise a *false* value is stored.

C / C++

Fortran

When the *is_fallback* argument is provided, a *true* value is assigned to the argument if a *successful groupprivate fallback situation* occurred for the *dynamic groupprivate block*, and otherwise a *false* value is assigned.

Fortran

Restrictions

Restrictions to the `omp_get_dyn_groupprivate_ptr` routine are as follows:

- If the storage referenced by the return value of the *routine* is freed, the effect is *unspecified*.
- When an access through a pointer that references a *dynamic groupprivate block* is out-of-bounds, the effect is *unspecified*.
- If the value of the *offset* argument is equal to or greater than the size of the *dynamic groupprivate block*, the effect is *unspecified*.
- When no corresponding *dynamic groupprivate block* exists, the *routine* returns *NULL*.
- If the call to the *routine* does not occur in a *region* that corresponds to a *groupprivate-instantiating construct*, the *routine* returns *NULL*.

33.13.2 `omp_get_dyn_groupprivate_size` Routine

Name: `omp_get_dyn_groupprivate_size`
Category: *function*

Properties: *dynamic-groupprivate-information-routine, groupprivate-information-routine, memory-management-routine, overloaded*

Return Type and Arguments

Name	Type	Properties
<return type>	<i>c_size_t</i>	<i>default</i>
<i>access_group</i>	<i>access</i>	<i>value, omp, optional</i>

1 **Prototypes**

```
2       └────────────────────────── C ───────────────────────────┘
3       | size_t omp_get_dyn_groupprivate_size(omp_access_t access_group);
4       └────────────────────────── C ───────────────────────────┘
5       └────────────────────────── C++ ─────────────────────────┘
6       | size_t omp_get_dyn_groupprivate_size(
7       |     omp_access_t access_group = omp_access_cgroup);
8       └────────────────────────── C++ ─────────────────────────┘
9       └────────────────────────── Fortran ─────────────────────┘
10       | integer (kind=c_size_t) function omp_get_dyn_groupprivate_size(&
11       |     access_group)
12       |     use, intrinsic :: iso_c_binding, only : c_size_t
13       |     integer (kind=omp_access_kind), value, optional :: access_group
14       └────────────────────────── Fortran ─────────────────────┘
```

9 **Effect**

10 The `omp_get_dyn_groupprivate_size` routine returns the size of the `dynamic`
11 `groupprivate block`, if it exists, and otherwise returns zero. If the `dynamic groupprivate block`
12 exists, this value is the result of the evaluation of the `size` argument that was specified in the
13 corresponding `dyn_groupprivate` clause that instantiated the `dynamic groupprivate block`.

14 **Restrictions**

15 Restrictions to the `omp_get_dyn_groupprivate_size` routine are as follows:

- 16 • If the call to the routine does not occur in a `region` that corresponds to a
17 `groupprivate-instantiating construct`, the routine returns zero.

18 **Cross References**

- 19 • `dyn_groupprivate` Clause, see [Section 13.9](#)
- 20 • `target` Construct, see [Section 21.8](#)
- 21 • `teams` Construct, see [Section 18.2](#)

22 **33.13.3 omp_get_groupprivate_limit Routine**

23 Name: <code>omp_get_groupprivate_limit</code> Category: <code>function</code>	Properties: <code>groupprivate-information-</code> <code>routine</code> , <code>memory-management-routine</code> , <code>overloaded</code>
---	---

1 **Return Type and Arguments**

Name	Type	Properties
<return type>	c_size_t	default
device_num	c_int	iso_c, value
access_group	access	value, omp, optional

3 **Prototypes**

C

4 **size_t omp_get_groupprivate_limit**(int device_num,
5 omp_access_t access_group);

C

C++

6 **size_t omp_get_groupprivate_limit**(int device_num,
7 omp_access_t access_group = omp_access_cgroup);

C++

Fortran

8 **integer** (kind=c_size_t) **function** **omp_get_groupprivate_limit**(&
9 device_num, access_group)
10 use, intrinsic :: iso_c_binding, only : c_size_t, c_int
11 **integer** (kind=c_int), value :: device_num
12 **integer** (kind=omp_access_kind), value, optional :: access_group

Fortran

13 **Effect**

14 The **omp_get_groupprivate_limit** routine returns the groupprivate absolute limit in bytes
15 of the groupprivate memory space that groupprivate variables and dynamic groupprivate blocks
16 with the access group type access_group may use in any region that corresponds to a
17 groupprivate-instantiating construct that targets the device device_num. The device device_num
18 may be any host device or non-host device.

19 **Restrictions**

20 Restrictions to device memory routines are as follows:

- 21 • The device_num argument must be a conforming device number.
- 22 • When called from within a target region, the effect is unspecified.

34 Lock Routines

This chapter describes general-purpose [lock routines](#) that can be used for synchronization via mutual exclusion. These [routines](#) with the [lock property](#) operate on OpenMP [locks](#) that are represented by [OpenMP lock variables](#). [OpenMP lock variables](#) must be accessed only through the [lock routines](#); [OpenMP programs](#) that otherwise access [OpenMP lock variables](#) are [non-conforming](#).

A [lock](#) can be in one of the following [lock states](#): *uninitialized*; *unlocked*; or *locked*. If a [lock](#) is in the [unlocked state](#), a [task](#) can acquire the [lock](#) by executing a [lock-acquiring routine](#), a [routine](#) that has the [lock-acquiring property](#), through which it changes the [lock state](#) to the [locked state](#). The [task](#) that acquires the [lock](#) is then said to *own* the [lock](#). A [task](#) that owns a [lock](#) can release it by executing a [lock-releasing routine](#), a [routine](#) that has the [lock-releasing property](#), through which it returns the [lock state](#) to the [unlocked state](#). An [OpenMP program](#) in which a [task](#) executes a [lock-releasing routine](#) on a [lock](#) that is owned by another [task](#) is [non-conforming](#).

OpenMP supports two types of [locks](#): [simple locks](#) and [nestable locks](#). A [nestable lock](#) can be acquired (i.e., set) multiple times by the same [task](#) before being released (i.e., unset); a [simple lock](#) cannot be acquired 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](#) ([routines](#) that have the [simple lock property](#)). [Nestable lock variables](#) are associated with [nestable locks](#) and can only be passed to [nestable lock routines](#) ([routines](#) that have the [nestable lock property](#)).

Each type of [lock](#) can also have a [synchronization hint](#) that contains information about the intended usage of the [lock](#) by the [OpenMP program](#). 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 [compliant implementation](#) can safely ignore the hint.

Constraints on the [lock 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 [lock routines](#) access an [OpenMP lock variable](#) such that they always read and update its most current value. An [OpenMP program](#) does not need to include explicit [flush directives](#) to ensure that the value of a [lock](#) is consistent among different [tasks](#).

Restrictions

Restrictions to OpenMP [lock routines](#) are as follows:

- The use of the same [lock](#) in different [contention groups](#) results in [unspecified behavior](#).

34.1 Lock Initializing Routines

Lock-initializing routines are routines with the lock-initializing property. These routines 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.

Restrictions

Restrictions to lock-initializing routines are as follows:

- A lock-initializing routine must not access a lock that is not in the uninitialized state.

34.1.1 omp_init_lock Routine

Name: <code>omp_init_lock</code> Category: <code>subroutine</code>	Properties: <code>all-contention-group-tasks-binding, lock-initializing, simple-lock</code>
---	---

Arguments

Name	Type	Properties
<i>svar</i>	lock	C/C++ pointer, omp

Prototypes

<div>C / C++</div> <div><code>void omp_init_lock(omp_lock_t *svar);</code></div>	<div>C / C++</div> <div></div>
<div>Fortran</div> <div><code>subroutine omp_init_lock(svar) integer (kind=omp_lock_kind) svar</code></div>	<div>Fortran</div> <div></div>

Effect

The `omp_init_lock` routine is a lock-initializing routine.

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.

Tool Callbacks

A thread dispatches a registered `lock_init` callback with `omp_sync_hint_none` as the *hint* argument and `ompt_mutex_lock` as the *kind* argument for each occurrence of a `lock-init` event in that thread. This callback occurs in the task that encounters the routine.

Cross References

- OpenMP `lock` Type, see Section 26.9.3
- `lock_init` Callback, see Section 40.7.9
- OMPT `mutex` Type, see Section 39.20









34.1.2 omp_init_nest_lock Routine

Name: <code>omp_init_nest_lock</code> Category: <code>subroutine</code>	Properties: all-contention-group- tasks-binding, lock-initializing, nestable-lock
--	--

Arguments

Name	Type	Properties
<code>nvar</code>	<code>nest_lock</code>	C/C++ pointer, <code>omp</code>

Prototypes

 C / C++ 
<code>void omp_init_nest_lock(omp_nest_lock_t *nvar);</code>
 C / C++ 
 Fortran 
<code>subroutine omp_init_nest_lock(nvar)</code>
<code>integer (kind=omp_nest_lock_kind) nvar</code>
 Fortran 

Effect

The `omp_init_nest_lock` routine is a lock-initializing routine.

Execution Model Events

The `nest-lock-init` event occurs in a `thread` that executes an `omp_init_nest_lock` region after initialization of the `lock`, but before it finishes the `region`.

Tool Callbacks

A `thread` dispatches a registered `lock_init` callback with `omp_sync_hint_none` as the `hint` argument and `ompt_mutex_nest_lock` as the `kind` argument for each occurrence of a `nest-lock-init` event in that `thread`. This `callback` occurs in the `task` that encounters the `routine`.

Cross References

- `lock_init` Callback, see [Section 40.7.9](#)
- OMPT `mutex` Type, see [Section 39.20](#)
- OpenMP `nest_lock` Type, see [Section 26.9.4](#)

34.1.3 omp_init_lock_with_hint Routine

Name: <code>omp_init_lock_with_hint</code> Category: <code>subroutine</code>	Properties: all-contention-group- tasks-binding, lock-initializing, simple- lock
---	---

Arguments

Name	Type	Properties
<i>svar</i>	lock	C/C++ pointer, omp
<i>hint</i>	sync_hint	omp

Prototypes

C / C++

```
void omp_init_lock_with_hint(omp_lock_t *svar,  
    omp_sync_hint_t hint);
```

C / C++

Fortran

```
subroutine omp_init_lock_with_hint(svar, hint)  
    integer (kind=omp_lock_kind) svar  
    integer (kind=omp_sync_hint_kind) hint
```

Fortran

Effect

The `omp_init_lock_with_hint` routine is a lock-initializing routine.

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.

Tool Callbacks

A thread dispatches a registered `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. This callback occurs in the task that encounters the routine.

Cross References

- OpenMP `lock` Type, see Section 26.9.3
- `lock_init` Callback, see Section 40.7.9
- OMPT `mutex` Type, see Section 39.20
- OpenMP `sync_hint` Type, see Section 26.9.5

34.1.4 `omp_init_nest_lock_with_hint` Routine

Name: `omp_init_nest_lock_with_hint`
Category: `subroutine`

Properties: all-contention-group-
tasks-binding, lock-initializing,
nestable-lock

1

2

3

4

6

7

9

0

1

2

4

5

9

- 20

24

25

27

28

- 29

34.2.1 omp_destroy_lock Routine

Name: <code>omp_destroy_lock</code> Category: <code>subroutine</code>	Properties: <code>all-contention-group-</code> <code>tasks-binding</code> , <code>lock-destroying</code> , <code>simple-</code> <code>lock</code>
--	--

Arguments

Name	Type	Properties
<code>svar</code>	<code>lock</code>	<code>C/C++ pointer</code> , <code>omp</code>

Prototypes

<code>C / C++</code>	<code>void omp_destroy_lock(omp_lock_t *svar);</code>
<code>C / C++</code>	
<code>Fortran</code>	<code>subroutine omp_destroy_lock(svar)</code> <code>integer (kind=omp_lock_kind) svar</code>
<code>Fortran</code>	

Effect

The `omp_destroy_lock` routine is a `lock-destroying` routine.

Execution Model Events

The `lock-destroy` event occurs in a `thread` that executes an `omp_destroy_lock` region before it finishes the region.

Tool Callbacks

A `thread` dispatches a registered `lock_destroy` callback with `ompt_mutex_lock` as the `kind` argument for each occurrence of a `lock-destroy` event in that `thread`. This `callback` occurs in the `task` that encounters the `routine`.

Cross References

- OpenMP `lock` Type, see [Section 26.9.3](#)
- `lock_destroy` Callback, see [Section 40.7.11](#)
- OMPT `mutex` Type, see [Section 39.20](#)

34.2.2 omp_destroy_nest_lock Routine

Name: <code>omp_destroy_nest_lock</code> Category: <code>subroutine</code>	Properties: <code>all-contention-group-</code> <code>tasks-binding</code> , <code>lock-destroying</code> , <code>nestable-lock</code>
---	--

Arguments

Name	Type	Properties
<i>nvar</i>	nest_lock	C/C++ pointer, omp

Prototypes

C / C++

void omp_destroy_nest_lock(omp_nest_lock_t *nvar) ;

C / C++

Fortran

subroutine omp_destroy_nest_lock(nvar)
integer (kind=omp_nest_lock_kind) nvar

Fortran

Effect

The `omp_destroy_nest_lock` routine is a lock-destroying routine.

Execution Model Events

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 `lock_destroy` callback with `ompt_mutex_nest_lock` as the *kind* argument for each occurrence of a *nest-lock-destroy event* in that thread. This occurs in the task that encounters the routine.

Cross References

- `lock_destroy` Callback, see [Section 40.7.11](#)
- OMPT `mutex` Type, see [Section 39.20](#)
- OpenMP `nest_lock` Type, see [Section 26.9.4](#)

34.3 Lock Acquiring Routines

Lock-acquiring routines are routines with the lock-acquiring property. These routines provide a means of setting locks. The *encountering task region* behaves as if it was suspended until the lock can be acquired by this task.

Note – The semantics of *lock-acquiring routine* are 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.

Restrictions

Restrictions to `lock-acquiring routines` are as follows:

- A `lock-acquiring routine` must not access a `lock` that is in the `uninitialized state`.









34.3.1 `omp_set_lock` Routine

Name: <code>omp_set_lock</code> Category: <code>subroutine</code>	Properties: <code>all-contention-group-tasks-binding, lock-acquiring, simple-lock</code>
--	---

Arguments

Name	Type	Properties
<i>svar</i>	<code>lock</code>	<code>C/C++ pointer, omp</code>

Prototypes

	
<code>void omp_set_lock(omp_lock_t *svar) ;</code>	
	
	
<code>subroutine omp_set_lock(svar) integer (kind=omp_lock_kind) svar</code>	
	

Effect

A `simple lock` is available when it is in the `unlocked state`. Ownership of the `lock` is granted to the `task` that executes the `routine`.

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 `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`.

Tool Callbacks

A `thread` dispatches a registered `mutex_acquire callback` for each occurrence of a `lock-acquire event` in that `thread`. A `thread` dispatches a registered `mutex_acquired callback` for each occurrence of a `lock-acquired event` in that `thread`. These `callbacks` occur in the `task` that encounters the `omp_set_lock routine` and their `kind` argument is `ompt_mutex_lock`.

Restrictions

Restrictions to the `omp_set_lock routine` are as follows:

- A `task` must not already own the `lock` that it accesses with a call to `omp_set_lock` (or deadlock will result).

Cross References

- OpenMP `lock` Type, see [Section 26.9.3](#)
- OMPT `mutex` Type, see [Section 39.20](#)
- `mutex_acquire` Callback, see [Section 40.7.8](#)
- `mutex_acquired` Callback, see [Section 40.7.12](#)

34.3.2 `omp_set_nest_lock` Routine

Name: <code>omp_set_nest_lock</code> Category: subroutine	Properties: all-contention-group-tasks-binding , lock-acquiring , nestable-lock
--	--

Arguments

Name	Type	Properties
<code>nvar</code>	<code>nest_lock</code>	C/C++ pointer , omp

Prototypes

C / C++

`void omp_set_nest_lock(omp_nest_lock_t *nvar) ;`

C / C++

Fortran

`subroutine omp_set_nest_lock(nvar)
integer (kind=omp_nest_lock_kind) nvar`

Fortran

Effect

A [nestable lock](#) is available if it is in the [unlocked state](#) 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 [nest-lock-acquire event](#) occurs in a [thread](#) that executes an [omp_set_nest_lock region](#) before the associated [lock](#) is requested. The [nest-lock-acquired event](#) occurs in a [thread](#) that executes an [omp_set_nest_lock region](#) if the [task](#) 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 [task](#) when it already owns the [lock](#) and executes an [omp_set_nest_lock region](#). The [nest-lock-owned event](#) occurs after the nesting count is incremented but before the [task](#) finishes the [region](#).

1 **Tool Callbacks**

2 A [thread](#) dispatches a registered [mutex_acquire](#) callback for each occurrence of a
3 *nest-lock-acquire event* in that [thread](#). A [thread](#) dispatches a registered [mutex_acquired](#)
4 [callback](#) for each occurrence of a *nest-lock-acquired event* in that [thread](#). A [thread](#) dispatches a
5 registered [nest_lock](#) callback with [ompt_scope_begin](#) as its *endpoint* argument for each
6 occurrence of a *nest-lock-owned event* in that [thread](#). These [callbacks](#) occur in the [task](#) that
7 encounters the [omp_set_nest_lock](#) routine and their *kind* argument is
8 [ompt_mutex_nest_lock](#).

9 **Cross References**

- 10 • OMPT **mutex** Type, see [Section 39.20](#)
- 11 • **mutex_acquire** Callback, see [Section 40.7.8](#)
- 12 • **mutex_acquired** Callback, see [Section 40.7.12](#)
- 13 • **nest_lock** Callback, see [Section 40.7.14](#)
- 14 • OpenMP **nest_lock** Type, see [Section 26.9.4](#)
- 15 • OMPT **scope_endpoint** Type, see [Section 39.27](#)

16 **34.4 Lock Releasing Routines**

17 [Lock-releasing routines](#) are [routines](#) with the [lock-releasing property](#). These [routines](#) provide a
18 means of unsetting [locks](#). If the effect of a [lock-releasing routine](#) changes the [lock state](#) to the
19 [unlocked state](#) and one or more [task regions](#) were effectively suspended because the [lock](#) was
20 unavailable, the effect is that one [task](#) is chosen and given ownership of the [lock](#).

21 **Restrictions**

22 Restrictions to [lock-releasing routines](#) are as follows:

- 23 • A [lock-releasing routine](#) must not access a [lock](#) that is not in the [locked state](#).
- 24 • A [lock-releasing routine](#) must not access a [lock](#) that is owned by a [task](#) other than the
25 [encountering task](#).

26 **34.4.1 omp_unset_lock Routine**

27

Name: omp_unset_lock	Properties: all-contention-group- tasks-binding , lock-releasing , simple- lock
Category: subroutine	

28 **Arguments**

29

Name	Type	Properties
<i>svar</i>	lock	C/C++ pointer , omp

1 **Prototypes**

C / C++

2 **void omp_unset_lock(omp_lock_t *svar);**

C / C++

Fortran

3 **subroutine omp_unset_lock(svar)**

4 **integer (kind=omp_lock_kind) svar**

Fortran

5 **Effect**

6 The **omp_unset_lock** routine changes the **lock state** to the **unlocked state**.

7 **Execution Model Events**

8 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**.

10 **Tool Callbacks**

11 A **thread** dispatches a registered **mutex_released** callback with **ompt_mutex_lock** as the **kind** argument for each occurrence of a **lock-release event** in that **thread**. This **callback** occurs in the **encountering task**.

14 **Cross References**

- 15 • OpenMP **lock** Type, see [Section 26.9.3](#)
- 16 • OMPT **mutex** Type, see [Section 39.20](#)
- 17 • **mutex_released** Callback, see [Section 40.7.13](#)

18 **34.4.2 omp_unset_nest_lock Routine**

Name: <code>omp_unset_nest_lock</code>	Properties: all-contention-group-tasks-binding , lock-releasing , nestable-lock
Category: subroutine	

20 **Arguments**

Name	Type	Properties
<code>nvar</code>	nest_lock	C/C++ pointer , omp

22 **Prototypes**

C / C++

23 **void omp_unset_nest_lock(omp_nest_lock_t *nvar);**

C / C++

Fortran

24 **subroutine omp_unset_nest_lock(nvar)**

25 **integer (kind=omp_nest_lock_kind) nvar**

Fortran

Effect

The `omp_unset_nest_lock` routine decrements the nesting count and, if the resulting nesting count is zero, changes the `lock state` to the `unlocked state`.

Execution Model Events

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 `mutex_released` callback with `ompt_mutex_nest_lock` as the *kind* argument for each occurrence of a *nest-lock-release event* in that `thread`. A `thread` dispatches a registered `nest_lock` callback with `ompt_scope_end` as its *endpoint* argument for each occurrence of a *nest-lock-held event* in that `thread`. These `callbacks` occur in the `encountering task`.

Cross References

- OMPT `mutex` Type, see [Section 39.20](#)
- `mutex_released` Callback, see [Section 40.7.13](#)
- `nest_lock` Callback, see [Section 40.7.14](#)
- OpenMP `nest_lock` Type, see [Section 26.9.4](#)
- OMPT `scope_endpoint` Type, see [Section 39.27](#)

34.5 Lock Testing Routines

`Lock-testing routines` are `routines` with the `lock-testing property`. These `routines` attempt to acquire a `lock` in the same manner as `lock-acquiring routines`, except that they do not suspend execution of the `encountering task`

Restrictions

Restrictions on `lock-testing routines` are as follows.

- A `lock-testing routine` must not access a `lock` that is in the `uninitialized state`.

34.5.1 omp_test_lock Routine

Name: <code>omp_test_lock</code>	Properties: <code>all-contention-group-tasks-binding, lock-testing, simple-lock</code>
Category: <code>function</code>	

Return Type and Arguments

Name	Type	Properties
<return type>	logical	<i>default</i>
<i>svar</i>	lock	C/C++ pointer, <code>omp</code>

1 **Prototypes**

C / C++

```
2       int omp_test_lock(omp_lock_t *svar);
```

C / C++

Fortran

```
3       logical function omp_test_lock(svar)
```

Fortran

```
4       integer (kind=omp_lock_kind) svar
```

5 **Effect**

6 The `omp_test_lock` routine returns *true* if it successfully acquires the `lock`; otherwise, it returns

7 *false*.

8 **Execution Model Events**

9 The *lock-test event* occurs in a `thread` that executes an `omp_test_lock` region before the

10 associated `lock` is tested. The *lock-test-acquired event* occurs in a `thread` that executes an

11 `omp_test_lock` region before it finishes the region if the associated `lock` was acquired.

12 **Tool Callbacks**

13 A `thread` dispatches a registered `mutex_acquire` callback for each occurrence of a *lock-test*

14 *event* in that `thread`. A `thread` dispatches a registered `mutex_acquired` callback for each

15 occurrence of a *lock-test-acquired event* in that `thread`. These *callbacks* occur in the *encountering*

16 *task* and their *kind* argument is `ompt_mutex_test_lock`.

17 **Restrictions**

18 Restrictions to `omp_test_lock` routines are as follows:

19

- 20 • An `omp_test_lock` routine must not access a `lock` that is already owned by the
- encountering task*.

21 **Cross References**

- 22
 - OpenMP `lock` Type, see [Section 26.9.3](#)
 - 23 • OMPT `mutex` Type, see [Section 39.20](#)
 - 24 • `mutex_acquire` Callback, see [Section 40.7.8](#)
 - 25 • `mutex_acquired` Callback, see [Section 40.7.12](#)

26 **34.5.2 omp_test_nest_lock Routine**

27

Name: <code>omp_test_nest_lock</code> Category: <code>function</code>	Properties: <code>all-contention-group-</code> <code>tasks-binding</code> , <code>lock-testing</code> , <code>nestable-</code> <code>lock</code>
--	---

1 **Return Type and Arguments**

Name	Type	Properties
<return type>	integer	default
nvar	nest_lock	C/C++ pointer, omp

3 **Prototypes**

		C / C++	
int omp_test_nest_lock(omp_nest_lock_t *nvar);			
		C / C++	
		Fortran	
integer function omp_test_nest_lock(nvar)			
integer (kind=omp_nest_lock_kind) nvar			
		Fortran	

7 **Effect**

8 The **omp_test_nest_lock** routine returns the new nesting count if it successfully sets the **lock**;
9 otherwise, it returns zero.

10 **Execution Model Events**

11 The *nest-lock-test* event occurs in a **thread** that executes an **omp_test_nest_lock** region
12 before the associated **lock** is tested. The *nest-lock-test-acquired* event occurs in a **thread** that
13 executes an **omp_test_nest_lock** region before it finishes the **region** if the associated **lock**
14 was acquired and the **thread** did not already own the **lock**. The *nest-lock-owned* event occurs in a
15 **thread** that executes an **omp_test_nest_lock** region before it finishes the **region** after the
16 nesting count is incremented if the **thread** already owned the **lock**.

17 **Tool Callbacks**

18 A **thread** dispatches a registered **mutex_acquire** callback for each occurrence of a *nest-lock-test*
19 **event** in that **thread**. A **thread** dispatches a registered **mutex_acquired** callback for each
20 occurrence of a *nest-lock-test-acquired* event in that **thread**. A **thread** dispatches a registered
21 **nest_lock** callback with **ompt_scope_begin** as its *endpoint* argument for each occurrence
22 of a *nest-lock-owned* event in that **thread**. These **callbacks** occur in the **encountering task** and their
23 *kind* argument is **ompt_mutex_test_nest_lock**.

24 **Cross References**

- 25 • OMPT **mutex** Type, see [Section 39.20](#)
- 26 • **mutex_acquire** Callback, see [Section 40.7.8](#)
- 27 • **mutex_acquired** Callback, see [Section 40.7.12](#)
- 28 • **nest_lock** Callback, see [Section 40.7.14](#)
- 29 • OpenMP **nest_lock** Type, see [Section 26.9.4](#)
- 30 • OMPT **scope_endpoint** Type, see [Section 39.27](#)

35 Thread Affinity Routines

This chapter describes routines that specify and obtain information about thread affinity policies, which govern the placement of threads in the execution environment of OpenMP programs.

35.1 omp_get_proc_bind Routine

Name: <code>omp_get_proc_bind</code> Category: function	Properties: ICV-retrieving
--	--

Return Type

Name	Type	Properties
<code><return type></code>	<code>proc_bind</code>	default

Prototypes

C / C++	<code>omp_proc_bind_t omp_get_proc_bind(void);</code>
C / C++	
Fortran	<code>integer (kind=omp_proc_bind_kind) function omp_get_proc_bind()</code>
Fortran	

Effect

The effect of this routine is to return the value of the first element of the *bind-var* ICV of the current task, which will be used for the subsequent nested `parallel` regions that do not specify a `proc_bind` clause. See [Section 18.1.3](#) for the rules that govern the thread affinity policy.

Cross References

- Controlling OpenMP Thread Affinity, see [Section 18.1.3](#)
- bind-var* ICV, see [Table 3.1](#)
- `parallel` Construct, see [Section 18.1](#)
- OpenMP `proc_bind` Type, see [Section 26.10.1](#)

35.2 omp_get_num_places Routine

Name: <code>omp_get_num_places</code>	Properties: <code>all-device-threads-binding</code>
Category: <code>function</code>	

Return Type

Name	Type	Properties
<code><return type></code>	integer	<code>default</code>

Prototypes

<code>int omp_get_num_places(void);</code>	C / C++
<code>integer function omp_get_num_places()</code>	Fortran

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 3.1](#)

35.3 omp_get_place_num_procs Routine

Name: <code>omp_get_place_num_procs</code>	Properties: <code>all-device-threads-binding</code> , <code>ICV-retrieving</code>
Category: <code>function</code>	

Return Type and Arguments

Name	Type	Properties
<code><return type></code>	integer	<code>default</code>
<code>place_num</code>	integer	<code>default</code>

Prototypes

<code>int omp_get_place_num_procs(int place_num);</code>	C / C++
<code>integer function omp_get_place_num_procs(place_num) integer place_num</code>	Fortran

Effect

The `omp_get_place_num_procs` routine returns the number of processors associated with the place numbered `place_num` as per the *place-partition-var* ICV. The routine returns zero when `place_num` is negative or is greater than or equal to the value returned by `omp_get_num_places`.

Cross References

- *place-partition-var* ICV, see Table 3.1
- `omp_get_num_places` Routine, see Section 35.2

35.4 omp_get_place_proc_ids Routine

Name: <code>omp_get_place_proc_ids</code>	Properties: all-device-threads-binding, ICV-retrieving
Category: subroutine	

Arguments

Name	Type	Properties
<code>place_num</code>	integer	<i>default</i>
<code>ids</code>	integer	pointer

Prototypes

C / C++
<code>void omp_get_place_proc_ids(int place_num, int *ids);</code>
C / C++
Fortran
<code>subroutine omp_get_place_proc_ids(place_num, ids)</code> <code>integer place_num, ids(*)</code>
Fortran

Effect

The `omp_get_place_proc_ids` routine returns the numerical identifiers of each processor associated with the place numbered `place_num` as per the *place-partition-var* ICV. 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 4.1.6](#)
- `omp_get_num_places` Routine, see [Section 35.2](#)
- `omp_get_place_num_procs` Routine, see [Section 35.3](#)





35.5 `omp_get_place_num` Routine

Name: <code>omp_get_place_num</code> Category: function	Properties: default
--	-------------------------------------

Return Type

Name	Type	Properties
<code><return type></code>	integer	default

Prototypes

 <code>int omp_get_place_num(void);</code>
 <code>integer function omp_get_place_num()</code>
 <code>integer function omp_get_place_num()</code>


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 zero 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` Routine, see [Section 35.2](#)

35.6 `omp_get_partition_num_places` Routine

Name: <code>omp_get_partition_num_places</code> Category: function	Properties: ICV-retrieving
---	--

Return Type

Name	Type	Properties
<code><return type></code>	integer	default

1 **Prototypes**

C / C++

```
2 | int omp_get_partition_num_places(void);
```

C / C++

Fortran

```
3 | integer function omp_get_partition_num_places()
```

Fortran

4 **Effect**

5 The `omp_get_partition_num_places` routine returns the number of `places` in the

6 *place-partition-var* ICV.

7 **Cross References**

- 8 • *place-partition-var* ICV, see [Table 3.1](#)

9 **35.7 omp_get_partition_place_nums Routine**

Name: <code>omp_get_partition_place_nums</code> Category: subroutine	Properties: ICV-retrieving
---	--

11 **Arguments**

Name	Type	Properties
<i>place_nums</i>	integer	pointer

13 **Prototypes**

C / C++

```
14 | void omp_get_partition_place_nums(int *place_nums);
```

C / C++

Fortran

```
15 | subroutine omp_get_partition_place_nums(place_nums)  
16 |     integer place_nums(*)
```

Fortran

17 **Effect**

18 The `omp_get_partition_place_nums` routine returns the list of `place numbers` that

19 correspond to the `places` in the *place-partition-var* ICV of the innermost `implicit task`. The array

20 must be sufficiently large to contain `omp_get_partition_num_places` integers; otherwise,

21 the behavior is [unspecified](#).

Cross References

- *place-partition-var* ICV, see Table 3.1
- `omp_get_partition_num_places` Routine, see Section 35.6

35.8 `omp_set_affinity_format` Routine

Name: <code>omp_set_affinity_format</code> Category: subroutine	Properties: ICV-modifying
--	---

Arguments

Name	Type	Properties
<i>format</i>	char	pointer , intent(in)

Prototypes

C / C++
<code>void omp_set_affinity_format(const char *format);</code>
C / C++
Fortran
<code>subroutine omp_set_affinity_format (format) character(len=*), intent(in) :: format</code>
Fortran

Effect

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. The value of the ICV is set by copying the character string specified by the *format* argument into the 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](#).

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](#).

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 4.3.5](#)
- `OMP_DISPLAY_AFFINITY`, see [Section 4.3.4](#)
- Controlling OpenMP Thread Affinity, see [Section 18.1.3](#)
- *affinity-format-var* ICV, see [Table 3.1](#)
- `parallel` Construct, see [Section 18.1](#)
- `teams` Construct, see [Section 18.2](#)

35.9 omp_get_affinity_format Routine

Name: <code>omp_get_affinity_format</code> Category: function	Properties: ICV-retrieving
--	--

Return Type and Arguments

Name	Type	Properties
<i><return type></i>	<code>size_t</code>	default
<i>buffer</i>	<code>char</code>	pointer , intent(out)
<i>size</i>	<code>size_t</code>	C/C++-only

Prototypes

C / C++
<code>size_t omp_get_affinity_format(char *buffer, size_t size);</code>
C / C++
Fortran
<code>integer function omp_get_affinity_format(buffer) character(len=*) , intent(out) :: buffer</code>
Fortran

Effect

C / C++
The <code>omp_get_affinity_format</code> routine returns the number of characters in the <i>affinity-format-var</i> ICV on the current device , excluding the terminating null byte (' <code>\0</code> ') and, if <i>size</i> is non-zero, writes the value of the <i>affinity-format-var</i> ICV on the current device to <i>buffer</i> followed by a null byte. If the return value is larger or equal to <i>size</i> , the affinity format specification is truncated, with the terminating null byte stored to <i>buffer</i> [<i>size</i> –1]. If <i>size</i> is zero, nothing is stored and <i>buffer</i> 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*.

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*.

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

- *affinity-format-var* ICV, see [Table 3.1](#)
- `parallel` Construct, see [Section 18.1](#)
- `target` Construct, see [Section 21.8](#)
- `teams` Construct, see [Section 18.2](#)

35.10 omp_display_affinity Routine

Name: <code>omp_display_affinity</code> Category: <i>subroutine</i>	Properties: <i>default</i>
--	-----------------------------------

Arguments

Name	Type	Properties
<i>format</i>	char	pointer, intent(in)

Prototypes

C / C++

```
void omp_display_affinity(const char *format);
```

C / C++

Fortran

```
subroutine omp_display_affinity(format)
  character(len=*), intent(in) :: format
```

Fortran

Effect

The `omp_display_affinity` routine prints the [thread affinity](#) information of the [encountering 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 3.1](#)
- [target](#) Construct, see [Section 21.8](#)

35.11 omp_capture_affinity Routine

Name: <code>omp_capture_affinity</code>	Properties: default
Category: function	

Return Type and Arguments

Name	Type	Properties
<i><return type></i>	<code>size_t</code>	default
<i>buffer</i>	<code>char</code>	pointer, intent(out)
<i>size</i>	<code>size_t</code>	C/C++-only
<i>format</i>	<code>char</code>	pointer, intent(in)

Prototypes

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

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 `encountering 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 `encountering 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 3.1](#)
- `target` Construct, see [Section 21.8](#)

36 Execution Control Routines

This chapter describes the [OpenMP API routines](#) that control the execution state of the OpenMP implementation and provide information about that state. These [routines](#) include:

- [Routines](#) that monitor and control [cancellation](#);
- [Resource-relinquishing routines](#) that free resources used by the [OpenMP program](#);
- [Routines](#) that support timing measurements of [OpenMP programs](#); and
- The environment display [routine](#) that displays the initial values of [ICVs](#).

36.1 omp_get_cancellation Routine

Name: <code>omp_get_cancellation</code>	Properties: ICV-retrieving
Category: function	

Return Type

Name	Type	Properties
<code><return type></code>	logical	default

Prototypes

C / C++	<code>int omp_get_cancellation(void);</code>
C / C++	
Fortran	<code>logical function omp_get_cancellation()</code>
Fortran	

Effect

The [omp_get_cancellation routine](#) returns the value of the [cancel-var ICV](#). Thus, it returns [true](#) if [cancellation](#) is enabled and otherwise it returns [false](#).

Cross References

- [cancel-var ICV](#), see [Table 3.1](#)

36.2 Resource Relinquishing Routines

This section describes routines that have the resource-relinquishing property. Each resource-relinquishing routine region implies a barrier. Each resource-relinquishing routine returns zero in case of success, and non-zero otherwise.

Tool Callbacks

If the tool is not allowed to interact with the specified device after encountering the resource-relinquishing routine, then the runtime must call the tool finalizer for that device.

Restrictions

Restrictions to resource-relinquishing routines are as follows:

- A resource-relinquishing routine region may not be nested in any explicit region.
- A resource-relinquishing 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.

36.2.1 omp_pause_resource Routine

Name: <code>omp_pause_resource</code> Category: <code>function</code>	Properties: <code>all-tasks-binding</code> , <code>resource-relinquishing</code>
--	--

Return Type and Arguments

Name	Type	Properties
<i><return type></i>	integer	<i>default</i>
<i>kind</i>	pause_resource	<i>default</i>
<i>device_num</i>	integer	<i>default</i>

Prototypes

C / C++

```
int omp_pause_resource(omp_pause_resource_t kind, int device_num);
```

C / C++

Fortran

```
integer function omp_pause_resource(kind, device_num)  
  integer (kind=omp_pause_resource_kind) kind  
  integer device_num
```

Fortran

Effect

The `omp_pause_resource` routine allows the runtime to relinquish resources used by OpenMP on the specified `device`. The `device_num` argument indicates the `device` that will be paused. If the `device number` has the value `omp_invalid_device`, runtime error termination is performed. The `binding task set` for an `omp_pause_resource` routine region is `all tasks` on the specified `device`. That is, this routines has the `all-device-tasks binding property`. If `omp_pause_stop_tool` is specified for a `non-host device`, the effect is the same as for `omp_pause_hard` and (unlike for the `host device`) does not shutdown the OMPT interface.

Restrictions

Restrictions to the `omp_pause_resource` routine are as follows:

- The `device_num` argument must be a `conforming device number`.

Cross References

- Predefined Identifiers, see [Section 26.1](#)
- OpenMP `pause_resource` Type, see [Section 26.11.1](#)









36.2.2 omp_pause_resource_all Routine

Name: <code>omp_pause_resource_all</code>	Properties: <code>all-tasks-binding</code> , <code>resource-relinquishing</code>
Category: <code>function</code>	

Return Type and Arguments

Name	Type	Properties
<code><return type></code>	<code>integer</code>	<code>default</code>
<code>kind</code>	<code>pause_resource</code>	<code>default</code>

Prototypes

 C / C++ 
<code>int omp_pause_resource_all(omp_pause_resource_t kind);</code>
 C / C++ 
 Fortran 
<code>integer function omp_pause_resource_all(kind)</code> <code>integer (kind=omp_pause_resource_kind) kind</code>
 Fortran 

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 `binding task set` for a `omp_pause_resource_all` routine region is `all tasks` in the OpenMP program. That is, this routine has the `all-tasks binding property`.

Cross References

- `omp_pause_resource` Routine, see [Section 36.2.1](#)
- OpenMP `pause_resource` Type, see [Section 26.11.1](#)

36.3 Timing Routines

This section describes [routines](#) that support a portable wall clock timer.





36.3.1 `omp_get_wtime` Routine

Name: <code>omp_get_wtime</code> Category: function	Properties: <i>default</i>
--	----------------------------

Return Type

Name	Type	Properties
<i><return type></i>	double	<i>default</i>

Prototypes

 <code>double omp_get_wtime(void);</code> 
 <code>double precision function omp_get_wtime()</code> 

Effect

The `omp_get_wtime` routine returns a value equal to the elapsed wall clock time in seconds since some *time-in-the-past*. The actual *time-in-the-past* is arbitrary, but it is guaranteed not to change during the execution of an [OpenMP program](#). The time returned is a *per-thread time*, so it is not required to be globally consistent across [all threads](#) that participate in an [OpenMP program](#).

36.3.2 `omp_get_wtick` Routine

Name: <code>omp_get_wtick</code> Category: function	Properties: <i>default</i>
--	----------------------------

Return Type

Name	Type	Properties
<i><return type></i>	double	<i>default</i>

1 **Prototypes**

		C / C++
2 double omp_get_wtick(void);		
		C / C++
		Fortran
3 double precision function omp_get_wtick()		
		Fortran

4 **Effect**

5 The `omp_get_wtick` routine returns the precision of the timer used by `omp_get_wtime` as a
6 value equal to the number of seconds between successive clock ticks. The return value of the
7 `omp_get_wtick` routine is not guaranteed to be consistent across any set of threads.

8 **Cross References**

- 9 • `omp_get_wtime` Routine, see [Section 36.3.1](#)

10 **36.4 omp_display_env Routine**

11 Name: <code>omp_display_env</code>	Properties: <i>default</i>
12 Category: subroutine	

13 **Arguments**

Name	Type	Properties
<i>verbose</i>	logical	intent(in)

14 **Prototypes**

		C / C++
15 void omp_display_env(int <i>verbose</i>);		
		C / C++
		Fortran
16 subroutine omp_display_env(<i>verbose</i>)		
17 logical, intent(in) :: <i>verbose</i>		
		Fortran

18 **Effect**

19 Each time that the `omp_display_env` routine is invoked, the runtime system prints the OpenMP
20 version number and the initial values of the ICVs associated with the [environment variables](#)
21 described in [Chapter 4](#). The displayed values are the values of the ICVs after they have been
22 modified according to the [environment variable](#) settings and before the execution of any [construct](#)
23 or [routine](#).

The display begins with "OPENMP DISPLAY ENVIRONMENT BEGIN", followed by the `__OPENMP` version macro (or the `openmp_version` predefined identifier for Fortran) and ICV values, in the format `NAME '=' VALUE`. `NAME` corresponds to the macro or environment variable name, 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 a comma-separated list of the devices on which the value of the ICV is applied. It is **host** if the device is the host device; **device** if the ICV applies to all non-host devices; **all** if the ICV has global scope or the value applies to the host device and all non-host devices; **dev**, a space, and the device number if it applies to a specific non-host devices. Instead of a single number a range can also be specified using the first and last device number separated by a hyphen. Whether ICVs with the same value are combined or displayed in multiple lines is implementation defined. The display is terminated with "OPENMP DISPLAY ENVIRONMENT END".

If the `verbose` argument evaluates to *false*, the runtime displays the OpenMP version number defined by the `__OPENMP` version macro (or the `openmp_version` predefined identifier for Fortran) value and the initial ICV values for the environment variables listed in Chapter 4. If the `verbose` argument evaluates to *true*, the runtime may also display the values of vendor-specific ICVs that may be modified by vendor-specific environment variables.

Example output:

```
OPENMP DISPLAY ENVIRONMENT BEGIN
__OPENMP='202411'
[dev 1] OMP_SCHEDULE='GUIDED,4'
[host] OMP_NUM_THREADS='4,3,2'
[device] OMP_NUM_THREADS='2'
[host, dev 2] OMP_DYNAMIC='TRUE'
[dev 2-3, dev 5] OMP_DYNAMIC='FALSE'
[all] OMP_WAIT_POLICY='ACTIVE'
[host] OMP_PLACES='{0:4},{4:4},{8:4},{12:4}'
...
OPENMP DISPLAY ENVIRONMENT END
```

Restrictions

Restrictions to the `omp_display_env` routine are as follows:

- When called from within a **target** region the effect is **unspecified**.

Cross References

- Predefined Identifiers, see Section 26.1

37 Tool Support Routines

This chapter describes the [OpenMP API routines](#) that support the use of OpenMP [tool](#) interfaces.

37.1 omp_control_tool Routine

Name: <code>omp_control_tool</code>	Properties: <i>default</i>
Category: <i>function</i>	

Return Type and Arguments

Name	Type	Properties
<i><return type></i>	<code>control_tool_result</code>	<i>default</i>
<i>command</i>	<code>control_tool</code>	<i>omp</i>
<i>modifier</i>	<code>integer</code>	<i>default</i>
<i>arg</i>	<code>void</code>	<i>C/C++-only, C/C++ pointer</i>

Prototypes

	C / C++	
<code>omp_control_tool_result_t</code>	<code>omp_control_tool(</code>	
<code>omp_control_tool_t</code>	<code>command, int modifier, void *arg);</code>	
	C / C++	
	Fortran	
<code>integer (kind=omp_control_tool_result_kind)</code>	<code>function &</code>	
<code>omp_control_tool(command, modifier)</code>		
<code>integer (kind=omp_control_tool_kind)</code>	<code>command</code>	
<code>integer</code>	<code>modifier</code>	
	Fortran	

Effect

An [OpenMP program](#) may use the [omp_control_tool routine](#) to pass commands to a [tool](#). An [OpenMP program](#) can use the [routine](#) to request: that a [tool](#) starts or restarts data collection when a code [region](#) of interest is encountered; that a [tool](#) pauses data collection when leaving the [region](#) of interest; that a [tool](#) flushes any data that it has collected so far; or that a [tool](#) ends data collection. Additionally, the [omp_control_tool routine](#) can be used to pass [tool](#)-specific commands to a particular [tool](#).

Any values for *modifier* and *arg* are **tool defined**.

If the OMPT interface state is **OMPT inactive**, the OpenMP implementation returns **omp_control_tool_notool**. If the OMPT interface state is **OMPT active**, but no **callback** is registered for the *tool-control event*, the OpenMP implementation returns **omp_control_tool_nocallback**. An OpenMP implementation may return other **implementation defined** negative values strictly smaller than -64; an **OpenMP program** may assume that any negative return value indicates that a **tool** has not received the command. A return value of **omp_control_tool_success** indicates that the **tool** has performed the specified command. A return value of **omp_control_tool_ignored** indicates that the **tool** has ignored the specified command. A **tool** may return other positive values strictly greater than 64 that are **tool defined**.

Execution Model Events

The *tool-control event* occurs in the **encountering thread** inside the corresponding **region**.

Tool Callbacks

A **thread** dispatches a registered **control_tool callback** for each occurrence of a *tool-control event*. The **callback** executes in the context of the call that occurs in the user program. The **callback** may return any **non-negative** value, which will be returned to the **OpenMP program** by the OpenMP implementation as the return value of the **omp_control_tool** call that triggered the **callback**.

Arguments passed to the **callback** are those passed by the user to **omp_control_tool**. If the call is made in Fortran, the **tool** will be passed **NULL** as the third argument to the **callback**. If any of the standard commands is presented to a **tool**, the **tool** will ignore the *modifier* and *arg* 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

- **control_tool** Callback, see [Section 40.8](#)
- OpenMP **control_tool** Type, see [Section 26.12.1](#)
- OpenMP **control_tool_result** Type, see [Section 26.12.2](#)
- OMPT Overview, see [Chapter 38](#)

1

Part V

2

OMPT

38 OMPT Overview

This chapter provides an overview of 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.

38.1 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 OMPT types. These definitions, which are listed throughout this and the immediately following chapters, 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 defined values.

The `ompt_start_tool` procedure is an external function with C linkage.

C / C++

38.2 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 OMPT-tool initializer of the tool, which enables the tool to prepare to monitor execution on the host device. 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 activities.

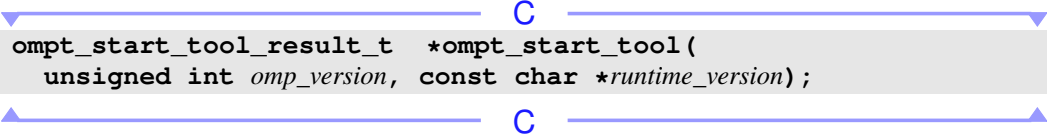
38.2.1 `ompt_start_tool` Procedure

Name: <code>ompt_start_tool</code> Category: <code>function</code>	Properties: C-only, OMPT
---	---------------------------------

1 **Return Type and Arguments**

Name	Type	Properties
<return type>	start_tool_result	pointer, OMPT
omp_version	integer	unsigned
runtime_version	char	intent(in), pointer

3 **Prototypes**

4  C

```
omp_start_tool_result_t *omp_start_tool(  
    unsigned int omp_version, const char *runtime_version);
```

6 **Semantics**

7 For a [tool](#) to use the [OMPT](#) interface that an OpenMP implementation provides, the [tool](#) must
8 define a globally-visible implementation of the [omp_start_tool](#) procedure. The [tool](#)
9 indicates that it will use the [OMPT](#) interface that an OpenMP implementation provides by returning
10 a [non-null pointer](#) to a [start_tool_result](#) OMPT type structure from the
11 [omp_start_tool](#) implementation that it provides. The [start_tool_result](#) structure
12 contains pointers to [initialize](#) and [finalize](#) callbacks as well as a [tool](#) data word that an
13 OpenMP implementation must pass by reference to these [callbacks](#). A [tool](#) may return [NULL](#) from
14 [omp_start_tool](#) to indicate that it will not use the [OMPT](#) interface in a particular execution.

15 A [tool](#) may use the *omp_version* argument to determine if it is compatible with the [OMPT](#) interface
16 that the OpenMP implementation provides. The *omp_version* argument is the value of the
17 `_OPENMP` version macro associated with the OpenMP implementation. This value identifies the
18 version that an implementation supports, which specifies the version of the [OMPT](#) interface that it
19 supports. The *runtime_version* argument is a version string that unambiguously identifies the
20 OpenMP implementation.

21 If a [tool](#) returns a [non-null pointer](#) to a [start_tool_result](#) OMPT type structure, an OpenMP
22 implementation will call the [OMPT-tool initializer](#) specified by the [initialize](#) field in this
23 [structure](#) before beginning execution of any [construct](#) or completing execution of any [routine](#); the
24 OpenMP implementation will call the [OMPT-tool finalizer](#) specified by the [finalize](#) field in this
25 [structure](#) when the OpenMP implementation shuts down.

26 **Restrictions**

27 Restrictions to [omp_start_tool](#) procedures are as follows:

- 28 • The *runtime_version* argument must be an immutable string that is defined for the lifetime of
29 a program execution.

Cross References

- **finalize** Callback, see [Section 40.1.2](#)
- **initialize** Callback, see [Section 40.1.1](#)
- OMPT **start_tool_result** Type, see [Section 39.30](#)

38.2.2 Determining Whether to Initialize a First-Party Tool

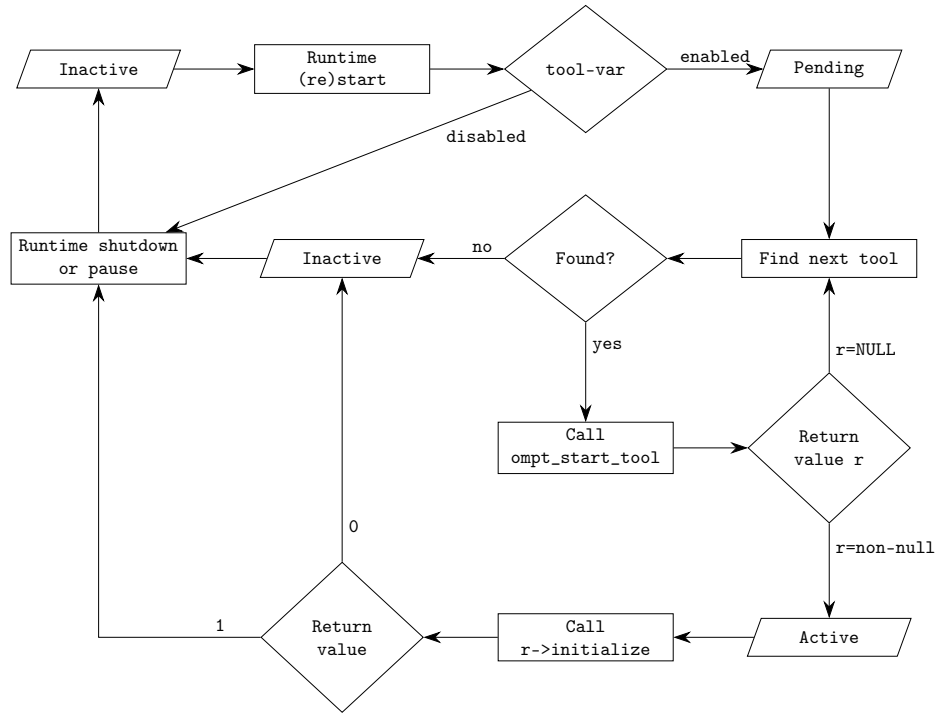


FIGURE 38.1: First-Party Tool Activation Flow Chart

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 `address space` of the program; 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` and that includes a definition of `ompt_start_tool`.

If the value of `tool-var` is *enabled*, the OpenMP implementation must check if a `tool` has provided an implementation of `ompt_start_tool`. The OpenMP implementation first checks if a `tool`-provided implementation of `ompt_start_tool` is available in the `address space`, either statically-linked into the `OpenMP program` or in a dynamically-linked library loaded in the `address space`. If multiple implementations of `ompt_start_tool` are available, the implementation will use the first `tool`-provided implementation of `ompt_start_tool` that it finds.

If the implementation does not find a `tool`-provided implementation of `ompt_start_tool` in the `address space`, it consults the `tool-libraries-var` ICV, which contains a (possibly empty) `list` of dynamically-linked libraries. As described in detail in [Section 4.5.2](#), the libraries in `tool-libraries-var` are then searched for the first usable implementation of `ompt_start_tool` that one of the libraries in the `list` provides.

If the implementation finds a `tool`-provided definition of `ompt_start_tool`, it invokes that `procedure`; if a `NULL` pointer is returned, the `OMPT interface state` remains `OMPT pending` and the implementation continues to look for implementations of `ompt_start_tool`; otherwise a `non-null pointer` to a `start_tool_result` `OMPT type structure` is returned, the `OMPT interface state` changes to `OMPT active` and the OpenMP implementation makes the `OMPT interface` available as the program executes. In this case, as the OpenMP implementation completes its initialization, it initializes the `OMPT interface`.

If no `tool` can be found, the `OMPT interface state` changes to `OMPT inactive`.

Cross References

- `tool-libraries-var` ICV, see [Table 3.1](#)
- `tool-var` ICV, see [Table 3.1](#)
- `ompt_start_tool` Procedure, see [Section 38.2.1](#)
- `OMPT start_tool_result` Type, see [Section 39.30](#)

38.2.3 Initializing a First-Party Tool

To initialize the `OMPT interface`, the OpenMP implementation invokes the `OMPT-tool initializer` that is specified in the `initialize` field of the `start_tool_result` structure that `ompt_start_tool` returns. This `initialize` callback is invoked prior to the occurrence of any OpenMP `event`.

An **initialize** callback uses the **entry point** specified in its *lookup* argument to look up pointers to **OMPT entry points** that the OpenMP implementation provides; this process is described in [Section 38.2.3.1](#). Typically, an **OMPT-tool initializer** obtains a pointer to the **set_callback** entry point and then uses it to perform **callback registration** for **events**, as described in [Section 38.2.4](#).

An **OMPT-tool initializer** may use the **enumerate_states** entry point to determine the **thread states** that an OpenMP implementation employs. Similarly, it may use the **enumerate_mutex_impls** entry point to determine the mutual exclusion implementations that the OpenMP implementation employs.

If an **OMPT-tool initializer** returns a non-zero value, the **OMPT interface state** remains **OMPT active** for the execution; otherwise, the **OMPT interface state** changes to **OMPT inactive**.

Cross References

- **enumerate_mutex_impls** Entry Point, see [Section 42.3](#)
- **enumerate_states** Entry Point, see [Section 42.2](#)
- Binding Entry Points, see [Section 38.2.3.1](#)
- **initialize** Callback, see [Section 40.1.1](#)
- **ompt_start_tool** Procedure, see [Section 38.2.1](#)
- **set_callback** Entry Point, see [Section 42.4](#)
- **OMPT start_tool_result** Type, see [Section 39.30](#)

38.2.3.1 Binding Entry Points

Routines that an OpenMP implementation provides to support **OMPT** are not defined as global symbols. Instead, they are defined as **runtime entry points** that a **tool** can only identify through the value returned in the *lookup* argument of the **initialize** callback. A **tool** can use this **function_lookup** entry point to obtain a pointer to each of the other **entry points** that an OpenMP implementation provides to support **OMPT**. Once a **tool** has obtained a **function_lookup** entry point, it may employ it at any point in the future.

For each **OMPT entry point** for the **host device**, [Table 38.1](#) provides the string name by which it is known and its associated type signature. Implementations can provide additional **implementation defined** names and corresponding **entry points**.

During initialization, a **tool** should look up each **entry point** by name and assign the **entry point** to a pointer that it maintains so it can later invoke that **entry point**. The **entry points** described in [Table 38.1](#) enable a **tool** to assess the **thread states** and mutual exclusion implementations that an implementation supports for **callback registration**, to inspect **registered callbacks**, to introspect OpenMP state associated with **threads**, and to use tracing to monitor computations that execute on **target devices**.

TABLE 38.1: OMPT Callback Interface Runtime Entry Point Names and Their Type Signatures

Entry Point String Name	OMPT Type
"ompt_enumerate_states"	enumerate_states
"ompt_enumerate_mutex_impls"	enumerate_mutex_impls
"ompt_set_callback"	set_callback
"ompt_get_callback"	get_callback
"ompt_get_thread_data"	get_thread_data
"ompt_get_num_places"	get_num_places
"ompt_get_place_proc_ids"	get_place_proc_ids
"ompt_get_place_num"	get_place_num
"ompt_get_partition_place_nums"	get_partition_place_nums
"ompt_get_proc_id"	get_proc_id
"ompt_get_state"	get_state
"ompt_get_parallel_info"	get_parallel_info
"ompt_get_task_info"	get_task_info
"ompt_get_task_memory"	get_task_memory
"ompt_get_num_devices"	get_num_devices
"ompt_get_num_procs"	get_num_procs
"ompt_get_target_info"	get_target_info
"ompt_get_unique_id"	get_unique_id
"ompt_finalize_tool"	finalize_tool

Cross References

- `enumerate_mutex_impls` Entry Point, see [Section 42.3](#)
- `enumerate_states` Entry Point, see [Section 42.2](#)
- `finalize_tool` Entry Point, see [Section 42.20](#)
- `function_lookup` Entry Point, see [Section 42.1](#)
- `get_callback` Entry Point, see [Section 42.5](#)
- `get_num_devices` Entry Point, see [Section 42.18](#)
- `get_num_places` Entry Point, see [Section 42.8](#)
- `get_num_procs` Entry Point, see [Section 42.7](#)
- `get_parallel_info` Entry Point, see [Section 42.14](#)
- `get_partition_place_nums` Entry Point, see [Section 42.11](#)
- `get_place_num` Entry Point, see [Section 42.10](#)
- `get_place_proc_ids` Entry Point, see [Section 42.9](#)

- `get_proc_id` Entry Point, see [Section 42.12](#)
- `get_state` Entry Point, see [Section 42.13](#)
- `get_target_info` Entry Point, see [Section 42.17](#)
- `get_task_info` Entry Point, see [Section 42.15](#)
- `get_task_memory` Entry Point, see [Section 42.16](#)
- `get_thread_data` Entry Point, see [Section 42.6](#)
- `get_unique_id` Entry Point, see [Section 42.19](#)
- `initialize` Callback, see [Section 40.1.1](#)
- `set_callback` Entry Point, see [Section 42.4](#)

38.2.4 Monitoring Activity on the Host with OMPT

To monitor the execution of an [OpenMP program](#) on the [host device](#), an [OMPT-tool initializer](#) must register to receive notification of [events](#) that occur as an [OpenMP program](#) executes. A [tool](#) can use the `set_callback` entry point to perform [callback registrations](#) for [events](#). The return codes for `set_callback` use the `set_result` OMPT type. If the `set_callback` entry point is called outside an `initialize` OMPT callback, [callback registration](#) may fail for supported [callbacks](#) with a return value of `ompt_set_error`. All [registered callbacks](#) and all [callbacks](#) returned by `get_callback` use the `callback` OMPT type as a dummy type signature.

For [callbacks](#) listed in [Table 38.2](#), `ompt_set_always` is the only registration return code that is allowed. An OpenMP implementation must guarantee that the [callback](#) will be invoked every time that a runtime [event](#) that is associated with it occurs. Support for such [callbacks](#) is required in a minimal implementation of the [OMPT](#) interface.

For any other [callbacks](#) not listed in [Table 38.2](#), the `set_callback` entry point may return any non-error code. Whether an OpenMP implementation invokes a registered [callback](#) never, sometimes, or always is [implementation defined](#). If registration for a [callback](#) allows a return code of `ompt_set_never`, support for invoking such a [callback](#) may not be present in a minimal implementation of the [OMPT](#) interface. The return code from [callback registration](#) indicates the [implementation defined](#) level of support for the [callback](#).

Two techniques reduce the size of the [OMPT](#) interface. First, in cases where [events](#) are naturally paired, for example the beginning and end of a [region](#), and the arguments needed by the [callback](#) at each [region endpoint](#) are identical, a [tool](#) registers a single [callback](#) for the pair of [events](#), with `ompt_scope_begin` or `ompt_scope_end` provided as an argument to identify for which [region endpoint](#) the [callback](#) is invoked. Second, when a class of [events](#) is amenable to uniform treatment, [OMPT](#) provides a single [callback](#) for that class of [events](#); for example, a `sync_region_wait` [callback](#) is used for multiple kinds of synchronization [regions](#), such as

TABLE 38.2: Callbacks for which `set_callback` Must Return `ompt_set_always`

Callback Name
<code>thread_begin</code>
<code>thread_end</code>
<code>parallel_begin</code>
<code>parallel_end</code>
<code>task_create</code>
<code>task_schedule</code>
<code>implicit_task</code>
<code>target_data_op_emi</code>
<code>target_emi</code>
<code>target_submit_emi</code>
<code>control_tool</code>
<code>device_initialize</code>
<code>device_finalize</code>
<code>device_load</code>
<code>device_unload</code>
<code>error</code>

`barrier`, `taskwait`, and `taskgroup` regions. Some `events`, for example, those that correspond to `sync_region_wait`, use both techniques.

Cross References

- `get_callback` Entry Point, see [Section 42.5](#)
- `initialize` Callback, see [Section 40.1.1](#)
- OMPT `scope_endpoint` Type, see [Section 39.27](#)
- `set_callback` Entry Point, see [Section 42.4](#)
- OMPT `set_result` Type, see [Section 39.28](#)

38.2.5 Tracing Activity on Target Devices

A `target device` may not initialize a full OpenMP runtime system. Without one, using a `tool` interface based on `callbacks` to monitor activity on a `device` may incur unacceptable overhead. Thus, OMPT defines a monitoring interface for tracing activity on `target devices`. This section details the use of that interface.

First, to prepare to trace `device` activity, a `tool` must register a `device_initialize` callback. A `tool` may also register a `device_load` callback to be notified when code is loaded onto a `target`

`device` or a `device_unload` callback to be notified when code is unloaded from a `target device`. A `tool` may also optionally register a `device_finalize` callback.

When an OpenMP implementation initializes a `target device`, it dispatches the `device_initialize` callback (the `device` initializer) of the `tool` on the `host device`. If the OpenMP implementation or `target device` does not support tracing, the OpenMP implementation passes `NULL` to the `device` initializer of the `tool` for its `lookup` argument; otherwise, the OpenMP implementation passes a pointer to a `device`-specific `function_lookup` entry point to the `device_initialize` callback of the `tool`.

If the `lookup` argument of the `device_initialize` of the `tool` is a `non-null pointer`, the `tool` may use it to determine the `entry points` in the tracing interface that are available for the `device` and may bind the returned function pointers to `tool variables`. Table 38.3 lists the names of `runtime entry points` that may be available for a `device`; an implementation may provide additional `implementation defined` names and corresponding `entry points`. The driver for the `device` provides the `entry points` that enable a `tool` to control the trace collection interface of the `device`. The `native trace format` that the interface uses may be `device-specific` and the available kinds of `trace records` are `implementation defined`.

Some `devices` may allow a `tool` to collect `trace records` in a `standard trace format` known as `OMPT trace records`. Each `OMPT trace record` serves as a substitute for an `OMPT callback` that is not appropriate to be dispatched on the `device`. The fields in each `trace record` type are defined in the description of the `callback` that the record represents. If this type of record is provided then the `function_lookup` entry point returns values for the `entry points` `set_trace_ompt` and `get_record_ompt`, which support collecting and decoding `OMPT` traces. If the `native trace format` for a `device` is the `OMPT` format then tracing can be controlled using the `entry points` for native or `OMPT` tracing.

The `tool` uses the `set_trace_native` and/or the `set_trace_ompt` runtime entry point to specify what types of `events` or activities to monitor on the `device`. The return codes for `set_trace_ompt` and `set_trace_native` use the `set_result` `OMPT` type. If the `set_trace_native` or the `set_trace_ompt` entry point is called outside a `device` initializer, registration of supported `callbacks` may fail with a return code of `ompt_set_error`.

After specifying the `events` or activities to monitor, the `tool` initiates tracing of `device` activity by invoking the `start_trace` entry point. Arguments to `start_trace` include two `tool callbacks` through which the OpenMP implementation can manage traces associated with the `device`. The `buffer_request` callback allocates a buffer in which `trace records` that correspond to `device` activity can be deposited. The `buffer_complete` callback processes a buffer of `trace records` from the `device`.

If the OpenMP implementation requires a trace buffer for `device` activity, it invokes the `tool-supplied callback` on the `host device` to request a new buffer. The OpenMP implementation then 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 device` operation that the `host device`

TABLE 38.3: OMPT Tracing Interface Runtime Entry Point Names and Their Type Signatures

Entry Point String Name	OMPT Type
"ompt_get_device_num_procs"	get_device_num_procs
"ompt_get_device_time"	get_device_time
"ompt_translate_time"	translate_time
"ompt_set_trace_ompt"	set_trace_ompt
"ompt_set_trace_native"	set_trace_native
"ompt_get_buffer_limits"	get_buffer_limits
"ompt_start_trace"	start_trace
"ompt_pause_trace"	pause_trace
"ompt_flush_trace"	flush_trace
"ompt_stop_trace"	stop_trace
"ompt_advance_buffer_cursor"	advance_buffer_cursor
"ompt_get_record_type"	get_record_type
"ompt_get_record_ompt"	get_record_ompt
"ompt_get_record_native"	get_record_native
"ompt_get_record_abstract"	get_record_abstract

initiated to cause these activities.

To correlate activities on the **host device** with activities on a **target device**, a **tool** can register a **target_submit_emi** callback. Before and after the **host device** initiates creation of an **initial task** on a **device** associated with a **structured block** for a **target construct**, the OpenMP implementation dispatches the **target_submit_emi** callback on the **host device** in the **thread** that is executing the **encountering task** of 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_complete** callback to process a non-empty sequence of **trace records** in a trace buffer that is associated with the **device**. The **buffer_complete** callback may return immediately, ignoring records in the trace buffer, or it may iterate through them using the **advance_buffer_cursor** entry point to inspect each **trace record**.

A **tool** may use the **get_record_type** entry point to inspect the type of the **trace record** at the current cursor position. Three entry points (**get_record_ompt**, **get_record_native**, and **get_record_abstract**) allow **tools** to inspect the contents of some or all **trace records** in a trace buffer. The **get_record_native** entry point uses the **native trace format** of the **device**. The **get_record_abstract** entry point decodes the contents of a **native trace record** and summarizes them as a **record_abstract** OMPT type record. The **get_record_ompt** entry point can only be used to retrieve **trace records** in **OMPT** format.

Once **device** tracing has been started, a **tool** may pause or resume **device** tracing at any time by

invoking `pause_trace` with an appropriate flag value as an argument. Further, a `tool` may invoke the `flush_trace` entry point for a `device` at any time between `device` initialization and finalization to cause the pending `trace records` for that `device` to be flushed.

At any time, a `tool` may use the `start_trace` entry point to start or the `stop_trace` entry point to stop `device` tracing. When `device` tracing is stopped, the OpenMP implementation eventually gathers all `trace records` already collected from `device` tracing and presents them to the `tool` using the buffer-completion `callback`.

An OpenMP implementation can be shut down while `device` tracing is in progress. When an OpenMP implementation is shut down, it finalizes each `device`. `Device` finalization occurs in three steps. First, the OpenMP implementation halts any tracing in progress for the `device`. Second, the OpenMP implementation flushes all `trace records` collected for the `device` and uses the `buffer_complete` `callback` associated with that `device` to present them to the `tool`. Finally, the OpenMP implementation dispatches any `device_finalize` `callback` registered for the `device`.

Cross References

- `advance_buffer_cursor` Entry Point, see [Section 43.11](#)
- `buffer_complete` Callback, see [Section 41.6](#)
- `buffer_request` Callback, see [Section 41.5](#)
- `device_finalize` Callback, see [Section 41.2](#)
- `device_initialize` Callback, see [Section 41.1](#)
- `device_load` Callback, see [Section 41.3](#)
- `device_unload` Callback, see [Section 41.4](#)
- `flush_trace` Entry Point, see [Section 43.9](#)
- `function_lookup` Entry Point, see [Section 42.1](#)
- `get_buffer_limits` Entry Point, see [Section 43.6](#)
- `get_device_num_procs` Entry Point, see [Section 43.1](#)
- `get_device_time` Entry Point, see [Section 43.2](#)
- `get_record_abstract` Entry Point, see [Section 43.15](#)
- `get_record_native` Entry Point, see [Section 43.14](#)
- `get_record_ompt` Entry Point, see [Section 43.13](#)
- `get_record_type` Entry Point, see [Section 43.12](#)
- `pause_trace` Entry Point, see [Section 43.8](#)
- OMPT `record_abstract` Type, see [Section 39.24](#)

- OMPT **set_result** Type, see [Section 39.28](#)
- **set_trace_native** Entry Point, see [Section 43.5](#)
- **set_trace_ompt** Entry Point, see [Section 43.4](#)
- **start_trace** Entry Point, see [Section 43.7](#)
- **stop_trace** Entry Point, see [Section 43.10](#)
- **translate_time** Entry Point, see [Section 43.3](#)

38.3 Finalizing a First-Party Tool

If the OMPT interface state is OMPT active, the OMPT-tool finalizer, which is a **finalize callback** and is specified by the **finalize** field in the **start_tool_result** OMPT type structure returned from the **ompt_start_tool** procedure, is called when the OpenMP implementation shuts down.

Cross References

- **finalize** Callback, see [Section 40.1.2](#)
- **ompt_start_tool** Procedure, see [Section 38.2.1](#)
- OMPT **start_tool_result** Type, see [Section 39.30](#)

39 OMPT Data Types

This chapter specifies OMPT types that the `omp-tools.h` C/C++ header file defines.

C / C++

39.1 OMPT Predefined Identifiers

Predefined Identifiers

Name	Value	Properties
<code>ompt_addr_none</code>	<code>~0</code>	<i>default</i>
<code>ompt_mutex_impl_none</code>	<code>0</code>	<i>default</i>

In addition to the predefined identifiers of OMPT type that are defined with their corresponding OMPT type, the OpenMP API includes the predefined identifiers shown above. The `ompt_addr_none` `void *` predefined identifier indicates that no address on the relevant device is available. The `ompt_mutex_impl_none` predefined identifier indicates an invalid mutex implementation.

C / C++

39.2 OMPT any_record_ompt Type

Name: <code>any_record_ompt</code> Properties: C/C++-only, OMPT	Base Type: <code>union</code>
--	--------------------------------------

Fields

Name	Type	Properties
<i>thread_begin</i>	<code>thread_begin</code>	C/C++-only
<i>parallel_begin</i>	<code>parallel_begin</code>	C/C++-only
<i>parallel_end</i>	<code>parallel_end</code>	C/C++-only
<i>work</i>	<code>work</code>	C/C++-only
<i>dispatch</i>	<code>dispatch</code>	C/C++-only
<i>task_create</i>	<code>task_create</code>	C/C++-only
<i>dependences</i>	<code>dependences</code>	C/C++-only
<i>task_dependence</i>	<code>task_dependence</code>	C/C++-only
<i>task_schedule</i>	<code>task_schedule</code>	C/C++-only
<i>implicit_task</i>	<code>implicit_task</code>	C/C++-only
<i>masked</i>	<code>masked</code>	C/C++-only
<i>sync_region</i>	<code>sync_region</code>	C/C++-only
<i>mutex_acquire</i>	<code>mutex_acquire</code>	C/C++-only
<i>mutex</i>	<code>mutex</code>	C/C++-only
<i>nest_lock</i>	<code>nest_lock</code>	C/C++-only
<i>flush</i>	<code>flush</code>	C/C++-only
<i>cancel</i>	<code>cancel</code>	C/C++-only
<i>target_emi</i>	<code>target_emi</code>	C/C++-only
<i>target_data_op_emi</i>	<code>target_data_op_emi</code>	C/C++-only
<i>target_map_emi</i>	<code>target_map_emi</code>	C/C++-only
<i>target_submit_emi</i>	<code>target_submit_emi</code>	C/C++-only
<i>control_tool</i>	<code>control_tool</code>	C/C++-only
<i>error</i>	<code>error</code>	C/C++-only

Type Definition

C / C++

```
typedef union ompt_any_record_ompt_t {  
    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;  
    ompt_record_task_dependence_t task_dependence;  
    ompt_record_task_schedule_t task_schedule;
```

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15

```
ompt_record_implicit_task_t implicit_task;  
ompt_record_masked_t masked;  
ompt_record_sync_region_t sync_region;  
ompt_record_mutex_acquire_t mutex_acquire;  
ompt_record_mutex_t mutex;  
ompt_record_nest_lock_t nest_lock;  
ompt_record_flush_t flush;  
ompt_record_cancel_t cancel;  
ompt_record_target_emi_t target_emi;  
ompt_record_target_data_op_emi_t target_data_op_emi;  
ompt_record_target_map_emi_t target_map_emi;  
ompt_record_target_submit_emi_t target_submit_emi;  
ompt_record_control_tool_t control_tool;  
ompt_record_error_t error;  
} ompt_any_record_ompt_t;
```

▲ C / C++ ▲

16
17
18
19
20
21
22
23

Additional information

The [union](#) also includes `target`, `target_data_op`, `target_kernel`, and `target_map` fields with corresponding [trace record OMPT types](#). These fields have been [deprecated](#).

Semantics

The [any_record_ompt](#) OMPT type is a union of all [standard trace format event-specific trace record OMPT types](#) that is the type of the `record` field of the `record_ompt` OMPT type.

Cross References

- OMPT `record_ompt` Type, see [Section 39.26](#)

24

39.3 OMPT buffer Type

25

Name: <code>buffer</code> Properties: C/C++-only , OMPT , opaque	Base Type: <code>void</code>
---	-------------------------------------

26

Type Definition

27

```
typedef void ompt_buffer_t;
```

▲ C / C++ ▲

28
29

Semantics

The [buffer](#) OMPT type represents a [handle](#) for a [device](#) buffer.

1
2
3
4
5
6
7
8
9
10
11
12
13
14

39.4 OMPT buffer_cursor Type

Name: <code>buffer_cursor</code> Properties: C/C++-only, OMPT, opaque	Base Type: <code>c_uint64_t</code>
--	------------------------------------

Type Definition

C / C++

```
typedef uint64_t ompt_buffer_cursor_t;
```

C / C++

Summary

The `buffer_cursor` OMPT type represents a `handle` for a position in a `device` buffer.

39.5 OMPT callback Type

Name: <code>callback</code> Category: <code>subroutine</code> pointer	Properties: C/C++-only, OMPT
--	------------------------------

Type Signature

C / C++

```
typedef void (*ompt_callback_t) (void);
```

C / C++

Semantics

Pointers to OMPT callbacks with different type signatures are passed to the `set_callback` entry point and returned by the `get_callback` entry point. For convenience, these entry points require all type signatures to be cast to the `callback` OMPT type.

39.6 OMPT callbacks Type

Name: callbacks Properties: C/C++-only, OMPT	Base Type: enumeration
---	-------------------------------

Values

Name	Value	Properties
ompt_callback_thread_begin	1	C-only, OMPT
ompt_callback_thread_end	2	C-only, OMPT
ompt_callback_parallel_begin	3	C-only, OMPT
ompt_callback_parallel_end	4	C-only, OMPT
ompt_callback_task_create	5	C-only, OMPT
ompt_callback_task_schedule	6	C-only, OMPT
ompt_callback_implicit_task	7	C-only, OMPT
ompt_callback_control_tool	11	C-only, OMPT
ompt_callback_device_initialize	12	C-only, OMPT
ompt_callback_device_finalize	13	C-only, OMPT
ompt_callback_device_load	14	C-only, OMPT
ompt_callback_device_unload	15	C-only, OMPT
ompt_callback_sync_region_wait	16	C-only, OMPT
ompt_callback_mutex_released	17	C-only, OMPT
ompt_callback_dependences	18	C-only, OMPT
ompt_callback_task_dependence	19	C-only, OMPT
ompt_callback_work	20	C-only, OMPT
ompt_callback_masked	21	C-only, OMPT
ompt_callback_sync_region	23	C-only, OMPT
ompt_callback_lock_init	24	C-only, OMPT
ompt_callback_lock_destroy	25	C-only, OMPT
ompt_callback_mutex_acquire	26	C-only, OMPT
ompt_callback_mutex_acquired	27	C-only, OMPT
ompt_callback_nest_lock	28	C-only, OMPT
ompt_callback_flush	29	C-only, OMPT
ompt_callback_cancel	30	C-only, OMPT
ompt_callback_reduction	31	C-only, OMPT
ompt_callback_dispatch	32	C-only, OMPT
ompt_callback_target_emi	33	C-only, OMPT
ompt_callback_target_data_op_emi	34	C-only, OMPT
ompt_callback_target_submit_emi	35	C-only, OMPT
ompt_callback_target_map_emi	36	C-only, OMPT
ompt_callback_error	37	C-only, OMPT

Type Definition

```

1  typedef enum omp_t_callbacks_t {
2      omp_callback_thread_begin          = 1,
3      omp_callback_thread_end            = 2,
4      omp_callback_parallel_begin        = 3,
5      omp_callback_parallel_end          = 4,
6      omp_callback_task_create           = 5,
7      omp_callback_task_schedule         = 6,
8      omp_callback_implicit_task         = 7,
9      omp_callback_control_tool          = 11,
10     omp_callback_device_initialize      = 12,
11     omp_callback_device_finalize       = 13,
12     omp_callback_device_load           = 14,
13     omp_callback_device_unload         = 15,
14     omp_callback_sync_region_wait      = 16,
15     omp_callback_mutex_released        = 17,
16     omp_callback_dependences           = 18,
17     omp_callback_task_dependence       = 19,
18     omp_callback_work                  = 20,
19     omp_callback_masked                 = 21,
20     omp_callback_sync_region            = 23,
21     omp_callback_lock_init              = 24,
22     omp_callback_lock_destroy           = 25,
23     omp_callback_mutex_acquire          = 26,
24     omp_callback_mutex_acquired         = 27,
25     omp_callback_nest_lock              = 28,
26     omp_callback_flush                  = 29,
27     omp_callback_cancel                 = 30,
28     omp_callback_reduction              = 31,
29     omp_callback_dispatch                = 32,
30     omp_callback_target_emi             = 33,
31     omp_callback_target_data_op_emi     = 34,
32     omp_callback_target_submit_emi      = 35,
33     omp_callback_target_map_emi         = 36,
34     omp_callback_error                  = 37
35 } omp_callbacks_t;

```

Additional information

The following instances and associated values of the **callbacks OMPT type** are also defined: **omp_callback_target**, with value 8; **omp_callback_target_data_op**, with value 9; **omp_callback_target_submit**, with value 10; and **omp_callback_target_map**, with value 22. These instances have been [deprecated](#).

Semantics

The `callbacks` OMPT type provides codes that identify OMPT callbacks when registering or querying them.

39.7 OMPT `cancel_flag` Type

Name: <code>cancel_flag</code> Properties: C/C++-only, OMPT	Base Type: enumeration
--	-------------------------------

Values

Name	Value	Properties
<code>ompt_cancel_parallel</code>	<code>0x01</code>	C/C++-only, OMPT
<code>ompt_cancel_sections</code>	<code>0x02</code>	C/C++-only, OMPT
<code>ompt_cancel_loop</code>	<code>0x04</code>	C/C++-only, OMPT
<code>ompt_cancel_taskgroup</code>	<code>0x08</code>	C/C++-only, OMPT
<code>ompt_cancel_activated</code>	<code>0x10</code>	C/C++-only, OMPT
<code>ompt_cancel_detected</code>	<code>0x20</code>	C/C++-only, OMPT
<code>ompt_cancel_discarded_task</code>	<code>0x40</code>	C/C++-only, OMPT

Type Definition

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++

Semantics

The `cancel_flag` OMPT type defines cancel flag values.

39.8 OMPT `data` Type

Name: <code>data</code> Properties: C/C++-only, OMPT	Base Type: union
---	-------------------------

Fields

Name	Type	Properties
<i>value</i>	c_uint64_t	<i>default</i>
<i>ptr</i>	void	C/C++-only, pointer

Predefined Identifiers

Name	Value	Properties
ompt_data_none	0	C/C++-only, OMPT

Type Definition

C / C++

```
typedef union ompt_data_t {  
    uint64_t value;  
    void *ptr;  
} ompt_data_t;
```

C / C++

Semantics

The **data OMPT type** represents data that is reserved for **tool** use. When an OpenMP implementation creates a **thread** or an instance of a **parallel region**, **teams region**, **task region**, or **device region**, it initializes the associated **data** object with the value **ompt_data_none**.

39.9 OMPT dependence Type

Name: dependence Properties: C/C++-only, OMPT	Base Type: structure
---	------------------------------------

Fields

Name	Type	Properties
<i>variable</i>	data	C/C++-only
<i>dependence_type</i>	dependence_type	C/C++-only

Type Definition

C / C++

```
typedef struct ompt_dependence_t {  
    ompt_data_t variable;  
    ompt_dependence_type_t dependence_type;  
} ompt_dependence_t;
```

C / C++

Semantics

The **dependence OMPT type** represents a **dependence** in a **structure** that holds information about a **depend** or **doacross** clause. For **task dependences**, the **ptr** field of its **variable** field points to the **storage location** of the **dependence**. For **doacross dependences**, the **value** field of the **variable** 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 the `variable` field is undefined and the `dependence_type` field contains a value that has the `_all_memory` suffix.

Cross References

- OMPT `data` Type, see [Section 39.8](#)
- OMPT `dependence_type` Type, see [Section 39.10](#)

39.10 OMPT `dependence_type` Type

Name: <code>dependence_type</code> Properties: C/C++-only , OMPT	Base Type: enumeration
---	--

Values

Name	Value	Properties
<code>ompt_dependence_type_in</code>	1	C/C++-only , OMPT
<code>ompt_dependence_type_out</code>	2	C/C++-only , OMPT
<code>ompt_dependence_type_inout</code>	3	C/C++-only , OMPT
<code>ompt_dependence_type_mutexinoutset</code>	4	C/C++-only , OMPT
<code>ompt_dependence_type_source</code>	5	C/C++-only , OMPT
<code>ompt_dependence_type_sink</code>	6	C/C++-only , OMPT
<code>ompt_dependence_type_inoutset</code>	7	C/C++-only , OMPT
<code>ompt_dependence_type_out_all_memory</code>	34	C/C++-only , OMPT
<code>ompt_dependence_type_inout_all_memory</code>	35	C/C++-only , OMPT

Type Definition

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++

1 **Semantics**

2 The `dependence_type` OMPT type defines task dependence type values. The
3 `ompt_dependence_type_in`, `ompt_dependence_type_out`,
4 `ompt_dependence_type_inout`, `ompt_dependence_type_mutexinoutset`,
5 `ompt_dependence_type_inoutset`, `ompt_dependence_type_out_all_memory`,
6 and `ompt_dependence_type_inout_all_memory` values represent the task dependence
7 type present in a `depend` clause while the `ompt_dependence_type_source` and
8 `ompt_dependence_type_sink` values represent the *dependence-type* present in a
9 `doacross` clause. The `ompt_dependence_type_out_all_memory` and
10 `ompt_dependence_type_inout_all_memory` represent task dependences for which the
11 `omp_all_memory` reserved locator is specified.

12 **39.11 OMPT device Type**

13

Name: <code>device</code> Properties: C/C++-only, OMPT, opaque	Base Type: <code>void</code>
---	------------------------------

14 **Type Definition**

15
16

16 **Semantics**

17 The `device` OMPT type represents a `device`.

18 **39.12 OMPT device_time Type**

19

Name: <code>device_time</code> Properties: C/C++-only, OMPT, opaque	Base Type: <code>c_uint64_t</code>
--	------------------------------------

20 **Predefined Identifiers**

21

Name	Value	Properties
<code>ompt_time_none</code>	0	C/C++-only, OMPT

22 **Type Definition**

23
24

24 **Semantics**

25 The `device_time` OMPT type represents raw `device` time values; `ompt_time_none`
26 represents an unknown or unspecified time.

39.13 OMPT dispatch Type

Name: <code>dispatch</code> Properties: C/C++-only, OMPT, overlapping-type-name	Base Type: <code>enumeration</code>
--	--

Values

Name	Value	Properties
<code>ompt_dispatch_iteration</code>	1	C/C++-only, OMPT
<code>ompt_dispatch_section</code>	2	C/C++-only, OMPT
<code>ompt_dispatch_ws_loop_chunk</code>	3	C/C++-only, OMPT
<code>ompt_dispatch_taskloop_chunk</code>	4	C/C++-only, OMPT
<code>ompt_dispatch_distribute_chunk</code>	5	C/C++-only, OMPT

Type Definition

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++

Semantics

The `dispatch` OMPT type defines the valid dispatch values.

39.14 OMPT dispatch_chunk Type

Name: <code>dispatch_chunk</code> Properties: C/C++-only, OMPT	Base Type: <code>structure</code>
---	--

Fields

Name	Type	Properties
<code>start</code>	<code>c_uint64_t</code>	<i>default</i>
<code>iterations</code>	<code>c_uint64_t</code>	<i>default</i>

Type Definition

C / C++

```
typedef struct ompt_dispatch_chunk_t {  
    uint64_t start;  
    uint64_t iterations;  
} ompt_dispatch_chunk_t;
```

C / C++

1 **Semantics**

2 The **dispatch_chunk** OMPT type represents **chunk** information for a dispatched **chunk**. The
3 **start** field specifies the first **logical iteration** of the **chunk** and the **iterations** field specifies
4 the number of **logical iterations** in the **chunk**. Whether the **chunk** of a **taskloop** region is
5 contiguous is **implementation defined**.

6 **39.15 OMPT frame Type**

7

Name: frame Properties: C/C++-only, OMPT	Base Type: structure
--	------------------------------------

8 **Fields**

9

Name	Type	Properties
<i>exit_frame</i>	data	C/C++-only, OMPT
<i>enter_frame</i>	data	C/C++-only, OMPT
<i>exit_frame_flags</i>	integer	<i>default</i>
<i>enter_frame_flags</i>	integer	<i>default</i>

10 **Type Definition**

11

C / C++

```
12      typedef struct ompt_frame_t {  
13          ompt_data_t exit_frame;  
14          ompt_data_t enter_frame;  
15          int exit_frame_flags;  
16          int enter_frame_flags;  
17      } ompt_frame_t;
```

C / C++

17 **Semantics**

18 The **frame** OMPT type describes **procedure frame** information for a **task**. Each **frame** object is
19 associated with the **task** to which the **procedure frames** belong. Every **task** that is not a **merged task**
20 with one or more **frames** on the stack of a **native thread**, whether an **initial task**, an **implicit task**, an
21 **explicit task**, or a **target task**, has an associated **frame** object.

22 The **exit_frame** field contains information to identify the first **procedure frame** executing the
23 **task region**. The **exit_frame** for the **frame** object associated with the **initial task** that is not
24 nested inside any OpenMP **construct** is **ompt_data_none**. The **enter_frame** field contains
25 information to identify the latest still active **procedure frame** executing the **task region** before
26 entering the OpenMP runtime implementation or before executing a different **task**. If a **task** with
27 **frames** on the stack is not executing **implementation code** in the OpenMP runtime, the value of
28 **enter_frame** for its associated **frame** object is **ompt_data_none**.

29 For the **frame** indicated by **exit_frame** (**enter_frame**), the **exit_frame_flags**
30 (**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 [frame_flag](#) OMPT type.

The lifetime of a [frame](#) 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 a [frame](#) object is passed to some [callbacks](#); a pointer to the [frame](#) object of a [task](#) can also be retrieved by a [tool](#) at any time, including in a [signal handler](#), by invoking the [get_task_info](#) entry point. A pointer to a [frame](#) 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 [frame](#) objects observed just prior to when their field values will be set or cleared.

Cross References

- OMPT [data](#) Type, see [Section 39.8](#)
- OMPT [frame_flag](#) Type, see [Section 39.16](#)
- [get_task_info](#) Entry Point, see [Section 42.15](#)

39.16 OMPT frame_flag Type

Name: frame_flag Properties: C/C++-only , OMPT	Base Type: enumeration
---	--

Values

Name	Value	Properties
ompt_frame_runtime	0x00	C/C++-only , OMPT
ompt_frame_application	0x01	C/C++-only , OMPT
ompt_frame_cfa	0x10	C/C++-only , OMPT
ompt_frame_framepointer	0x20	C/C++-only , OMPT
ompt_frame_stackaddress	0x30	C/C++-only , OMPT

Type Definition

C / C++

```
typedef enum ompt_frame_flag_t {  
    ompt_frame_runtime      = 0x00,  
    ompt_frame_application  = 0x01,  
    ompt_frame_cfa          = 0x10,  
    ompt_frame_framepointer = 0x20,  
};
```

```
1  ompt_frame_stackaddress = 0x30
2  } ompt_frame_flag_t;
```

C / C++

Semantics

The **frame_flag** OMPT type defines frame information flags. The **ompt_frame_runtime** value indicates that a frame address is a procedure frame in the OpenMP runtime implementation. The **ompt_frame_application** value indicates that a frame address is a procedure frame in the OpenMP program. Higher order bits indicate the specific information for a particular frame pointer. The **ompt_frame_cfa** value indicates that a frame address specifies a canonical frame address. The **ompt_frame_framepointer** value indicates that a frame address provides the value of the frame pointer register. The **ompt_frame_stackaddress** value indicates that a frame address specifies a pointer address that is contained in the current stack frame.

39.17 OMPT hwid Type

Name: hwid Properties: C/C++-only, OMPT	Base Type: c_uint64_t
--	-----------------------

Predefined Identifiers

Name	Value	Properties
ompt_hwid_none	0	C/C++-only, OMPT

Type Definition

C / C++

```
typedef uint64_t ompt_hwid_t;
```

C / C++

Semantics

The **hwid** OMPT type is a handle for a hardware identifier for a target device; **ompt_hwid_none** represents an unknown or unspecified hardware identifier. If no specific value for the **hwid** field is associated with an instance of the **record_abstract** OMPT type then the value of **hwid** is **ompt_hwid_none**.

Cross References

- OMPT **record_abstract** Type, see [Section 39.24](#)

39.18 OMPT id Type

Name: id Properties: C/C++-only, OMPT	Base Type: c_uint64_t
--	-----------------------

1 **Predefined Identifiers**

2

Name	Value	Properties
<code>ompt_id_none</code>	0	C/C++-only, OMPT

3 **Type Definition**

4

<code>typedef uint64_t ompt_id_t;</code>	C / C++
--	---------

5 **Semantics**

6 The `id` OMPT type is used to provide various identifiers to `tools`; `ompt_id_none` is used when
7 the specific ID is unknown or unavailable. When tracing asynchronous activity on `devices`,
8 identifiers enable `tools` to correlate `device regions` and operations that the `host device` initiates with
9 associated activities on a `target device`. In addition, OMPT provides identifiers to refer to `parallel`
10 `regions` and `tasks` that execute on a `device`.

11 **Restrictions**

12 Restrictions to the `id` OMPT type are as follows:

- 13
- 14 Identifiers created on each `device` must be unique from the time an OpenMP implementation
15 is initialized until it is shut down. Identifiers for each `device region` and target data operation
16 instance that the `host device` initiates must be unique over time on the `host device`. Identifiers
17 for instances of `parallel regions` and `task regions` that execute on a `device` must be unique over
time within that `device`.

18 **39.19 OMPT `interface_fn` Type**

19

Name: <code>interface_fn</code>	Properties: C/C++-only, OMPT
Category: <code>subroutine</code> pointer	

20 **Type Signature**

21

<code>typedef void (*ompt_interface_fn_t) (void);</code>	C / C++
--	---------

22 **Semantics**

23 The `interface_fn` OMPT type serves as a generic function pointer that the
24 `function_lookup` entry point returns to provide access to a `tool` to `entry points` by name.

39.20 OMPT mutex Type

Name: mutex Properties: C/C++-only, OMPT, overlapping-type-name	Base Type: enumeration
--	-------------------------------

Values

Name	Value	Properties
ompt_mutex_lock	1	C/C++-only, OMPT
ompt_mutex_test_lock	2	C/C++-only, OMPT
ompt_mutex_nest_lock	3	C/C++-only, OMPT
ompt_mutex_test_nest_lock	4	C/C++-only, OMPT
ompt_mutex_critical	5	C/C++-only, OMPT
ompt_mutex_atomic	6	C/C++-only, OMPT
ompt_mutex_ordered	7	C/C++-only, OMPT

Type Definition

C / C++
<pre>typedef enum ompt_mutex_t { ompt_mutex_lock = 1, ompt_mutex_test_lock = 2, ompt_mutex_nest_lock = 3, ompt_mutex_test_nest_lock = 4, ompt_mutex_critical = 5, ompt_mutex_atomic = 6, ompt_mutex_ordered = 7 } ompt_mutex_t;</pre>
C / C++

Semantics

The **mutex** OMPT type defines the valid mutex values.

39.21 OMPT native_mon_flag Type

Name: native_mon_flag Properties: C/C++-only, OMPT	Base Type: enumeration
---	-------------------------------

1 **Values**

Name	Value	Properties
ompt_native_data_motion_explicit	0x01	C/C++-only, OMPT
ompt_native_data_motion_implicit	0x02	C/C++-only, OMPT
ompt_native_kernel_invocation	0x04	C/C++-only, OMPT
ompt_native_kernel_execution	0x08	C/C++-only, OMPT
ompt_native_driver	0x10	C/C++-only, OMPT
ompt_native_runtime	0x20	C/C++-only, OMPT
ompt_native_overhead	0x40	C/C++-only, OMPT
ompt_native_idleness	0x80	C/C++-only, OMPT

3 **Type Definition**

```
4 typedef enum ompt_native_mon_flag_t {  
5     ompt_native_data_motion_explicit = 0x01,  
6     ompt_native_data_motion_implicit = 0x02,  
7     ompt_native_kernel_invocation    = 0x04,  
8     ompt_native_kernel_execution     = 0x08,  
9     ompt_native_driver               = 0x10,  
10    ompt_native_runtime               = 0x20,  
11    ompt_native_overhead              = 0x40,  
12    ompt_native_idleness              = 0x80  
13 } ompt_native_mon_flag_t;
```

14 **Semantics**

15 The `native_mon_flag` OMPT type defines the valid native monitoring flag values.

16 **39.22 OMPT parallel_flag Type**

Name: parallel_flag Properties: C/C++-only, OMPT	Base Type: enumeration
---	-------------------------------

18 **Values**

Name	Value	Properties
ompt_parallel_invoker_program	0x00000001	C/C++-only, OMPT
ompt_parallel_invoker_runtime	0x00000002	C/C++-only, OMPT
ompt_parallel_league	0x40000000	C/C++-only, OMPT
ompt_parallel_team	0x80000000	C/C++-only, OMPT

1 **Type Definition**

C / C++

```
2      typedef enum ompt_parallel_flag_t {
3          ompt_parallel_invoker_program = 0x00000001,
4          ompt_parallel_invoker_runtime = 0x00000002,
5          ompt_parallel_league            = 0x40000000,
6          ompt_parallel_team             = 0x80000000
7      } ompt_parallel_flag_t;
```

C / C++

8 **Semantics**

9 The **parallel_flag** OMPT type defines valid invoker values, which indicate how the code that
10 implements the associated **structured block** of the region is invoked or encountered. The
11 **ompt_parallel_invoker_program** value indicates that the **encountering thread** for a
12 **parallel** or **teams** region executes code to implement its associated **structured block** as if
13 directly invoked or encountered in application code. The
14 **ompt_parallel_invoker_runtime** value indicates that the **encountering thread** for a
15 **parallel** or **teams** region invokes the code that implements its associated **structured block**
16 from the runtime. The **ompt_parallel_league** value indicates that the **callback** is invoked
17 due to the creation of a **league** of **teams** by a **teams** construct. The **ompt_parallel_team**
18 value indicates that the **callback** is invoked due to the creation of a **team** of **threads** by a **parallel**
19 construct.

20 **39.23 OMPT record Type**

Name: record	Base Type: enumeration
Properties: C/C++-only, OMPT	

22 **Values**

Name	Value	Properties
ompt_record_ompt	1	C/C++-only, OMPT
ompt_record_native	2	C/C++-only, OMPT
ompt_record_invalid	3	C/C++-only, OMPT

24 **Type Definition**

C / C++

```
25      typedef enum ompt_record_t {
26          ompt_record_ompt        = 1,
27          ompt_record_native      = 2,
28          ompt_record_invalid = 3
29      } ompt_record_t;
```

C / C++

Semantics

The **record** OMPT type indicates the integer codes that identify OMPT trace record formats.

39.24 OMPT record_abstract Type

Name: record_abstract Properties: C/C++-only, OMPT	Base Type: structure
--	-----------------------------

Fields

Name	Type	Properties
<i>rclass</i>	record_native	C/C++-only, OMPT
<i>type</i>	char	common-field, intent(in), pointer
<i>start_time</i>	device_time	C/C++-only, OMPT
<i>end_time</i>	device_time	C/C++-only, OMPT
<i>hwid</i>	hwid	C/C++-only, OMPT

Type Definition

```
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

The **record_abstract** OMPT type is an abstract trace record format that summarizes native trace records. It contains information that a tool can use to process a native trace record that it may not fully understand. The **rclass** field indicates that the trace record is informational or that it represents an event; this information can help a tool determine how to present the trace record. The **type** field points to a statically-allocated, immutable character string that provides a meaningful name that a tool can use to describe the event. 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 trace record then the value of **hwid** is **ompt_hwid_none**.

Cross References

- OMPT `device_time` Type, see [Section 39.12](#)
- OMPT `hwid` Type, see [Section 39.17](#)
- OMPT `record_native` Type, see [Section 39.25](#)

39.25 OMPT `record_native` Type

Name: <code>record_native</code> Properties: C/C++-only , OMPT	Base Type: enumeration
---	---

Values

Name	Value	Properties
<code>ompt_record_native_info</code>	1	C/C++-only , OMPT
<code>ompt_record_native_event</code>	2	C/C++-only , OMPT

Type Definition

C / C++

```
typedef enum ompt_record_native_t {  
    ompt_record_native_info = 1,  
    ompt_record_native_event = 2  
} ompt_record_native_t;
```

C / C++

Semantics

The `record_native` OMPT type indicates the integer codes that identify OMPT native trace record contents.

39.26 OMPT `record_ompt` Type

Name: <code>record_ompt</code> Properties: C/C++-only , OMPT	Base Type: structure
---	---

Fields

Name	Type	Properties
<i>type</i>	callbacks	C/C++-only , common-field , OMPT
<i>time</i>	<code>device_time</code>	C/C++-only , OMPT
<i>thread_id</i>	<code>id</code>	C/C++-only , OMPT
<i>target_id</i>	<code>id</code>	C/C++-only , OMPT
<i>record</i>	<code>any_record_ompt</code>	C/C++-only , OMPT

Type Definition

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;
    ompt_any_record_ompt_t record;
} ompt_record_ompt_t;
```

C / C++

Semantics

The `record_ompt` OMPT type provides a complete `trace record` by specifying the common fields of the `standard trace format` along with a field that is an instance of the `any_record_ompt OMPT type`. The `type` field specifies the type of `trace record` that the `structure` provides. According to the type, `event`-specific information is stored in the matching `record` field.

Restrictions

Restrictions to the `record_ompt` OMPT type are as follows:

- If `type` is `ompt_callback_thread_end` then the value of `record` is undefined.

Cross References

- OMPT `any_record_ompt` Type, see [Section 39.2](#)
- OMPT `callbacks` Type, see [Section 39.6](#)
- OMPT `device_time` Type, see [Section 39.12](#)
- OMPT `id` Type, see [Section 39.18](#)

39.27 OMPT `scope_endpoint` Type

Name: <code>scope_endpoint</code> Properties: <code>C/C++-only</code> , <code>OMPT</code>	Base Type: <code>enumeration</code>
--	--

Values

Name	Value	Properties
<code>ompt_scope_begin</code>	1	<code>C/C++-only</code> , <code>OMPT</code>
<code>ompt_scope_end</code>	2	<code>C/C++-only</code> , <code>OMPT</code>
<code>ompt_scope_beginend</code>	3	<code>C/C++-only</code> , <code>OMPT</code>

1 **Type Definition**

C / C++

```
2      typedef enum ompt_scope_endpoint_t {  
3          ompt_scope_begin        = 1,  
4          ompt_scope_end         = 2,  
5          ompt_scope_beginend = 3  
6      } ompt_scope_endpoint_t;
```

C / C++

7 **Summary**

8 The **scope_endpoint** OMPT type defines valid **region endpoint** values.

9 **39.28 OMPT set_result Type**

Name: set_result Properties: C/C++-only, OMPT	Base Type: enumeration
---	-------------------------------

11 **Values**

Name	Value	Properties
ompt_set_error	0	C/C++-only, OMPT
ompt_set_never	1	C/C++-only, OMPT
ompt_set_impossible	2	C/C++-only, OMPT
ompt_set_sometimes	3	C/C++-only, OMPT
ompt_set_sometimes_paired	4	C/C++-only, OMPT
ompt_set_always	5	C/C++-only, OMPT

13 **Type Definition**

C / C++

```
14     typedef enum ompt_set_result_t {  
15        ompt_set_error                = 0,  
16        ompt_set_never                = 1,  
17        ompt_set_impossible          = 2,  
18        ompt_set_sometimes            = 3,  
19        ompt_set_sometimes_paired = 4,  
20        ompt_set_always               = 5  
21     } ompt_set_result_t;
```

C / C++

1 **Summary**

2 The `set_result` OMPT type corresponds to values that the `set_callback`,
3 `set_trace_ompt` and `set_trace_native` entry points return. Its values indicate several
4 possible outcomes. The `ompt_set_error` value indicates that the associated call failed.
5 Otherwise, the value indicates when an `event` may occur and, when appropriate, `callback dispatch`
6 leads to the invocation of the `callback`. The `ompt_set_never` value indicates that the `event` will
7 never occur or that the `callback` will never be invoked at runtime. The `ompt_set_impossible`
8 value indicates that the `event` may occur but that tracing of it is not possible. The
9 `ompt_set_sometimes` value indicates that the `event` may occur and, for an `implementation`
10 `defined` subset of associated `event` occurrences, will be traced or the `callback` will be invoked at
11 runtime. The `ompt_set_sometimes_paired` value indicates the same result as
12 `ompt_set_sometimes` and, in addition, that a `callback` with an `endpoint` value of
13 `ompt_scope_begin` will be invoked if and only if the same `callback` with an `endpoint` value of
14 `ompt_scope_end` will also be invoked sometime in the future. The `ompt_set_always` value
15 indicates that, whenever an associated `event` occurs, it will be traced or the `callback` will be invoked.

16 **Cross References**

- 17 • OMPT `scope_endpoint` Type, see [Section 39.27](#)
- 18 • `set_callback` Entry Point, see [Section 42.4](#)
- 19 • `set_trace_native` Entry Point, see [Section 43.5](#)
- 20 • `set_trace_ompt` Entry Point, see [Section 43.4](#)

21 **39.29 OMPT severity Type**

22 Name: severity Properties: C/C++-only, OMPT	Base Type: enumeration
---	--

23 **Values**

24 Name	Value	Properties
<code>ompt_warning</code>	1	C/C++-only, OMPT
<code>ompt_fatal</code>	2	C/C++-only, OMPT

25 **Type Definition**

C / C++

```
26      typedef enum ompt_severity_t {  
27          ompt_warning = 1,  
28          ompt_fatal   = 2  
29      } ompt_severity_t;
```

C / C++

Semantics

The **severity** OMPT type defines severity values.

39.30 OMPT start_tool_result Type

Name: <code>start_tool_result</code> Properties: C/C++-only, OMPT	Base Type: <code>structure</code>
--	--

Fields

Name	Type	Properties
<i>initialize</i>	initialize	C/C++-only, OMPT
<i>finalize</i>	finalize	C/C++-only, OMPT
<i>tool_data</i>	data	C/C++-only, OMPT

Type Definition

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++

Semantics

The **ompt_start_tool** procedure returns a pointer to a `structure` of the **start_tool_result** OMPT type, which provides pointers to the tool's **initialize** and **finalize** callbacks as well as a **data** object for use by the **tool**.

Restrictions

Restrictions to the **start_tool_result** OMPT type are as follows:

- The **initialize** and **finalize** callback pointer values in a **start_tool_result** structure that **ompt_start_tool** returns must be **non-null** values.

Cross References

- OMPT **data** Type, see [Section 39.8](#)
- **finalize** Callback, see [Section 40.1.2](#)
- **initialize** Callback, see [Section 40.1.1](#)
- **ompt_start_tool** Procedure, see [Section 38.2.1](#)

39.31 OMPT state Type

Name: <code>state</code> Properties: <code>C/C++-only</code> , <code>OMPT</code>	Base Type: <code>enumeration</code>
---	--

Values

Name	Value	Properties
<code>ompt_state_work_serial</code>	<code>0x000</code>	<code>C/C++-only</code> , <code>OMPT</code>
<code>ompt_state_work_parallel</code>	<code>0x001</code>	<code>C/C++-only</code> , <code>OMPT</code>
<code>ompt_state_work_reduction</code>	<code>0x002</code>	<code>C/C++-only</code> , <code>OMPT</code>
<code>ompt_state_work_free_agent</code>	<code>0x003</code>	<code>C/C++-only</code> , <code>OMPT</code>
<code>ompt_state_work_induction</code>	<code>0x004</code>	<code>C/C++-only</code> , <code>OMPT</code>
<code>ompt_state_wait_barrier_implicit_parallel</code>	<code>0x011</code>	<code>C/C++-only</code> , <code>OMPT</code>
<code>ompt_state_wait_barrier_implicit_workshare</code>	<code>0x012</code>	<code>C/C++-only</code> , <code>OMPT</code>
<code>ompt_state_wait_barrier_explicit</code>	<code>0x014</code>	<code>C/C++-only</code> , <code>OMPT</code>
<code>ompt_state_wait_barrier_implementation</code>	<code>0x015</code>	<code>C/C++-only</code> , <code>OMPT</code>
<code>ompt_state_wait_barrier_teams</code>	<code>0x016</code>	<code>C/C++-only</code> , <code>OMPT</code>
<code>ompt_state_wait_taskwait</code>	<code>0x020</code>	<code>C/C++-only</code> , <code>OMPT</code>
<code>ompt_state_wait_taskgroup</code>	<code>0x021</code>	<code>C/C++-only</code> , <code>OMPT</code>
<code>ompt_state_wait_mutex</code>	<code>0x040</code>	<code>C/C++-only</code> , <code>OMPT</code>
<code>ompt_state_wait_lock</code>	<code>0x041</code>	<code>C/C++-only</code> , <code>OMPT</code>
<code>ompt_state_wait_critical</code>	<code>0x042</code>	<code>C/C++-only</code> , <code>OMPT</code>
<code>ompt_state_wait_atomic</code>	<code>0x043</code>	<code>C/C++-only</code> , <code>OMPT</code>
<code>ompt_state_wait_ordered</code>	<code>0x044</code>	<code>C/C++-only</code> , <code>OMPT</code>
<code>ompt_state_wait_target</code>	<code>0x080</code>	<code>C/C++-only</code> , <code>OMPT</code>
<code>ompt_state_wait_target_map</code>	<code>0x081</code>	<code>C/C++-only</code> , <code>OMPT</code>
<code>ompt_state_wait_target_update</code>	<code>0x082</code>	<code>C/C++-only</code> , <code>OMPT</code>
<code>ompt_state_idle</code>	<code>0x100</code>	<code>C/C++-only</code> , <code>OMPT</code>
<code>ompt_state_overhead</code>	<code>0x101</code>	<code>C/C++-only</code> , <code>OMPT</code>
<code>ompt_state_undefined</code>	<code>0x102</code>	<code>C/C++-only</code> , <code>OMPT</code>

Type Definition

C / C++

```
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_work_induction = 0x004,  
    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,  
}
```

```

1      ompt_state_wait_barrier_teams          = 0x016,
2      ompt_state_wait_taskwait              = 0x020,
3      ompt_state_wait_taskgroup             = 0x021,
4      ompt_state_wait_mutex                 = 0x040,
5      ompt_state_wait_lock                  = 0x041,
6      ompt_state_wait_critical               = 0x042,
7      ompt_state_wait_atomic                = 0x043,
8      ompt_state_wait_ordered               = 0x044,
9      ompt_state_wait_target                = 0x080,
10     ompt_state_wait_target_map             = 0x081,
11     ompt_state_wait_target_update          = 0x082,
12     ompt_state_idle                       = 0x100,
13     ompt_state_overhead                    = 0x101,
14     ompt_state_undefined                   = 0x102
15 } ompt_state_t;

```

C / C++

Semantics

The **state OMPT type** defines **thread states** that indicate the current activity of a **thread**. 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**. A **thread state** must be one of the values of the **state OMPT type** or an **implementation defined** state value of 0x200 (512) or higher that extends the OMPT type.

A **tool** can query the OpenMP **thread state** at any time. If a **tool** queries the **thread state** of a **native thread** that is not associated with OpenMP then the implementation reports the state as **ompt_state_undefined**.

The **ompt_state_work_serial** value indicates that the **thread** is executing code outside all **parallel regions**. The **ompt_state_work_parallel** value indicates that the **thread** is executing code within the scope of a **parallel region**. The **ompt_state_work_reduction** value 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 **ompt_state_work_free_agent** value indicates that the **thread** is executing code within the scope of a **task** while not being assigned to the **current team** of that **task**. The **ompt_state_wait_barrier_implicit_parallel** value indicates that the **thread** is waiting at the **implicit barrier** at the end of a **parallel region**. The **ompt_state_wait_barrier_implicit_workshare** value indicates that the **thread** is waiting at an **implicit barrier** at the end of a **worksharing construct**. The **ompt_state_wait_barrier_explicit** value indicates that the **thread** is waiting in an **explicit barrier region**. The **ompt_state_wait_barrier_implementation** value indicates that the **thread** is waiting in a **barrier** that the OpenMP specification does not require but the implementation introduces. The **ompt_state_wait_barrier_teams** value indicates

1 that the `thread` is waiting at a `barrier` at the end of a `teams region`. The value
2 `ompt_state_wait_taskwait` indicates that the `thread` is waiting at a `taskwait` construct.
3 The `ompt_state_wait_taskgroup` value indicates that the `thread` is waiting at the end of a
4 `taskgroup` construct. The `ompt_state_wait_mutex` value indicates that the `thread` is
5 waiting for a mutex of an unspecified type. The `ompt_state_wait_lock` value indicates that
6 the `thread` is waiting for a `lock` or `nestable lock`. The `ompt_state_wait_critical` value
7 indicates that the `thread` is waiting to enter a `critical region`. The
8 `ompt_state_wait_atomic` value indicates that the `thread` is waiting to enter an `atomic`
9 `region`. The `ompt_state_wait_ordered` value indicates that the `thread` is waiting to enter an
10 `ordered region`. The `ompt_state_wait_target` value indicates that the `thread` is waiting
11 for a `target region` to complete. The `ompt_state_wait_target_map` value indicates that
12 the `thread` is waiting for a `mapping operation` to complete. An implementation may report
13 `ompt_state_wait_target` for `target_data` constructs. The
14 `ompt_state_wait_target_update` value indicates that the `thread` is waiting for a
15 `target_update` operation to complete. An implementation may report
16 `ompt_state_wait_target` for `target_update` constructs. The `ompt_state_idle`
17 value indicates that the `native thread` is an `idle thread`, that is, it is an `unassigned thread` that is not a
18 `free-agent thread`. The `ompt_state_overhead` value indicates that the `thread` is in the
19 overhead state at any point while executing within the OpenMP runtime, except while waiting at a
20 synchronization point. The `ompt_state_undefined` value indicates that the `native thread` is
21 not created by the OpenMP implementation.

22 39.32 OMPT subvolume Type

23	Name: <code>subvolume</code> Properties: C/C++-only, OMPT	Base Type: <code>structure</code>
----	--	--

24 Fields

	Name	Type	Properties
	<i>base</i>	<code>c_ptr</code>	C/C++-only, in- tent(in), value
	<i>size</i>	<code>c_uint64_t</code>	value
	<i>num_dims</i>	<code>c_uint64_t</code>	value, positive
25	<i>volume</i>	<code>c_uint64_t</code>	C/C++-only, in- tent(in), pointer
	<i>offsets</i>	<code>c_uint64_t</code>	C/C++-only, in- tent(in), pointer
	<i>dimensions</i>	<code>c_uint64_t</code>	C/C++-only, in- tent(in), pointer

26 Type Definition

C / C++

```
1 typedef struct ompt_subvolume_t {
2     const void *base;
3     uint64_t size;
4     uint64_t num_dims;
5     const uint64_t *volume;
6     const uint64_t *offsets;
7     const uint64_t *dimensions;
8 } ompt_subvolume_t;
```

C / C++

Semantics

The **subvolume** OMPT type represents a rectangular subvolume used in a [rectangular-memory-copying routine](#).

Cross References

- Memory Copying Routines, see [Section 31.7](#)

39.33 OMPT sync_region Type

Name: **sync_region**

Properties: C/C++-only, OMPT, overlapping-type-name

Base Type: [enumeration](#)

Values

Name	Value	Properties
ompt_sync_region_barrier_explicit	3	C/C++-only, OMPT
ompt_sync_region_barrier_implementation	4	C/C++-only, OMPT
ompt_sync_region_taskwait	5	C/C++-only, OMPT
ompt_sync_region_taskgroup	6	C/C++-only, OMPT
ompt_sync_region_reduction	7	C/C++-only, OMPT
ompt_sync_region_barrier_implicit_workshare	8	C/C++-only, OMPT
ompt_sync_region_barrier_implicit_parallel	9	C/C++-only, OMPT
ompt_sync_region_barrier_teams	10	C/C++-only, OMPT

Type Definition

C / C++

```
19 typedef enum ompt_sync_region_t {
20     ompt_sync_region_barrier_explicit          = 3,
21     ompt_sync_region_barrier_implementation    = 4,
22     ompt_sync_region_taskwait                  = 5,
23     ompt_sync_region_taskgroup                  = 6,
```

```

1      ompt_sync_region_reduction          = 7,
2      ompt_sync_region_barrier_implicit_workshare = 8,
3      ompt_sync_region_barrier_implicit_parallel = 9,
4      ompt_sync_region_barrier_teams        = 10
5  } ompt_sync_region_t;

```

C / C++

Semantics

The **sync_region** OMPT type defines the valid synchronization **region** values.

39.34 OMPT target Type

Name: target Properties: C/C++-only, OMPT	Base Type: enumeration
--	-------------------------------

Values

Name	Value	Properties
ompt_target	1	C/C++-only, OMPT
ompt_target_enter_data	2	C/C++-only, OMPT
ompt_target_exit_data	3	C/C++-only, OMPT
ompt_target_update	4	C/C++-only, OMPT
ompt_target_nowait	9	C/C++-only, OMPT
ompt_target_enter_data_nowait	10	C/C++-only, OMPT
ompt_target_exit_data_nowait	11	C/C++-only, OMPT
ompt_target_update_nowait	12	C/C++-only, OMPT

Type Definition

C / C++

```

13  typedef enum ompt_target_t {
14      ompt_target          = 1,
15      ompt_target_enter_data = 2,
16      ompt_target_exit_data = 3,
17      ompt_target_update    = 4,
18      ompt_target_nowait    = 9,
19      ompt_target_enter_data_nowait = 10,
20      ompt_target_exit_data_nowait = 11,
21      ompt_target_update_nowait    = 12
22  } ompt_target_t;

```

C / C++

Semantics

The **target** OMPT type defines valid values to identify **device constructs**.

39.35 OMPT target_data_op Type

Name: <code>target_data_op</code> Properties: C/C++-only, OMPT	Base Type: enumeration
---	--

Values

Name	Value	Properties
<code>ompt_target_data_alloc</code>	1	C/C++-only, OMPT
<code>ompt_target_data_delete</code>	4	C/C++-only, OMPT
<code>ompt_target_data_associate</code>	5	C/C++-only, OMPT
<code>ompt_target_data_disassociate</code>	6	C/C++-only, OMPT
<code>ompt_target_data_transfer</code>	7	C/C++-only, OMPT
<code>ompt_target_data_memset</code>	8	C/C++-only, OMPT
<code>ompt_target_data_transfer_rect</code>	9	C/C++-only, OMPT
<code>ompt_target_data_alloc_async</code>	17	C/C++-only, OMPT
<code>ompt_target_data_delete_async</code>	20	C/C++-only, OMPT
<code>ompt_target_data_transfer_async</code>	23	C/C++-only, OMPT
<code>ompt_target_data_memset_async</code>	24	C/C++-only, OMPT
<code>ompt_target_data_transfer_rect_async</code>	25	C/C++-only, OMPT

Type Definition

C / C++

```
typedef enum ompt_target_data_op_t {  
    ompt_target_data_alloc          = 1,  
    ompt_target_data_delete         = 4,  
    ompt_target_data_associate      = 5,  
    ompt_target_data_disassociate   = 6,  
    ompt_target_data_transfer       = 7,  
    ompt_target_data_memset         = 8,  
    ompt_target_data_transfer_rect  = 9,  
    ompt_target_data_alloc_async    = 17,  
    ompt_target_data_delete_async   = 20,  
    ompt_target_data_transfer_async = 23,  
    ompt_target_data_memset_async   = 24,  
    ompt_target_data_transfer_rect_async = 25  
} ompt_target_data_op_t;
```

C / C++

Additional information

The following instances and associated values of the [target_data_op](#) OMPT type are also defined: `ompt_target_data_transfer_to_device`, with value 2;
`ompt_target_data_transfer_from_device`, with value 3;
`ompt_target_data_transfer_to_device_async`, with value 18; and

`ompt_target_data_transfer_from_device`, with value 19. These instances have been deprecated.

Semantics

The `target_data_op` OMPT type indicates the kind of target data operation for `target_data_op_emi` callbacks, which can be allocate (`ompt_target_data_alloc` and `ompt_target_data_alloc_async`); delete (`ompt_target_data_delete` and `ompt_target_data_delete_async`); associate (`ompt_target_data_associate`); disassociate (`ompt_target_data_disassociate`); transfer (`ompt_target_data_transfer`, `ompt_target_data_transfer_async`, `ompt_target_data_transfer_rect`, and `ompt_target_data_transfer_rect_async`); or memset (`ompt_target_data_memset` and `ompt_target_data_memset_async`), where the values that end with `_async` correspond to asynchronous data operations.

39.36 OMPT target_map_flag Type

Name: <code>target_map_flag</code> Properties: C/C++-only, OMPT	Base Type: enumeration
--	--

Values

Name	Value	Properties
<code>ompt_target_map_flag_to</code>	<code>0x01</code>	C/C++-only, OMPT
<code>ompt_target_map_flag_from</code>	<code>0x02</code>	C/C++-only, OMPT
<code>ompt_target_map_flag_alloc</code>	<code>0x04</code>	C/C++-only, OMPT
<code>ompt_target_map_flag_release</code>	<code>0x08</code>	C/C++-only, OMPT
<code>ompt_target_map_flag_delete</code>	<code>0x10</code>	C/C++-only, OMPT
<code>ompt_target_map_flag_implicit</code>	<code>0x20</code>	C/C++-only, OMPT
<code>ompt_target_map_flag_always</code>	<code>0x40</code>	C/C++-only, OMPT
<code>ompt_target_map_flag_present</code>	<code>0x80</code>	C/C++-only, OMPT
<code>ompt_target_map_flag_close</code>	<code>0x100</code>	C/C++-only, OMPT
<code>ompt_target_map_flag_shared</code>	<code>0x200</code>	C/C++-only, OMPT

Type Definition

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 `target_map_flag` OMPT type defines the valid map flag values. The `ompt_target_map_flag_to`, `ompt_target_map_flag_from`, `ompt_target_map_flag_alloc`, and `ompt_target_map_flag_release` values are set when the mapping operations have the corresponding *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` values are set. The `ompt_target_map_flag_implicit` value is set if the mapping operations correspond to implicitly determined data-mapping attributes. The `ompt_target_map_flag_delete`, `ompt_target_map_flag_always`, `ompt_target_map_flag_present`, and

`ompt_target_map_flag_close`, values are set if the `mapping operations` are specified with the corresponding `map-type-modifier` modifiers. The `ompt_target_map_flag_shared` value is set if the `original storage` and `corresponding storage` are shared for the `mapping operation`.

39.37 OMPT `task_flag` Type

Name: <code>task_flag</code> Properties: C/C++-only, OMPT	Base Type: <code>enumeration</code>
--	--

Values

Name	Value	Properties
<code>ompt_task_initial</code>	<code>0x00000001</code>	C/C++-only, OMPT
<code>ompt_task_implicit</code>	<code>0x00000002</code>	C/C++-only, OMPT
<code>ompt_task_explicit</code>	<code>0x00000004</code>	C/C++-only, OMPT
<code>ompt_task_target</code>	<code>0x00000008</code>	C/C++-only, OMPT
<code>ompt_task_taskwait</code>	<code>0x00000010</code>	C/C++-only, OMPT
<code>ompt_task_importing</code>	<code>0x02000000</code>	C/C++-only, OMPT
<code>ompt_task_exporting</code>	<code>0x04000000</code>	C/C++-only, OMPT
<code>ompt_task_undeffered</code>	<code>0x08000000</code>	C/C++-only, OMPT
<code>ompt_task_untied</code>	<code>0x10000000</code>	C/C++-only, OMPT
<code>ompt_task_final</code>	<code>0x20000000</code>	C/C++-only, OMPT
<code>ompt_task_mergeable</code>	<code>0x40000000</code>	C/C++-only, OMPT
<code>ompt_task_merged</code>	<code>0x80000000</code>	C/C++-only, OMPT

Type Definition

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_importing    = 0x02000000,  
    ompt_task_exporting    = 0x04000000,  
    ompt_task_undeffered   = 0x08000000,  
    ompt_task_untied       = 0x10000000,  
    ompt_task_final        = 0x20000000,  
    ompt_task_mergeable    = 0x40000000,  
    ompt_task_merged       = 0x80000000,  
};
```

```
1      ompt_task_taskwait    = 0x00000010,  
2      ompt_task_importing  = 0x02000000,  
3      ompt_task_exporting  = 0x04000000,  
4      ompt_task_underruned = 0x08000000,  
5      ompt_task_untied     = 0x10000000,  
6      ompt_task_final      = 0x20000000,  
7      ompt_task_mergeable  = 0x40000000,  
8      ompt_task_merged     = 0x80000000  
9  } ompt_task_flag_t;
```

C / C++

Semantics

The `task_flag` OMPT type defines valid `task` values. The least significant byte provides information about the general classification of the `task`. The other bits represent its properties.

39.38 OMPT `task_status` Type

Name: <code>task_status</code> Properties: C/C++-only, OMPT	Base Type: <code>enumeration</code>
--	-------------------------------------

Values

Name	Value	Properties
<code>ompt_task_complete</code>	1	C/C++-only, OMPT
<code>ompt_task_yield</code>	2	C/C++-only, OMPT
<code>ompt_task_cancel</code>	3	C/C++-only, OMPT
<code>ompt_task_detach</code>	4	C/C++-only, OMPT
<code>ompt_task_early_fulfill</code>	5	C/C++-only, OMPT
<code>ompt_task_late_fulfill</code>	6	C/C++-only, OMPT
<code>ompt_task_switch</code>	7	C/C++-only, OMPT
<code>ompt_taskwait_complete</code>	8	C/C++-only, OMPT

Type Definition

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 `task_status` OMPT type indicates the reason that a `task` was switched when it reached a `task scheduling point`. The `ompt_task_complete` value indicates that the `task` that encountered the `task scheduling point` completed execution of its associated `structured block` and any associated `allow-completion event` was fulfilled. The `ompt_task_yield` value indicates that the `task` encountered a `taskyield` construct. The `ompt_task_cancel` value indicates that the `task` was canceled when it encountered an active `cancellation point`. The `ompt_task_detach` value 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 `ompt_task_early_fulfill` value indicates that the `allow-completion event` of the `task` was fulfilled before the `task` completed execution of the associated `structured block`. The `ompt_task_late_fulfill` value indicates that the `allow-completion event` of the `task` was fulfilled after the `task` completed execution of the associated `structured block`. The `ompt_taskwait_complete` value indicates completion of the `dependent task` that results from a `taskwait` construct with one or more `depend` clauses. The `ompt_task_switch` value is used for all other cases that a `task` was switched.

39.39 OMPT thread Type

Name: <code>thread</code> Properties: C/C++-only, OMPT	Base Type: <code>enumeration</code>
---	--

Values

Name	Value	Properties
<code>ompt_thread_initial</code>	1	C/C++-only, OMPT
<code>ompt_thread_worker</code>	2	C/C++-only, OMPT
<code>ompt_thread_other</code>	3	C/C++-only, OMPT
<code>ompt_thread_unknown</code>	4	C/C++-only, OMPT

Type Definition

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

The `thread` OMPT type defines the valid `thread` type values. 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`.



39.40 OMPT wait_id Type

Name: <code>wait_id</code> Properties: C/C++-only, OMPT	Base Type: <code>c_uint64_t</code>
--	---

Predefined Identifiers

Name	Value	Properties
<code>ompt_wait_id_none</code>	0	C/C++-only, OMPT

Type Definition


<code>typedef uint64_t ompt_wait_id_t;</code>


Semantics

The `wait_id` OMPT type describes `wait identifiers` for a `thread`; each `thread` maintains one of these `wait identifiers`. 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 the *name* argument of a critical section, or a `lock`, or 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`.

39.41 OMPT work Type

Name: <code>work</code> Properties: C/C++-only, OMPT, overlapping-type-name	Base Type: <code>enumeration</code>
--	--

Values

Name	Value	Properties
ompt_work_loop	1	C/C++-only, OMPT
ompt_work_sections	2	C/C++-only, OMPT
ompt_work_single_executor	3	C/C++-only, OMPT
ompt_work_single_other	4	C/C++-only, OMPT
ompt_work_workshare	5	C/C++-only, OMPT
ompt_work_distribute	6	C/C++-only, OMPT
ompt_work_taskloop	7	C/C++-only, OMPT
ompt_work_scope	8	C/C++-only, OMPT
ompt_work_workdistribute	9	C/C++-only, OMPT
ompt_work_loop_static	10	C/C++-only, OMPT
ompt_work_loop_dynamic	11	C/C++-only, OMPT
ompt_work_loop_guided	12	C/C++-only, OMPT
ompt_work_loop_other	13	C/C++-only, OMPT

Type Definition

C / C++

```
typedef enum ompt_work_t {  
    ompt_work_loop = 1,  
    ompt_work_sections = 2,  
    ompt_work_single_executor = 3,  
    ompt_work_single_other = 4,  
    ompt_work_workshare = 5,  
    ompt_work_distribute = 6,  
    ompt_work_taskloop = 7,  
    ompt_work_scope = 8,  
    ompt_work_workdistribute = 9,  
    ompt_work_loop_static = 10,  
    ompt_work_loop_dynamic = 11,  
    ompt_work_loop_guided = 12,  
    ompt_work_loop_other = 13  
} ompt_work_t;
```

C / C++

Semantics

The **work OMPT type** defines the valid work values.

40 General Callbacks and Trace Records

This chapter describes general **OMPT callbacks** that an **OMPT tool** may register and that are called during the runtime of an **OpenMP program**. The C/C++ header file (**omp-tools.h**) provides the types that this chapter defines. **Tool** implementations of **callbacks** are not required to be **async signal safe**.

Several **OMPT callbacks** include a *codeptr_ra* argument that relates the implementation of an OpenMP **region** to its source code. If a **routine** implements the **region** associated with a **callback** then *codeptr_ra* contains the return address of the call to that **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**.

Several **OMPT callbacks** have a *flags* argument; the meaning and valid values for that argument is described with the **callback**. Some **callbacks** have an *encountering_task_frame* argument that points to the **frame** object that is associated with the **encountering task**. The behavior for accessing the **frame** object after the **callback** returns is unspecified. Some **callbacks** have a *tool_data* argument that is a pointer to the **tool_data** field in the **start_tool_result** structure that **omp_start_tool** returned. Some **callbacks** have a *parallel_data* argument; the binding of these arguments is the **parallel** or **teams region** that is beginning or ending or the current **parallel region** for **callbacks** that are dispatched during the execution of one. Some **callbacks** have an *encountering_task_data* argument; the binding of these arguments is the **encountering task**. Some **callbacks** have an *endpoint* argument that indicates whether the **callback** signals that a **region** begins or ends. Some **callbacks** have a *wait_id* argument, which indicates the object being awaited.

Several **OMPT callbacks** have a *task_data* argument; unless otherwise specified, the binding of these arguments is the **encountering task** of the **event** for which the implementation dispatches the **callback**. For some of those **callbacks**, OpenMP semantics imply that this **task** to which the *task_data* argument binds is the **implicit task** that executes the **structured block** of the binding **parallel region** or **teams region**.

An implementation may also provide a trace of **events** per **device**. Along with the **callbacks**, this chapter also defines standard **trace records**. For these **trace records**, unless otherwise specified, **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**. If the **target_emi** callback is dispatched, the **target_id** used in any **trace records** associated with the **device region** is given by the **value** field of the *target_data* **data** object that is set in the **callback**.

Restrictions

Restrictions to OpenMP **tool callbacks** are as follows:

- **Tool callbacks** may not use **directives** or call any **routines**.
- **Tool callbacks** must exit by either returning to the caller or aborting.

40.1 Initialization and Finalization Callbacks

This section describes **callbacks** that are called to initialize and to finalize **tools** and when **native threads** are initialized and finalized.

40.1.1 initialize Callback

Name: initialize	Properties: C/C++-only, OMPT
Category: function	

Return Type and Arguments

Name	Type	Properties
<i><return type></i>	integer	<i>default</i>
<i>lookup</i>	function_lookup	OMPT
<i>initial_device_num</i>	integer	<i>default</i>
<i>tool_data</i>	data	OMPT, pointer

Type Signature

C / C++

```
typedef int (*ompt_initialize_t) (ompt_function_lookup_t lookup,  
    int initial_device_num, ompt_data_t *tool_data);
```

C / C++

Semantics

A **tool** provides an **initialize** callback, which has the **initialize** OMPT type, in the non-null pointer to a **start_tool_result** OMPT type structure that its implementation of **ompt_start_tool** returns. An OpenMP implementation must call this OMPT-tool initializer after fully initializing itself but before beginning execution of any **construct** or **routine**. An **initialize** callback returns a non-zero value if it succeeds; otherwise, the OMPT interface state changes to OMPT inactive as described in Section 38.2.3.

The *lookup* argument of an **initialize** callback is a pointer to a runtime entry point that a **tool** must use to obtain pointers to the other entry points in the OMPT interface. The *initial_device_num* argument provides the value that a call to **omp_get_initial_device** would return.

C / C++

A callback of **initialize** OMPT type is a callback of type **ompt_initialize_t**.

C / C++

Cross References

- OMPT `data` Type, see [Section 39.8](#)
- `omp_get_initial_device` Routine, see [Section 30.10](#)
- `ompt_start_tool` Procedure, see [Section 38.2.1](#)
- OMPT `start_tool_result` Type, see [Section 39.30](#)

40.1.2 finalize Callback

Name: <code>finalize</code> Category: subroutine	Properties: C/C++-only , OMPT
---	---

Arguments

Name	Type	Properties
<code>tool_data</code>	data	OMPT , pointer

Type Signature

C / C++

```
typedef void (*ompt_finalize_t) (ompt_data_t *tool_data);
```

C / C++

Semantics

A `tool` provides a `finalize` callback, which has the `finalize` OMPT type, in the non-null pointer to a `start_tool_result` OMPT type structure that its implementation of `ompt_start_tool` returns. An OpenMP implementation must call this OMPT-tool finalizer after the last OMPT event as the OpenMP implementation shuts down.

C / C++

A callback of `finalize` OMPT type is a callback of type `ompt_finalize_t`.

C / C++

Cross References

- OMPT `data` Type, see [Section 39.8](#)
- `ompt_start_tool` Procedure, see [Section 38.2.1](#)
- OMPT `start_tool_result` Type, see [Section 39.30](#)

40.1.3 thread_begin Callback

Name: <code>thread_begin</code> Category: subroutine	Properties: C/C++-only , OMPT
---	---

Arguments

Name	Type	Properties
<i>thread_type</i>	thread	OMPT
<i>thread_data</i>	data	OMPT, pointer, untraced-argument

Type Signature

C / C++

```
typedef void (*ompt_callback_thread_begin_t) (  
    ompt_thread_t thread_type, ompt_data_t *thread_data);
```

Trace Record

C / C++

```
typedef struct ompt_record_thread_begin_t {  
    ompt_thread_t thread_type;  
} ompt_record_thread_begin_t;
```

Semantics

A tool provides a **thread_begin** callback, which has the **thread_begin** OMPT type, that the OpenMP implementation dispatches when **native threads** are created. The *thread_type* argument indicates the type of the new **thread**: initial, worker, other, or unknown. The binding of the *thread_data* argument is the new **thread**.

Cross References

- OMPT **data** Type, see [Section 39.8](#)
- OMPT **thread** Type, see [Section 39.39](#)

40.1.4 thread_end Callback

Name: thread_end	Properties: C/C++-only, OMPT
Category: subroutine	

Arguments

Name	Type	Properties
<i>thread_data</i>	data	OMPT, pointer

Type Signature

C / C++

```
typedef void (*ompt_callback_thread_end_t) (  
    ompt_data_t *thread_data);
```

1 **Semantics**

2 A [tool](#) provides a [thread_end](#) callback, which has the [thread_end](#) OMPT type, that the
3 OpenMP implementation dispatches when [native threads](#) are destroyed. The binding of the
4 *thread_data* argument is the [thread](#) that will be destroyed.

5 **Cross References**

- 6 • OMPT [data](#) Type, see [Section 39.8](#)

7 **40.2 error Callback**

8

Name: error	Properties: C/C++-only , OMPT
Category: subroutine	

9 **Arguments**

10

Name	Type	Properties
<i>severity</i>	severity	OMPT
<i>message</i>	char	intent(in) , pointer
<i>length</i>	size_t	default
<i>codeptr_ra</i>	void	intent(in) , pointer

11 **Type Signature**

12 C / C++
13

```
typedef void (*ompt_callback_error_t) (ompt_severity_t severity,  
const char *message, size_t length, const void *codeptr_ra);
```


14 C / C++

14 **Trace Record**

15 C / C++
16

```
typedef struct ompt_record_error_t {  
17       ompt_severity_t severity;  
18       const char *message;  
19       size_t length;  
20       const void *codeptr_ra;  
} ompt_record_error_t;
```


21 C / C++

21 **Semantics**

22 A [tool](#) provides an [error](#) callback, which has the [error](#) OMPT type, that the OpenMP
23 implementation dispatches when an [error directive](#) is encountered for which the *action-time*
24 argument of the [at clause](#) is specified as [execution](#). The *severity* argument passes the specified
25 severity level. The *message* argument passes the C string from the [message clause](#). The *length*
26 argument provides the length of the C string.

Cross References

- **error** Directive, see [Section 16.1](#)
- OMPT **severity** Type, see [Section 39.29](#)

40.3 Parallelism Generation Callback Signatures

This section describes [callbacks](#) that are related to [constructs](#) for generating and controlling parallelism.

40.3.1 parallel_begin Callback

Name: <code>parallel_begin</code>	Properties: C/C++-only , OMPT
Category: subroutine	

Arguments

Name	Type	Properties
<i>encountering_task_data</i>	data	OMPT , pointer
<i>encountering_task_frame</i>	frame	intent(in) , OMPT , pointer , untraced-argument
<i>parallel_data</i>	data	OMPT , pointer
<i>requested_parallelism</i>	integer	unsigned
<i>flags</i>	integer	default
<i>codeptr_ra</i>	void	intent(in) , pointer

Type Signature

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++

1 **Semantics**

2 A [tool](#) provides a [parallel_begin](#) callback, which has the [parallel_begin](#) OMPT type,
3 that the OpenMP implementation dispatches when a [parallel](#) or [teams](#) region starts. The
4 *requested_parallelism* argument indicates the number of [threads](#) or [teams](#) that the user requested.
5 The *flags* argument indicates whether the code for the [region](#) is inlined into the application or
6 invoked by the runtime and also whether the [region](#) is a [parallel](#) or [teams](#) region. Valid values
7 for *flags* are a disjunction of elements in the [parallel_flag](#) OMPT type.

8 **Cross References**

- 9 • OMPT **data** Type, see [Section 39.8](#)
- 10 • OMPT **frame** Type, see [Section 39.15](#)
- 11 • OMPT **id** Type, see [Section 39.18](#)
- 12 • **parallel** Construct, see [Section 18.1](#)
- 13 • OMPT **parallel_flag** Type, see [Section 39.22](#)
- 14 • **teams** Construct, see [Section 18.2](#)

15 **40.3.2 parallel_end Callback**

16 Name: <code>parallel_end</code>	Properties: C/C++-only, OMPT
Category: subroutine	

17 **Arguments**

Name	Type	Properties
<i>parallel_data</i>	data	OMPT , pointer
<i>encountering_task_data</i>	data	OMPT , pointer
<i>flags</i>	integer	default
<i>codeptr_ra</i>	void	intent(in) , pointer

19 **Type Signature**

20

C / C++

```
typedef void (*ompt_callback_parallel_end_t) (  
21       ompt_data_t *parallel_data, ompt_data_t *encountering_task_data,  
22       int flags, const void *codeptr_ra);
```

C / C++

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;
```

C / C++

Semantics

A tool provides a `parallel_end` callback, which has the `parallel_end` OMPT type, that the OpenMP implementation dispatches when a `parallel` or `teams` region ends. 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 `parallel_flag` OMPT type.

Cross References

- OMPT `data` Type, see [Section 39.8](#)
- OMPT `id` Type, see [Section 39.18](#)
- `parallel` Construct, see [Section 18.1](#)
- OMPT `parallel_flag` Type, see [Section 39.22](#)
- `teams` Construct, see [Section 18.2](#)

40.3.3 masked Callback

Name: <code>masked</code> Category: <code>subroutine</code>	Properties: <code>C/C++-only</code> , <code>OMPT</code>
--	---

Arguments

Name	Type	Properties
<i>endpoint</i>	<code>scope_endpoint</code>	<code>OMPT</code>
<i>parallel_data</i>	<code>data</code>	<code>OMPT</code> , <code>pointer</code>
<i>task_data</i>	<code>data</code>	<code>OMPT</code> , <code>pointer</code>
<i>codeptr_ra</i>	<code>void</code>	<code>intent(in)</code> , <code>pointer</code>

Type Signature

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++

Semantics

A [tool](#) provides a **masked callback**, which has the **masked OMPT type**, that the OpenMP implementation dispatches for **masked regions**.

Cross References

- OMPT **data** Type, see [Section 39.8](#)
- **masked** Construct, see [Section 18.5](#)
- OMPT **id** Type, see [Section 39.18](#)
- OMPT **scope_endpoint** Type, see [Section 39.27](#)

40.4 Work Distribution Callback Signatures

This section describes [callbacks](#) that are related to [work-distribution constructs](#).

40.4.1 work Callback

Name: <code>work</code> Category: subroutine	Properties: C/C++-only , OMPT , overlapping-type-name
---	--

Arguments

Name	Type	Properties
<i>work_type</i>	work	OMPT , overlapping-type-name
<i>endpoint</i>	scope_endpoint	OMPT
<i>parallel_data</i>	data	OMPT , pointer
<i>task_data</i>	data	OMPT , pointer
<i>count</i>	c_uint64_t	default
<i>codeptr_ra</i>	void	intent(in) , pointer

Type Signature

C / C++

```
typedef void (*ompt_callback_work_t) (ompt_work_t work_type,  
    ompt_scope_endpoint_t endpoint, ompt_data_t *parallel_data,  
    ompt_data_t *task_data, uint64_t count, const void *codeptr_ra);
```

C / C++

Trace Record

C / C++

```
typedef struct ompt_record_work_t {  
    ompt_work_t work_type;  
    ompt_scope_endpoint_t endpoint;  
    ompt_id_t parallel_id;  
    ompt_id_t task_id;  
    uint64_t count;  
    const void *codeptr_ra;  
} ompt_record_work_t;
```

C / C++

Semantics

A tool provides a **work** callback, which has the **work** OMPT type, that the OpenMP implementation dispatches for **worksharing regions** and **taskloop regions**. The *work_type* argument indicates the kind of **region**. The *count* argument is a measure of the quantity of work involved in the **construct**. For a **worksharing-loop construct** or **taskloop construct**, *count* represents the number of **collapsed iterations**. For a **sections construct**, *count* represents the number of sections. For a **workshare** or **workdistribute construct**, *count* represents the **units of work**, as defined by the **workshare** or **workdistribute construct**. For a **single** or **scope construct**, *count* is always 1. When the *endpoint* argument signals the end of a **region**, a *count* value of 0 indicates that the actual *count* value is not available.

Cross References

- OMPT **data** Type, see [Section 39.8](#)
- Work-Distribution Constructs, see [Chapter 19](#)
- OMPT **id** Type, see [Section 39.18](#)
- OMPT **scope_endpoint** Type, see [Section 39.27](#)
- **taskloop** Construct, see [Section 20.2](#)
- OMPT **work** Type, see [Section 39.41](#)

40.4.2 dispatch Callback

Name: <code>dispatch</code> Category: <code>subroutine</code>	Properties: <code>C/C++-only</code> , <code>OMPT</code> , <code>overlapping-type-name</code>
--	--

Arguments

Name	Type	Properties
<code>parallel_data</code>	<code>data</code>	<code>OMPT</code> , <code>pointer</code>
<code>task_data</code>	<code>data</code>	<code>OMPT</code> , <code>pointer</code>
<code>kind</code>	<code>dispatch</code>	<code>OMPT</code> , <code>overlapping-type-name</code>
<code>instance</code>	<code>data</code>	<code>OMPT</code>

Type Signature

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_id_t instance;  
} ompt_record_dispatch_t;
```

C / C++

Semantics

A `tool` provides a `dispatch` callback, which has the `dispatch` OMPT type (which has an `overlapping type name` with the `dispatch` OMPT type that applies to the `kind` argument of the `callback`), that the OpenMP implementation dispatches when a `thread` begins to execute a section or a `collapsed iteration`. The `kind` argument indicates whether a `collapsed 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 `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 `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 `dispatch_chunk` that contains the information for the `chunk`.

Cross References

- OMPT **data** Type, see [Section 39.8](#)
- OMPT **dispatch** Type, see [Section 39.13](#)
- OMPT **dispatch_chunk** Type, see [Section 39.14](#)
- Worksharing-Loop Constructs, see [Section 19.6](#)
- OMPT **id** Type, see [Section 39.18](#)
- **sections** Construct, see [Section 19.3](#)
- **taskloop** Construct, see [Section 20.2](#)

40.5 Tasking Callback Signatures

This section describes [callbacks](#) that are related to [tasks](#).

40.5.1 task_create Callback

Name: <code>task_create</code> Category: subroutine	Properties: C/C++-only , OMPT
--	--

Arguments

Name	Type	Properties
<i>encountering_task_data</i>	data	OMPT , pointer
<i>encountering_task_frame</i>	frame	intent(in) , OMPT , pointer , untraced-argument
<i>new_task_data</i>	data	OMPT , pointer
<i>flags</i>	integer	default
<i>has_dependences</i>	integer	default
<i>codeptr_ra</i>	void	intent(in) , pointer

Type Signature

C / C++

```
typedef void (*ompt_callback_task_create_t) (  
    ompt_data_t *encountering_task_data,  
    const ompt_frame_t *encountering_task_frame,  
    ompt_data_t *new_task_data, int flags, int has_dependences,  
    const void *codeptr_ra);
```

C / C++

Trace Record

C / C++

```
typedef struct ompt_record_task_create_t {
    ompt_id_t encountering_task_id;
    ompt_id_t new_task_id;
    int flags;
    int has_dependences;
    const void *codeptr_ra;
} ompt_record_task_create_t;
```

C / C++

Semantics

A [tool](#) provides a **task_create** callback, which has the **task_create** OMPT type, that the OpenMP implementation dispatches when [task regions](#) are generated. The binding of the *new_task_data* argument is the [generated task](#). The *flags* argument indicates the kind of [task](#) ([explicit task](#) or [target task](#)) that is generated. Values for *flags* are a disjunction of elements in the **task_flag** OMPT type. The *has_dependences* argument is *true* if the [generated task](#) has [dependences](#) and *false* otherwise.

Cross References

- OMPT **data** Type, see [Section 39.8](#)
- OMPT **frame** Type, see [Section 39.15](#)
- Initial Task, see [Section 20.13](#)
- OMPT **id** Type, see [Section 39.18](#)
- **task** Construct, see [Section 20.1](#)
- OMPT **task_flag** Type, see [Section 39.37](#)

40.5.2 task_schedule Callback

Name: <code>task_schedule</code>	Properties: C/C++-only, OMPT
Category: subroutine	

Arguments

Name	Type	Properties
<i>prior_task_data</i>	data	OMPT, pointer
<i>prior_task_status</i>	task_status	OMPT
<i>next_task_data</i>	data	OMPT, pointer

1 **Type Signature**

C / C++

```
2       typedef void (*ompt_callback_task_schedule_t) (  
3               ompt_data_t *prior_task_data,  
4               ompt_task_status_t prior_task_status,  
5               ompt_data_t *next_task_data);
```

C / C++

6 **Trace Record**

C / C++

```
7       typedef struct ompt_record_task_schedule_t {  
8               ompt_id_t prior_task_id;  
9               ompt_task_status_t prior_task_status;  
10              ompt_id_t next_task_id;  
11      } ompt_record_task_schedule_t;
```

C / C++

12 **Semantics**

13 A [tool](#) provides a **task_schedule** callback, which has the **task_schedule** OMPT type, that
14 the OpenMP implementation dispatches when **task** scheduling decisions are made. The binding of
15 the *prior_task_data* argument is the **task** that arrived at the **task scheduling point**. This argument
16 can be **NULL** if no **task** was active when the next **task** is scheduled. The *prior_task_status*
17 argument indicates the status of that prior **task**. The binding of the *next_task_data* argument is the
18 **task** that is resumed at the **task scheduling point**. This argument is **NULL** if the **callback** is
19 dispatched for a *task-fulfill event* or if the **callback** signals completion of a **taskwait** construct.
20 This argument can be **NULL** if no **task** was active when the prior **task** was scheduled.

21 **Cross References**

- 22 • OMPT **data** Type, see [Section 39.8](#)
- 23 • Task Scheduling, see [Section 20.14](#)
- 24 • OMPT **id** Type, see [Section 39.18](#)
- 25 • OMPT **task_status** Type, see [Section 39.38](#)

26 **40.5.3 implicit_task Callback**

27 Name: **implicit_task**
 Category: [subroutine](#)

Properties: C/C++-only, OMPT

Arguments

Name	Type	Properties
<i>endpoint</i>	scope_endpoint	OMPT
<i>parallel_data</i>	data	OMPT, pointer
<i>task_data</i>	data	OMPT, pointer
<i>actual_parallelism</i>	integer	unsigned
<i>index</i>	integer	unsigned
<i>flags</i>	integer	default

Type Signature

```
C / C++
typedef void (*ompt_callback_implicit_task_t) (
    ompt_scope_endpoint_t endpoint, ompt_data_t *parallel_data,
    ompt_data_t *task_data, unsigned int actual_parallelism,
    unsigned int index, int flags);
```

Trace Record

```
C / C++
typedef struct ompt_record_implicit_task_t {
    ompt_scope_endpoint_t endpoint;
    ompt_id_t parallel_id;
    ompt_id_t task_id;
    unsigned int actual_parallelism;
    unsigned int index;
    int flags;
} ompt_record_implicit_task_t;
```

Semantics

A tool provides an **implicit_task** callback, which has the **implicit_task** OMPT type, that the OpenMP implementation dispatches when **initial tasks** and **implicit tasks** are generated and completed. The *flags* argument, which has the **task_flag** OMPT type, indicates the kind of **task** (**initial task** or **implicit task**). For the *implicit-task-end* and the *initial-task-end* events, the *parallel_data* argument is **NULL**.

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 region** 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**.

Cross References

- OMPT **data** Type, see [Section 39.8](#)
- OMPT **id** Type, see [Section 39.18](#)
- **parallel** Construct, see [Section 18.1](#)
- OMPT **scope_endpoint** Type, see [Section 39.27](#)
- OMPT **task_flag** Type, see [Section 39.37](#)
- **teams** Construct, see [Section 18.2](#)

40.6 cancel Callback

Name: <code>cancel</code>	Properties: C/C++-only, OMPT
Category: subroutine	

Arguments

Name	Type	Properties
<i>task_data</i>	data	OMPT, pointer
<i>flags</i>	integer	<i>default</i>
<i>codeptr_ra</i>	void	intent(in), pointer

Type Signature

C / C++

```
typedef void (*ompt_callback_cancel_t) (ompt_data_t *task_data,  
    int flags, const void *codeptr_ra);
```

C / C++

Trace Record

C / C++

```
typedef struct ompt_record_cancel_t {  
    ompt_id_t task_id;  
    int flags;  
    const void *codeptr_ra;  
} ompt_record_cancel_t;
```

C / C++

Semantics

A [tool](#) provides a [cancel callback](#), which has the [cancel OMPT type](#), that the OpenMP implementation dispatches when *cancellation*, *cancel* and *discarded-task events* occur. The *flags* argument, which is defined by the [cancel_flag OMPT type](#), indicates whether *cancellation* is activated by the [encountering task](#) or detected as being activated by another [task](#). The [construct](#) that is being canceled is also described in the *flags* argument. When several [constructs](#) are detected as being concurrently canceled, each corresponding bit in the argument will be set.

Cross References

- OMPT `cancel_flag` Type, see [Section 39.7](#)
- OMPT `data` Type, see [Section 39.8](#)
- OMPT `id` Type, see [Section 39.18](#)

40.7 Synchronization Callback Signatures

This section describes [callbacks](#) that are related to [synchronization constructs](#) and [clauses](#).

40.7.1 dependences Callback

Name: <code>dependences</code>	Properties: C/C++-only, OMPT
Category: subroutine	

Arguments

Name	Type	Properties
<i>task_data</i>	data	OMPT, pointer
<i>deps</i>	dependence	intent(in), pointer
<i>ndeps</i>	integer	default

Type Signature

C / C++

```
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_dependences_t {  
    ompt_id_t task_id;  
    ompt_dependence_t dep;  
    int ndeps;  
} ompt_record_dependences_t;
```

C / C++

Semantics

A [tool](#) provides a [dependences](#) callback, which has the [dependences](#) OMPT type, that the OpenMP implementation dispatches when [tasks](#) are generated and when [ordered constructs](#) are encountered. 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 [device construct](#), the [depend object](#) in an asynchronous [routine](#), or the [encountering task](#) for a [doacross](#) clause of the [ordered construct](#). The *deps* argument points to an array of [structures](#) of [dependence](#) OMPT type that

represent [dependences](#) of the [generated task](#) or the *iteration-specifier* of the **doacross** clause. [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.

When the implementation logs [dependences](#) [trace records](#) for a given [event](#), the **ndeps** field determines the number of [trace records](#) that are logged, one for each [dependence](#). The **dep** field in a given [trace record](#) denotes a [structure](#) of [dependence OMPT type](#) that represents the [dependence](#).

Cross References

- OMPT **data** Type, see [Section 39.8](#)
- **depend** Clause, see [Section 23.9.5](#)
- OMPT **dependence** Type, see [Section 39.9](#)
- OMPT **id** Type, see [Section 39.18](#)
- Stand-alone **ordered** Construct, see [Section 23.10.1](#)

40.7.2 task_dependence Callback

Name: <code>task_dependence</code>	Properties: C/C++-only , OMPT
Category: subroutine	

Arguments

Name	Type	Properties
<i>src_task_data</i>	data	OMPT , pointer
<i>sink_task_data</i>	data	OMPT , pointer

Type Signature

C / C++

```
typedef void (*ompt_callback_task_dependence_t) (
    ompt_data_t *src_task_data, ompt_data_t *sink_task_data);
```

C / C++

Trace Record

C / C++

```
typedef struct ompt_record_task_dependence_t {
    ompt_id_t src_task_id;
    ompt_id_t sink_task_id;
} ompt_record_task_dependence_t;
```

C / C++

Semantics

A [tool](#) provides a [task_dependence](#) callback, which has the [task_dependence](#) OMPT type, that the OpenMP implementation dispatches when it encounters an unfulfilled [task dependence](#). The binding of the *src_task_data* argument is an uncompleted [antecedent task](#). The binding of the *sink_task_data* argument is a corresponding [dependent task](#).

Cross References

- OMPT **data** Type, see [Section 39.8](#)
- **depend** Clause, see [Section 23.9.5](#)
- OMPT **id** Type, see [Section 39.18](#)

40.7.3 OMPT sync_region Type

Name: <code>sync_region</code>	Properties: C/C++-only, OMPT, overlapping-type-name
Category: subroutine pointer	

Arguments

Name	Type	Properties
<i>kind</i>	<code>sync_region</code>	OMPT
<i>endpoint</i>	<code>scope_endpoint</code>	OMPT
<i>parallel_data</i>	<code>data</code>	OMPT, pointer
<i>task_data</i>	<code>data</code>	OMPT, pointer
<i>codeptr_ra</i>	<code>void</code>	intent(in), pointer

Type Signature

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++

1 **Semantics**

2 Callbacks that have the **sync_region** OMPT type are synchronizing-region callbacks, which
3 each have the synchronizing-region property. A tool provides these callbacks to mark the beginning
4 and end of regions that have synchronizing semantics. The *kind* argument, which has the
5 **sync_region** OMPT type, indicates the kind of synchronization.

6 **Cross References**

- 7 • OMPT **data** Type, see [Section 39.8](#)
- 8 • OMPT **id** Type, see [Section 39.18](#)
- 9 • OMPT **scope_endpoint** Type, see [Section 39.27](#)
- 10 • OMPT **sync_region** Type, see [Section 39.33](#)

11 **40.7.4 sync_region Callback**

Name: sync_region Category: subroutine	Properties: C/C++-only, common- type-callback, synchronizing-region, OMPT
---	--

13 **Type Signature**

14 **sync_region**

15 **Semantics**

16 A tool provides a **sync_region** callback, which has the **sync_region** OMPT type, that the
17 OpenMP implementation dispatches when barrier regions, **taskwait** regions, and **taskgroup**
18 regions begin and end. For the *implicit-barrier-end* event at the end of a parallel region,
19 *parallel_data* argument is **NULL**.

20 **Cross References**

- 21 • **barrier** Construct, see [Section 23.3.1](#)
- 22 • Implicit Barriers, see [Section 23.3.2](#)
- 23 • OMPT **sync_region** Type, see [Section 40.7.3](#)
- 24 • **taskgroup** Construct, see [Section 23.4](#)
- 25 • **taskwait** Construct, see [Section 23.5](#)

26 **40.7.5 sync_region_wait Callback**

Name: sync_region_wait Category: subroutine	Properties: C/C++-only, common- type-callback, synchronizing-region, OMPT
--	--

1 **Type Signature**

2 [sync_region](#)

3 **Semantics**

4 A [tool](#) provides a [sync_region_wait](#) callback, which has the [sync_region](#) OMPT type,
5 that the OpenMP implementation dispatches when waiting begins and ends for [barrier regions](#),
6 [taskwait regions](#), and [taskgroup regions](#). For the *implicit-barrier-wait-begin* and
7 *implicit-barrier-wait-end* [events](#) at the end of a [parallel region](#), whether *parallel_data* is [NULL](#) or
8 is the current [parallel region](#) is [implementation defined](#).

9 **Cross References**

- 10 • [barrier](#) Construct, see [Section 23.3.1](#)
- 11 • Implicit Barriers, see [Section 23.3.2](#)
- 12 • OMPT [sync_region](#) Type, see [Section 40.7.3](#)
- 13 • [taskgroup](#) Construct, see [Section 23.4](#)
- 14 • [taskwait](#) Construct, see [Section 23.5](#)

15 **40.7.6 reduction Callback**

Name: reduction Category: subroutine	Properties: C/C++-only, common- type-callback, synchronizing-region, OMPT
---	--

17 **Type Signature**

18 [sync_region](#)

19 **Semantics**

20 A [tool](#) provides a [reduction](#) callback, which is a [synchronizing-region](#) callback, that the
21 OpenMP implementation dispatches when it performs reductions.

22 **Cross References**

- 23 • Properties Common to All Reduction Clauses, see [Section 8.6](#)
- 24 • OMPT [sync_region](#) Type, see [Section 40.7.3](#)

25 **40.7.7 OMPT mutex_acquire Type**

Name: mutex_acquire Category: subroutine pointer	Properties: C/C++-only, OMPT
---	-------------------------------------

Arguments

Name	Type	Properties
<i>kind</i>	mutex	OMPT, overlapping-type-name
<i>hint</i>	integer	unsigned
<i>impl</i>	integer	unsigned
<i>wait_id</i>	wait_id	OMPT
<i>codeptr_ra</i>	void	intent(in), pointer

Type Signature

C / C++

```
typedef void (*ompt_callback_mutex_acquire_t) (ompt_mutex_t kind,  
    unsigned int hint, unsigned int impl, ompt_wait_id_t wait_id,  
    const void *codeptr_ra);
```

Trace Record

C / C++

```
typedef struct ompt_record_mutex_acquire_t {  
    ompt_mutex_t kind;  
    unsigned int hint;  
    unsigned int impl;  
    ompt_wait_id_t wait_id;  
    const void *codeptr_ra;  
} ompt_record_mutex_acquire_t;
```

Semantics

Callbacks that have the **mutex_acquire** OMPT type are **mutex-acquiring callbacks**, which each have the **mutex-acquiring** property. A tool provides these **callbacks** to monitor the beginning of **regions** associated with **mutual-exclusion** constructs, **lock-initializing routines** and **lock-acquiring routines**. The *kind* argument, which has the **mutex** OMPT type, 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.

Cross References

- OMPT **mutex** Type, see [Section 39.20](#)
- OMPT **wait_id** Type, see [Section 39.40](#)

40.7.8 mutex_acquire Callback

Name: <code>mutex_acquire</code> Category: subroutine	Properties: C/C++-only , common-type-callback , mutex-acquiring , OMPT
--	---

Type Signature

[mutex_acquire](#)

Semantics

A [tool](#) provides a [mutex_acquire](#) callback, which has the [mutex_acquire](#) OMPT type, that the OpenMP implementation dispatches when [regions](#) associated with [mutual-exclusion constructs](#), [lock-acquiring routines](#) and [lock-testing routines](#) are begun.

Cross References

- OMPT [mutex_acquire](#) Type, see [Section 40.7.7](#)

40.7.9 lock_init Callback

Name: <code>lock_init</code> Category: subroutine	Properties: C/C++-only , common-type-callback , mutex-acquiring , OMPT
--	---

Type Signature

[mutex_acquire](#)

Semantics

A [tool](#) provides a [lock_init](#) callback, which has the [mutex_acquire](#) OMPT type, that the OpenMP implementation dispatches when [lock-initializing routines](#) are executed.

Cross References

- OMPT [mutex_acquire](#) Type, see [Section 40.7.7](#)

40.7.10 OMPT mutex Type

Name: <code>mutex</code> Category: subroutine pointer	Properties: C/C++-only , OMPT , overlapping-type-name
--	--

Arguments

Name	Type	Properties
<i>kind</i>	mutex	OMPT , overlapping-type-name
<i>wait_id</i>	wait_id	OMPT
<i>codeptr_ra</i>	void	intent(in) , pointer

1 **Type Signature**

C / C++

2 **typedef void (*ompt_callback_mutex_t) (ompt_mutex_t kind,**
3 **ompt_wait_id_t wait_id, const void *codeptr_ra);**

C / C++

4 **Trace Record**

C / C++

5 **typedef struct ompt_record_mutex_t {**
6 **ompt_mutex_t kind;**
7 **ompt_wait_id_t wait_id;**
8 **const void *codeptr_ra;**
9 **} ompt_record_mutex_t;**

C / C++

10 **Semantics**

11 Callbacks that have the **mutex** OMPT type are **mutex-execution callbacks**, which each have the **mutex-execution property**. A **tool** provides these **callbacks** to monitor the execution of a **lock-destroying routine** or the beginning or completion of execution of either the **structured block** associated with a **mutual-exclusion construct**, or the **region** guarded by a **lock-acquiring routine** or **lock-testing routine** paired with a **lock-releasing routine**. The *kind* argument, which has the **mutex** OMPT type, indicates the kind of mutual exclusion **event**.

17 **Cross References**

- Lock Acquiring Routines, see [Section 34.3](#)
- Lock Destroying Routines, see [Section 34.2](#)
- Lock Releasing Routines, see [Section 34.4](#)
- Lock Testing Routines, see [Section 34.5](#)
- OMPT **mutex** Type, see [Section 39.20](#)
- OMPT **wait_id** Type, see [Section 39.40](#)

24 **40.7.11 lock_destroy Callback**

25 **Name:** `lock_destroy`
 Category: [subroutine](#)

Properties: C/C++-only, common-type-callback, mutex-execution, OMPT

26 **Type Signature**

27 **mutex**

1 **Semantics**

2 A [tool](#) provides a [lock_destroy](#) callback, which has the [mutex](#) OMPT type, that the OpenMP
3 implementation dispatches when it executes a [lock-destroying routine](#).

4 **Cross References**

- 5 • Lock Destroying Routines, see [Section 34.2](#)
6 • OMPT [mutex](#) Type, see [Section 40.7.10](#)

7 **40.7.12 mutex_acquired Callback**

8

Name: mutex_acquired Category: subroutine	Properties: C/C++-only , common-type-callback , mutex-execution , OMPT
--	---

9 **Type Signature**

10 [mutex](#)

11 **Semantics**

12 A [tool](#) provides a [mutex_acquired](#) callback, which has the [mutex](#) OMPT type, that the
13 OpenMP implementation dispatches when the [structured block](#) associated with a [mutual-exclusion](#)
14 [construct](#) begins execution or when a [region](#) guarded by a [lock-acquiring routine](#) or [lock-testing](#)
15 [routine](#) begins execution.

16 **Cross References**

- 17 • Lock Acquiring Routines, see [Section 34.3](#)
18 • Lock Testing Routines, see [Section 34.5](#)
19 • OMPT [mutex](#) Type, see [Section 40.7.10](#)

20 **40.7.13 mutex_released Callback**

21

Name: mutex_released Category: subroutine	Properties: C/C++-only , common-type-callback , mutex-execution , OMPT
--	---

22 **Type Signature**

23 [mutex](#)

24 **Semantics**

25 A [tool](#) provides a [mutex_released](#) callback, which has the [mutex](#) OMPT type, that the
26 OpenMP implementation dispatches when the [structured block](#) associated with a [mutual-exclusion](#)
27 [construct](#) completes execution or, similarly, when a [region](#) that a [lock-releasing routine](#) guards
28 completes execution.

29 **Cross References**

- 30 • Lock Releasing Routines, see [Section 34.4](#)
31 • OMPT [mutex](#) Type, see [Section 40.7.10](#)

40.7.14 nest_lock Callback

Name: nest_lock Category: subroutine	Properties: C/C++-only , OMPT
--	---

Arguments

Name	Type	Properties
<i>endpoint</i>	scope_endpoint	OMPT
<i>wait_id</i>	wait_id	OMPT
<i>codeptr_ra</i>	void	intent(in) , pointer

Type Signature

C / C++
<pre>typedef void (*ompt_callback_nest_lock_t) (ompt_scope_endpoint_t endpoint, ompt_wait_id_t wait_id, const void *codeptr_ra);</pre>
C / C++

Trace Record

C / C++
<pre>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;</pre>
C / C++

Semantics

A [tool](#) provides a [nest_lock](#) callback, which has the [nest_lock](#) OMPT type, that the OpenMP implementation dispatches when a [thread](#) that owns a [nestable](#) lock invokes a [routine](#) that alters the nesting count of the [lock](#) but does not relinquish its ownership.

Cross References

- OMPT [scope_endpoint](#) Type, see [Section 39.27](#)
- OMPT [wait_id](#) Type, see [Section 39.40](#)

40.7.15 flush Callback

Name: flush Category: subroutine	Properties: C/C++-only , OMPT
--	---

Arguments

Name	Type	Properties
<i>thread_data</i>	data	OMPT, pointer, untraced-argument
<i>codeptr_ra</i>	void	intent(in), pointer

Type Signature

C / C++

```
typedef void (*ompt_callback_flush_t) (ompt_data_t *thread_data,  
    const void *codeptr_ra);
```

Trace Record

C / C++

```
typedef struct ompt_record_flush_t {  
    const void *codeptr_ra;  
} ompt_record_flush_t;
```

Semantics

A [tool](#) provides a **flush** callback, which has the **flush** OMPT type, that the OpenMP implementation dispatches when it encounters a **flush** construct. The binding of the *thread_data* argument is the [encountering thread](#).

Cross References

- OMPT **data** Type, see [Section 39.8](#)
- **flush** Construct, see [Section 23.8.6](#)

40.8 control_tool Callback

Name: <code>control_tool</code> Category: function	Properties: C/C++-only, OMPT
---	-------------------------------------

Return Type and Arguments

Name	Type	Properties
<return type>	integer	default
<i>command</i>	c_uint64_t	default
<i>modifier</i>	c_uint64_t	default
<i>arg</i>	c_ptr	iso_c, value, untraced-argument
<i>codeptr_ra</i>	void	intent(in), pointer

Type Signature

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

A **tool** provides a **control_tool** callback, which has the **control_tool** OMPT type, that the OpenMP implementation uses to dispatch *tool-control events*. This **callback** may return any **non-negative** value, which will be returned to the **OpenMP program** as the return value of the **omp_control_tool** call that triggered the **callback**.

The *command* argument passes a command from an **OpenMP program** to a **tool**. Standard values for *command* are defined by the **control_tool** OpenMP type. The *modifier* argument passes a command modifier from an **OpenMP program** to a **tool**. The *command* and *modifier* arguments may have **tool-defined** 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 **OpenMP program** to exchange arbitrary state. The *arg* argument may be **NULL**.

Restrictions

Restrictions on **control_tool** callbacks are as follows:

- **Tool-defined** values for *command* must be greater than or equal to 64 and less than or equal to 2147483647 (**INT32_MAX**).
- **Tool-defined** values for *modifier* must be **non-negative** and less than or equal to 2147483647 (**INT32_MAX**).

Cross References

- OpenMP **control_tool** Type, see [Section 26.12.1](#)
- **omp_control_tool** Routine, see [Section 37.1](#)

41 Device Callbacks and Tracing

This chapter describes [device-tracing callbacks](#), which have the [device-tracing](#) property. An [OMPT tool](#) may register these [callbacks](#) to monitor and to trace [events](#) that involve [device](#) execution. The C/C++ header file (`omp-tools.h`) also provides the types that this chapter defines.

41.1 `device_initialize` Callback

Name: <code>device_initialize</code> Category: subroutine	Properties: C/C++-only , device-tracing , OMPT
--	---

Arguments

Name	Type	Properties
<i>device_num</i>	integer	default
<i>type</i>	char	intent(in) , pointer
<i>device</i>	device	OMPT , opaque , pointer
<i>lookup</i>	function_lookup	OMPT
<i>documentation</i>	char	intent(in) , pointer

Type Signature

C / C++

```
typedef void (*ompt_callback_device_initialize_t) (  
    int device_num, const char *type, ompt_device_t *device,  
    ompt_function_lookup_t lookup, const char *documentation);
```

C / C++

Semantics

A [tool](#) provides [device_initialize](#) callbacks, which have the [device_initialize](#) [OMPT type](#), that the OpenMP implementation can use to initialize asynchronous collection of traces for [devices](#). The OpenMP implementation dispatches this [callback](#) after OpenMP is initialized for the [device](#) but before execution of any [construct](#) is started on the [device](#).

A [device_initialize](#) [callback](#) must fulfill several duties. First, the *type* argument should be used to determine if any special knowledge about the hardware or software of a [device](#) is employed. Second, the *lookup* argument should be used to look up pointers to [device-tracing entry points](#) for the [device](#). Finally, these [entry points](#) should be used to set up tracing for the [device](#). Initialization of tracing for a [target device](#) is described in [Section 38.2.5](#).

The *device_num* argument indicates the [device number](#) of the [device](#) that is being initialized. The *type* argument is a C string that indicates the type of the [device](#). A [device](#) type string is a semicolon-separated character string that includes, at a minimum, the vendor and model name of the [device](#). These names may be followed by a semicolon-separated sequence of characteristics of the hardware or software of the [device](#).

The *device* argument is a pointer to an [OpenMP object](#) that represents the [target device](#) instance. [Device-tracing entry points](#) use this pointer to identify the [device](#) that is being addressed. The *lookup* argument points to a [function_lookup entry point](#) that a [tool](#) must use to obtain pointers to other [device-tracing entry points](#). If a [device](#) does not support tracing then *lookup* is [NULL](#). The *documentation* argument is a C string that describes how to use these [entry points](#). This documentation string may be a pointer to external documentation, or it may be inline descriptions that include names and type signatures for any [device-specific entry points](#) that are available through the [function_lookup entry point](#) along with descriptions of how to use them to control monitoring and analysis of [device](#) traces.

The *type* and *documentation* arguments are immutable strings that are defined for the lifetime of program execution.

Cross References

- OMPT **device** Type, see [Section 39.11](#)
- **function_lookup** Entry Point, see [Section 42.1](#)

41.2 device_finalize Callback

Name: <code>device_finalize</code>	Properties: C/C++-only , device-tracing , OMPT
Category: subroutine	

Arguments

Name	Type	Properties
<i>device_num</i>	integer	default

Type Signature

C / C++

```
typedef void (*ompt_callback_device_finalize_t) (int device_num);
```

C / C++

Semantics

A [tool](#) provides [device_finalize](#) callbacks, which have the [device_finalize](#) OMPT [type](#), that the OpenMP implementation can use to finalize asynchronous collection of traces for [devices](#). The OpenMP implementation dispatches this [callback](#) immediately prior to finalizing the [device](#) that the *device_num* argument identifies. Prior to dispatching a [device_finalize](#) [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_complete](#) callbacks as needed prior to the dispatch of the [device_finalize](#) callback.

Cross References

- [buffer_complete](#) Callback, see [Section 41.6](#)

41.3 device_load Callback

Name: device_load	Properties: C/C++-only , device-tracing , OMPT
Category: subroutine	

Arguments

Name	Type	Properties
<i>device_num</i>	integer	default
<i>filename</i>	char	intent(in) , pointer
<i>offset_in_file</i>	c_int64_t	iso_c , value
<i>vma_in_file</i>	c_ptr	iso_c , value
<i>bytes</i>	c_size_t	iso_c , value
<i>host_addr</i>	c_ptr	iso_c , value
<i>device_addr</i>	c_ptr	iso_c , value
<i>module_id</i>	c_uint64_t	default

Type Signature

[C / C++](#)

```
typedef void (*ompt_callback_device_load_t) (int device_num,  
      const char *filename, int64_t offset_in_file, void *vma_in_file,  
      size_t bytes, void *host_addr, void *device_addr,  
      uint64_t module_id);
```

[C / C++](#)

Semantics

A [tool](#) provides a [device_load](#) callback, which has the [device_load](#) OMPT type, that the OpenMP implementation can use to indicate that it has just loaded code onto the specified [device](#). The *device_num* argument indicates the [device number](#) of the [device](#) that is being loaded. 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. The *vma_in_file* argument indicates a virtual address in *filename* at which the code can be found. If no virtual address in the file is available then [ompt_addr_none](#) is used. The *bytes* argument indicates the size of the [device](#) code object in bytes.

The *host_addr* argument indicates the address at which a copy of the [device](#) code is available in host [memory](#). The *device_addr* argument indicates the address at which the [device](#) code has been loaded in [device memory](#). Both *host_addr* and *device_addr* will be [ompt_addr_none](#) when no

code address is available for the relevant [device](#). The *module_id* argument is an identifier that is associated with the [device](#) code object.

41.4 device_unload Callback

Name: <code>device_unload</code> Category: subroutine	Properties: C/C++-only , device-tracing , OMPT
--	---

Arguments

Name	Type	Properties
<i>device_num</i>	integer	default
<i>module_id</i>	<code>c_uint64_t</code>	default

Type Signature

C / C++
<pre>typedef void (*ompt_callback_device_unload_t) (int device_num, uint64_t module_id);</pre>
C / C++

Semantics

A [tool](#) provides a [device_unload](#) callback, which has the [device_unload](#) OMPT type, that the OpenMP implementation can use to indicate that it is about to unload code from the specified [device](#). The *device_num* argument indicates the [device number](#) of the [device](#) that is being unloaded. The *module_id* argument is an identifier that is associated with the [device](#) code object.

41.5 buffer_request Callback

Name: <code>buffer_request</code> Category: subroutine	Properties: C/C++-only , device-tracing , OMPT
---	---

Arguments

Name	Type	Properties
<i>device_num</i>	integer	default
<i>buffer</i>	buffer	pointer-to-pointer
<i>bytes</i>	<code>size_t</code>	pointer

Type Signature

C / C++
<pre>typedef void (*ompt_callback_buffer_request_t) (int device_num, ompt_buffer_t **buffer, size_t *bytes);</pre>
C / C++

Semantics

A [tool](#) provides a [buffer_request](#) callback, which has the [buffer_request](#) OMPT type, that the OpenMP implementation dispatches to request a buffer in which to store [trace records](#) for the [device](#) specified by the *device* argument. The [callback](#) sets the location to which the *buffer*

argument points to point to the location of the provided buffer. On entry to the [callback](#), the location to which the *bytes* argument points holds the minimum size of the buffer in bytes that the implementation requests; the implementation must ensure that this size does not exceed the recommended buffer size returned by the [get_buffer_limits](#) entry point for that *device*. A buffer request [callback](#) may set the location to which *bytes* points to 0 if it does not provide a buffer. If a [callback](#) sets that location to a value less than the minimum requested buffer size, further recording of [events](#) for the *device* may be disabled until the next invocation of the [start_trace](#) entry point. This action causes the implementation to drop any [trace records](#) for the *device* until recording is restarted.

Cross References

- OMPT **buffer** Type, see [Section 39.3](#)
- [get_buffer_limits](#) Entry Point, see [Section 43.6](#)

41.6 `buffer_complete` Callback

Name: <code>buffer_complete</code> Category: subroutine	Properties: C/C++-only, device-tracing, OMPT
--	---

Arguments

Name	Type	Properties
<i>device_num</i>	integer	default
<i>buffer</i>	buffer	pointer
<i>bytes</i>	size_t	default
<i>begin</i>	buffer_cursor	OMPT, opaque
<i>buffer_owned</i>	integer	default

Type Signature

<div>C / C++</div> <pre>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);</pre> <div>C / C++</div>

Semantics

A [tool](#) provides a [buffer_complete](#) callback, which has the [buffer_complete](#) OMPT type, that the OpenMP implementation dispatches to indicate that it will not record any more [trace records](#) in the buffer at the location to which the *buffer* argument points. The implementation guarantees that all [trace records](#) in the buffer, which was previously allocated by a [buffer_request](#) callback, are valid. The *device* argument specifies the [device](#) for which the [trace records](#) were gathered. The *bytes* argument indicates the full size of the buffer. The *begin* argument is an [OpenMP object](#) that indicates the position of the beginning of the first [trace record](#)

1 in the buffer. The *buffer_owned* argument is 1 if the data to which *buffer* points can be deleted by
2 the [callback](#) and 0 otherwise. If multiple [devices](#) accumulate [events](#) into a single buffer, this
3 [callback](#) may be invoked with a pointer to one or more [trace records](#) in a shared buffer with
4 *buffer_owned* equal to zero.

5 Typically, a [tool](#) will iterate through the [trace records](#) in the buffer and process them. The OpenMP
6 implementation makes these [callbacks](#) on a [native thread](#) that is not an [OpenMP thread](#) so these
7 [buffer_complete](#) callbacks are not required to be [async signal safe](#).

8 **Restrictions**

9 Restrictions on [buffer_complete](#) callbacks are as follows:

- 10
 - The [callback](#) must not delete the buffer if *buffer_owned* is zero.

11 **Cross References**

- 12
 - OMPT **buffer** Type, see [Section 39.3](#)
 - 13
 - OMPT **buffer_cursor** Type, see [Section 39.4](#)

14 **41.7 target_data_op_emi Callback**

15	Name: <code>target_data_op_emi</code> Category: subroutine	Properties: C/C++-only , device-tracing , OMPT
----	---	---

16 **Arguments**

17

Name	Type	Properties
<i>endpoint</i>	scope_endpoint	OMPT, untraced-argument
<i>target_task_data</i>	data	OMPT, pointer, untraced-argument
<i>target_data</i>	data	OMPT, pointer, untraced-argument
<i>host_op_id</i>	id	OMPT, pointer
<i>optype</i>	target_data_op	OMPT
<i>dev1_addr</i>	c_ptr	iso_c, value
<i>dev1_device_num</i>	integer	default
<i>dev2_addr</i>	c_ptr	iso_c, value
<i>dev2_device_num</i>	integer	default
<i>bytes</i>	size_t	default
<i>codeptr_ra</i>	void	intent(in), pointer

Type Signature

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);
```

C / C++

Trace Record

C / C++

```
typedef struct ompt_record_target_data_op_emi_t {  
    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;  
    ompt_device_time_t end_time;  
    const void *codeptr_ra;  
} ompt_record_target_data_op_emi_t;
```

C / C++

Additional information

The [target_data_op](#) callback may also be used. This [callback](#) has identical arguments to the [target_data_op_emi](#) callback except that the *endpoint* and *target_task_data* arguments are omitted and the *target_data* argument is replaced by the *target_id* argument, which has the **id OMPT type**, and the *host_op_id* argument is not a pointer and is provided by the implementation. If this [callback](#) is registered, it is dispatched for the *target_data_op_end*, *target-data-allocation-end*, *target-data-free-begin*, *target-data-associate*, *target-global-data-op*, and *target-data-disassociate* [events](#). This [callback](#) has been [deprecated](#). In addition to the standard [trace record OMPT type](#) name, the **target_data_op** name may be used to specify a [trace record OMPT type](#) with identical fields. This [OMPT type](#) name has been [deprecated](#).

Semantics

A [tool](#) provides a [target_data_op_emi](#) callback, which has the [target_data_op_emi OMPT type](#), that the OpenMP implementation dispatches when a [device memory](#) is allocated or freed, as well as when data is copied to or from a [device](#).

Note – An OpenMP implementation may aggregate [variables](#) and data operations upon them. For instance, an implementation may synthesize a composite to represent multiple [scalar variables](#) and then allocate, free, or copy this composite as a whole rather than performing data operations on each one individually. Thus, the implementation may not dispatch [callbacks](#) for separate data operations on each [variable](#).

The binding of the *target_task_data* argument is the [target task region](#). The binding of the *target_data* argument is the [device region](#). The *host_op_id* argument points to a [tool](#)-controlled integer value that identifies a data operation for 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 41.1 or NULL if the table indicates none for [device memory routines](#) that solely operate on device memory. For [rectangular-memory-copying routines](#) this argument points to a structure of [subvolume OMPT type](#) that describes a rectangular subvolume of a multi-dimensional array *src*, in the [device data environment](#) of [device](#) *dev1_device_num*. The address *src* of the array is referenced as *base* in the [subvolume OMPT type](#). The *dev1_device_num* argument indicates the [device number](#) on the [device](#) given by Table 41.1. The *dev2_addr* argument indicates the data address on the [device](#) given by Table 41.1. For [rectangular-memory-copying routines](#) this argument points to a structure of [subvolume OMPT type](#) that describes a rectangular subvolume of a multi-dimensional array *dst*, in the [device data environment](#) of [device](#) *dev2_device_num*. The address *dst* of the array is referenced as *base* in the [subvolume OMPT type](#). The *dev2_device_num* argument indicates the [device number](#) on the [device](#) given by Table 41.1. 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 the data in bytes.

If [set_trace_ompt](#) has configured the implementation to trace data operations to [device memory](#) then the implementation will log a [target_data_op_emi](#) trace record in a trace. The fields in the record are as follows:

- The [host_op_id](#) field contains an identifier of a data operation for a [target device](#); if the corresponding [target_data_op_emi](#) callback was dispatched, this identifier is the [tool](#)-controlled integer value to which the *host_op_id* argument of the [callback](#) points so that a [tool](#) may correlate the [trace record](#) with the [callback](#), and otherwise the [host_op_id](#) field contains an implementation-controlled identifier;
- The [optype](#), [dev1_addr](#), [dev1_device_num](#), [dev2_addr](#), [dev2_device_num](#), [bytes](#), and [codeptr_ra](#) fields contain the same values as the [callback](#);
- The time when the data operation began execution for the [device](#) is recorded in the [time](#) field of an enclosing [trace record](#) of [record_ompt](#) OMPT type; and
- The time when the data operation completed execution for the [device](#) is recorded in the [end_time](#) field.

TABLE 41.1: Association of dev1 and dev2 arguments for target data operations

Data op	dev1	dev2
allocate	host/none	device
transfer	<i>from</i> device	<i>to</i> device
delete	host/none	device
associate	host	device
disassociate	host	device
memset	none	device

Restrictions

Restrictions to `target_data_op_emi` callbacks are as follows:

- The deprecated `target_data_op` callback must not be registered if a `target_data_op_emi` callbacks is registered.

Cross References

- OMPT `data` Type, see [Section 39.8](#)
- OMPT `device_time` Type, see [Section 39.12](#)
- OMPT `id` Type, see [Section 39.18](#)
- `map` Clause, see [Section 11.3](#)
- OMPT `scope_endpoint` Type, see [Section 39.27](#)
- OMPT `target_data_op` Type, see [Section 39.35](#)

41.8 target_emi Callback

Name: <code>target_emi</code> Category: subroutine	Properties: C/C++-only , device-tracing , OMPT
---	---

Arguments

Name	Type	Properties
<i>kind</i>	target	OMPT
<i>endpoint</i>	scope_endpoint	OMPT
<i>device_num</i>	integer	default
<i>task_data</i>	data	OMPT , pointer
<i>target_task_data</i>	data	OMPT , pointer , untraced-argument
<i>target_data</i>	data	OMPT , pointer
<i>codeptr_ra</i>	void	intent(in) , pointer

Type Signature

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);
```

C / C++

Trace Record

C / C++

```
typedef struct ompt_record_target_emi_t {  
ompt_target_t kind;  
ompt_scope_endpoint_t endpoint;  
int device_num;  
ompt_id_t task_id;  
ompt_id_t target_id;  
const void *codeptr_ra;  
} ompt_record_target_emi_t;
```

C / C++

Additional information

The **target** callback may also be used. This callback has identical arguments to the **target_emi** callback except that the *target_task_data* argument is omitted and the *target_data* argument is replaced by the *target_id* argument, which has the **id OMPT type**. If this callback is registered, it is dispatched for the *target-begin*, *target-end*, *target-enter-data-begin*, *target-enter-data-end*, *target-exit-data-begin*, *target-exit-data-end*, *target-update-begin*, and *target-update-end* events. This callback has been deprecated. In addition to the standard **trace record OMPT type** name, the **target** name may be used to specify a **trace record OMPT type** with identical fields. This OMPT type name has been deprecated.

Semantics

A tool provides a **target_emi** callback, which has the **target_emi** OMPT type, that the OpenMP implementation dispatches when a thread begins to execute a device construct. The *kind* argument indicates the kind of device region. The *device_num* argument specifies the device number of the target device associated with the region. The binding of the *task_data* argument is the encountering task. The binding of the *target_task_data* argument is the target task. If a device region does not have a target task or if the target task is a merged task, this argument is NULL. The binding of the *target_data* argument is the device region.

Restrictions

Restrictions to **target_emi** callbacks are as follows:

- The deprecated **target** callback must not be registered if a **target_emi** callback is registered.

Cross References

- OMPT `data` Type, see [Section 39.8](#)
- OMPT `id` Type, see [Section 39.18](#)
- OMPT `scope_endpoint` Type, see [Section 39.27](#)
- `target` Construct, see [Section 21.8](#)
- OMPT `target` Type, see [Section 39.34](#)
- `target_data` Construct, see [Section 21.7](#)
- `target_enter_data` Construct, see [Section 21.5](#)
- `target_exit_data` Construct, see [Section 21.6](#)
- `target_update` Construct, see [Section 21.9](#)

41.9 target_map_emi Callback

Name: <code>target_map_emi</code> Category: subroutine	Properties: C/C++-only , device-tracing , OMPT
---	---

Arguments

Name	Type	Properties
<i>target_data</i>	<code>data</code>	OMPT , pointer
<i>nitems</i>	<code>integer</code>	unsigned
<i>host_addr</i>	<code>void</code>	pointer-to-pointer
<i>device_addr</i>	<code>void</code>	pointer-to-pointer
<i>bytes</i>	<code>size_t</code>	pointer
<i>mapping_flags</i>	<code>integer</code>	unsigned , pointer
<i>codeptr_ra</i>	<code>void</code>	intent(in) , pointer

Type Signature

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);
```

C / C++

Trace Record

C / C++

```
typedef struct ompt_record_target_map_emi_t {
    ompt_id_t target_id;
    unsigned int nitems;
    void **host_addr;
    void **device_addr;
    size_t *bytes;
    unsigned int *mapping_flags;
    const void *codeptr_ra;
} ompt_record_target_map_emi_t;
```

C / C++

Additional information

The **target_map** callback may also be used. This callback has identical arguments to the **target_map_emi** callback except that the *target_data* argument is replaced by the *target_id* argument, which has the **id OMPT type**. If this callback is registered, it is dispatched for any *target-map* events. This callback has been deprecated. In addition to the standard **trace record OMPT type** name, the **target_map** name may be used to specify a **trace record OMPT type** with identical fields. This **OMPT type** name has been deprecated.

Semantics

A tool provides a **target_map_emi** callback, which has the **target_map_emi OMPT type**, that the OpenMP implementation dispatches to indicate data mapping relationships. The implementation may report mappings associated with multiple **map clauses** that appear on the same **construct** with a single callback to report the effect of all mappings or multiple callbacks with each reporting a subset of the mappings. Further, the implementation may omit mappings that it determines are unnecessary. If the implementation issues multiple **target_map_emi** callbacks, these callbacks may be interleaved with **target_data_op_emi** callbacks that report data operations associated with the mappings.

The binding of the *target_data* argument is the **device region**. The *nitems* argument indicates the number of data mappings that the callback reports. The *host_addr* argument indicates an array of host addresses. The *device_addr* argument indicates an array of device addresses. The *bytes* argument indicates an array of sizes of data. The *mapping_flags* argument indicates the kind of mapping operations, which may result from explicit **map clauses** or the implicit data-mapping rules (see [Section 11.1](#)). Flags for the mapping operations include one or more values specified by the **target_map_flag** type.

Restrictions

Restrictions to **target_map_emi** callbacks are as follows:

- The deprecated **target_map** callback must not be registered if a **target_map_emi** callback is registered.

Cross References

- OMPT `data` Type, see [Section 39.8](#)
- OMPT `id` Type, see [Section 39.18](#)
- `map` Clause, see [Section 11.3](#)
- `target_data_op_emi` Callback, see [Section 41.7](#)
- OMPT `target_map_flag` Type, see [Section 39.36](#)

41.10 `target_submit_emi` Callback

Name: <code>target_submit_emi</code> Category: subroutine	Properties: C/C++-only , device-tracing , OMPT
--	---

Arguments

Name	Type	Properties
<i>endpoint</i>	<code>scope_endpoint</code>	OMPT , untraced-argument
<i>target_data</i>	<code>data</code>	OMPT , pointer , untraced-argument
<i>host_op_id</i>	<code>id</code>	OMPT , pointer
<i>requested_num_teams</i>	<code>integer</code>	unsigned

Type Signature

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);
```

C / C++

Trace Record

C / C++

```
typedef struct ompt_record_target_submit_emi_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_submit_emi_t;
```

C / C++

Additional information

The `target_submit` callback may also be used. This callback has identical arguments to the `target_submit_emi` callback except that the `endpoint` argument is omitted and the `target_data` argument is replaced by the `target_id` argument, which has the `id` OMPT type, and the `host_op_id` argument is not a pointer and is provided by the implementation. If this callback is registered, it is dispatched for any `target-submit-begin` events. This callback has been deprecated. In addition to the standard `trace record` OMPT type name, the `target_kernel` name may be used to specify a `trace record` OMPT type with identical fields. This OMPT type name has been deprecated.

Semantics

A tool provides a `target_submit_emi` callback, which has the `target_submit_emi` OMPT type, that the OpenMP implementation dispatches before and after a `target` task initiates creation of an `initial task` on a `device`. The binding of the `target_data` argument is the `device region`. The `host_op_id` argument points to a tool-controlled integer value that identifies an `initial task` on a `target device`. The `requested_num_teams` argument is the number of `teams` that the `device construct` requested to execute the `region`. The actual number of `teams` that execute the `region` may be smaller and generally will not be known until the `region` begins to execute on the `device`.

If `set_trace_ompt` has configured the implementation to trace `device region` execution for a `device` then the implementation will log a `target_submit_emi` trace record. The fields in the record are as follows:

- The `host_op_id` field contains an identifier that identifies the `initial task` on the `device`; if the corresponding `target_submit_emi` callback was dispatched, this identifier is the tool-controlled integer value to which the `host_op_id` argument of the callback points so that a tool may correlate the `trace record` with the callback, and otherwise the `host_op_id` field contains an implementation-controlled identifier;
- The `requested_num_teams` field contains the number of `teams` that the `device construct` requested to execute the `device region`;
- The `granted_num_teams` field contains the number of `teams` that the `device` actually used to execute the `device region`;
- The time when the `initial task` began execution on the `device` is recorded in the `time` field of an enclosing `trace record` of `record_ompt` OMPT type; and
- The time when the `initial task` completed execution on the `device` is recorded in the `end_time` field.

Restrictions

Restrictions to `target_submit_emi` callbacks are as follows:

- The deprecated `target_submit` callback must not be registered if a `target_submit_emi` callback is registered.

Cross References

- OMPT **data** Type, see [Section 39.8](#)
- OMPT **device_time** Type, see [Section 39.12](#)
- OMPT **id** Type, see [Section 39.18](#)
- OMPT **scope_endpoint** Type, see [Section 39.27](#)
- **target** Construct, see [Section 21.8](#)

42 General Entry Points

OMPT supports two principal sets of runtime entry points for tools. For both sets, entry points should not be global symbols since tools cannot rely on the visibility of such symbols. This chapter defines the first set, which enables a tool to register callbacks for events and to inspect the state of threads while executing in a callback or a signal handler. The `omp-tools.h` C/C++ header file provides the definitions of the types that are specified throughout this chapter.

OMPT also supports entry points for two classes of lookup entry points. The first class of lookup entry points contains a single member that is provided through the `initialize` callback: a `function_lookup` entry point that returns pointers to the set of entry points that are defined in this chapter. The second class of lookup entry points includes a unique lookup entry point for each kind of device that can return pointers to entry points in a device's OMPT tracing interface.

The binding thread set for each OMPT entry point is the encountering thread unless otherwise specified. The binding task set is the task executing on the encountering thread.

Several entry points are `async-signal-safe` entry points, which means they each have the `async-signal-safe` property, which implies that they are `async signal safe`.

Restrictions

Restrictions on OMPT runtime entry points are as follows:

- 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.
- Device entry points must not be called after a `device-finalize` event for that device.

42.1 function_lookup Entry Point

Name: <code>function_lookup</code> Category: <code>function</code>	Properties: <code>C/C++-only</code> , <code>OMPT</code>
---	--

Return Type and Arguments

Name	Type	Properties
<code><return type></code>	<code>interface_fn</code>	<code>default</code>
<code>interface_function_name</code>	<code>char</code>	<code>intent(in)</code> , <code>pointer</code>

1 **Type Signature**

 C / C++

2 **typedef omp_t_interface_fn_t** (*omp_t_function_lookup_t) (
3 **const char** *interface_function_name);

 C / C++

4 **Semantics**

5 The **function_lookup** entry point, which has the **function_lookup** OMPT type, enables
6 tools to look up pointers to **OMPT entry points** by name. When an OpenMP implementation
7 invokes the **initialize** callback to configure the **OMPT callback** interface, it provides an **entry**
8 **point** that provides pointers to other **entry points** that implement **routines** that are part of the **OMPT**
9 **callback** interface. Alternatively, when it invokes a **device_initialize** callback to configure
10 the **OMPT** tracing interface for a **device**, it provides an **entry point** that provides pointers to **entry**
11 **points** that implement tracing control **routines** appropriate for that **device**.

12 For these **entry points**, the *interface_function_name* argument is a C string that represents the name
13 of the **entry point** to look up. If the name is unknown to the implementation, the **entry point** returns
14 **NULL**. In a **compliant implementation**, the **entry point** that is provided by the **initialize**
15 **callback** returns a valid function pointer for any **entry point** name listed in Table 38.1. Similarly, in
16 a **compliant implementation**, the **entry point** that is provided by the **device_initialize**
17 **callback** returns non-**NULL** function pointers for any **entry point** name listed in Table 38.3, except
18 for **set_trace_omp** and **get_record_omp**, as described in Section 38.2.5.

19 **Cross References**

- 20 • **device_initialize** Callback, see Section 41.1
- 21 • Binding Entry Points, see Section 38.2.3.1
- 22 • Tracing Activity on Target Devices, see Section 38.2.5
- 23 • **initialize** Callback, see Section 40.1.1
- 24 • **OMPT interface_fn** Type, see Section 39.19

25 **42.2 enumerate_states Entry Point**

26 Name: enumerate_states	Properties: C/C++-only, OMPT
Category: function	

1 **Return Type and Arguments**

Name	Type	Properties
<i><return type></i>	integer	<i>default</i>
<i>current_state</i>	integer	<i>default</i>
<i>next_state</i>	integer	pointer
<i>next_state_name</i>	const char	intent(out), pointer-to-pointer

3 **Type Signature**

C / C++

```
typedef int (*ompt_enumerate_states_t) (int current_state,  
int *next_state, const char **next_state_name);
```

C / C++

6 **Semantics**

7 An OpenMP implementation may support only a subset of the [thread states](#) that the **state OMPT**
8 [type](#) defines. An OpenMP implementation may also support [implementation defined](#) states. The
9 **enumerate_states** entry point, which has the **enumerate_states** OMPT type, is the
10 entry point that enables a [tool](#) to enumerate the supported [thread states](#).

11 When a supported [thread state](#) is passed as *current_state*, the [entry point](#) assigns the next [thread](#)
12 [state](#) in the enumeration to the [variable](#) passed by reference in *next_state* and assigns the name
13 associated with that state to the character pointer passed by reference in *next_state_name*; the
14 returned string is immutable and defined for the lifetime of program execution. Whenever one or
15 more states are left in the enumeration, the **enumerate_states** entry point returns 1. When the
16 last state in the enumeration is passed as *current_state*, **enumerate_states** returns 0, which
17 indicates that the enumeration is complete.

18 To begin enumerating the supported states, a [tool](#) should pass **ompt_state_undefined** as
19 *current_state*. Subsequent invocations of **enumerate_states** should pass the value assigned to
20 the variable that was passed by reference in *next_state* to the previous call. The
21 **ompt_state_undefined** value is returned to indicate an invalid [thread state](#).

22 **Cross References**

- OMPT **state** Type, see [Section 39.31](#)

24 **42.3 enumerate_mutex_impls Entry Point**

Name: <code>enumerate_mutex_impls</code> Category: function	Properties: C/C++-only , OMPT
--	--

1 **Return Type and Arguments**

Name	Type	Properties
<i><return type></i>	integer	<i>default</i>
<i>current_impl</i>	integer	<i>default</i>
<i>next_impl</i>	integer	pointer
<i>next_impl_name</i>	const char	intent(out), pointer-to-pointer

3 **Type Signature**

4 **C / C++**
5 **typedef int (*ompt_enumerate_mutex_impls_t) (int current_impl,**
6 **int *next_impl, const char **next_impl_name);**
7 **C / C++**

6 **Semantics**

7 Mutual exclusion for **locks**, **critical regions**, and **atomic regions** may be implemented in
8 several ways. The **enumerate_mutex_impls** entry point, which has the
9 **enumerate_mutex_impls** OMPT type, enables a **tool** to enumerate the supported mutual
10 exclusion implementations.

11 When a supported mutex implementation is passed as *current_impl*, the **entry point** assigns the next
12 mutex implementation in the **enumeration** to the **variable** passed by reference in *next_impl* and
13 assigns the name associated with that mutex implementation to the character pointer passed by
14 reference in *next_impl_name*; the returned string is immutable and defined for the lifetime of
15 program execution. Whenever one or more mutex implementations are left in the **enumeration**, the
16 **enumerate_mutex_impls** entry point returns 1. When the last mutex implementation in the
17 **enumeration** is passed as *current_impl*, the **entry point** returns 0, which indicates that the
18 **enumeration** is complete.

19 To begin enumerating the supported mutex implementations, a **tool** should pass
20 **ompt_mutex_impl_none** as *current_impl*. Subsequent invocations of
21 **enumerate_mutex_impls** should pass the value assigned to the variable that was passed by
22 reference in *next_impl* to the previous call. The value **ompt_mutex_impl_none** is returned to
23 indicate an invalid mutex implementation.

24 **42.4 set_callback Entry Point**

Name: set_callback Category: function	Properties: C/C++-only, OMPT
--	-------------------------------------

1 **Return Type and Arguments**

Name	Type	Properties
<return type>	set_result	default
event	callbacks	OMPT
callback	callback	OMPT

3 **Type Signature**

C / C++

```
typedef ompt_set_result_t (*ompt_set_callback_t) (  
    ompt_callbacks_t event, ompt_callback_t callback);
```

C / C++

6 **Semantics**

7 OpenMP implementations can use [callbacks](#) to indicate the occurrence of [events](#) during the
8 execution of an [OpenMP program](#). The [set_callback](#) entry point, which has the
9 [set_callback](#) OMPT type, enables a [tool](#) to register the [callback](#) indicated by the *callback*
10 argument for the [event](#) indicated by the *event* argument on the [current device](#). The return value of
11 [set_callback](#) indicates the outcome of registering the [callback](#) and may be any value in the
12 [set_result](#) OMPT type except [ompt_set_impossible](#). If *callback* is [NULL](#) then
13 [callbacks](#) associated with *event* are disabled. If [callbacks](#) are successfully disabled then
14 [ompt_set_always](#) is returned.

15 **Restrictions**

16 Restrictions on the [set_callback](#) entry point are as follows:

- The type signature for *callback* must match the type signature appropriate for the [event](#).

18 **Cross References**

- OMPT [callback](#) Type, see [Section 39.5](#)
- OMPT [callbacks](#) Type, see [Section 39.6](#)
- Monitoring Activity on the Host with OMPT, see [Section 38.2.4](#)
- OMPT [set_result](#) Type, see [Section 39.28](#)


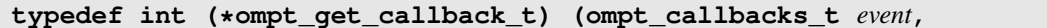


23 **42.5 get_callback Entry Point**

Name: get_callback Category: function	Properties: C/C++-only , OMPT
--	---

1 **Return Type and Arguments**

Name	Type	Properties
<return type>	integer	default
event	callbacks	OMPT
callback	callback	OMPT, pointer

3 **Type Signature**

4  C / C++ 
5 **typedef int (*ompt_get_callback_t) (ompt_callbacks_t event,**
6 **ompt_callback_t *callback);**
7  C / C++ 

6 **Semantics**

7 The **get_callback** entry point, which has the **get_callback** OMPT type, enables a tool to
8 retrieve a pointer to a registered **callback** (if any) that an OpenMP implementation invokes when a
9 host **event** occurs. If the **callback** that is registered for the **event** that is specified by the **event**
10 argument is not **NULL**, the pointer to the **callback** is assigned to the **variable** passed by reference in
11 **callback** and **get_callback** returns 1; otherwise, it returns 0. If **get_callback** returns 0, the
12 value of the **variable** passed by reference as **callback** is **undefined**.

13 **Restrictions**

14 Restrictions on the **get_callback** entry point are as follows:

- 15
 - The **callback** argument must not be **NULL** and must point to valid storage.

16 **Cross References**

- 17
 - OMPT **callback** Type, see [Section 39.5](#)
 - OMPT **callbacks** Type, see [Section 39.6](#)
 - **set_callback** Entry Point, see [Section 42.4](#)





20 **42.6 get_thread_data Entry Point**

Name: get_thread_data	Properties: async-signal-safe , C/C++-only , OMPT
Category: function	

22 **Return Type**

Name	Type	Properties
<return type>	data	pointer

24 **Type Signature**

25  C / C++ 
typedef ompt_data_t *(*ompt_get_thread_data_t) (void);
26  C / C++ 

Semantics

Each `thread` can have an associated `thread` data object of `data OMPT` type. The `get_thread_data` entry point, which has the `get_thread_data OMPT` type, enables a `tool` to retrieve a pointer to the `thread` data object, if any, that is associated with the `encountering thread`. A `tool` may use a pointer to a `thread`'s data object that `get_thread_data` retrieves to inspect or to modify the value of the data object. When a `thread` is created, its data object is initialized with the value `ompt_data_none`.

Cross References

- OMPT `data` Type, see [Section 39.8](#)

42.7 `get_num_procs` Entry Point

Name: <code>get_num_procs</code> Category: function	Properties: all-device-threads-binding , async-signal-safe , C/C++-only , OMPT
--	---

Return Type

Name	Type	Properties
<i><return type></i>	integer	<i>default</i>

Type Signature

C / C++

```
typedef int (*ompt_get_num_procs_t) (void);
```

C / C++

Semantics

The `get_num_procs` entry point, which has the `get_num_procs OMPT` type, enables a `tool` to retrieve the number of `processors` that are available on the `host device` at the time the `entry point` 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. The `binding thread set` of this `entry point` is `all threads` on the `host device`.

42.8 `get_num_places` Entry Point

Name: <code>get_num_places</code> Category: function	Properties: all-device-threads-binding , async-signal-safe , C/C++-only , OMPT
---	---

Return Type

Name	Type	Properties
<i><return type></i>	integer	<i>default</i>

1 **Type Signature**

C / C++

```
2       typedef int (*ompt_get_num_places_t) (void);
```

C / C++

3 **Semantics**

4 The `get_num_places` entry point, which has the `get_num_places` OMPT type, enables a
5 [tool](#) to retrieve the number of `places` in the `place list`. This value is equal to the number of `places` in
6 the `place-partition-var` ICV in the execution environment of the `initial task`. The `binding thread set`
7 of this `entry point` is `all threads` on the `host device`.

8 **Cross References**

- 9
 - `OMP_PLACES`, see [Section 4.1.6](#)
 - `place-partition-var` ICV, see [Table 3.1](#)

11 **42.9 get_place_proc_ids Entry Point**

Name: <code>get_place_proc_ids</code> Category: function	Properties: all-device-threads-binding , C/C++-only , OMPT
---	--

13 **Return Type and Arguments**

Name	Type	Properties
<code><return type></code>	integer	default
<code>place_num</code>	integer	default
<code>ids_size</code>	integer	default
<code>ids</code>	integer	pointer

15 **Type Signature**

C / C++

```
16      typedef int (*ompt_get_place_proc_ids_t) (int place_num,  
17      int ids_size, int *ids);
```

C / C++

18 **Semantics**

19 The `get_place_proc_ids` entry point, which has the `get_place_proc_ids` OMPT type,
20 enables a [tool](#) to retrieve the numerical identifiers of each `processor` that is associated with the `place`
21 specified by the `place_num` argument. The `ids` argument is an array in which the `entry point` can
22 return a vector of `processor` identifiers in the specified `place`; these identifiers are `non-negative`, and
23 their meaning is `implementation defined`. The `ids_size` argument indicates the size of the result
24 array that is specified by `ids`. The `binding thread set` of this `entry point` is `all threads` on the `device`.

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 [entry point](#) returns are [undefined](#). The [entry point](#) always returns the number of numerical identifiers of the [processors](#) that are available to the execution environment in the specified [place](#).

42.10 `get_place_num` Entry Point

Name: <code>get_place_num</code> Category: function	Properties: async-signal-safe , C/C++-only , OMPT
--	--

Return Type

Name	Type	Properties
<code><return type></code>	integer	default

Type Signature

C / C++
<code>typedef int (*ompt_get_place_num_t) (void);</code>
C / C++

Semantics

When the [encountering thread](#) is bound to a [place](#), the `get_place_num` entry point, which has the `get_place_num` OMPT type, enables a [tool](#) to retrieve the [place number](#) associated with the [thread](#). The returned value is between zero and one less than the value returned by `get_num_places`, inclusive. When the [encountering thread](#) is not bound to a [place](#), the [entry point](#) returns `-1`.

42.11 `get_partition_place_nums` Entry Point

Name: <code>get_partition_place_nums</code> Category: function	Properties: async-signal-safe , C/C++-only , OMPT
---	--

Return Type and Arguments

Name	Type	Properties
<code><return type></code>	integer	default
<code>place_nums_size</code>	integer	default
<code>place_nums</code>	integer	pointer

Type Signature

C / C++
<code>typedef int (*ompt_get_partition_place_nums_t) (int place_nums_size, int *place_nums);</code>
C / C++

1 **Semantics**

2 The `get_partition_place_nums` entry point, which has the
3 [get_partition_place_nums](#) OMPT type, enables a [tool](#) to retrieve a list of [place numbers](#)
4 that correspond to the [places](#) in the [place-partition-var](#) ICV of the innermost [implicit task](#). The
5 [place_nums](#) argument is an array in which the [entry point](#) can return a vector of [place](#) identifiers.
6 The [place_nums_size](#) argument indicates the size of that array.

7 If the [place_nums](#) array of size [place_nums_size](#) is large enough to contain all identifiers then they
8 are returned in [place_nums](#) and their order in the array is [implementation defined](#). Otherwise, if the
9 [place_nums](#) array is too small, the values in [place_nums](#) when the [entry point](#) returns are
10 [undefined](#). The [entry point](#) always returns the number of [places](#) in the [place-partition-var](#) ICV of
11 the innermost [implicit task](#).

12 **Cross References**

- 13 • [OMP_PLACES](#), see [Section 4.1.6](#)
14 • [place-partition-var](#) ICV, see [Table 3.1](#)

15



42.12 `get_proc_id` Entry Point

Name: <code>get_proc_id</code> Category: function	Properties: async-signal-safe , C/C++-only , OMPT
--	--

17 **Return Type**

Name	Type	Properties
<code><return type></code>	integer	default

19 **Type Signature**


<code>typedef int (*ompt_get_proc_id_t) (void);</code>


21 The [get_proc_id](#) entry point, which has the [get_proc_id](#) OMPT type, enables a [tool](#) to
22 retrieve the numerical identifier of the [processor](#) of the [encountering thread](#). A defined numerical
23 identifier is [non-negative](#), and its meaning is [implementation defined](#). A negative number indicates
24 a failure to retrieve the numerical identifier.

25

42.13 `get_state` Entry Point

Name: <code>get_state</code> Category: function	Properties: async-signal-safe , C/C++-only , OMPT
--	--

1 **Return Type and Arguments**

Name	Type	Properties
<return type>	integer	default
wait_id	wait_id	OMPT, pointer

3 **Type Signature**

4 **typedef int (*ompt_get_state_t) (ompt_wait_id_t *wait_id);**

5 **Semantics**

6 Each [thread](#) has an associated state and a [wait identifier](#). If the [thread state](#) indicates that the [thread](#)
7 is waiting for mutual exclusion then its [wait identifier](#) contains a [handle](#) that indicates the data
8 object upon which the [thread](#) is waiting. The [get_state](#) entry point, which has the [get_state](#)
9 [OMPT type](#), enables a [tool](#) to retrieve the state and the [wait identifier](#) of the [encountering thread](#).
10 The returned value may be any one of the states predefined by the [state OMPT type](#) or a value
11 that represents an [implementation defined](#) state. The [tool](#) may obtain a string representation for
12 each state with the [enumerate_states](#) entry point. If the returned state indicates that the
13 [thread](#) is waiting for a [lock](#), [nestable lock](#), [critical region](#), [atomic region](#), or [ordered](#)
14 [region](#) and the [wait identifier](#) passed as the [wait_id](#) argument is not [NULL](#) then the value of the [wait](#)
15 [identifier](#) is assigned to that argument, which is a pointer to a [handle](#). If the returned state is not one
16 of the specified wait states then the value of that [handle](#) is undefined after the call.

17 **Restrictions**

18 Restrictions on the [get_state](#) entry point are as follows:

- 19
 - The [wait_id](#) argument must be a reference to a [variable](#) of the [wait_id OMPT type](#) or
- 20 [NULL](#).

21 **Cross References**

- 22
 - [enumerate_states](#) Entry Point, see [Section 42.2](#)
- 23
 - OMPT [state](#) Type, see [Section 39.31](#)
- 24
 - OMPT [wait_id](#) Type, see [Section 39.40](#)

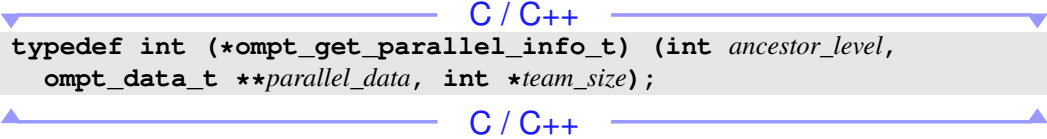
25 **42.14 get_parallel_info Entry Point**

Name: get_parallel_info Category: function	Properties: async-signal-safe , C/C++-only , OMPT
---	---

1 **Return Type and Arguments**

Name	Type	Properties
<i><return type></i>	integer	<i>default</i>
<i>ancestor_level</i>	integer	<i>default</i>
<i>parallel_data</i>	data	OMPT, pointer-to-pointer
<i>team_size</i>	integer	pointer

3 **Type Signature**

4  C / C++
5 **typedef int (*ompt_get_parallel_info_t) (int ancestor_level,
ompt_data_t **parallel_data, int *team_size);**

6 **Semantics**

7 During execution, an OpenMP program may employ nested parallel regions. The
8 **get_parallel_info** entry point, which has the **get_parallel_info** OMPT type, enables
9 a tool to retrieve information about the current parallel region and any enclosing parallel regions for
10 the current execution context.

11 The *ancestor_level* argument specifies the parallel region of interest by its ancestor level. Ancestor
12 level 0 refers to the innermost parallel region; information about enclosing parallel regions may be
13 obtained using larger values for *ancestor_level*. Information about a parallel region may not be
14 available if the ancestor level is 0; otherwise it must be available if a parallel region exists at the
15 specified ancestor level. The entry point returns 2 if a parallel region exists at the specified ancestor
16 level and the information is available, 1 if a parallel region exists at the specified ancestor level but
17 the information is currently unavailable, and 0 otherwise. The *parallel_data* argument returns the
18 parallel data if the argument is not NULL. The *team_size* argument returns the team size if the
19 argument is not NULL. If no parallel region exists at the specified ancestor level or the information
20 is unavailable then the values of variables passed by reference to the entry point are undefined when
21 **get_parallel_info** returns.

22 A tool may use the pointer to the data object of a parallel region that it obtains from this entry point
23 to inspect or to modify the value of the data object. When a parallel region is created, its data object
24 will be initialized with the value **ompt_data_none**. Between a *parallel-begin* event and an
25 *implicit-task-begin* event, a call to **get_parallel_info** with an *ancestor_level* value of 0 may
26 return information about the outer team or the new team. If a thread is in the
27 **ompt_state_wait_barrier_implicit_parallel** state then a call to
28 **get_parallel_info** may return a pointer to a copy of the specified parallel region's
29 *parallel_data* rather than a pointer to the data word for the region itself. This convention enables
30 the primary thread for a parallel region to free storage for the region immediately after the region
31 ends, yet avoid having some other thread in the team that is executing the region potentially
32 reference the *parallel_data* object for the region after it has been freed.

33 If **get_parallel_info** returns two then the entry point has the following effects:

- If a **non-null value** was passed for *parallel_data*, the value returned in *parallel_data* is a pointer to a data word that is associated with the **parallel region** at the specified level; and
- If a **non-null value** was passed for *team_size*, the value returned in the integer to which *team_size* points is the number of **threads** in the **team** that is associated with the **parallel region**.

Restrictions

Restrictions on the **get_parallel_info** entry point are as follows:

- While the *ancestor_level* argument is passed by value, all other arguments must be valid pointers to **variables** of the specified types or **NULL**.

Cross References

- OMPT **data** Type, see [Section 39.8](#)
- OMPT **state** Type, see [Section 39.31](#)

42.15 get_task_info Entry Point

Name: <code>get_task_info</code>	Properties: <code>async-signal-safe</code> , <code>C/C++-only</code> , <code>OMPT</code>
Category: <code>function</code>	

Return Type and Arguments

Name	Type	Properties
<i><return type></i>	integer	<i>default</i>
<i>ancestor_level</i>	integer	<i>default</i>
<i>flags</i>	integer	<i>pointer</i>
<i>task_data</i>	data	<i>OMPT, pointer-to-pointer</i>
<i>task_frame</i>	frame	<i>OMPT, pointer-to-pointer</i>
<i>parallel_data</i>	data	<i>OMPT, pointer-to-pointer</i>
<i>thread_num</i>	integer	<i>pointer</i>

Type Signature

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 **routines**. The **get_task_info** entry point, which has the **get_task_info** OMPT type, enables a **tool** to retrieve information about any **task** on the stack of the **encountering thread**.

The *ancestor_level* argument specifies the **task region** of interest by its ancestor level. Ancestor level 0 refers to the **encountering 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 a **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. The *flags* argument returns the **task** type if the argument is not **NULL**. The *task_data* argument returns the **task** data if the argument is not **NULL**. The *task_frame* argument returns the **task frame** pointer if the argument is not **NULL**. The *parallel_data* argument returns the parallel data if the argument is not **NULL**. The *thread_num* argument returns the **thread number** if the argument is not **NULL**. 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 **get_task_info** returns.

A **tool** may use a pointer to a data object for a **task** or **parallel region** that it obtains from **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**.

If **get_task_info** returns 2 then the **entry point** has the following effects:

- If a **non-null value** was passed for *flags* then the value returned in the integer to which *flags* points represents the type of the **task** at the specified level; possible **task** types include **initial task**, **implicit task**, **explicit task**, and **target task**;
- If a **non-null value** was passed for *task_data* then the value that is returned in the object to which it points is a pointer to a data word that is associated with the **task** at the specified level;
- If a **non-null value** was passed for *task_frame* then the value that is returned in the object to which *task_frame* points is a pointer to the **frame OMPT type structure** that is associated with the **task** at the specified level;
- If a **non-null value** was passed for *parallel_data* then the value that is returned in the object to which *parallel_data* points is a pointer to a data word that is associated with the **parallel region** that contains the **task** at the specified level or, if the **task** at the specified level is an **initial task**, **NULL**; and
- If a **non-null value** was passed for *thread_num*, then the value that is returned in the object to which *thread_num* points indicates the number of the **thread** in the **parallel region** that is executing the **task** at the specified level.

1 **Restrictions**

2 Restrictions on the `get_task_info` entry point are as follows:

- 3 • While the *ancestor_level* argument is passed by value, all other arguments must be valid
4 pointers to *variables* of the specified types or `NULL`.

5 **Cross References**

- 6 • OMPT `data` Type, see [Section 39.8](#)
7 • OMPT `frame` Type, see [Section 39.15](#)
8 • OMPT `task_flag` Type, see [Section 39.37](#)

9 **42.16 get_task_memory Entry Point**

Name: <code>get_task_memory</code> Category: function	Properties: async-signal-safe , C/C++-only , OMPT
--	--

11 **Return Type and Arguments**

Name	Type	Properties
<i><return type></i>	integer	default
<i>addr</i>	void	pointer-to-pointer
<i>size</i>	<code>size_t</code>	pointer
<i>block</i>	integer	default

13 **Type Signature**

```
14 C / C++  
15 typedef int (*ompt_get_task_memory_t) (void **addr, size_t *size,  
16     int block);  
C / C++
```

16 **Semantics**

17 During execution, a *thread* may be executing a *task*. The OpenMP implementation must preserve
18 the *data environment* from the generation of the *task* for its execution. The `get_task_memory`
19 *entry point*, which has the `get_task_memory` OMPT type, enables a *tool* to retrieve information
20 about *memory* ranges that store the *data environment* for the *encountering task*. Multiple *memory*
21 ranges may be used to store these data. The *addr* argument is a pointer to a void pointer return
22 value to provide the start address of a *memory* range. The *size* argument is a pointer to a size type
23 return value to provide the size of the *memory* range. The *block* argument, which is an integer
24 value to specify the *memory* block of interest, supports iteration over the *memory* ranges. The
25 `get_task_memory` entry point returns one if more *memory* ranges are available, and zero
26 otherwise. If no *memory* is used for a *task*, *size* is set to zero. In this case, the value to which *addr*
27 points is *undefined*.

42.17 get_target_info Entry Point

Name: <code>get_target_info</code> Category: <code>function</code>	Properties: <code>async-signal-safe</code> , <code>C/C++-only</code> , <code>OMPT</code>
---	---

Return Type and Arguments

Name	Type	Properties
<code><return type></code>	<code>integer</code>	<code>default</code>
<code>device_num</code>	<code>c_uint64_t</code>	<code>pointer</code>
<code>target_id</code>	<code>id</code>	<code>OMPT</code> , <code>pointer</code>
<code>host_op_id</code>	<code>id</code>	<code>OMPT</code> , <code>pointer-to-pointer</code>

Type Signature

```
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 `get_target_info` entry point, which has the `get_target_info` OMPT type, enables a tool to retrieve identifiers that specify the current `target` region and target operation ID of the `encountering thread`, if any. This entry point returns one if the `encountering thread` is in a `target` region and zero otherwise. If the entry point returns zero then the values of the `variables` passed by reference as its arguments are `undefined`. If the `encountering thread` is in a `target` region then `get_target_info` returns information about the `current device`, active `target` region, and active host operation, if any. In this case, the `device_num` argument returns the `device number` of the `target` region and the `target_id` argument returns the `target` region identifier. If the `encountering thread` is in the process of initiating an operation on a `target device` (for example, copying data to or from a `device`) then `host_op_id` returns the identifier for the operation; otherwise, `host_op_id` returns `ompt_id_none`.

Restrictions

Restrictions on the `get_target_info` entry point are as follows:

- All arguments must be valid pointers to `variables` of the specified types.

Cross References

- OMPT `id` Type, see [Section 39.18](#)

42.18 get_num_devices Entry Point



Name: <code>get_num_devices</code> Category: <code>function</code>	Properties: <code>async-signal-safe</code> , <code>C/C++-only</code> , <code>OMPT</code>
---	---

1 **Return Type**

2

Name	Type	Properties
<i><return type></i>	integer	<i>default</i>

3 **Type Signature**

4  **typedef int (*ompt_get_num_devices_t) (void);** 

5 **Semantics**

6 The `get_num_devices` entry point, which has the `get_num_devices` OMPT type, is the
7 entry point that enables a tool to retrieve the number of devices available to an OpenMP program.

8 **42.19 get_unique_id Entry Point**

9



Name: <code>get_unique_id</code> Category: <code>function</code>	Properties: <code>async-signal-safe</code> , <code>C/C++-only</code> , <code>OMPT</code>
---	---

10 **Return Type**

11

Name	Type	Properties
<i><return type></i>	<code>c_uint64_t</code>	<i>default</i>

12 **Type Signature**

13  **typedef uint64_t (*ompt_get_unique_id_t) (void);** 

14 **Semantics**



15 The `get_unique_id` entry point, which has the `get_unique_id` OMPT type, enables a tool
16 to retrieve a number that is unique for the duration of an OpenMP program. Successive invocations
17 may not result in consecutive or even increasing numbers.

18 **42.20 finalize_tool Entry Point**

19

Name: <code>finalize_tool</code> Category: <code>subroutine</code>	Properties: <code>C/C++-only</code> , <code>OMPT</code>
---	--

20 **Type Signature**

21  **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 **finalize_tool** entry point, which has the **finalize_tool** OMPT type. Upon completion of **finalize_tool**, no OMPT callbacks are dispatched. The entry point detaches the **tool** from the runtime, unregisters all callbacks and invalidates all OMPT entry points passed to the **tool** by **function_lookup**. Upon completion of **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 on the **finalize_tool** entry point are as follows:

- The entry point must not be called from inside an **explicit region**.
- As **finalize_tool** 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 entry point has completed results in **unspecified behavior**.

43 Device Tracing Entry Points

The second set of [OMPT 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 device](#) to process these [trace records](#). This chapter defines that set of [entry points](#).

Several [OMPT entry points](#) have a *device* argument. This argument is a pointer to an [OpenMP object](#) that represents the [target device](#). [Callbacks](#) in the [device](#) tracing interface use a pointer to this [device](#) object to identify the [device](#) being addressed.

43.1 get_device_num_procs Entry Point

Name: <code>get_device_num_procs</code> Category: function	Properties: C/C++-only , OMPT
---	--

Return Type and Arguments

Name	Type	Properties
<i><return type></i>	integer	default
<i>device</i>	device	OMPT , pointer

Type Signature

C / C++

```
typedef int (*ompt_get_device_num_procs_t) (  
    ompt_device_t *device);
```

C / C++

Semantics

The [get_device_num_procs](#) entry point, which has the [get_device_num_procs](#) [OMPT type](#), enables a [tool](#) to retrieve the number of [processors](#) that are available on the [device](#) at the time the [entry point](#) 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.

Cross References

- [OMPT device](#) Type, see [Section 39.11](#)

43.2 get_device_time Entry Point

Name: <code>get_device_time</code> Category: function	Properties: C/C++-only , OMPT
--	---

Return Type and Arguments

Name	Type	Properties
<return type>	<code>device_time</code>	default
<i>device</i>	<code>device</code>	OMPT , pointer

Type Signature

[C / C++](#)

```
typedef ompt_device_time_t (*ompt_get_device_time_t) (  
    ompt_device_t *device);
```

[C / C++](#)

Semantics

[Host devices](#) and [target devices](#) are typically distinct and run independently. If the [host device](#) and any [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 [get_device_time](#) entry point, which has the [get_device_time](#) OMPT type, enables a [tool](#) to retrieve the current time on the [device](#) specified by the *device* argument. A [tool](#) can use the information retrieved by [get_device_time](#) to align time stamps from different [devices](#).

Cross References

- OMPT `device` Type, see [Section 39.11](#)
- OMPT `device_time` Type, see [Section 39.12](#)

43.3 translate_time Entry Point

Name: <code>translate_time</code> Category: function	Properties: C/C++-only , OMPT
---	---

Return Type and Arguments

Name	Type	Properties
<return type>	<code>double</code>	default
<i>device</i>	<code>device</code>	OMPT , pointer
<i>time</i>	<code>device_time</code>	OMPT

Type Signature

[C / C++](#)

```
typedef double (*ompt_translate_time_t) (ompt_device_t *device,  
    ompt_device_time_t time);
```

[C / C++](#)

Semantics

The `translate_time` entry point, which has the `translate_time` OMPT type, enables a tool to translate a time value, specified by the `time` argument, obtained from the `device` specified by the `device` argument 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`.

Cross References

- OMPT `device` Type, see [Section 39.11](#)
- OMPT `device_time` Type, see [Section 39.12](#)
- `omp_get_wtime` Routine, see [Section 36.3.1](#)

43.4 set_trace_ompt Entry Point

Name: <code>set_trace_ompt</code> Category: function	Properties: C/C++-only , OMPT
---	--

Return Type and Arguments

Name	Type	Properties
<i><return type></i>	<code>set_result</code>	default
<i>device</i>	<code>device</code>	OMPT , pointer
<i>enable</i>	<code>integer</code>	OMPT , unsigned
<i>etype</i>	<code>integer</code>	OMPT , unsigned

Type Signature

C / C++

```
typedef ompt_set_result_t (*ompt_set_trace_ompt_t) (  
    ompt_device_t *device, unsigned int enable, unsigned int etype);
```

C / C++

Semantics

A tool uses the `set_trace_ompt` entry point, which has the `set_trace_ompt` OMPT type, to enable or to disable the recording of standard `trace records` for one or more types of `events` that the `etype` argument indicates. If the value of `etype` is zero then the invocation applies to all `events`. If `etype` is `positive` then it applies to the `event` in the `callbacks` OMPT type that matches that value. The `enable` argument indicates whether tracing should be enabled or disabled for the `events` that `etype` specifies; a `positive` value indicates that recording should be enabled while a value of zero indicates that recording should be disabled. If `etype` specifies any of the `events` that correspond to the `target_data_op_emi` or `target_submit_emi` callbacks then tracing, if supported, is enabled or disabled for those `events` when they occur on the `host device`. If `etype` specifies any other `events` then tracing, if supported, is enabled or disabled for those `events` when they occur on the specified `target device`. The return value of `set_trace_ompt` indicates the outcome of

enabling or disabling the recording of the [trace records](#) and can be any value in the [set_result OMPT type](#) except [ompt_set_sometimes_paired](#).

Cross References

- OMPT **callbacks** Type, see [Section 39.6](#)
- OMPT **device** Type, see [Section 39.11](#)
- Tracing Activity on Target Devices, see [Section 38.2.5](#)
- OMPT **set_result** Type, see [Section 39.28](#)

43.5 set_trace_native Entry Point

Name: <code>set_trace_native</code> Category: function	Properties: C/C++-only , OMPT
---	--

Return Type and Arguments

Name	Type	Properties
<i><return type></i>	set_result	default
<i>device</i>	device	OMPT , pointer
<i>enable</i>	integer	default
<i>flags</i>	integer	default

Type Signature

C / C++

```
typedef ompt_set_result_t (*ompt_set_trace_native_t) (  
    ompt_device_t *device, int enable, int flags);
```

C / C++

Semantics

A [tool](#) uses the [set_trace_native](#) entry point, which has the [set_trace_native OMPT type](#), to enable or to disable the recording of native [trace records](#). 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 of the [native_mon_flag OMPT type](#). The return value of [set_trace_native](#) indicates the outcome of enabling or disabling the recording of the [trace records](#) and can be any value in the [set_result OMPT type](#) except [ompt_set_sometimes_paired](#).

This interface is designed for use by a [tool](#) that cannot directly use native control [procedures](#) for the [device](#). If a [tool](#) can directly use the native control [procedures](#) then it can invoke them directly using pointers that the [function_lookup](#) entry point associated with the [device](#) provides and that are described in the *documentation* string that is provided to its [device_initialize](#) callback.

Cross References

- OMPT `device` Type, see [Section 39.11](#)
- Tracing Activity on Target Devices, see [Section 38.2.5](#)
- OMPT `native_mon_flag` Type, see [Section 39.21](#)
- OMPT `set_result` Type, see [Section 39.28](#)

43.6 `get_buffer_limits` Entry Point

Name: <code>get_buffer_limits</code> Category: subroutine	Properties: C/C++-only , OMPT
--	---

Arguments

Name	Type	Properties
<i>device</i>	device	OMPT , pointer
<i>max_concurrent_allocs</i>	integer	pointer
<i>recommended_bytes</i>	size_t	pointer

Type Signature

C / C++

```
typedef void (*ompt_get_buffer_limits_t) (ompt_device_t *device,  
int *max_concurrent_allocs, size_t *recommended_bytes);
```

C / C++

Semantics

The `get_buffer_limits` entry point, which has the `get_buffer_limits` OMPT type, enables a [tool](#) to retrieve the maximum number of concurrent buffer allocations and the recommended size of any buffer allocation that will be requested of the [tool](#) for a specified [device](#). The *max_concurrent_allocs* points to a location in which the [entry point](#) returns the maximum number of buffer allocations that the implementation may request for tracing activity on the [target device](#) without the implementation performing [callback dispatch](#) of [buffer_complete callbacks](#) with its *buffer_owned* argument set to a non-zero value for any of the buffers. The *recommended_bytes* argument points to a location in which the [entry point](#) returns the recommended buffer size of the buffer to be returned by the [tool](#) when the implementation dispatches a [buffer_request callback](#) for the [target device](#).

A [tool](#) may use this [entry point](#) prior to a call to the `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 that this [entry point](#) returns remain the same on each successive invocation unless the `stop_trace` entry point is called for the same [target device](#) between the successive invocations.

Cross References

- `buffer_complete` Callback, see [Section 41.6](#)
- `buffer_request` Callback, see [Section 41.5](#)
- OMPT `device` Type, see [Section 39.11](#)
- `start_trace` Entry Point, see [Section 43.7](#)
- `stop_trace` Entry Point, see [Section 43.10](#)

43.7 start_trace Entry Point

Name: <code>start_trace</code> Category: function	Properties: C/C++-only , OMPT
--	---

Return Type and Arguments

Name	Type	Properties
<i><return type></i>	integer	default
<i>device</i>	device	OMPT , pointer
<i>request</i>	<code>buffer_request</code>	OMPT , procedure
<i>complete</i>	<code>buffer_complete</code>	OMPT , procedure

Type Signature

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

The `start_trace` entry point, which has the `start_trace` OMPT type, enables a tool to start tracing of activity on a specified device. The *request* argument specifies a callback that supplies a buffer in which a device can deposit events. The *complete* argument specifies a callback that the OpenMP implementation invokes to empty a buffer that contains trace records.

Under normal operating conditions, every event buffer that a tool callback provides for a device is returned to the tool before the OpenMP runtime shuts down. If an exceptional condition terminates execution of an OpenMP program, the runtime may not return buffers provided for the device. An invocation of `start_trace` returns one if the entry point succeeds and zero otherwise.

Cross References

- `buffer_complete` Callback, see [Section 41.6](#)
- `buffer_request` Callback, see [Section 41.5](#)
- OMPT `device` Type, see [Section 39.11](#)

43.8 pause_trace Entry Point

Name: <code>pause_trace</code> Category: function	Properties: C/C++-only , OMPT
--	--

Return Type and Arguments

Name	Type	Properties
<return type>	integer	default
<i>device</i>	device	OMPT , pointer
<i>begin_pause</i>	integer	default

Type Signature

<div>C / C++</div> <pre>typedef int (*ompt_pause_trace_t) (ompt_device_t *device, int begin_pause);</pre> <div>C / C++</div>
--

Semantics

The [pause_trace](#) entry point, which has the [pause_trace](#) OMPT type, enables a [tool](#) to pause or to resume tracing on a [device](#). 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. An invocation of [pause_trace](#) returns one if it succeeds and zero otherwise. Redundant pause or resume commands are idempotent and will return the same value as the prior command.

Cross References

- OMPT `device` Type, see [Section 39.11](#)

43.9 flush_trace Entry Point

Name: <code>flush_trace</code> Category: function	Properties: C/C++-only , OMPT
--	--

Return Type and Arguments

Name	Type	Properties
<return type>	integer	default
<i>device</i>	device	OMPT , pointer

Type Signature

<div>C / C++</div> <pre>typedef int (*ompt_flush_trace_t) (ompt_device_t *device);</pre> <div>C / C++</div>

1 **Semantics**

2 The **flush_trace** entry point, which has the **flush_trace** OMPT type, enables a tool to
3 cause the OpenMP implementation to issue a sequence of zero or more **buffer_complete**
4 **callbacks** to deliver all **trace records** that have been collected prior to the flush for the specified
5 **device**. An invocation of **flush_trace** returns one if the **entry point** succeeds and zero
6 otherwise.

7 **Cross References**

- 8
 - OMPT **device** Type, see [Section 39.11](#)

9

43.10 stop_trace Entry Point

10

Name: <code>stop_trace</code> Category: function	Properties: C/C++-only , OMPT
---	--

11 **Return Type and Arguments**

12

Name	Type	Properties
<i><return type></i>	integer	default
<i>device</i>	device	OMPT , pointer

13 **Type Signature**

14

▼

C / C++

▼

typedef int (*ompt_stop_trace_t) (ompt_device_t *device);

▲

C / C++

▲

15 **Semantics**

16 The **stop_trace** entry point, which has the **stop_trace** OMPT type, provides a superset of
17 the functionality of the **flush_trace** entry point. Specifically, the **stop_trace** entry point
18 stops tracing for the specified **device** in addition to flushing pending trace records. An invocation of
19 **stop_trace** returns one if the **entry point** succeeds and zero otherwise.

20 **Cross References**

- 21
 - OMPT **device** Type, see [Section 39.11](#)
 - **flush_trace** Entry Point, see [Section 43.9](#)

23

43.11 advance_buffer_cursor Entry Point

24

Name: <code>advance_buffer_cursor</code> Category: function	Properties: C/C++-only , OMPT
--	--

1 **Return Type and Arguments**

Name	Type	Properties
<return type>	integer	<i>default</i>
<i>device</i>	device	OMPT, pointer
<i>buffer</i>	buffer	OMPT, pointer
<i>size</i>	size_t	<i>default</i>
<i>current</i>	buffer_cursor	OMPT, opaque
<i>next</i>	buffer_cursor	OMPT, opaque, pointer

3 **Type Signature**

```
4       C / C++
5       typedef int (*ompt_advance_buffer_cursor_t) (
6       ompt_device_t *device, ompt_buffer_t *buffer, size_t size,
7       ompt_buffer_cursor_t current, ompt_buffer_cursor_t *next);
8       C / C++
```

7 **Semantics**

8 The [advance_buffer_cursor](#) entry point, which has the [advance_buffer_cursor](#)
9 OMPT type, enables a [tool](#) to advance the trace buffer pointer for the buffer that the *buffer*
10 argument indicates to the next [trace record](#). The *size* argument indicates the size of *buffer* in bytes.
11 The *current* argument is an [OpenMP object](#) that indicates the current position, while the *next*
12 argument returns an [OpenMP object](#) with the next value. An invocation of
13 [advance_buffer_cursor](#) returns *true* if the advance is successful and the next position in the
14 buffer is valid. Otherwise, it returns *false*.

15 **Cross References**

- 16 • OMPT **buffer** Type, see [Section 39.3](#)
- 17 • OMPT **buffer_cursor** Type, see [Section 39.4](#)
- 18 • OMPT **device** Type, see [Section 39.11](#)

19 **43.12 get_record_type Entry Point**

Name: get_record_type Category: function	Properties: C/C++-only , OMPT
--	---

21 **Return Type and Arguments**

Name	Type	Properties
<return type>	record	<i>default</i>
<i>buffer</i>	buffer	OMPT, pointer
<i>current</i>	buffer_cursor	OMPT

1

2

4

5

6

- 7

8

9

0

1

1

2

2

Semantics

The `get_record_ompt` entry point, which has the `get_record_ompt` OMPT type, enables a tool to obtain a pointer to an OMPT trace record from a trace buffer associated with a device. The pointer may point to storage in the buffer indicated by the `buffer` argument or it may point to a trace record in thread-local storage in which the information extracted from a trace record was assembled. The information available for an event depends upon its type. The `current` argument is an OpenMP object that indicates the position from which to extract the trace record. The return value of the `record_ompt` OMPT type includes a field of the `any_record_ompt` OMPT type, which is a union that can represent information for any OMPT trace record type. Another call to the entry point may overwrite the contents of the fields in a trace record returned by a prior invocation.

Cross References

- OMPT `any_record_ompt` Type, see Section 39.2
- OMPT `buffer` Type, see Section 39.3
- OMPT `buffer_cursor` Type, see Section 39.4
- OMPT `device` Type, see Section 39.11
- OMPT `record_ompt` Type, see Section 39.26

43.14 get_record_native Entry Point

Name: <code>get_record_native</code> Category: function	Properties: C/C++-only, OMPT
--	-------------------------------------

Return Type and Arguments

Name	Type	Properties
<i><return type></i>	<code>c_ptr</code>	<i>default</i>
<i>buffer</i>	<code>buffer</code>	OMPT, pointer
<i>current</i>	<code>buffer_cursor</code>	OMPT, opaque
<i>host_op_id</i>	<code>id</code>	OMPT, pointer

Type Signature

C / C++

```
typedef void (*ompt_get_record_native_t) (  
    ompt_buffer_t *buffer, ompt_buffer_cursor_t current,  
    ompt_id_t *host_op_id);
```

C / C++

1 **Semantics**

2 The `get_record_native` entry point, which has the `get_record_native` OMPT type,
3 enables a `tool` to obtain a pointer to a `native trace record` from a trace buffer associated with a
4 `device`. The pointer may point to storage in the buffer indicated by the `buffer` argument or it may
5 point to a `trace record` in `thread`-local storage in which the information extracted from a `trace record`
6 was assembled. The information available for a native `event` depends upon its type. The `current`
7 argument is an `OpenMP object` that indicates the position from which to extract the `trace record`. If
8 the `entry point` returns a `non-null pointer` result, it will also set the object to which the `host_op_id`
9 argument points to a host-side identifier for the operation that is associated with the `trace record` on
10 the `target device` and was created when the operation was initiated by the `host device`. Another call
11 to the `entry point` may overwrite the contents of the fields in a `trace record` returned by a prior
12 invocation.

13 **Cross References**

- 14 • OMPT `buffer` Type, see [Section 39.3](#)
- 15 • OMPT `buffer_cursor` Type, see [Section 39.4](#)
- 16 • OMPT `id` Type, see [Section 39.18](#)

17

43.15 `get_record_abstract` Entry Point

Name: <code>get_record_abstract</code> Category: <code>function</code>	Properties: <code>C/C++-only</code> , <code>OMPT</code>
---	---

19 **Return Type and Arguments**

Name	Type	Properties
<code><return type></code>	<code>record_abstract</code>	<code>pointer</code>
<code>native_record</code>	<code>void</code>	<code>pointer</code>

21 **Type Signature**

22

C / C++

```
typedef ompt_record_abstract_t *  
(*ompt_get_record_abstract_t) (void *native_record);
```

C / C++

24 **Semantics**

25 An OpenMP implementation may execute on a `device` that logs `trace records` in a `native trace`
26 `format` that a `tool` cannot interpret directly. The `get_record_abstract` entry point, which has
27 the `get_record_abstract` OMPT type, enables a `tool` to translate a `native trace record` to
28 which the `native_record` argument points into a standard form.

29 **Cross References**

- 30 • OMPT `record_abstract` Type, see [Section 39.24](#)

1

Part VI

2

OMPD

44 OMPD Overview

This chapter provides an overview of **OMPD**, which is an interface for **third-party tools**, such as a debugger. **Third-party tools** exist in separate processes from the **OpenMP program**. To provide **OMPD** support, an OpenMP implementation must provide an **OMPD library** that the **third-party tool** can load. An OpenMP implementation does not need to maintain any extra information to support **OMPD** inquiries from **third-party tools** unless it is explicitly instructed to do so.

OMPD allows **third-party tools** to inspect the OpenMP state of a live **OpenMP program** or core file in an implementation-agnostic manner. Thus, a **third-party tool** that uses **OMPD** should work with any **compliant implementation**. An OpenMP implementation 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 an **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**, which exports the **OMPD API** and which the **third-party tool** uses to determine OpenMP information about the **OpenMP program**. The **OMPD library** must look up 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 44.3.2](#) and [Section 44.3.3](#). 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 [Chapter 48](#). Each of these symbols corresponds to an **event** type. The OpenMP runtime must ensure that control passes through the appropriate named location when **events** occur. If the **third-party tool** requires notification of an **event**, it can plant a breakpoint at the matching

location. The location can, but may not, be a function. It can, for example, simply be a label. However, the names of the locations must have external **C** linkage.

44.1 OMPD Interfaces Definitions

C / C++

A **compliant implementation** must supply a set of definitions for the **OMP third-party tool callback** signatures, **third-party tool** interface **routines** and the special data types of their parameters and return values. These definitions, which are listed throughout the **OMP** chapters, 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 defined** values. The **ompd_dll_locations** variable and all **OMP third-party tool** interface **routines** are external symbols with **C** linkage.

C / C++

44.2 Thread and Signal Safety

The **OMP** library does not need to be reentrant. The **third-party 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 **third-party tool**.

44.3 Activating a Third-Party Tool

The **third-party tool** and the **OpenMP** program exist as separate processes. Thus, OMPD requires coordination between the OpenMP runtime and the **third-party tool**.

44.3.1 Enabling Runtime Support for OMPD

In order to support **third-party tools**, the OpenMP runtime may need to collect and to store information that it may not otherwise maintain. The OpenMP runtime collects whatever information is necessary to support **OMP** if the **debug-var ICV** is set to **enabled**.

Cross References

- *debug-var* ICV, see [Table 3.1](#)

44.3.2 ompd_dll_locations

Format

C

```
extern const char **ompd_dll_locations;
```

C

Semantics

An OpenMP runtime may have more than one **OMPD library**. The **third-party tool** must be able to locate the right library to use for the program that it is examining. The **ompd_dll_locations** global **variable** points to the locations of **OMPD libraries** that are compatible with the OpenMP implementation. The OpenMP runtime system must provide this public **variable**, which is an **argv**-style vector of pathname string pointers that provide the names of the 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 should 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** 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.

44.3.3 ompd_dll_locations_valid Breakpoint

Format

```
void ompd_dll_locations_valid(void);
```

Semantics

Since **ompd_dll_locations** may not be a static **variable**, it may require runtime initialization. The OpenMP runtime notifies **third-party tools** that **ompd_dll_locations** is valid by having execution pass through a location that the symbol **ompd_dll_locations_valid** identifies. If **ompd_dll_locations** is **NULL**, a **third-party tool** can place a breakpoint at **ompd_dll_locations_valid** to be notified that **ompd_dll_locations** is initialized. In practice, the symbol **ompd_dll_locations_valid** may not be a function; instead, it may be a labeled machine instruction through which execution passes once the vector is valid.

45 OMPD Data Types

This chapter defines OMPD types, which support interactions with the OMPD library and provide information about the device architecture.

45.1 OMPD addr Type

Name: <code>addr</code> Properties: C/C++-only, OMPD	Base Type: <code>c_uint64_t</code>
---	---

Type Definition

C / C++

```
typedef uint64_t ompd_addr_t;
```

C / C++

Semantics

The `addr` OMPD type represents an address in an OpenMP process as an unsigned integer.

45.2 OMPD address Type

Name: <code>address</code> Properties: C/C++-only, OMPD	Base Type: <code>structure</code>
--	--

Fields

Name	Type	Properties
<i>segment</i>	<code>seg</code>	C/C++-only, OMPD
<i>address</i>	<code>addr</code>	C/C++-only, OMPD

Type Definition

C / C++

```
typedef struct ompd_address_t {  
    ompd_seg_t segment;  
    ompd_addr_t address;  
} ompd_address_t;
```

C / C++

Semantics

The `address` type is a `structure` that OMPD uses to specify addresses, which may or may not be segmented. For non-segmented architectures, `ompd_segment_none` is used in the `segment` field of the `address` OMPD type.

Cross References

- OMPd `addr` Type, see [Section 45.1](#)
- OMPd `seg` Type, see [Section 45.10](#)

45.3 OMPd `address_space_context` Type

Name: <code>address_space_context</code> Properties: C/C++-only, handle , OMPd	Base Type: <code>aspace_cont</code>
---	-------------------------------------

Type Definition

C / C++

```
typedef struct _ompd_aspace_cont ompd_address_space_context_t;
```

C / C++

Semantics

A [third-party tool](#) uses the `address_space_context` OMPd type, which represents [handles](#) that are opaque to the OMPd library and that define an [address space context](#) uniquely, to identify the [address space](#) of the [OpenMP process](#) that it is monitoring.

45.4 OMPd `callbacks` Type

Name: <code>callbacks</code> Properties: C/C++-only, OMPd	Base Type: structure
--	--------------------------------------

Fields

Name	Type	Properties
<i>alloc_memory</i>	<code>memory_alloc</code>	C-only , OMPd
<i>free_memory</i>	<code>memory_free</code>	C-only , OMPd
<i>print_string</i>	<code>print_string</code>	C-only , OMPd
<i>sizeof_type</i>	<code>sizeof</code>	C-only , OMPd
<i>symbol_addr_lookup</i>	<code>symbol_addr</code>	C-only , OMPd
<i>read_memory</i>	<code>memory_read</code>	C-only , OMPd
<i>write_memory</i>	<code>memory_write</code>	C-only , OMPd
<i>read_string</i>	<code>memory_read</code>	C-only , OMPd
<i>device_to_host</i>	<code>device_host</code>	C-only , OMPd
<i>host_to_device</i>	<code>device_host</code>	C-only , OMPd
<i>get_thread_context_for_thread_id</i>	<code>get_thread_context_for_thread_id</code>	C-only , OMPd

Type Definition

C / C++

```
typedef struct ompd_callbacks_t {  
    ompd_callback_memory_alloc_fn_t alloc_memory;  
    ompd_callback_memory_free_fn_t free_memory;  
    ompd_callback_print_string_fn_t print_string;  
    ompd_callback_sizeof_fn_t sizeof_type;  
    ompd_callback_symbol_addr_fn_t symbol_addr_lookup;  
    ompd_callback_memory_read_fn_t read_memory;  
    ompd_callback_memory_write_fn_t write_memory;  
    ompd_callback_memory_read_fn_t read_string;  
    ompd_callback_device_host_fn_t device_to_host;  
    ompd_callback_device_host_fn_t host_to_device;  
    ompd_callback_get_thread_context_for_thread_id_fn_t  
        get_thread_context_for_thread_id;  
} ompd_callbacks_t;
```

C / C++

Semantics

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. The set of callbacks that the OMPD library must use is collected in an instance of the callbacks OMPD type that is passed to the OMPD library as an argument to `ompd_initialize`. Each field points to a procedure that the OMPD library must use to interact with the OpenMP program or for memory operations.

The `alloc_memory` and `free_memory` fields are pointers to `alloc_memory` and `free_memory` callbacks, which the OMPD library uses to allocate and to release dynamic memory. The `print_string` field points to a `print_string` callback 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 `sizeof_type` callback 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 `device_to_host` callback 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 `host_to_device` callback to which the `host_to_device` field points.

The `symbol_addr_lookup` field points to a `symbol_addr_lookup` callback, which 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 `read_memory`, `read_string` and `write_memory` callbacks 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 `get_thread_context_for_thread_id` callback that the OMPD library can use to obtain a native thread context that corresponds to a native thread identifier.

Cross References

- `alloc_memory` Callback, see [Section 46.1.1](#)
- `device_to_host` Callback, see [Section 46.4.2](#)
- `free_memory` Callback, see [Section 46.1.2](#)
- `get_thread_context_for_thread_id` Callback, see [Section 46.3.1](#)
- `host_to_device` Callback, see [Section 46.4.3](#)
- `ompd_initialize` Routine, see [Section 47.1.1](#)
- `print_string` Callback, see [Section 46.5](#)
- `read_memory` Callback, see [Section 46.2.2.1](#)
- `read_string` Callback, see [Section 46.2.2.2](#)
- `sizeof_type` Callback, see [Section 46.3.2](#)
- `symbol_addr_lookup` Callback, see [Section 46.2.1](#)
- `write_memory` Callback, see [Section 46.2.3](#)

45.5 OMPD device Type

Name: <code>device</code> Properties: <code>C/C++-only</code> , <code>OMPD</code>	Base Type: <code>c_uint64_t</code>
--	------------------------------------

Type Definition

C / C++

```
typedef uint64_t ompd_device_t;
```

C / C++

Semantics

The **device** OMPD type provides information about OpenMP devices. 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 OMPD. Instead, a device identifier is passed across the interface via its device kind, its size in bytes and a pointer to where it is stored. The OMPD library and the third-party tool use the device 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 device 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/>.

45.6 OMPD device_type_sizes Type

Name: <code>device_type_sizes</code> Properties: C/C++-only, OMPD	Base Type: <code>structure</code>
--	--

Fields

Name	Type	Properties
<code>sizeof_char</code>	<code>c_uint8_t</code>	C/C++-only, OMPD
<code>sizeof_short</code>	<code>c_uint8_t</code>	C/C++-only, OMPD
<code>sizeof_int</code>	<code>c_uint8_t</code>	C/C++-only, OMPD
<code>sizeof_long</code>	<code>c_uint8_t</code>	C/C++-only, OMPD
<code>sizeof_long_long</code>	<code>c_uint8_t</code>	C/C++-only, OMPD
<code>sizeof_pointer</code>	<code>c_uint8_t</code>	C/C++-only, OMPD

Type Definition

C / C++

```
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;
```

C / C++

Semantics

The **device_type_sizes** OMPD type is used in OMPD callbacks through which the OMPD library can interrogate a 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 **device_type_sizes** 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

- `sizeof_type` Callback, see [Section 46.3.2](#)

45.7 OMPD `frame_info` Type

Name: <code>frame_info</code> Properties: C/C++-only, OMPD	Base Type: <code>structure</code>
---	-----------------------------------

Fields

Name	Type	Properties
<code>frame_address</code>	<code>address</code>	C/C++-only, OMPD
<code>frame_flag</code>	<code>word</code>	C/C++-only, OMPD

Type Definition

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 `frame_info` OMPD type is a `structure` type that OMPD uses to specify `frame` information. The `frame_address` field of `frame_info` identifies a `frame`. The `frame_flag` field of `frame_info` indicates what type of information is provided in `frame_address`. The values and meaning are the same as are defined for the `frame_flag` OMPT type.

Cross References

- OMPD `address` Type, see [Section 45.2](#)
- OMPT `frame_flag` Type, see [Section 39.16](#)
- OMPD `word` Type, see [Section 45.17](#)

45.8 OMPD `icv_id` Type

Name: <code>icv_id</code> Properties: C/C++-only, OMPD	Base Type: <code>c_uint64_t</code>
---	------------------------------------

Predefined Identifiers

Name	Value	Properties
<code>ompd_icv_undefined</code>	0	C/C++-only, OMPD

Type Definition

```
typedef uint64_t ompd_icv_id_t;
```

Semantics

The `icv_id` OMPD type identifies ICVs.

45.9 OMPD rc Type

Name: <code>rc</code> Properties: C/C++-only, OMPD	Base Type: enumeration
---	------------------------

Values

Name	Value	Properties
<code>ompd_rc_ok</code>	0	C-only, OMPD
<code>ompd_rc_unavailable</code>	1	C-only, OMPD
<code>ompd_rc_stale_handle</code>	2	C-only, OMPD
<code>ompd_rc_bad_input</code>	3	C-only, OMPD
<code>ompd_rc_error</code>	4	C-only, OMPD
<code>ompd_rc_unsupported</code>	5	C-only, OMPD
<code>ompd_rc_needs_state_tracking</code>	6	C-only, OMPD
<code>ompd_rc_incompatible</code>	7	C-only, OMPD
<code>ompd_rc_device_read_error</code>	8	C-only, OMPD
<code>ompd_rc_device_write_error</code>	9	C-only, OMPD
<code>ompd_rc_nomem</code>	10	C-only, OMPD
<code>ompd_rc_incomplete</code>	11	C-only, OMPD
<code>ompd_rc_callback_error</code>	12	C-only, OMPD
<code>ompd_rc_incompatible_handle</code>	13	C-only, OMPD

Type Definition

```
typedef enum ompd_rc_t {
    ompd_rc_ok = 0,
    ompd_rc_unavailable = 1,
    ompd_rc_stale_handle = 2,
    ompd_rc_bad_input = 3,
    ompd_rc_error = 4,
    ompd_rc_unsupported = 5,
    ompd_rc_needs_state_tracking = 6,
    ompd_rc_incompatible = 7,
    ompd_rc_device_read_error = 8,
    ompd_rc_device_write_error = 9,
```

```

1      ompd_rc_nomem           = 10,
2      ompd_rc_incomplete     = 11,
3      ompd_rc_callback_error = 12,
4      ompd_rc_incompatible_handle = 13
5  } ompd_rc_t;

```

C / C++

The **rc** OMPD type is the return code type of OMPD routines and OMPD callbacks. The values of the **rc** OMPD type and their semantics are defined as follows:

- **ompd_rc_ok**: The routine or callback procedure was successful;
- **ompd_rc_unavailable**: Information was not available for the specified context;
- **ompd_rc_stale_handle**: The specified handle was not valid;
- **ompd_rc_bad_input**: The arguments (other than handles) are invalid;
- **ompd_rc_error**: A fatal error occurred;
- **ompd_rc_unsupported**: The requested routine or callback is not supported for the specified arguments;
- **ompd_rc_needs_state_tracking**: The state tracking operation failed because state tracking was not enabled;
- **ompd_rc_incompatible**: The selected OMPD library was incompatible with the OpenMP program or was incapable of handling it;
- **ompd_rc_device_read_error**: A read operation failed on the device;
- **ompd_rc_device_write_error**: A write operation failed on the device;
- **ompd_rc_nomem**: A memory allocation failed;
- **ompd_rc_incomplete**: The information provided on return was incomplete, while the arguments were set to valid values;
- **ompd_rc_callback_error**: The callback interface or one of the required callback procedures provided by the third-party tool was invalid; and
- **ompd_rc_incompatible_handle**: The specified handle was incompatible with the routine or callback.

45.10 OMPD seg Type



Name: seg Properties: C/C++-only, OMPD	Base Type: c_uint64_t
---	------------------------------

1 **Predefined Identifiers**

2

Name	Value	Properties
<code>ompd_segment_none</code>	0	C/C++-only, OMPD

3 **Type Definition**

4  `typedef uint64_t ompd_seg_t;` 

5 **Semantics**

6 The **seg** OMPD type represents a **segment** value as an unsigned integer.

7 **45.11 OMPD scope Type**

8



Name: <code>scope</code> Properties: C/C++-only, OMPD	Base Type: enumeration
--	-------------------------------

9 **Values**

10

Name	Value	Properties
<code>ompd_scope_global</code>	1	C-only, OMPD
<code>ompd_scope_address_space</code>	2	C-only, OMPD
<code>ompd_scope_thread</code>	3	C-only, OMPD
<code>ompd_scope_parallel</code>	4	C-only, OMPD
<code>ompd_scope_implicit_task</code>	5	C-only, OMPD
<code>ompd_scope_task</code>	6	C-only, OMPD
<code>ompd_scope_teams</code>	7	C-only, OMPD
<code>ompd_scope_target</code>	8	C-only, OMPD

11 **Type Definition**

12  `typedef enum ompd_scope_t {` 
13 `ompd_scope_global = 1,`
14 `ompd_scope_address_space = 2,`
15 `ompd_scope_thread = 3,`
16 `ompd_scope_parallel = 4,`
17 `ompd_scope_implicit_task = 5,`
18 `ompd_scope_task = 6,`
19 `ompd_scope_teams = 7,`
20 `ompd_scope_target = 8`
21 `} ompd_scope_t;`

Semantics

The **scope** OMPD type is used for **scope handles** to identify OpenMP scopes, including those related to **parallel regions** and **tasks**. When used in an **OMP routine** or **OMP callback procedure**, the **scope** OMPD type and the **OMP handle** must match according to Table 45.1.





TABLE 45.1: Mapping of Scope Type and OMPD Handles

Scope types	Handles
<code>ompd_scope_global</code>	Address space handle for the host device
<code>ompd_scope_address_space</code>	Any address space handle
<code>ompd_scope_thread</code>	Any native thread handle
<code>ompd_scope_parallel</code>	Any parallel handle
<code>ompd_scope_implicit_task</code>	Task handle for an implicit task
<code>ompd_scope_teams</code>	Parallel handle for an implicit parallel region generated from a teams construct
<code>ompd_scope_target</code>	Parallel handle for an implicit parallel region generated from a target construct
<code>ompd_scope_task</code>	Any task handle

45.12 OMPD size Type

Name: <code>size</code> Properties: C/C++-only, OMPD	Base Type: <code>c_uint64_t</code>
---	---

Type Definition

 C / C++ 
<code>typedef uint64_t ompd_size_t;</code>
 C / C++ 

The **size** OMPD type specifies the number of bytes in opaque data objects that are passed across the **OMP** API.

45.13 OMPD team_generator Type

Name: <code>team_generator</code> Properties: C/C++-only, OMPD	Base Type: <code>enumeration</code>
---	--

1 **Values**

Name	Value	Properties
<code>ompd_generator_program</code>	0	C-only, OMPD
<code>ompd_generator_parallel</code>	1	C-only, OMPD
<code>ompd_generator_teams</code>	2	C-only, OMPD
<code>ompd_generator_target</code>	3	C-only, OMPD

3 **Type Definition**

4 **C / C++**

```
typedef enum ompd_team_generator_t {  
5     ompd_generator_program = 0,  
6     ompd_generator_parallel = 1,  
7     ompd_generator_teams = 2,  
8     ompd_generator_target = 3  
9 } ompd_team_generator_t;
```

6 **C / C++**

10 **Semantics**

11 The `team_generator` OMPD type represents the value of the *team-generator-var* ICV. The
12 `ompd_generator_program` value indicates that the *team* is the *initial team* created at the start
13 of the OpenMP program. The `ompd_generator_parallel`, `ompd_generator_teams`,
14 and `ompd_generator_target` values indicate that the *team* was created by an encountered
15 *parallel* construct, *teams* construct, or *target* construct, respectively.

16 **45.14 OMPD thread_context Type**

Name: <code>thread_context</code> Properties: C/C++-only, handle, OMPD	Base Type: <code>thread_cont</code>
---	--

18 **Type Definition**

19 **C / C++**

```
typedef struct _ompd_thread_cont ompd_thread_context_t;
```

20 **C / C++**

21 **Semantics**

22 A *third-party tool* uses the `thread_context` OMPD type, which represents *handles* that are
23 opaque to the OMPD library and that uniquely identify a *native thread* of the OpenMP process that
it is monitoring.

45.15 OMPD thread_id Type

Name: thread_id Properties: C/C++-only, OMPD	Base Type: c_uint64_t
--	------------------------------

Type Definition

C / C++

```
typedef uint64_t ompd_thread_id_t;
```

C / C++

Semantics

The **thread_id** OMPD type provides information about **native threads**. 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 OMPD. Instead, a **native thread identifier** is passed across the interface via the **thread_id** OMPD type, its size in bytes and a pointer to where it is stored. The OMPD library and the third-party tool use the **thread_id** OMPD type 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 the **thread_id** OMPD type 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/>.

45.16 OMPD wait_id Type

Name: wait_id Properties: C/C++-only, OMPD	Base Type: c_uint64_t
--	------------------------------

Type Definition

C / C++

```
typedef uint64_t ompd_wait_id_t;
```

C / C++

Semantics

A **variable** of **wait_id** OMPD type identifies the object on which a **thread** waits. The values and meaning of **wait_id** are the same as those defined for the **wait_id** OMPT type.

Cross References

- OMPT **wait_id** Type, see [Section 39.40](#)

45.17 OMPD word Type

Name: word Properties: C/C++-only, OMPD	Base Type: c_int64_t
--	-----------------------------

Type Definition

C / C++

```
typedef int64_t ompd_word_t;
```

C / C++

Semantics

The **word OMPD type** represents a data word from the OpenMP runtime as a signed integer.

45.18 OMPD Handle Types

The **OMP library** defines **handles**, which have **OMP types** that are **handle types** (i.e., they have the **handle property**). These **handles** are used to refer to **address spaces**, **threads**, **parallel regions** and **tasks** and are managed by the OpenMP runtime. The internal **structures** that these **handles** represent are opaque to the **third-party tool**. Defining externally visible type names in this way introduces type safety to the interface and helps to catch instances where incorrect **handles** are passed by a **third-party tool** to the **OMP library**. The **structures** do not need to be defined; instead, the **OMP library** must cast incoming (pointers to) **handles** to the appropriate internal, private types.

Each **OMP routine** or **OMP callback procedure** that applies to a particular **address space**, **thread**, **parallel region** or **task** must explicitly specify a corresponding **handle**. A **handle** remains constant and valid while the associated entity is managed by the OpenMP runtime or until it is released with the corresponding **OMP routine** for releasing **handles** of that type. If a **third-party tool** receives notification of the end of the lifetime of a managed entity (see **Chapter 48**) or it releases the **handle**, the **handle** may no longer be referenced.

45.18.1 OMPD address_space_handle Type

Name: address_space_handle Properties: C/C++-only, handle, OMPD	Base Type: aspace_handle
--	---------------------------------

Type Definition

C / C++

```
typedef struct _ompd_aspace_handle ompd_address_space_handle_t;
```

C / C++

Semantics

The **address_space_handle OMPD type** is used for **handles** that represent **address spaces**.

1 **45.18.2 OMPD parallel_handle Type**

2

Name: parallel_handle Properties: C/C++-only, handle, OMPD	Base Type: parallel_handle
---	-----------------------------------

3 **Type Definition**

4

C / C++

typedef struct _ompd_parallel_handle ompd_parallel_handle_t;

C / C++

5 **Semantics**

6 The **parallel_handle** OMPD type is used for **parallel handles** that represent **parallel regions**.

7 **45.18.3 OMPD task_handle Type**

8

Name: task_handle Properties: C/C++-only, handle, OMPD	Base Type: task_handle
---	-------------------------------

9 **Type Definition**

10

C / C++

typedef struct _ompd_task_handle ompd_task_handle_t;

C / C++

11 **Semantics**

12 The **task_handle** OMPD type is used for **handles** that represent **tasks**.

13 **45.18.4 OMPD thread_handle Type**

14

Name: thread_handle Properties: C/C++-only, handle, OMPD	Base Type: thread_handle
---	---------------------------------

15 **Type Definition**

16

C / C++

typedef struct _ompd_thread_handle ompd_thread_handle_t;

C / C++

17 **Semantics**

18 The **thread_handle** OMPD type is used for **handles** that represent **threads**.

46 OMPD 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 be able to extract information from the OpenMP process that the **third-party tool** is examining. The process on which the **tool** is operating may be either a live process or a core file, and a **thread** may be either a live **thread** in a live process or a **thread** in a core file. To enable the **OMPd library** to extract state information from a process or core file, the **tool** must supply the **OMPd library** with **callbacks** to inquire about the size of primitive types in the **device** of the 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 **OMPd callbacks** in response to calls to the **OMPd library**. The names of the **OMPd callbacks** correspond to the names of the fields of the **callbacks OMPD type**.

Restrictions

The following restrictions apply to all **OMPd callbacks**:

- Unless explicitly specified otherwise, all **OMPd callbacks** must return **ompd_rc_ok** or **ompd_rc_stale_handle**.

46.1 Memory Management of OMPD Library

A **tool** provides **alloc_memory** and **free_memory callbacks** to obtain and to release heap **memory**. This mechanism ensures that the **OMPd library** does not interfere with any custom **memory** management scheme that the **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 such that any of its definitions of **new** and **delete** do not interfere with any that the **tool** defines. In some cases, the **OMPd library** must allocate **memory** to return results to the **tool**. The **tool** then owns this **memory** and has the responsibility to release it. Thus, the **OMPd library** and the **tool** must use the same **memory** manager.

The **OMPd library** creates **OMPd handles**, which are opaque to **tools** and may have a complex internal structure. A **tool** cannot determine if the **handle** pointers that **OMPd** returns correspond to discrete heap allocations. Thus, the **tool** must not simply deallocate a **handle** by passing an address that it receives from the **OMPd library** to its own **memory** manager. Instead, **OMPd** includes **routines** that the **tool** must use when it no longer needs a **handle**.

A [tool](#) creates [tool contexts](#) and passes them to the [OMPD library](#). The [OMPD library](#) does not release [tool contexts](#); instead the [tool](#) releases them after it releases any [handles](#) that may reference the [tool contexts](#).

Cross References

- [alloc_memory](#) Callback, see [Section 46.1.1](#)
- [free_memory](#) Callback, see [Section 46.1.2](#)

46.1.1 alloc_memory Callback

Name: alloc_memory Category: function	Properties: C-only , OMPD
--	---

Return Type and Arguments

Name	Type	Properties
<i><return type></i>	rc	default
<i>nbytes</i>	size	default
<i>ptr</i>	void	pointer-to-pointer

Type Signature

<pre>typedef ompd_rc_t (*ompd_callback_memory_alloc_fn_t) (ompd_size_t nbytes, void **ptr);</pre>	C
--	-------------------

Semantics

A [tool](#) provides an [alloc_memory](#) callback, which has the [memory_alloc](#) OMPD type, that the [OMPD library](#) may call to allocate [memory](#). 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](#) but is not guaranteed to be set to zero.

Cross References

- OMPD [rc](#) Type, see [Section 45.9](#)
- OMPD [size](#) Type, see [Section 45.12](#)

46.1.2 free_memory Callback

Name: free_memory Category: function	Properties: C-only , OMPD
---	---

1 **Return Type and Arguments**

Name	Type	Properties
<return type>	rc	default
ptr	void	pointer

3 **Type Signature**

4  C 
typedef ompd_rc_t (*ompd_callback_memory_free_fn_t) (void *ptr);

5 **Semantics**

6 A tool provides a **free_memory** callback, which has the **memory_free** OMPD type, that the
7 OMPd library may call to deallocate **memory** that was obtained from a prior call to the
8 **alloc_memory** callback. The *ptr* argument is the address of the block to be deallocated.

9 **Cross References**

- **alloc_memory** Callback, see [Section 46.1.1](#)
- OMPD **rc** Type, see [Section 45.9](#)

12 **46.2 Accessing Program or Runtime Memory**

13 The **OMPd library** cannot directly read from or write to **memory** of the **OpenMP program**. Instead
14 the **OMPd library** must use **callbacks** into the **third-party tool** that perform the operation.

15 **46.2.1 symbol_addr_lookup Callback**

Name: symbol_addr_lookup Category: function	Properties: C-only, OMPD
---	---------------------------------

17 **Return Type and Arguments**

Name	Type	Properties
<return type>	rc	default
address_space_context	address_space_context	pointer
thread_context	thread_context	pointer
symbol_name	char	intent(in), pointer
symbol_addr	address	pointer
file_name	char	intent(in), pointer

Type Signature

```
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

A [tool](#) provides a [symbol_addr_lookup](#) callback, which has the [symbol_addr](#) OMPD type, that the [OMP library](#) may call to look up the address of the symbol provided in the *symbol_name* argument within the [address space](#) specified by the *address_space_context* argument. This argument provides the [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 [tool](#) uses the *symbol_name* argument that the [OMP library](#) supplies verbatim. In particular, no name mangling, demangling or other transformations are performed before 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](#), [procedure](#), or untyped label. The symbol can have local, global, or weak binding. The [callback](#) returns the address of the symbol in the location to which *symbol_addr* points.

The *file_name* argument is an optional input argument 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 lookup order as when *file_name* is [NULL](#).

In addition to the general return codes for [OMP callbacks](#), [symbol_addr_lookup](#) callbacks may also return the following return codes:

- [ompd_rc_error](#) if the symbol that the *symbol_name* argument specifies is not found; or
- [ompd_rc_bad_input](#) if no symbol name is provided.

Restrictions

Restrictions on `symbol_addr_lookup` callbacks are as follows:

- The `address_space_context` argument must be a [non-null value](#).
- 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 [procedure](#).

Cross References

- OMPD `address` Type, see [Section 45.2](#)
- OMPD `address_space_context` Type, see [Section 45.3](#)
- OMPD `rc` Type, see [Section 45.9](#)
- OMPD `thread_context` Type, see [Section 45.14](#)

46.2.2 OMPD `memory_read` Type

Name: <code>memory_read</code>	Properties: C-only, OMPD
Category: function pointer	

Return Type and Arguments

Name	Type	Properties
<i><return type></i>	<code>rc</code>	default
<i>address_space_context</i>	<code>address_space_context</code>	pointer
<i>thread_context</i>	<code>thread_context</code>	pointer
<i>addr</i>	<code>address</code>	intent(in), pointer
<i>nbytes</i>	<code>size</code>	default
<i>buffer</i>	<code>void</code>	pointer

Type Signature

C

```
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);
```

C

Callbacks that have the `memory_read` OMPD type are [memory-reading callbacks](#), which each have the [memory-reading property](#). A [tool](#) provides these [callbacks](#) to read [memory](#) from an [OpenMP program](#). The `thread_context` argument of this type should be `NULL` for global [memory](#) accesses. If it is a [non-null value](#), the `thread_context` argument identifies the [native thread context](#) for the [memory](#) access for the purpose of accessing [thread](#) local storage. The data are returned through the `buffer` argument, which is allocated and owned by the [OMP library](#). The contents of

the buffer are unstructured, raw bytes. The [OMPD library](#) must use the [device_to_host callback](#) to perform any transformations such as byte-swapping that may be necessary.

In addition to the general return codes for [OMPD callbacks](#), [memory-reading callbacks](#) may also return the following return code:

- [ompd_rc_error](#) if unallocated [memory](#) is reached while reading *nbytes*.

Cross References

- OMPD [address](#) Type, see [Section 45.2](#)
- OMPD [address_space_context](#) Type, see [Section 45.3](#)
- [device_to_host](#) Callback, see [Section 46.4.2](#)
- OMPD [rc](#) Type, see [Section 45.9](#)
- OMPD [size](#) Type, see [Section 45.12](#)
- OMPD [thread_context](#) Type, see [Section 45.14](#)

46.2.2.1 read_memory Callback

Name: read_memory Category: function	Properties: C-only, common-type-callback, memory-reading, OMPD
---	---

Type Signature

[memory_read](#)

Semantics

A [tool](#) provides a [read_memory](#) callback, which is a [memory-reading callback](#), that the [OMPD library](#) may call to copy a block of data from *addr* within the [address space](#) given by *address_space_context* to the [tool buffer](#).

Cross References

- OMPD [address](#) Type, see [Section 45.2](#)
- OMPD [address_space_context](#) Type, see [Section 45.3](#)
- OMPD [memory_read](#) Type, see [Section 46.2.2](#)

46.2.2.2 read_string Callback

Name: read_string Category: function	Properties: C-only, common-type-callback, memory-reading, OMPD
---	---

Type Signature

[memory_read](#)

Semantics

A `tool` provides a `read_string` callback, which is a `memory-reading callback`, that the `OMPD library` may call to copy a string to which `addr` points, including the terminating null byte (`'\0'`), to the `tool buffer`. At most `nbytes` bytes are copied. If a null byte is not among the first `nbytes` bytes, the string placed in `buffer` is not null-terminated.

In addition to the general return codes for `memory-reading callbacks`, `read_string` callbacks may also return the following return code:

- `ompd_rc_incomplete` if no terminating null byte is found while reading `nbytes` using the `read_string` callback.

Cross References

- OMPD `rc` Type, see [Section 45.9](#)
- OMPD `size` Type, see [Section 45.12](#)

46.2.3 write_memory Callback

Name: <code>write_memory</code> Category: <code>function</code>	Properties: <code>C-only</code> , <code>OMPD</code>
--	--

Return Type and Arguments

Name	Type	Properties
<code><return type></code>	<code>rc</code>	<code>default</code>
<code>address_space_context</code>	<code>address_space_context</code>	<code>pointer</code>
<code>thread_context</code>	<code>thread_context</code>	<code>pointer</code>
<code>addr</code>	<code>address</code>	<code>intent(in)</code> , <code>pointer</code>
<code>nbytes</code>	<code>size</code>	<code>default</code>
<code>buffer</code>	<code>void</code>	<code>pointer</code>

Type Signature

C

```
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, void *buffer);
```

C

1 **Semantics**

2 A [tool](#) provides a [write_memory](#) callback, which has the [memory_write](#) OMPD type, that the
3 [OMP library](#) may call to have the [tool](#) write a block of data to a location within an [address space](#)
4 from a provided buffer. The address to which the data are to be written in the [OpenMP program](#)
5 that [address_space_context](#) specifies is given by *addr*. The *nbytes* argument is the number of bytes
6 to be transferred. The *thread_context* argument for global [memory](#) accesses should be `NULL`. If it
7 is a [non-null value](#), then *thread_context* identifies the [native thread context](#) for the [memory](#) access
8 for the purpose of accessing [thread](#) local storage.

9 The data to be written are passed through *buffer*, which is allocated and owned by the [OMP](#)
10 [library](#). The contents of the buffer are unstructured, raw bytes. The [OMP library](#) must use the
11 [host_to_device](#) callback to perform any transformations such as byte-swapping that may be
12 necessary to render the data into a form that is compatible with the OpenMP runtime.

13 In addition to the general return codes for [OMP callbacks](#), [write_memory](#) callbacks may also
14 return the following return codes:

- 15
 - [ompd_rc_error](#) if unallocated [memory](#) is reached while writing *nbytes*.

16 **Cross References**

- 17
 - OMPD `address` Type, see [Section 45.2](#)
 - 18 • OMPD `address_space_context` Type, see [Section 45.3](#)
 - 19 • `host_to_device` Callback, see [Section 46.4.3](#)
 - 20 • OMPD `rc` Type, see [Section 45.9](#)
 - 21 • OMPD `size` Type, see [Section 45.12](#)
 - 22 • OMPD `thread_context` Type, see [Section 45.14](#)

23

46.3 Context Management and Navigation

24 **Summary**

25 A [tool](#) provides [callbacks](#) to manage and to navigate [tool context](#) relationships.

26

46.3.1 `get_thread_context_for_thread_id` Callback

27

Name: <code>get_thread_context_for_thread_id</code> Category: function	Properties: C-only , OMP
--	--

Return Type and Arguments

Name	Type	Properties
<return type>	rc	default
address_space_context	address_space_context	opaque, pointer
kind	thread_id	default
sizeof_thread_id	size	default
thread_id	void	intent(in), pointer
thread_context	thread_context	pointer-to-pointer

Type Signature

```
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

A tool provides a `get_thread_context_for_thread_id` callback, which has the `get_thread_context_for_thread_id` OMPD type, that the OMPD library may call to map 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.

The `address_space_context` argument is an opaque handle that the tool provides to reference an address space. The `kind`, `sizeof_thread_id`, and `thread_id` arguments represent a native thread identifier. On return, the `thread_context` argument provides a handle that maps a native thread identifier to a tool native thread context.

In addition to the general return codes for OMPD callbacks, `get_thread_context_for_thread_id` callbacks 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 `get_thread_context_for_thread_id` callbacks are as follows:

- The provided `thread_context` must be valid until the OMPD library returns from the tool procedure.

Cross References

- OMP `address_space_context` Type, see [Section 45.3](#)
- OMP `rc` Type, see [Section 45.9](#)
- OMP `size` Type, see [Section 45.12](#)
- OMP `thread_context` Type, see [Section 45.14](#)
- OMP `thread_id` Type, see [Section 45.15](#)

46.3.2 `sizeof_type` Callback

Name: <code>sizeof_type</code> Category: function	Properties: C-only, OMP
--	-------------------------

Return Type and Arguments

Name	Type	Properties
<i><return type></i>	<code>rc</code>	<i>default</i>
<i>address_space_context</i>	<code>address_space_context</code>	pointer
<i>sizes</i>	<code>device_type_sizes</code>	pointer

Type Signature

C

```
typedef ompd_rc_t (*ompd_callback_sizeof_fn_t) (  
    ompd_address_space_context_t *address_space_context,  
    ompd_device_type_sizes_t *sizes);
```

C

Semantics

A [tool](#) provides a [sizeof_type](#) callback, which has the [sizeof](#) OMP type, that the OMP library may call to query the sizes of the basic primitive types for the [address space](#) that the *address_space_context* argument specifies in the location to which *sizes* points.

Cross References

- OMP `address_space_context` Type, see [Section 45.3](#)
- OMP `device_type_sizes` Type, see [Section 45.6](#)
- OMP `rc` Type, see [Section 45.9](#)

46.4 Device Translating Callbacks

Summary

A [tool](#) provides [device-translating callbacks](#), which have the [device-translating property](#), to perform any necessary translations between [devices](#) on which the [tool](#) and [OMPD library](#) run and on which the [OpenMP program](#) runs.

46.4.1 OMPD device_host Type

Name: <code>device_host</code>	Properties: C-only, OMPD
Category: function pointer	

Return Type and Arguments

Name	Type	Properties
<i><return type></i>	rc	<i>default</i>
<i>address_space_context</i>	<code>address_space_context</code>	pointer
<i>input</i>	<code>void</code>	intent(in) , pointer
<i>unit_size</i>	<code>size</code>	<i>default</i>
<i>count</i>	<code>size</code>	<i>default</i>
<i>output</i>	<code>void</code>	pointer

Type Signature

```
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 [tool](#) and [OMPD library](#) use and the ones that the [OpenMP program](#) uses. The `device_host` OMPD type is the type signature of the `device_to_host` and `host_to_device` callbacks that the [tool](#) provides to convert data between formats.

The `address_space_context` argument specifies the [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.

1 **Cross References**

- 2 • OMPD `address_space_context` Type, see [Section 45.3](#)
- 3 • `device_to_host` Callback, see [Section 46.4.2](#)
- 4 • `host_to_device` Callback, see [Section 46.4.3](#)
- 5 • OMPD `rc` Type, see [Section 45.9](#)
- 6 • OMPD `size` Type, see [Section 45.12](#)

7 **46.4.2 device_to_host Callback**

8

Name: <code>device_to_host</code> Category: function	Properties: C-only, common-type-callback, device-translating, OMPD
---	---

9 **Type Signature**

10 [device_host](#)

11 **Semantics**

12 The [device_to_host](#) is the [device-translating callback](#) that translates data that is read from the

13 [OpenMP program](#).

14 **Cross References**

- 15 • OMPD `device_host` Type, see [Section 46.4.1](#)

16 **46.4.3 host_to_device Callback**

17

Name: <code>host_to_device</code> Category: function	Properties: C-only, common-type-callback, device-translating, OMPD
---	---

18 **Type Signature**

19 [device_host](#)

20 **Semantics**

21 The [host_to_device](#) is the [device-translating callback](#) that translates data that is to be written

22 to the [OpenMP program](#).

23 **Cross References**

- 24 • OMPD `device_host` Type, see [Section 46.4.1](#)

46.5 print_string Callback

Name: <code>print_string</code> Category: function	Properties: C-only , OMPD
---	--

Return Type and Arguments

Name	Type	Properties
<i><return type></i>	<code>rc</code>	default
<i>string</i>	<code>char</code>	intent(in) , pointer
<i>category</i>	<code>integer</code>	default

Type Signature

C

```
typedef ompd_rc_t (*ompd_callback_print_string_fn_t) (  
    const char *string, int category);
```

C

Semantics

A [tool](#) provides a [print_string](#) callback, which has the [print_string](#) OMPD type, that the [OMPD library](#) may call to emit output, such as logging or debug information. The [tool](#) may set the [print_string](#) callback to `NULL` to prevent the [OMPD library](#) from emitting output. The [OMPD library](#) may not write to file descriptors that it did not open. The *string* argument is the null-terminated string to be printed; no conversion or formatting is performed on the string. The *category* argument is the [implementation defined](#) category of the string to be printed.

Cross References

- OMPD `rc` Type, see [Section 45.9](#)

47 OMPD Routines

This chapter defines the [OMP](#) routines, which are [routines](#) that have the [OMP](#) property and, thus, are provided by the [OMP](#) library to be used by [third-party tools](#). Some [OMP](#) 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.

47.1 OMPD Library Initialization and Finalization

The [OMP](#) 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 [OMP](#) API that the library supports by calling [ompd_get_api_version](#). If the [tool](#) supports the version that [ompd_get_api_version](#) returns, the [tool](#) starts the initialization by calling [ompd_initialize](#) using the version of the [OMP](#) 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.

Cross References

- [ompd_get_api_version](#) Routine, see [Section 47.1.2](#)
- [ompd_initialize](#) Routine, see [Section 47.1.1](#)

47.1.1 ompd_initialize Routine

Name: ompd_initialize Category: function	Properties: C-only , OMP
---	--

Return Type and Arguments

Name	Type	Properties
<i><return type></i>	rc	default
<i>api_version</i>	word	default
<i>callbacks</i>	callbacks	intent(in) , pointer

Prototypes

C

```
ompd_rc_t ompd_initialize(ompd_word_t api_version,  
    const ompd_callbacks_t *callbacks);
```

C

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** 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.

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 **callbacks** in the *callbacks* input argument, which enables the **OMPD library** to allocate and to deallocate memory in the **address space** of the **tool**, 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**.

This **routine** returns **ompd_rc_bad_input** if invalid **callbacks** are provided. In addition to the return codes permitted for all **OMPD routines**, this **routine** may return **ompd_rc_unsupported** if the requested API version cannot be provided.

Cross References

- **OMPD callbacks** Type, see [Section 45.4](#)
- **ompd_get_api_version** Routine, see [Section 47.1.2](#)
- **OMPD rc** Type, see [Section 45.9](#)
- **OMPD word** Type, see [Section 45.17](#)

47.1.2 ompd_get_api_version Routine

Name: <code>ompd_get_api_version</code>	Properties: C-only, OMPD
Category: <code>function</code>	

Return Type and Arguments

Name	Type	Properties
<i><return type></i>	rc	<i>default</i>
<i>api_version</i>	word	<i>pointer</i>

Prototypes

C

`ompd_rc_t ompd_get_api_version(ompd_word_t *api_version);`

C

1 **Semantics**

2 The [tool](#) may call the [ompd_get_api_version](#) routine to obtain the latest OMPD API version
3 number of the OMPD library. The OMPD API version number is equal to the value of the
4 `_OPENMP` macro defined in the associated OpenMP implementation, if the C preprocessor is
5 supported. If the associated OpenMP implementation compiles Fortran codes without the use of a
6 C preprocessor, the OMPD API version number is equal to the value of the [openmp_version](#)
7 [predefined identifier](#). The latest version number is returned into the location to which the *version*
8 argument points.

9 **Cross References**

- 10 • [ompd_initialize](#) Routine, see [Section 47.1.1](#)
11 • OMPD `rc` Type, see [Section 45.9](#)
12 • OMPD `word` Type, see [Section 45.17](#)

13 **47.1.3 ompd_get_version_string Routine**

14

Name: <code>ompd_get_version_string</code> Category: function	Properties: C-only , OMP
--	---

15 **Return Type and Arguments**

16

Name	Type	Properties
<i><return type></i>	<code>rc</code>	default
<i>string</i>	<code>const char</code>	intent(out) , pointer-to-pointer

17 **Prototypes**

18

```
▼ C ▼  
| ompd_rc_t ompd_get_version_string(const char **string); |  
▲ C ▲
```

19 **Semantics**

20 The [ompd_get_version_string](#) routine returns a pointer to a descriptive version string of
21 the OMPD library vendor, implementation, internal version, date, or any other information that may
22 be useful to a [tool](#) user or vendor. An implementation should provide a different string for every
23 change to its source code or build that could be visible to the OMPD user.

24 A pointer to a descriptive version string is placed into the location to which the *string* output
25 argument points. The OMPD library owns the string that the OMPD library returns; the [tool](#) must
26 not modify or release this string. The string remains valid for as long as the library is loaded. The
27 [ompd_get_version_string](#) routine may be called before [ompd_initialize](#).
28 Accordingly, the OMPD library must not use heap or stack memory for the string.

29 The signatures of [ompd_get_api_version](#) and [ompd_get_version_string](#) are
30 guaranteed not to change in future versions of OMPD. In contrast, the type definitions and

1 prototypes in the rest of [OMPD](#) do not carry the same guarantee. Therefore a [tool](#) that uses [OMPD](#)
2 should check the version of the loaded [OMPD library](#) before it calls any other [OMPD routine](#).

3 **Cross References**

- 4 • [OMPD `address_space_handle` Type](#), see [Section 45.18.1](#)
- 5 • [`ompd_get_api_version` Routine](#), see [Section 47.1.2](#)
- 6 • [OMPD `rc` Type](#), see [Section 45.9](#)

7 **47.1.4 `ompd_finalize` Routine**

8

Name: <code>ompd_finalize</code> Category: function	Properties: C-only , OMPD
--	--

9 **Return Type**

10

Name	Type	Properties
<code><return type></code>	<code>rc</code>	default

11 **Prototypes**

12

C	<code>ompd_rc_t ompd_finalize(void);</code>	C
-------------------	---	-------------------

13 **Semantics**

14 When the [tool](#) is finished with the [OMPD library](#), it should call [`ompd_finalize`](#) before it
15 unloads the library. The call to the [`ompd_finalize` routine](#) must be the last [OMPD](#) call that the
16 [tool](#) makes before it unloads the library. This [routine](#) allows the [OMPD library](#) to free any resources
17 that it may be holding. The [OMPD library](#) may implement a *finalizer* section, which executes as the
18 library is unloaded and therefore after the call to [`ompd_finalize`](#). During finalization, the
19 [OMPD library](#) may use the [callbacks](#) that the [tool](#) provided earlier during the call to
20 [`ompd_initialize`](#). In addition to the return codes permitted for all [OMPD routines](#), this
21 [routine](#) returns [`ompd_rc_unsupported`](#) if the [OMPD library](#) is not initialized.

22 **Cross References**

- 23 • [OMPD `rc` Type](#), see [Section 45.9](#)

24 **47.2 Process Initialization and Finalization**

25 **47.2.1 `ompd_process_initialize` Routine**

26

Name: <code>ompd_process_initialize</code> Category: function	Properties: C-only , OMPD
--	--

1 **Return Type and Arguments**

Name	Type	Properties
<i><return type></i>	rc	<i>default</i>
<i>context</i>	address_space_context	opaque, pointer
<i>host_handle</i>	address_space_handle	opaque, pointer-to-pointer

3 **Prototypes**

4 **C**

```
ompd_rc_t ompd_process_initialize(  
5     ompd_address_space_context_t *context,  
6     ompd_address_space_handle_t **host_handle);
```

7 **C**

7 **Semantics**

8 A [tool](#) calls [ompd_process_initialize](#) to obtain an [address space handle](#) for the [host device](#)
9 when it initializes a session on an [OpenMP process](#) or core file. On return from
10 [ompd_process_initialize](#), the [tool](#) owns the [address space handle](#), which it must release
11 with [ompd_rel_address_space_handle](#). The initialization function must be called before
12 any [OMPD](#) operations are performed on the OpenMP process or core file. This [routine](#) allows the
13 [OMPD library](#) to confirm that it can handle the [OpenMP process](#) or core file that *context* identifies.

14 The *context* argument is an opaque [handle](#) that the [tool](#) provides to address an [address space](#) from
15 the [host device](#). On return, the *host_handle* argument provides an opaque [handle](#) to the [tool](#) for this
16 [address space](#), which the [tool](#) must release when it is no longer needed.

17 In addition to the return codes permitted for all [OMPD routines](#), this [routine](#) returns
18 [ompd_rc_incompatible](#) if the [OMPD library](#) is incompatible with the runtime library loaded
19 in the process.

20 **Cross References**

- 21 • [OMPD address_space_context](#) Type, see [Section 45.3](#)
- 22 • [OMPD address_space_handle](#) Type, see [Section 45.18.1](#)
- 23 • [ompd_rel_address_space_handle](#) Routine, see [Section 47.8.1](#)
- 24 • [OMPD rc](#) Type, see [Section 45.9](#)

25 **47.2.2 ompd_device_initialize Routine**

Name: ompd_device_initialize Category: function	Properties: C-only , OMPD
--	---

1

2

3

4
5
6
7
8

9

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. 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.

In addition to the return codes permitted for all `OMPD` routines, this routine may return `ompd_rc_unsupported` if the `OMPD` library has no support for the specific `device`.

20

- OMPD **address_space_context** Type, see [Section 45.3](#)
- OMPD **address_space_handle** Type, see [Section 45.18.1](#)
- OMPD **device** Type, see [Section 45.5](#)
- OMPD **rc** Type, see [Section 45.9](#)
- OMPD **size** Type, see [Section 45.12](#)

47.2.3 ompd_get_device_thread_id_kinds Routine

Name: <code>ompd_get_device_thread_id_kinds</code> Category: function	Properties: C-only , OMPD
---	---

Return Type and Arguments

Name	Type	Properties
<i><return type></i>	<code>rc</code>	default
<i>device_handle</i>	<code>address_space_handle</code>	opaque , pointer
<i>kinds</i>	<code>thread_id</code>	pointer-to-pointer
<i>thread_id_sizes</i>	<code>size</code>	pointer-to-pointer
<i>count</i>	<code>integer</code>	pointer

Prototypes

C

```
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);
```

C

Semantics

The `ompd_get_device_thread_id_kinds` routine 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.

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.

Cross References

- [OMPD address_space_handle](#) Type, see [Section 45.18.1](#)
- `ompd_device_initialize` Routine, see [Section 47.2.2](#)
- `ompd_process_initialize` Routine, see [Section 47.2.1](#)
- [OMPD rc](#) Type, see [Section 45.9](#)

- OMPD `size` Type, see [Section 45.12](#)
- OMPD `thread_id` Type, see [Section 45.15](#)

47.3 Address Space Information



47.3.1 `ompd_get_omp_version` Routine

Name: <code>ompd_get_omp_version</code>	Properties: C-only, OMPD
Category: function	

Return Type and Arguments

Name	Type	Properties
<i><return type></i>	<code>rc</code>	default
<i>address_space</i>	<code>address_space_handle</code>	opaque , pointer
<i>omp_version</i>	<code>word</code>	pointer

Prototypes


<code>ompd_rc_t ompd_get_omp_version(ompd_address_space_handle_t *address_space, ompd_word_t *omp_version);</code>


Semantics

The [tool](#) may call the [ompd_get_omp_version](#) routine to obtain the version of the OpenMP API that is associated with the [address space](#) *address_space*. The *address_space* argument is an opaque [handle](#) that the [tool](#) provides to reference the [address space](#) of the process or [device](#). Upon return, the *omp_version* argument contains the version of the OpenMP runtime in the `_OPENMP` version macro format.

Cross References

- OMPD `address_space_handle` Type, see [Section 45.18.1](#)
- OMPD `rc` Type, see [Section 45.9](#)
- OMPD `word` Type, see [Section 45.17](#)

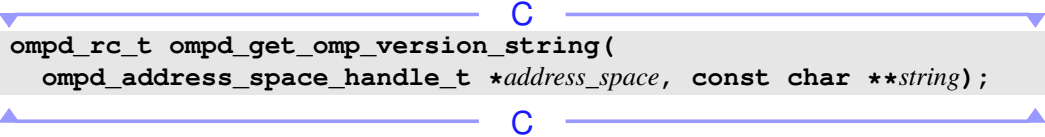
47.3.2 `ompd_get_omp_version_string` Routine

Name: <code>ompd_get_omp_version_string</code>	Properties: C-only, OMPD
Category: function	

1 **Return Type and Arguments**

Name	Type	Properties
<i><return type></i>	rc	<i>default</i>
<i>address_space</i>	address_space_handle	opaque, pointer
<i>string</i>	const char	intent(out), pointer-to-pointer

3 **Prototypes**

4  **C**

```
ompd_rc_t ompd_get_omp_version_string(  
    ompd_address_space_handle_t *address_space, const char **string);
```

6 **Semantics**

7 The [ompd_get_omp_version_string](#) routine returns a descriptive string for the OpenMP
8 API version that is associated with an [address space](#). The *address_space* argument is an opaque
9 [handle](#) that the [tool](#) provides to reference the [address space](#) of a process or [device](#). A pointer to a
10 descriptive version string is placed into the location to which the *string* output argument points.
11 After returning from the [routine](#), the [tool](#) owns the string. The [OMPD library](#) must use the memory
12 allocation [callback](#) that the [tool](#) provides to allocate the string storage. The [tool](#) is responsible for
13 releasing the [memory](#).

14 **Cross References**

- [OMPD Handle Types](#), see [Section 45.18](#)
- [OMPD rc Type](#), see [Section 45.9](#)

17 **47.4 Thread Handle Routines**

18 **47.4.1 ompd_get_thread_in_parallel Routine**

Name: ompd_get_thread_in_parallel Category: function	Properties: C-only , OMPD
---	---

20 **Return Type and Arguments**

Name	Type	Properties
<i><return type></i>	rc	<i>default</i>
<i>parallel_handle</i>	parallel_handle	opaque, pointer
<i>thread_num</i>	integer	<i>default</i>
<i>thread_handle</i>	thread_handle	opaque, pointer-to-pointer

1 **Prototypes**

C

```
2       ompd_rc_t ompd_get_thread_in_parallel(  
3           ompd_parallel_handle_t *parallel_handle, int thread_num,  
4           ompd_thread_handle_t **thread_handle);
```

C

5 **Semantics**

6 The **ompd_get_thread_in_parallel** routine enables a tool to obtain handles for OpenMP
7 threads that are associated with a parallel region. A successful invocation of
8 **ompd_get_thread_in_parallel** returns a pointer to a native thread handle in the location
9 to which *thread_handle* points. This routine yields meaningful results only if all OpenMP threads
10 in the team that is executing the parallel region are stopped.

11 The *parallel_handle* argument is an opaque handle for a parallel region and selects the parallel
12 region on which to operate. The *thread_num* argument represents the thread number and selects the
13 thread, the handle for which is to be returned. On return, the *thread_handle* argument is a handle
14 for the selected thread.

15 This routine returns **ompd_rc_bad_input** if the *thread_num* argument is greater than or equal
16 to the *team-size-var* ICV or negative, in which case the value returned in *thread_handle* is invalid.

17 **Cross References**

- 18 • **ompd_get_icv_from_scope** Routine, see Section 47.11.2
- 19 • OMPD **parallel_handle** Type, see Section 45.18.2
- 20 • OMPD **rc** Type, see Section 45.9
- 21 • OMPD **thread_handle** Type, see Section 45.18.4

22 **47.4.2 ompd_get_thread_handle Routine**

Name: ompd_get_thread_handle	Properties: C-only, OMPD
Category: function	

24 **Return Type and Arguments**

Name	Type	Properties
<return type>	rc	default
<i>handle</i>	address_space_handle	pointer
<i>kind</i>	thread_id	default
<i>sizeof_thread_id</i>	size	default
<i>thread_id</i>	void	intent(in), pointer
<i>thread_handle</i>	thread_handle	pointer-to-pointer

1 **Prototypes**

C

```
2       ompd_rc_t ompd_get_thread_handle(  
3           ompd_address_space_handle_t *handle, ompd_thread_id_t kind,  
4           ompd_size_t sizeof_thread_id, const void *thread_id,  
5           ompd_thread_handle_t **thread_handle);
```

C

6 **Semantics**

7 The **ompd_get_thread_handle** routine maps a [native thread](#) to a [native thread handle](#).
8 Further, the routine determines if the [native thread identifier](#) to which *thread_id* points represents an
9 [OpenMP thread](#). If so, the routine returns **ompd_rc_ok** and the location to which *thread_handle*
10 points is set to the [native thread handle](#) for the [native thread](#) to which the [OpenMP thread](#) is mapped.

11 The *handle* argument is a [handle](#) that the [tool](#) provides to reference an [address space](#). The *kind*,
12 *sizeof_thread_id*, and *thread_id* arguments represent a [native thread identifier](#). On return, the
13 *thread_handle* argument provides a [handle](#) to the [native thread](#) within the provided [address space](#).

14 The [native thread identifier](#) to which *thread_id* points must be valid for the duration of the call to
15 the routine. If the OMPD library must retain the [native thread identifier](#), it must copy it.

16 This routine returns **ompd_rc_bad_input** if a different value in *sizeof_thread_id* is expected
17 for a [thread](#) kind of *kind*. In addition to the return codes permitted for all [OMP routines](#), this
18 routine returns **ompd_rc_unsupported** if the *kind* of [thread](#) is not supported and it returns
19 **ompd_rc_unavailable** if the [native thread](#) is not an [OpenMP thread](#).

20 **Cross References**

- 21 • OMPD **address_space_handle** Type, see [Section 45.18.1](#)
- 22 • OMPD **rc** Type, see [Section 45.9](#)
- 23 • OMPD **size** Type, see [Section 45.12](#)
- 24 • OMPD **thread_handle** Type, see [Section 45.18.4](#)
- 25 • OMPD **thread_id** Type, see [Section 45.15](#)

26 **47.4.3 ompd_get_thread_id Routine**

27 Name: <code>ompd_get_thread_id</code> Category: function	Properties: C-only , OMP
--	---

1 **Return Type and Arguments**

Name	Type	Properties
<return type>	rc	default
thread_handle	thread_handle	pointer
kind	thread_id	default
sizeof_thread_id	size	default
thread_id	void	pointer

3 **Prototypes**

4 **C**

```
ompd_rc_t ompd_get_thread_id(ompd_thread_handle_t *thread_handle,  
5 ompd_thread_id_t kind, ompd_size_t sizeof_thread_id,  
6 void *thread_id);
```

C

7 **Semantics**

8 The `ompd_get_thread_id` routine maps a native thread handle to a native thread identifier.
9 This routine yields meaningful results only if the referenced OpenMP thread is stopped. The
10 `thread_handle` argument is a native thread handle. The `kind` argument represents the native thread
11 identifier. The `sizeof_thread_id` argument represents the size of the native thread identifier. On
12 return, the `thread_id` argument is a buffer that represents a native thread identifier.

13 This routine returns `ompd_rc_bad_input` if a different value in `sizeof_thread_id` is expected
14 for a native thread kind of `kind`. In addition to the return codes permitted for all OMPD routines,
15 this routine returns `ompd_rc_unsupported` if the `kind` of native thread is not supported.

16 **Cross References**

- 17 • OMPD `rc` Type, see [Section 45.9](#)
- 18 • OMPD `size` Type, see [Section 45.12](#)
- 19 • OMPD `thread_handle` Type, see [Section 45.18.4](#)
- 20 • OMPD `thread_id` Type, see [Section 45.15](#)

21 **47.4.4 ompd_get_device_from_thread Routine**

Name: <code>ompd_get_device_from_thread</code>	Properties: C-only, OMPD
Category: function	

23 **Return Type and Arguments**

Name	Type	Properties
<return type>	rc	default
thread_handle	thread_handle	pointer
device	address_space_handle	pointer-to-pointer

1 **Prototypes**

C

```
2       ompd_rc_t ompd_get_device_from_thread(  
3           ompd_thread_handle_t *thread_handle,  
4           ompd_address_space_handle_t **device);
```

C

5 **Semantics**

6 The `ompd_get_device_from_thread` routine obtains a pointer to the `address space handle`
7 for a `device` on which an `OpenMP thread` is executing. The returned pointer will be the same as the
8 `address space handle` pointer that was previously returned by a call to
9 `ompd_process_initialize` (for a `host device`) or a call to `ompd_device_initialize`
10 (for a `non-host device`). This routine yields meaningful results only if the referenced `OpenMP`
11 `thread` is stopped.

12 The `thread_handle` argument is a pointer to a `native thread handle` that represents a `native thread` to
13 which an `OpenMP thread` is mapped. On return, the `device` argument is the address of a pointer to
14 an `address space handle`.

15 **Cross References**

- 16 • OMPD `address_space_handle` Type, see [Section 45.18.1](#)
- 17 • OMPD `rc` Type, see [Section 45.9](#)
- 18 • OMPD `thread_handle` Type, see [Section 45.18.4](#)

19 **47.5 Parallel Region Handle Routines**
20 **47.5.1 ompd_get_curr_parallel_handle Routine**

21 Name: <code>ompd_get_curr_parallel_handle</code> Category: <code>function</code>	Properties: <code>C-only</code> , <code>OMPD</code>
--	---

22 **Return Type and Arguments**

Name	Type	Properties
<return type>	<code>rc</code>	<i>default</i>
<i>thread_handle</i>	<code>thread_handle</code>	<code>opaque</code> , <code>pointer</code>
<i>parallel_handle</i>	<code>parallel_handle</code>	<code>opaque</code> , <code>pointer-to-pointer</code>

24 **Prototypes**

C

```
25       ompd_rc_t ompd_get_curr_parallel_handle(  
26           ompd_thread_handle_t *thread_handle,  
27           ompd_parallel_handle_t **parallel_handle);
```

C

Semantics

The `ompd_get_curr_parallel_handle` routine enables a `tool` to obtain a pointer to the `parallel handle` for the innermost `parallel region` that is associated with an `OpenMP thread`. This routine 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`.

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.

In addition to the return codes permitted for all `OMPD routines`, this routine returns `ompd_rc_unavailable` if the `thread` is not currently part of a `team`.

Cross References

- `ompd_rel_parallel_handle` Routine, see [Section 47.8.2](#)
- `OMPD parallel_handle` Type, see [Section 45.18.2](#)
- `OMPD rc` Type, see [Section 45.9](#)
- `OMPD thread_handle` Type, see [Section 45.18.4](#)

47.5.2 ompd_get_enclosing_parallel_handle Routine

Name: <code>ompd_get_enclosing_parallel_handle</code> Category: function	Properties: C-only , OMPD
--	--

Return Type and Arguments

Name	Type	Properties
<return type>	<code>rc</code>	default
<i>parallel_handle</i>	<code>parallel_handle</code>	opaque , pointer
<i>enclosing_parallel_handle</i>	<code>parallel_handle</code>	opaque , pointer-to-pointer

Prototypes

C

```
ompd_rc_t ompd_get_enclosing_parallel_handle(  
    ompd_parallel_handle_t *parallel_handle,  
    ompd_parallel_handle_t **enclosing_parallel_handle);
```

C

1 **Semantics**

2 The `ompd_get_enclosing_parallel_handle` routine enables a `tool` to obtain a pointer to
3 the `parallel handle` for the `parallel region` that encloses the `parallel region` that `parallel_handle`
4 specifies. This routine yields meaningful results only if at least one `thread` in the `team` that is
5 executing the `parallel region` is stopped. A pointer to the `parallel handle` for the enclosing `region` is
6 returned in the location to which `enclosing_parallel_handle` points. After a call to this routine, the
7 `tool` owns the `handle`; the `tool` must release the `handle` with `ompd_rel_parallel_handle`
8 when it is no longer required. The `parallel_handle` argument is an opaque `handle` for a `parallel`
9 `region` that selects the `parallel region` on which to operate.

10 In addition to the return codes permitted for all OMPD routines, this routine returns
11 `ompd_rc_unavailable` if no enclosing `parallel region` exists.

12 **Cross References**

- 13 • `ompd_rel_parallel_handle` Routine, see [Section 47.8.2](#)
14 • OMPD `parallel_handle` Type, see [Section 45.18.2](#)
15 • OMPD `rc` Type, see [Section 45.9](#)

16 **47.5.3 ompd_get_task_parallel_handle Routine**

17

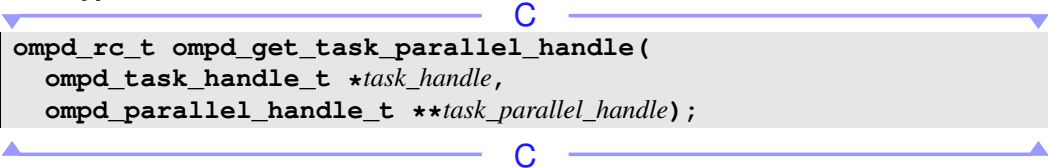
Name: <code>ompd_get_task_parallel_handle</code> Category: <code>function</code>	Properties: <code>C-only</code> , OMPD
---	--

18 **Return Type and Arguments**

19

Name	Type	Properties
<code><return type></code>	<code>rc</code>	<code>default</code>
<code>task_handle</code>	<code>task_handle</code>	<code>pointer</code>
<code>task_parallel_handle</code>	<code>parallel_handle</code>	<code>pointer-to-pointer</code>

20 **Prototypes**

21  `C`
22

```
ompd_rc_t ompd_get_task_parallel_handle(  
23           ompd_task_handle_t *task_handle,  
24           ompd_parallel_handle_t **task_parallel_handle);
```

24 **Semantics**

25 The `ompd_get_task_parallel_handle` routine enables a `tool` to obtain a pointer to the
26 `parallel handle` for the `parallel region` that encloses the `task region` that `task_handle` specifies. This
27 routine yields meaningful results only if at least one `thread` in the `team` that is executing the `parallel`
28 `region` is stopped. A pointer to the `parallel handle` is returned in the location to which
29 `task_parallel_handle` points. The `tool` owns that `parallel handle`, which it must release with
30 `ompd_rel_parallel_handle`.

Cross References

- `ompd_rel_parallel_handle` Routine, see [Section 47.8.2](#)
- OMPD `parallel_handle` Type, see [Section 45.18.2](#)
- OMPD `rc` Type, see [Section 45.9](#)
- OMPD `task_handle` Type, see [Section 45.18.3](#)

47.6 Task Handle Routines

47.6.1 `ompd_get_curr_task_handle` Routine

Name: <code>ompd_get_curr_task_handle</code> Category: function	Properties: C-only, OMPD
--	--------------------------

Return Type and Arguments

Name	Type	Properties
<return type>	rc	default
<i>thread_handle</i>	thread_handle	opaque, pointer
<i>task_handle</i>	task_handle	opaque, pointer-to-pointer

Prototypes

C

```
ompd_rc_t ompd_get_curr_task_handle(  
    ompd_thread_handle_t *thread_handle,  
    ompd_task_handle_t **task_handle);
```

C

Semantics

The `ompd_get_curr_task_handle` routine obtains a pointer to the task handle for the current task region that is associated with an OpenMP thread. This routine 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`. 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. In addition to the return codes permitted for all OMPD routines, this routine returns `ompd_rc_unavailable` if the thread is currently not executing a task.

Cross References

- `ompd_rel_task_handle` Routine, see [Section 47.8.3](#)
- OMPD `rc` Type, see [Section 45.9](#)
- OMPD `task_handle` Type, see [Section 45.18.3](#)
- OMPD `thread_handle` Type, see [Section 45.18.4](#)

47.6.2 `ompd_get_generating_task_handle` Routine

Name: <code>ompd_get_generating_task_handle</code> Category: function	Properties: C-only , OMPD
---	---

Return Type and Arguments

Name	Type	Properties
<i><return type></i>	<code>rc</code>	default
<i>task_handle</i>	<code>task_handle</code>	pointer
<i>generating_task_handle</i>	<code>task_handle</code>	pointer-to-pointer

Prototypes

C

```
ompd_rc_t ompd_get_generating_task_handle(  
    ompd_task_handle_t *task_handle,  
    ompd_task_handle_t **generating_task_handle);
```

C

Semantics

The `ompd_get_generating_task_handle` routine obtains a pointer to the `task handle` of the `generating task region`. The `generating task` is the `task` that was active when the `task` specified by `task_handle` was created. This routine 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`. On return, the `generating_task_handle` argument points to a location that points to a `handle` for the `generating task`. In addition to the return codes permitted for all OMPD routines, this routine returns `ompd_rc_unavailable` if no `generating task region` exists.

Cross References

- `ompd_rel_task_handle` Routine, see [Section 47.8.3](#)
- OMPD `rc` Type, see [Section 45.9](#)
- OMPD `task_handle` Type, see [Section 45.18.3](#)

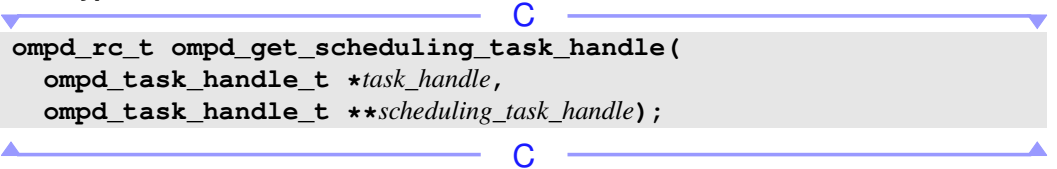
47.6.3 ompd_get_scheduling_task_handle Routine

Name: ompd_get_scheduling_task_handle Category: function	Properties: C-only , OMPD
--	---

Return Type and Arguments

Name	Type	Properties
<return type>	rc	default
task_handle	task_handle	pointer
scheduling_task_handle	task_handle	pointer-to-pointer

Prototypes

 C	
<pre>ompd_rc_t ompd_get_scheduling_task_handle(ompd_task_handle_t *task_handle, ompd_task_handle_t **scheduling_task_handle);</pre>	
C	

Semantics

The [ompd_get_scheduling_task_handle](#) routine 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 [routine](#) yields meaningful results only if the [thread](#) that is executing the [task](#) that *task_handle* specifies is stopped while executing the [task](#). 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](#). This [task handle](#) must be released with [ompd_rel_task_handle](#). In addition to the return codes permitted for all [OMPD routines](#), this [routine](#) returns [ompd_rc_unavailable](#) if no scheduling [task](#) exists.

Cross References

- [ompd_rel_task_handle](#) Routine, see [Section 47.8.3](#)
- [OMPD rc](#) Type, see [Section 45.9](#)
- [OMPD task_handle](#) Type, see [Section 45.18.3](#)

47.6.4 ompd_get_task_in_parallel Routine

Name: ompd_get_task_in_parallel Category: function	Properties: C-only , OMPD
---	---

1 **Return Type and Arguments**

Name	Type	Properties
<return type>	rc	default
parallel_handle	parallel_handle	opaque, pointer
thread_num	integer	default
task_handle	task_handle	opaque, pointer-to-pointer

3 **Prototypes**

4 **C**

```
5 ompd_rc_t ompd_get_task_in_parallel(  
6     ompd_parallel_handle_t *parallel_handle, int thread_num,  
7     ompd_task_handle_t **task_handle);
```

8 **C**

7 **Semantics**

8 The **ompd_get_task_in_parallel** routine obtains handles for the implicit tasks that are
9 associated with a parallel region. A successful invocation of **ompd_get_task_in_parallel**
10 returns a pointer to a task handle in the location to which *task_handle* points. This routine yields
11 meaningful results only if all OpenMP threads in the parallel region are stopped. The
12 *parallel_handle* argument is an opaque handle that selects the parallel region on which to operate.
13 The *thread_num* argument selects the implicit task of the team to be returned. The *thread_num*
14 argument is equal to the *thread-num-var* ICV value of the selected implicit task. This routine
15 returns **ompd_rc_bad_input** if the *thread_num* argument is greater than or equal to the
16 *team-size-var* ICV or negative.

17 **Cross References**

- 18 • **ompd_get_icv_from_scope** Routine, see [Section 47.11.2](#)
- 19 • OMPD **parallel_handle** Type, see [Section 45.18.2](#)
- 20 • OMPD **rc** Type, see [Section 45.9](#)
- 21 • OMPD **task_handle** Type, see [Section 45.18.3](#)

22 **47.6.5 ompd_get_task_function Routine**

Name: ompd_get_task_function	Properties: C-only, OMPD
Category: function	

24 **Return Type and Arguments**

Name	Type	Properties
<return type>	rc	default
task_handle	task_handle	opaque, pointer
entry_point	address	pointer

1 **Prototypes**

C

```
2       ompd_rc_t ompd_get_task_function(ompd_task_handle_t *task_handle,  
3       ompd_address_t *entry_point) ;
```

C

4 **Semantics**

5 The **ompd_get_task_function** routine returns the entry point of the code that corresponds to
6 the body of code that the **task** executes. This routine returns meaningful results only if the **thread**
7 that is executing the **task** that *task_handle* specifies is stopped while executing the **task**. That
8 argument is an opaque **handle** that selects the **task** on which to operate. On return, the *entry_point*
9 argument is set to an address that describes the beginning of application code that executes the **task**
10 **region**.

11 **Cross References**

- 12 • OMPD **address** Type, see [Section 45.2](#)
- 13 • OMPD **rc** Type, see [Section 45.9](#)
- 14 • OMPD **task_handle** Type, see [Section 45.18.3](#)

15 **47.6.6 ompd_get_task_frame Routine**

16 Name: ompd_get_task_frame	Properties: C-only, OMPD
17 Category: function	

17 **Return Type and Arguments**

Name	Type	Properties
<return type>	rc	default
<i>task_handle</i>	task_handle	pointer
<i>exit_frame</i>	frame_info	pointer
<i>enter_frame</i>	frame_info	pointer

19 **Prototypes**

C

```
20      ompd_rc_t ompd_get_task_frame(ompd_task_handle_t *task_handle,  
21      ompd_frame_info_t *exit_frame, ompd_frame_info_t *enter_frame) ;
```

C

22 **Semantics**

23 The **ompd_get_task_frame** routine extracts the **frame** pointers of a **task**. An OpenMP
24 implementation maintains an object of **frame OMP** type for every **implicit task** and **explicit task**.
25 The **ompd_get_task_frame** routine extracts the **enter_frame** and **exit_frame** fields of
26 the **frame** object of the **task** that *task_handle* identifies. This routine yields meaningful results only
27 if the **thread** that is executing the **task** that *task_handle* specifies is stopped while executing the **task**.

On return, the `exit_frame` argument points to a `frame_info` object that has the `frame` information with the same semantics as the `exit_frame` field in the `frame` object that is associated with the specified `task`. On return, the `enter_frame` argument points to a `frame_info` object that has the `frame` information with the same semantics as the `enter_frame` field in the `frame` object that is associated with the specified `task`.

Cross References

- OMPD `address` Type, see [Section 45.2](#)
- OMPT `frame` Type, see [Section 39.15](#)
- OMPD `frame_info` Type, see [Section 45.7](#)
- OMPD `rc` Type, see [Section 45.9](#)
- OMPD `task_handle` Type, see [Section 45.18.3](#)

47.7 Handle Comparing Routines

This section describes [handle-comparing routines](#), which are [routines](#) that have the [handle-comparing property](#) and, thus, enable the comparison of two [handles](#). The internal structure of [handles](#) is opaque to [tools](#). While [tools](#) can easily compare pointers to [handles](#), they cannot determine whether [handles](#) at two different addresses refer to the same underlying context and instead must use a [handle-comparing routine](#).

On success, a [handle-comparing routine](#) returns, in the location to which its `cmp_value` argument points, a signed integer value that indicates how the underlying contexts compare. A value less than, equal to, or greater than 0 indicates that the context to which `<handle-type>_handle_1` corresponds is, respectively, less than, equal to, or greater than that to which `<handle-type>_handle_2` corresponds. The `<handle-type>_handle_1` and `<handle-type>_handle_2` arguments are [handles](#) that correspond to the type of [handle](#) that the [routine](#) compares. In each [handle-comparing routine](#), `<handle-type>` is replaced with the name of the type of [handle](#) that the [routine](#) compares. For all types of [handles](#), the means by which two [handles](#) are ordered is [implementation defined](#).

47.7.1 ompd_parallel_handle_compare Routine

Name: <code>ompd_parallel_handle_compare</code> Category: function	Properties: C-only , handle-comparing , OMPD
---	---

Return Type and Arguments

Name	Type	Properties
<code><return type></code>	<code>rc</code>	default
<code>parallel_handle_1</code>	<code>parallel_handle</code>	opaque , pointer
<code>parallel_handle_2</code>	<code>parallel_handle</code>	opaque , pointer
<code>cmp_value</code>	<code>integer</code>	pointer

1 **Prototypes**

C

```
2       ompd_rc_t ompd_parallel_handle_compare(  
3           ompd_parallel_handle_t *parallel_handle_1,  
4           ompd_parallel_handle_t *parallel_handle_2, int *cmp_value);
```

C

5 **Semantics**

6 The **ompd_parallel_handle_compare** routine compares two **parallel handles**. The
7 *parallel_handle_1* and *parallel_handle_2* arguments are **parallel handles** that correspond to **parallel**
8 **regions**.

9 **Cross References**

- 10 • OMPD **parallel_handle** Type, see [Section 45.18.2](#)
- 11 • OMPD **rc** Type, see [Section 45.9](#)

12 **47.7.2 ompd_task_handle_compare Routine**

13 Name: ompd_task_handle_compare	Properties: C-only, handle-comparing,
Category: function	OMPD

14 **Return Type and Arguments**

Name	Type	Properties
<return type>	rc	default
<i>task_handle_1</i>	task_handle	opaque , pointer
<i>task_handle_2</i>	task_handle	opaque , pointer
<i>cmp_value</i>	integer	pointer

16 **Prototypes**

C

```
17       ompd_rc_t ompd_task_handle_compare(  
18           ompd_task_handle_t *task_handle_1,  
19           ompd_task_handle_t *task_handle_2, int *cmp_value);
```

C

20 **Semantics**

21 The **ompd_task_handle_compare** routine compares two **task handles**. The *task_handle_1*
22 and *task_handle_2* arguments are **task handles** that correspond to **tasks**.

23 **Cross References**

- 24 • OMPD **rc** Type, see [Section 45.9](#)
- 25 • OMPD **task_handle** Type, see [Section 45.18.3](#)

47.7.3 ompd_thread_handle_compare Routine

Name: <code>ompd_thread_handle_compare</code> Category: function	Properties: C-only , handle-comparing , OMPD
---	--

Return Type and Arguments

Name	Type	Properties
<i><return type></i>	<code>rc</code>	default
<i>thread_handle_1</i>	<code>thread_handle</code>	opaque , pointer
<i>thread_handle_2</i>	<code>thread_handle</code>	opaque , pointer
<i>cmp_value</i>	<code>integer</code>	pointer

Prototypes

C

```
ompd_rc_t ompd_thread_handle_compare(  
    ompd_thread_handle_t *thread_handle_1,  
    ompd_thread_handle_t *thread_handle_2, int *cmp_value);
```

C

Semantics

The `ompd_thread_handle_compare` routine compares two [native thread handles](#). The *thread_handle_1* and *thread_handle_2* arguments are [native thread handles](#) that correspond to [native threads](#).

Cross References

- OMPD `rc` Type, see [Section 45.9](#)
- OMPD `thread_handle` Type, see [Section 45.18.4](#)

47.8 Handle Releasing Routines

This section describes [handle-releasing routines](#), which are [routines](#) that have the [handle-releasing property](#) and, thus, release a [handle](#) owned by a [tool](#). When a [tool](#) finishes with a [handle](#) that a *handle* argument identifies, it should release it with the corresponding [handle-releasing routine](#) so the [OMPD library](#) can release any resources that it has related to the corresponding context.

Restrictions

Restrictions to [handle-releasing routines](#) are as follows:

- A context must not be used after its corresponding [handle](#) is released.

47.8.1 ompd_rel_address_space_handle Routine

Name: <code>ompd_rel_address_space_handle</code>	Properties: C-only, handle-releasing, OMPD
Category: function	

Return Type and Arguments

Name	Type	Properties
<i><return type></i>	<code>rc</code>	<i>default</i>
<i>handle</i>	<code>address_space_handle</code>	opaque , pointer

Prototypes

<code>ompd_rc_t ompd_rel_address_space_handle(ompd_address_space_handle_t *handle);</code>

Semantics

A [tool](#) calls `ompd_rel_address_space_handle` to release an [address space handle](#).

Cross References

- OMPD `address_space_handle` Type, see [Section 45.18.1](#)
- OMPD `rc` Type, see [Section 45.9](#)

47.8.2 ompd_rel_parallel_handle Routine

Name: <code>ompd_rel_parallel_handle</code>	Properties: C-only, handle-releasing, OMPD
Category: function	

Return Type and Arguments

Name	Type	Properties
<i><return type></i>	<code>rc</code>	<i>default</i>
<i>parallel_handle</i>	<code>parallel_handle</code>	opaque , pointer

Prototypes

<code>ompd_rc_t ompd_rel_parallel_handle(ompd_parallel_handle_t *parallel_handle);</code>
--

Semantics

A [tool](#) calls `ompd_rel_parallel_handle` to release a [parallel handle](#).

Cross References

- OMPD `parallel_handle` Type, see [Section 45.18.2](#)
- OMPD `rc` Type, see [Section 45.9](#)

47.8.3 `ompd_rel_task_handle` Routine

Name: <code>ompd_rel_task_handle</code> Category: function	Properties: C-only , handle-releasing , OMPD
---	---

Return Type and Arguments

Name	Type	Properties
<i><return type></i>	<code>rc</code>	default
<i>task_handle</i>	<code>task_handle</code>	opaque , pointer

Prototypes

C

```
ompd_rc_t ompd_rel_task_handle(ompd_task_handle_t *task_handle);
```

C

Semantics

A [tool](#) calls `ompd_rel_task_handle` to release a [task handle](#).

Cross References

- OMPD `rc` Type, see [Section 45.9](#)
- OMPD `task_handle` Type, see [Section 45.18.3](#)

47.8.4 `ompd_rel_thread_handle` Routine

Name: <code>ompd_rel_thread_handle</code> Category: function	Properties: C-only , handle-releasing , OMPD
---	---

Return Type and Arguments

Name	Type	Properties
<i><return type></i>	<code>rc</code>	default
<i>thread_handle</i>	<code>thread_handle</code>	opaque , pointer

Prototypes

C

```
ompd_rc_t ompd_rel_thread_handle(  
    ompd_thread_handle_t *thread_handle);
```

C

Semantics

A **tool** calls **ompd_rel_thread_handle** to release a **native thread handle**.

Cross References

- OMPD **rc** Type, see [Section 45.9](#)
- OMPD **thread_handle** Type, see [Section 45.18.4](#)

47.9 Querying Thread States

47.9.1 ompd_enumerate_states Routine

Name: <code>ompd_enumerate_states</code> Category: function	Properties: C-only , OMPD
--	--

Return Type and Arguments

Name	Type	Properties
<i><return type></i>	<code>rc</code>	default
<i>address_space_handle</i>	<code>address_space_handle</code>	opaque , pointer
<i>current_state</i>	<code>word</code>	default
<i>next_state</i>	<code>word</code>	pointer
<i>next_state_name</i>	<code>const char</code>	intent(out) , pointer-to-pointer
<i>more_enums</i>	<code>word</code>	pointer

Prototypes

C

```
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);
```

C

Semantics

An OpenMP implementation may support only a subset of the states that the **state** OMPT type defines. In addition, an OpenMP implementation may support **implementation defined** states. The **ompd_enumerate_states** routine enumerates the **thread states** that an OpenMP implementation supports.

When the *current_state* argument is a **thread state** that an OpenMP implementation supports, the **routine** 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 **tool** owns the *next_state_name* string. The **OMPD library** allocates storage for the string with the **alloc_memory** callback that the **tool** provides. The **tool** is responsible for releasing the storage.

1 On return, the location to which the *more_enums* argument points has the value 1 whenever one or
2 more states are left in the enumeration. On return, the location to which the *more_enums* argument
3 points has the value 0 when *current_state* is the last state in the enumeration.

4 The *address_space_handle* argument identifies the [address space](#). The *current_state* argument must
5 be a [thread state](#) that the OpenMP implementation supports. To begin enumerating the supported
6 states, a [tool](#) should pass [ompt_state_undefined](#) as the value of *current_state*. Subsequent
7 calls to [ompd_enumerate_states](#) by the [tool](#) should pass the value that the [routine](#) returned in
8 the *next_state* argument. This [routine](#) returns [ompd_rc_bad_input](#) if an unknown value is
9 provided in *current_state*.

10 **Cross References**

- 11 • OMPD [address_space_handle](#) Type, see [Section 45.18.1](#)
- 12 • OMPD [rc](#) Type, see [Section 45.9](#)
- 13 • OMPT [state](#) Type, see [Section 39.31](#)
- 14 • OMPD [word](#) Type, see [Section 45.17](#)

15 **47.9.2 ompd_get_state Routine**

16	Name: ompd_get_state Category: function	Properties: C-only , OMPD
----	--	---

17 **Return Type and Arguments**

18	Name	Type	Properties
	<i><return type></i>	rc	default
	<i>thread_handle</i>	thread_handle	opaque , pointer
	<i>state</i>	word	pointer
	<i>wait_id</i>	wait_id	pointer

19 **Prototypes**

```
20 C  
21 ompd\_rc\_t ompd\_get\_state(ompd\_thread\_handle\_t *thread\_handle,  
    ompd\_word\_t *state, ompd\_wait\_id\_t *wait\_id);  
C
```

22 **Semantics**

23 The [ompd_get_state](#) routine returns the state of an [OpenMP thread](#). This [routine](#) yields
24 meaningful results only if the referenced [thread](#) is stopped. The *thread_handle* argument identifies
25 the [thread](#). The *state* argument represents the state of that [thread](#) as represented by a value that
26 [ompd_enumerate_states](#) returns. On return, if the *wait_id* argument is a [non-null value](#) then
27 it points to a [handle](#) that corresponds to the *wait_id* wait identifier of the [thread](#). If the [thread state](#)
28 is not one of the specified wait states, the value to which *wait_id* points is undefined.

Cross References

- `ompd_enumerate_states` Routine, see [Section 47.9.1](#)
- OMPD `rc` Type, see [Section 45.9](#)
- OMPD `thread_handle` Type, see [Section 45.18.4](#)
- OMPD `wait_id` Type, see [Section 45.16](#)
- OMPD `word` Type, see [Section 45.17](#)

47.10 Display Control Variables

47.10.1 `ompd_get_display_control_vars` Routine

Name: <code>ompd_get_display_control_vars</code> Category: function	Properties: C-only, OMPD
--	--------------------------

Return Type and Arguments

Name	Type	Properties
<i><return type></i>	<code>rc</code>	default
<i>address_space_handle</i>	<code>address_space_handle</code>	opaque , pointer
<i>control_vars</i>	<code>const char * const *</code>	intent(out) , pointer

Prototypes

C

```
ompd_rc_t ompd_get_display_control_vars(  
    ompd_address_space_handle_t *address_space_handle,  
    const char * const * *control_vars);
```

C

Semantics

The `ompd_get_display_control_vars` routine returns a list of OpenMP control variables as a `NULL`-terminated vector of null-terminated strings of name/value pairs. These control variables have user-controllable settings and are important to the operation or performance of an OpenMP runtime system. The control variables that this interface exposes include all [OpenMP environment variables](#), settings that may come from vendor or platform-specific [environment variables](#), and other settings that affect the operation or functioning of an OpenMP runtime. The format of the strings is `NAME '=' VALUE`. `NAME` corresponds to the control variable name, optionally prepended with a bracketed `DEVICE`. `VALUE` corresponds to the value of the control variable.

On return, the [tool](#) owns the vector and the strings. The [OMP library](#) must satisfy the termination constraints; it may use static or dynamic [memory](#) for the vector and/or the strings and is unconstrained in how it arranges them in [memory](#). If it uses dynamic [memory](#) then the [OMP](#)

library must use the `alloc_memory` callback that the `tool` provides. The `tool` must use the `ompd_rel_display_control_vars` routine to release the vector and the strings.

The `address_space_handle` argument identifies the `address space`. On return, the `control_vars` argument points to the vector of display control variables.

Cross References

- OMPD `address_space_handle` Type, see [Section 45.18.1](#)
- `ompd_initialize` Routine, see [Section 47.1.1](#)
- `ompd_rel_display_control_vars` Routine, see [Section 47.10.2](#)
- OMPD `rc` Type, see [Section 45.9](#)

47.10.2 ompd_rel_display_control_vars Routine

Name: <code>ompd_rel_display_control_vars</code> Category: <code>function</code>	Properties: <code>C-only</code> , <code>OMPD</code>
---	---

Return Type and Arguments

Name	Type	Properties
<return type>	<code>rc</code>	<i>default</i>
<code>control_vars</code>	<code>const char * const *</code>	<i>pointer</i>

Prototypes

C

```
ompd_rc_t ompd_rel_display_control_vars(  
    const char * const * control_vars);
```

C

Semantics

After a `tool` calls `ompd_get_display_control_vars`, it owns the vector and strings that it acquires. The `tool` must call `ompd_rel_display_control_vars` to release them. The `control_vars` argument is the vector of display control variables to be released.

Cross References

- `ompd_get_display_control_vars` Routine, see [Section 47.10.1](#)
- OMPD `rc` Type, see [Section 45.9](#)

47.11 Accessing Scope-Specific Information

47.11.1 ompd_enumerate_icvs Routine

Name: <code>ompd_enumerate_icvs</code> Category: <code>function</code>	Properties: C-only, OMPD
---	--------------------------

Return Type and Arguments

Name	Type	Properties
<code><return type></code>	<code>rc</code>	<code>default</code>
<code>handle</code>	<code>address_space_handle</code>	<code>opaque</code> , <code>pointer</code>
<code>current</code>	<code>icv_id</code>	<code>default</code>
<code>next_id</code>	<code>icv_id</code>	<code>pointer</code>
<code>next_icv_name</code>	<code>const char</code>	<code>intent(out)</code> , <code>pointer-to-pointer</code>
<code>next_scope</code>	<code>scope</code>	<code>pointer</code>
<code>more</code>	<code>integer</code>	<code>pointer</code>

Prototypes

```
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

The `ompd_enumerate_icvs` routine enables a `tool` to enumerate the `ICVs` that an OpenMP implementation supports and their related scopes. An OpenMP implementation must support all `ICVs` listed in [Section 3.1](#). An OpenMP implementation may support additional `implementation defined ICVs`. An implementation may store `ICVs` in a different scope than [Section 3.1](#) indicates.

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 `tool` owns the `next_icv_name` string. The `OMPd` library uses the `alloc_memory` 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 `ICVs` are left in the enumeration. On return, that location has the value 0 when `current` is the last `ICV` in the enumeration. 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 `routine` 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 value to which the *next_scope* argument points identifies 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. This routine returns **ompd_rc_bad_input** if an unknown value is provided in *current*.

Cross References

- OMPD **address_space_handle** Type, see [Section 45.18.1](#)
- OMPD **icv_id** Type, see [Section 45.8](#)
- OMPD **rc** Type, see [Section 45.9](#)
- OMPD **scope** Type, see [Section 45.11](#)

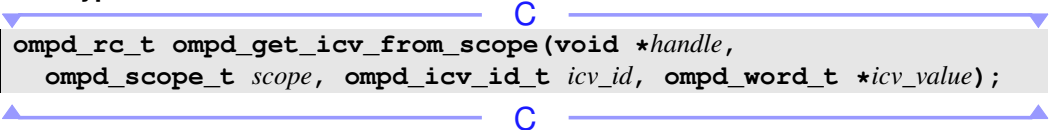

47.11.2 ompd_get_icv_from_scope Routine

Name: ompd_get_icv_from_scope Category: function	Properties: C-only, OMPD
---	---------------------------------

Return Type and Arguments

Name	Type	Properties
<return type>	rc	<i>default</i>
<i>handle</i>	void	opaque, pointer
<i>scope</i>	scope	<i>default</i>
<i>icv_id</i>	icv_id	<i>default</i>
<i>icv_value</i>	word	pointer

Prototypes


<pre>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);</pre>


Summary

The **ompd_get_icv_from_scope** routine returns the value of an **ICV**. 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**.

This routine returns **ompd_rc_bad_input** if an unknown value is provided in *icv_id*. In addition to the return codes permitted for all **OMP** routines, this routine returns **ompd_rc_incomplete** if only the first item of the **ICV** is returned in the integer (e.g., if *nthreads-var* has more than one **list item**). Further, it returns **ompd_rc_incompatible** if the **ICV** cannot be represented as an integer or if the scope of the *handle* matches neither the scope as defined in [Section 45.8](#) nor the scope for *icv_id* as identified by **ompd_enumerate_icvs**.

Cross References

- OMPD Handle Types, see [Section 45.18](#)
- OMPD `icv_id` Type, see [Section 45.8](#)
- `ompd_enumerate_icvs` Routine, see [Section 47.11.1](#)
- OMPD `rc` Type, see [Section 45.9](#)
- OMPD `scope` Type, see [Section 45.11](#)
- OMPD `word` Type, see [Section 45.17](#)

47.11.3 `ompd_get_icv_string_from_scope` Routine

Name: <code>ompd_get_icv_string_from_scope</code> Category: function	Properties: C-only , OMPD
---	---

Return Type and Arguments

Name	Type	Properties
<i><return type></i>	<code>rc</code>	default
<i>handle</i>	<code>void</code>	opaque , pointer
<i>scope</i>	<code>scope</code>	default
<i>icv_id</i>	<code>icv_id</code>	default
<i>icv_string</i>	<code>const char</code>	intent(out) , pointer-to-pointer

Prototypes

C

```
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);
```

C

Semantics

The `ompd_get_icv_string_from_scope` routine returns the value of an `ICV`. 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 `tool` owns the string. The `OMPD library` allocates the string storage with the `alloc_memory` callback that the `tool` provides. The `tool` is responsible for releasing the `memory`.

This routine returns `ompd_rc_bad_input` if an unknown value is provided in *icv_id*. In addition to the return codes permitted for all `OMPD routines`, this routine returns `ompd_rc_incompatible` if the scope of the *handle* does not match the *scope* as defined in

Section 45.8 or if it does not match the scope for *icv_id* as identified by `ompd_enumerate_icvs`.

Cross References

- OMPD Handle Types, see Section 45.18
- OMPD `icv_id` Type, see Section 45.8
- `ompd_enumerate_icvs` Routine, see Section 47.11.1
- OMPD `rc` Type, see Section 45.9
- OMPD `scope` Type, see Section 45.11

47.11.4 ompd_get_tool_data Routine

Name: <code>ompd_get_tool_data</code> Category: <code>function</code>	Properties: C-only, OMPD
--	--------------------------

Return Type and Arguments

Name	Type	Properties
<return type>	<code>rc</code>	<i>default</i>
<i>handle</i>	<code>void</code>	<i>opaque, pointer</i>
<i>scope</i>	<code>scope</code>	<i>default</i>
<i>value</i>	<code>word</code>	<i>pointer</i>
<i>ptr</i>	<code>address</code>	<i>pointer</i>

Prototypes

C

```
ompd_rc_t ompd_get_tool_data(void *handle, ompd_scope_t scope,
                             ompd_word_t *value, ompd_address_t *ptr);
```

C

Semantics

The `ompd_get_tool_data` routine provides access to the OMPT tool data stored for each scope. 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 *data OMPT type* stored for the selected scope. On return, the *ptr* argument points to the *ptr* field of the *data OMPT type* stored for the selected scope. In addition to the return codes permitted for all OMPD routines, this routine returns `ompd_rc_unsupported` if the runtime library does not support OMPT.

Cross References

- OMPD **address** Type, see [Section 45.2](#)
- OMPT **data** Type, see [Section 39.8](#)
- OMPD Handle Types, see [Section 45.18](#)
- OMPD **rc** Type, see [Section 45.9](#)
- OMPD **scope** Type, see [Section 45.11](#)
- OMPD **word** Type, see [Section 45.17](#)

48 OMPD Breakpoint Symbol Names

The OpenMP implementation must define several symbols through which execution must pass when particular [events](#) occur and data collection for OMPD is enabled. A [third-party tool](#) can enable notification of an [event](#) by setting a breakpoint at the address of the symbol.

OMP 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 OMPD symbol conceptually has a function type signature, it may not be a function. It may be a labeled location.

48.1 ompd_bp_thread_begin Breakpoint

Format

```
void ompd_bp_thread_begin(void);
```

Semantics

When starting a [native thread](#) that will be used as an [OpenMP thread](#), the implementation must execute [ompd_bp_thread_begin](#). Thus, 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](#).

48.2 ompd_bp_thread_end Breakpoint

Format

```
void ompd_bp_thread_end(void);
```

Semantics

When terminating an [OpenMP thread](#) or a [native thread](#) that has been used as an [OpenMP thread](#), the implementation must execute [ompd_bp_thread_end](#). Thus, 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 by calling [ompd_rel_thread_handle](#).

Cross References

- [ompd_rel_thread_handle](#) Routine, see [Section 47.8.4](#)

48.3 ompd_bp_device_begin Breakpoint

Format

```
void ompd_bp_device_begin(void);
```

Semantics

When initializing a [device](#) for execution of [target regions](#), the implementation must execute [ompd_bp_device_begin](#). Thus, the OpenMP implementation must execute [ompd_bp_device_begin](#) at every *device-initialize* event. This execution occurs before the work associated with any OpenMP [region](#) executes on the [device](#).

Cross References

- Device Initialization, see [Section 21.4](#)
- **target** Construct, see [Section 21.8](#)

48.4 ompd_bp_device_end Breakpoint

Format

```
void ompd_bp_device_end(void);
```

Semantics

When terminating use of a [device](#), the implementation must execute [ompd_bp_device_end](#). Thus, the OpenMP implementation must execute [ompd_bp_device_end](#) at every *device-finalize* event. This execution occurs after the [device](#) executes all OpenMP [regions](#). After execution of [ompd_bp_device_end](#), any *address_space_handle* that was acquired for this [device](#) is invalid and should be released by calling [ompd_rel_address_space_handle](#).

Cross References

- Device Initialization, see [Section 21.4](#)
- [ompd_rel_address_space_handle](#) Routine, see [Section 47.8.1](#)

48.5 ompd_bp_parallel_begin Breakpoint

Format

```
void ompd_bp_parallel_begin(void);
```

Semantics

Before starting execution of a **parallel region**, the implementation must execute **ompd_bp_parallel_begin**. Thus, the OpenMP implementation must execute **ompd_bp_parallel_begin** at every *parallel-begin event*. When the implementation reaches **ompd_bp_parallel_begin**, the **binding region** for **ompd_get_curr_parallel_handle** is the **parallel region** that is beginning and the **binding task set** for **ompd_get_curr_task_handle** is the **encountering task** for the **parallel construct**.

Cross References

- **ompd_get_curr_parallel_handle** Routine, see [Section 47.5.1](#)
- **ompd_get_curr_task_handle** Routine, see [Section 47.6.1](#)
- **parallel** Construct, see [Section 18.1](#)

48.6 ompd_bp_parallel_end Breakpoint

Format

```
void ompd_bp_parallel_end(void);
```

Semantics

After finishing execution of a **parallel region**, the implementation must execute **ompd_bp_parallel_end**. Thus, the OpenMP implementation must execute **ompd_bp_parallel_end** at every *parallel-end event*. When the implementation reaches **ompd_bp_parallel_end**, the **binding region** for **ompd_get_curr_parallel_handle** is the **parallel region** that is ending and the **binding task set** for **ompd_get_curr_task_handle** is the **encountering task** for 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 by calling **ompd_rel_parallel_handle**.

Cross References

- **ompd_get_curr_parallel_handle** Routine, see [Section 47.5.1](#)
- **ompd_get_curr_task_handle** Routine, see [Section 47.6.1](#)
- **ompd_rel_parallel_handle** Routine, see [Section 47.8.2](#)
- **parallel** Construct, see [Section 18.1](#)

48.7 ompd_bp_teams_begin Breakpoint

Format

```
void ompd_bp_teams_begin(void);
```

Semantics

Before starting execution of a **teams** region, the implementation must execute **ompd_bp_teams_begin**. Thus, the OpenMP implementation must execute **ompd_bp_teams_begin** at every *teams-begin* event. When the implementation reaches **ompd_bp_teams_begin**, the **binding** region for **ompd_get_curr_parallel_handle** is the **teams** region that is beginning and the **binding** task set for **ompd_get_curr_task_handle** is the **encountering** task for the **teams** construct.

Cross References

- **ompd_get_curr_parallel_handle** Routine, see [Section 47.5.1](#)
- **ompd_get_curr_task_handle** Routine, see [Section 47.6.1](#)
- **teams** Construct, see [Section 18.2](#)

48.8 ompd_bp_teams_end Breakpoint

Format

```
void ompd_bp_teams_end(void);
```

Semantics

After finishing execution of a **teams** region, the implementation must execute **ompd_bp_teams_end**. Thus, the OpenMP implementation must execute **ompd_bp_teams_end** at every *teams-end* event. When the implementation reaches **ompd_bp_teams_end**, the **binding** region for **ompd_get_curr_parallel_handle** is the **teams** region that is ending and the **binding** task set for **ompd_get_curr_task_handle** is the **encountering** task for 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 by calling **ompd_rel_parallel_handle**.

Cross References

- **ompd_get_curr_parallel_handle** Routine, see [Section 47.5.1](#)
- **ompd_get_curr_task_handle** Routine, see [Section 47.6.1](#)
- **ompd_rel_parallel_handle** Routine, see [Section 47.8.2](#)
- **teams** Construct, see [Section 18.2](#)

48.9 ompd_bp_task_begin Breakpoint

Format

```
void ompd_bp_task_begin(void);
```

Semantics

Before starting execution of a [task region](#), the implementation must execute [ompd_bp_task_begin](#). Thus, 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. When the implementation reaches [ompd_bp_task_begin](#), the [binding task set](#) for [ompd_get_curr_task_handle](#) is the [task](#) that is scheduled to execute.

Cross References

- [ompd_get_curr_task_handle](#) Routine, see [Section 47.6.1](#)

48.10 ompd_bp_task_end Breakpoint

Format

```
void ompd_bp_task_end(void);
```

Semantics

After finishing execution of a [task region](#), the implementation must execute [ompd_bp_task_end](#). Thus, the OpenMP implementation must execute [ompd_bp_task_end](#) immediately after completion of a [structured block](#) that is associated with a non-merged task. When the implementation reaches [ompd_bp_task_end](#), the [binding task set](#) 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 by calling [ompd_rel_task_handle](#).

Cross References

- [ompd_get_curr_task_handle](#) Routine, see [Section 47.6.1](#)
- [ompd_rel_task_handle](#) Routine, see [Section 47.8.3](#)

48.11 ompd_bp_target_begin Breakpoint

Format

```
void ompd_bp_target_begin(void);
```

Semantics

Before starting execution of a **target region**, the implementation must execute **ompd_bp_target_begin**. Thus, 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**. When the implementation reaches **ompd_bp_target_begin**, the *binding region* for **ompd_get_curr_parallel_handle** is the **target region** that is beginning and the *binding task set* for **ompd_get_curr_task_handle** is the *initial task* on the *device*.

Cross References

- **ompd_get_curr_parallel_handle** Routine, see [Section 47.5.1](#)
- **ompd_get_curr_task_handle** Routine, see [Section 47.6.1](#)
- **target** Construct, see [Section 21.8](#)

48.12 ompd_bp_target_end Breakpoint

Format

```
void ompd_bp_target_end(void);
```

Semantics

After finishing execution of a **target region**, the implementation must execute **ompd_bp_target_end**. Thus, 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**. When the implementation reaches **ompd_bp_target_end**, the *binding region* for **ompd_get_curr_parallel_handle** is the **target region** that is ending and the *binding task set* 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 by calling **ompd_rel_parallel_handle**.

Cross References

- **ompd_get_curr_parallel_handle** Routine, see [Section 47.5.1](#)
- **ompd_get_curr_task_handle** Routine, see [Section 47.6.1](#)
- **ompd_rel_parallel_handle** Routine, see [Section 47.8.2](#)
- **target** Construct, see [Section 21.8](#)

1

Part VII

2

Appendices

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:

- **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 1.3.1](#)).
- **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 1.3.2](#)).

Chapter 2:

- **Processor:** A hardware unit that is [implementation defined](#) (see [Chapter 2](#)).
- **Device:** An [implementation defined](#) logical execution engine (see [Chapter 2](#)).
- **Device pointer:** An [implementation defined handle](#) that refers to a [device address](#) (see [Chapter 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 [Chapter 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 [Chapter 2](#)).

Chapter 3:

- **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 3.2](#)).

Chapter 4:

- **OMP_DYNAMIC environment variable:** If the value is neither **true** nor **false**, the behavior of the program is **implementation defined** (see Section 4.1.2).
- **OMP_NUM_THREADS environment variable:** If any value of the specified **list** 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 4.1.3).
- **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 4.1.4).
- **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 4.1.5).
- **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 4.1.6).
- **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 4.1.7).
- **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 4.3.1).
- **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 4.3.2).
- **OMP_WAIT_POLICY environment variable:** The details of the **active** and **passive** behaviors are **implementation defined** (see Section 4.3.3).
- **OMP_DISPLAY_AFFINITY environment variable:** For all values of the **environment**

- variable other than **true** or **false**, the display action is **implementation defined** (see Section 4.3.4).
- **OMP_AFFINITY_FORMAT environment variable**: Additional **implementation defined** field types can be added (see Section 4.3.5).
 - **OMP_CANCELLATION environment variable**: If the value is set to neither **true** nor **false**, the behavior of the program is **implementation defined** (see Section 4.3.6).
 - **OMP_TARGET_OFFLOAD environment variable**: The support of **disabled** is **implementation defined** (see Section 4.3.9).
 - **OMP_THREADS_RESERVE environment variable**: If the requested values are greater than **OMP_THREAD_LIMIT**, the behavior of the program is **implementation defined** (see Section 4.3.10).
 - **OMP_TOOL_LIBRARIES environment variable**: Whether the value of the **environment variable** is case sensitive is **implementation defined** (see Section 4.5.2).
 - **OMP_TOOL_VERBOSE_INIT environment variable**: Support for logging to **stdout** or **stderr** is **implementation defined**. Whether the value of the **environment variable** is case sensitive when it is treated as a filename is **implementation defined**. The format and detail of the log is **implementation defined** (see Section 4.5.3).
 - **OMP_DEBUG environment variable**: If the value is neither **disabled** nor **enabled**, the behavior is **implementation defined** (see Section 4.6.1).
 - **OMP_NUM_TEAMS environment variable**: If the value is not a **positive** integer or is greater than the number of **teams** that an implementation can support, the behavior of the program is **implementation defined** (see Section 4.2.1).
 - **OMP_TEAMS_THREAD_LIMIT environment variable**: If the value is not a **positive** integer or is greater than the number of **threads** that an implementation can support, the behavior of the program is **implementation defined** (see Section 4.2.2).

Chapter 5:

C / C++

- A pragma **directive** that uses **omp** as the first processing token is **implementation defined** (see Chapter 5).
- The attribute namespace of an attribute specifier or the optional namespace qualifier within a **sequence** attribute that uses **omp** is **implementation defined** (see Chapter 5).

C / C++

C++

- Whether a **throw** executed inside a **region** that arises from an **exception-aborting directive** results in **runtime error termination** is **implementation defined** (see Chapter 5).

C++

Fortran

- Any **directive** that uses **omx** or **omp_x** in the sentinel is **implementation defined** (see Chapter 5).

Fortran

Chapter 6:

- **Collapsed loops**: The particular integer type used to compute the **iteration count** for the collapsed loop is **implementation defined** (see Section 6.4.3).

Chapter 7:

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 7.1.2).

Fortran

- **is_device_ptr clause**: Support for pointers created outside of the OpenMP **device** memory routines is **implementation defined** (see Section 7.3.6).

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 7.3.8 and Section 7.3.9).

Fortran

Chapter 8:

- None.

Chapter 9:

Fortran

- **threadprivate directive**: If the conditions for values of data in the **threadprivate** memories 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 9.1).
- **enter clause**: Whether a device-specific version of a procedure is created is **implementation defined** when the **clause** is in a different **compilation unit** than the definition of the procedure but both are in the same **translation unit** (see Section 9.4).

Fortran

Chapter 10:

- None.

Chapter 11:

- None.

Chapter 12:

- None.

Chapter 13:

- **Memory spaces:** The actual storage resources that each `memory space` defined in Table 13.1 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 13.1).
- **Memory allocators:** The minimum size for partitioning allocated memory over storage resources is `implementation defined`. The default value for the `omp_atk_pool_size` allocator trait (see Table 13.2) 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 13.3) are `implementation defined` (see Section 13.2).
- **dyn_groupprivate clause:** If the *fallback-mode* of the *fallback-modifier* is specified as `default_mem` and the implementation is unable to instantiate a `dynamic groupprivate` block from the `groupprivate memory space` for an `access group` of tasks, the `memory space` used for that `access group` as a fallback is `implementation defined` (see Section 13.9).

Chapter 14:

- **aligned clause:** If the *alignment modifier* is not specified, the default alignments for SIMD instructions on the target platforms are `implementation defined` (see Section 14.2).

Chapter 15:

- **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 15.1).
- **Metadirectives:** The number of times that each expression of the `context selector` of a `when` clause is evaluated is `implementation defined` (see Section 15.4.1).
- **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 15.6).

- **declare_simd directive:** If a **SIMD** version is created and the **simdlen** clause is not specified, the number of concurrent arguments for the **procedure** is **implementation defined** (see [Section 15.8](#)).
- **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 15.9](#)).

Chapter 16:

- **requires directive:** Support for any feature specified by a **requirement clause** on a **requires directive** is **implementation defined** (see [Section 16.5](#)).

Chapter 17:

- **flatten construct:** When applied to a **non-rectangular loops**, the number of inserted empty iterations is **implementation defined** (see [Section 17.3](#)).
- **flatten construct:** The particular integer type used to compute the **iteration count** for the flattened loop is **implementation defined** (see [Section 17.3](#)).
- **stripe construct:** If a generated **offsetting loop** and a generated **grid loop** are associated with the same **construct**, the **grid loops** may execute additional empty **logical iterations**. The number of empty **logical iterations** is **implementation defined** (see [Section 17.8](#)).
- **tile construct:** If a generated **grid loop** and a generated **tile loop** are associated with the same **construct**, the **tile loops** may execute additional empty **logical iterations**. The number of empty **logical iterations** is **implementation defined** (see [Section 17.9](#)).
- **tile construct:** If the number of **logical iterations** of an **affected loop** is not a multiple of the corresponding number of tiles specified by the **sizes clause** with the **grid sizes-selector**, the number of **logical iterations** of the **tile loop** may vary by one (see [Section 17.9](#)).
- **unroll construct:** If no **clauses** are specified, if and how the loop is unrolled is **implementation defined**. If the **partial clause** is specified without an **unroll-factor** argument then the unroll factor is a **positive** integer that is **implementation defined** (see [Section 17.10](#)).

Chapter 18:

- **Default safesync for non-host devices:** Unless indicated otherwise by a **device_safesync requirement clause**, if the **parallel construct** is encountered on a **non-host device** then the default behavior is as if the **safesync clause** appears on the **directive** with a **width** value that is **implementation defined** (see [Section 18.1](#)).
- **Dynamic adjustment of threads:** Providing the ability to adjust the number of **threads** dynamically is **implementation defined** (see [Section 18.1.1](#)).
- **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 18.1.2).

- **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 ET is not zero, which sets have AT positions and which sets have BT positions is **implementation defined**. Further, the positions assigned to the groups that are assigned sets with BT positions to make the number of positions assigned to each group AT is **implementation defined**. The determination of whether the **thread affinity** request can be fulfilled is **implementation defined**. If the **thread affinity** request cannot be fulfilled, then the **thread affinity** of threads in the **team** is **implementation defined** (see Section 18.1.3).
- **teams construct**: The number of **teams** that are created is **implementation defined**, but it is greater than or equal to the lower bound and less than or equal to the upper bound values of the **num_teams** clause if specified. If the **num_teams** clause is not specified, the number of **teams** is less than or equal to the value of the *ntteams-var* ICV if its value is **positive**. Otherwise it is an **implementation defined positive** value (see Section 18.2).
- **simd construct**: The number of iterations that are executed concurrently at any given time is **implementation defined** (see Section 18.4).
- **simklen clause**: For a given *type* that is specified by the *scaled-modifier*, $VL(type)$ is an **implementation defined** value, which may only be determined at runtime, that corresponds to the number of **SIMD lanes** that can be used to process *type* data elements in parallel.

Chapter 19:

- **single construct**: The method of choosing a **thread** to execute the **structured block** each time the **team** encounters the **construct** is **implementation defined** (see Section 19.1).
- **sections construct**: The method of scheduling the **structured block sequences** among threads in the **team** is **implementation defined** (see Section 19.3).
- **Worksharing-loop construct**: The schedule that is used is **implementation defined** if the **schedule** clause is not specified or if the specified schedule has the kind **auto**. The value of *simd_width* for the **simd** schedule modifier is **implementation defined** (see Section 19.6).
- **distribute construct**: If no **dist_schedule** clause is specified then the schedule for the **distribute** construct is **implementation defined** (see Section 19.7).

Chapter 20:

- **taskloop construct**: The number of **logical iterations** assigned to a **task** created from a **taskloop** construct is **implementation defined**, unless the **grainsize** or **num_tasks** clause is specified (see Section 20.2).

- **taskloop construct:** For **firstprivate variables** of class type, the number of invocations of copy constructors to perform the initialization is **implementation defined** (see [Section 20.2](#)).
- **taskgraph construct:** Whether **foreign tasks** are recorded or not in a **taskgraph record** and the manner in which they are executed during a **replay execution** if they are recorded is **implementation defined** (see [Section 20.3](#)).

Chapter 21:

- **thread_limit clause:** The maximum number of **threads** that participate in executing tasks in the **contention group** that each **team** initiates is **implementation defined** if no **thread_limit clause** is specified on the **construct**. Otherwise, it has the **implementation defined** upper bound of the **teams-thread-limit-var ICV**, if the value of this **ICV** is **positive** (see [Section 21.3](#)).
- **target construct:** If a **device clause** is specified with the **ancestor device-modifier**, whether a **storage block** on the **encountering device** that has no **corresponding storage** on the specified **device** may be mapped is **implementation defined** (see [Section 21.8](#)).

Chapter 22:

- **prefer-type modifier:** The supported **preference specifications** are **implementation defined**, including the supported **foreign runtime identifiers**, which may be non-standard names compatible with the **modifier**. The default **preference specification** when the implementation supports multiple values is **implementation defined** (see [Section 22.1.3](#)).

Chapter 23:

- **atomic construct:** A **compliant implementation** may enforce exclusive access between **atomic regions** that update different **storage locations**. The circumstances under which this occurs are **implementation defined**. If the **storage location** designated by x is not size-aligned (that is, if the byte alignment of x is not a multiple of the size of x), then the behavior of the **atomic region** is **implementation defined** (see [Section 23.8.5](#)).

Chapter 24:

- None.

Chapter 25:

- None.

Chapter 26:

- **Runtime routines:** Routine names that begin with the `ompx` prefix are [implementation defined](#) extensions to the OpenMP Runtime API (see [Chapter 26](#)).

C / C++

- **Runtime library definitions:** The types for the `allocator_handle`, `event_handle`, `interop_fr`, `memspace_handle` and `interop` OpenMP types are [implementation defined](#). The value of the `omp_invalid_device` predefined identifier is [implementation defined](#). The value of the `omp_unassigned_thread` predefined identifier is [implementation defined](#) (see [Chapter 26](#)).

C / C++

Fortran

- **Runtime library definitions:** Whether the deprecated include file `omp_lib.h` or the module `omp_lib` (or both) is provided is [implementation defined](#). Whether the `omp_lib.h` file provides derived-type definitions or those [routines](#) that require an explicit interface is [implementation defined](#). Whether any of the [OpenMP API routines](#) that take an argument are extended with a generic interface so arguments of different **KIND** type can be accommodated is [implementation defined](#). The values of the `omp_default_device` and `omp_invalid_device` predefined identifiers are [implementation defined](#) (see [Chapter 26](#)).

Fortran

- **Routine arguments:** The behavior is [implementation defined](#) if a [routine](#) argument is specified with a value that does not conform to the constraints that are implied by the [properties](#) of the argument (see [Section 26.3](#)).
- **Interoperability objects:** [Implementation defined properties](#) may use [non-negative](#) values for [properties](#) associated with an [interoperability object](#) (see [Section 26.7](#)).

Chapter 27:

- **`omp_set_schedule` routine:** For any [implementation defined schedule types](#), the values and associated meanings of the second argument are [implementation defined](#) (see [Section 27.9](#)).
- **`omp_get_schedule` routine:** The value returned by the second argument is [implementation defined](#) for any [schedule types](#) other than `omp_sched_static`, `omp_sched_dynamic` and `omp_sched_guided` (see [Section 27.10](#)).
- **`omp_get_supported_active_levels` routine:** The number of [active levels](#) supported by the implementation is [implementation defined](#), but must be [positive](#) (see [Section 27.11](#)).

- **omp_set_max_active_levels routine:** If the argument is a negative integer then the behavior is [implementation defined](#). If the argument is less than the *active-levels-var* ICV, the *max-active-levels-var* ICV is set to an [implementation defined](#) value between the value of the argument and the value of *active-levels-var*, inclusive (see [Section 27.12](#)).

Chapter 28:

- **omp_set_num_teams routine:** If the argument does not evaluate to a [positive](#) integer, the behavior of this routine is [implementation defined](#) (see [Section 28.2](#)).
- **omp_set_teams_thread_limit routine:** If the argument is not a [positive](#) integer, the behavior is [implementation defined](#) (see [Section 28.6](#)).

Chapter 29:

- None.

Chapter 30:

- None.

Chapter 31:

- **Rectangular-memory-copying routine:** The maximum number of dimensions supported is [implementation defined](#), but must be at least three (see [Section 31.7](#)).

Chapter 32:

- None.

Chapter 33:

- None.

Chapter 34:

- **Lock routines:** If a [lock](#) contains a [synchronization hint](#), the effect of the hint is [implementation defined](#) (see [Chapter 34](#)).

Chapter 35:

- **omp_get_place_proc_ids routine:** The meaning of the [non-negative](#) numerical identifiers returned by the **omp_get_place_proc_ids** routine is [implementation defined](#). The order of the numerical identifiers returned in the array *ids* is [implementation defined](#) (see [Section 35.4](#)).
- **omp_set_affinity_format routine:** 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 and the effect of this routine are [implementation defined](#) (see [Section 35.8](#)).

- **omp_get_affinity_format routine:** 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** (see [Section 35.9](#)).
- **omp_display_affinity routine:** If the *format* argument does not conform to the specified format then the result is **implementation defined** (see [Section 35.10](#)).
- **omp_capture_affinity routine:** If the *format* argument does not conform to the specified format then the result is **implementation defined** (see [Section 35.11](#)).

Chapter 36:

- **omp_display_env routine:** Whether **ICVs** with the same value are combined or displayed in multiple lines is **implementation defined** (see [Section 36.4](#)).

Chapter 37:

- None.

Chapter 38:

- **Tool callbacks:** If a **tool** attempts to register a **callback** not listed in [Table 38.2](#), whether the **registered callback** may never, sometimes or always invoke this **callback** for the associated events is **implementation defined** (see [Section 38.2.4](#)).
- **Device tracing:** Whether a **target device** supports tracing or not is **implementation defined**. If a **target device** does not support tracing, a **NULL** may be supplied for the *lookup* function to the **device** initializer of a **tool** (see [Section 38.2.5](#)).
- **set_trace_ompt and get_record_ompt entry points:** Whether a **device-specific** tracing interface defines this **entry point**, indicating that it can collect traces in **standard trace format**, is **implementation defined**. The kinds of **trace records** available for a **device** is **implementation defined** (see [Section 38.2.5](#)).

Chapter 39:

- **dispatch_chunk OMPT type:** Whether the **chunk** of a **taskloop** region is contiguous is **implementation defined** (see [Section 39.14](#)).
- **record_abstract OMPT type:** The meaning of a **hwid** value for a **device** is **implementation defined** (see [Section 39.24](#)).
- **state OMPT type:** The set of **OMPT thread states** supported is **implementation defined** (see [Section 39.31](#)).

Chapter 40:

- **sync_region_wait callback:** For the *implicit-barrier-wait-begin* and *implicit-barrier-wait-end* **events** at the end of a **parallel region**, whether the *parallel_data*

argument is [NULL](#) or points to the parallel data of the current [parallel region](#) is [implementation defined](#) (see [Section 40.7.5](#)).

Chapter 41:

- **[target_data_op_emi callbacks](#)**: Whether *dev1_addr* or *dev2_addr* points to an intermediate buffer in some operations is [implementation defined](#) (see [Section 41.7](#)).

Chapter 42:

- **[get_place_proc_ids entry point](#)**: The meaning of the numerical identifiers returned is [implementation defined](#). The order of *ids* returned in the array is [implementation defined](#) (see [Section 42.9](#)).
- **[get_partition_place_nums entry point](#)**: The order of the identifiers returned in the *place_nums* array is [implementation defined](#) (see [Section 42.11](#)).
- **[get_proc_id entry point](#)**: The meaning of the numerical identifier returned is [implementation defined](#) (see [Section 42.12](#)).

Chapter 43:

- None.

Chapter 44:

- None.

Chapter 45:

- None.

Chapter 46:

- **[print_string callback](#)**: The value of the *category* argument is [implementation defined](#) (see [Section 46.5](#)).

Chapter 47:

- **[handle-comparing routines](#)**: For all types of [handles](#), the means by which two [handles](#) are ordered is [implementation defined](#) (see [Section 47.7](#)).

Chapter 48:

- None.

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 6.1:

- A form for a **conditional-update-capture structured block** without a capture statement was deprecated for all base languages (see [Section 6.3.3](#)).

The following features were deprecated in Version 6.0:

Fortran

- Omitting the optional **white space** to separate adjacent keywords in the *directive-name* in free source form and fixed source form **directives** is deprecated (see [Section 5.1.1](#) and [Section 5.1.2](#)).

Fortran

- The syntax of the **declare_reduction** directive that specifies the **combiner expression** in the **directive** argument was deprecated (see [Section 8.15](#)).
- The Fortran include file **omp_lib.h** has been deprecated (see [Chapter 26](#)).
- The **target**, **target_data_op**, **target_submit** and **target_map** values of the **callbacks** OMPT types and the associated **trace record** OMPT type names were deprecated (see [Section 39.6](#)).
- 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 the **target_data_op** OMPT type were deprecated (see [Section 39.35](#)).
- The **target_data_op**, **target**, **target_map** and **target_submit** callbacks and the associated **trace record** OMPT type names were deprecated (see [Section 41.7](#), [Section 41.8](#), [Section 41.9](#) and [Section 41.10](#)).

B.2 Version 6.0 to 6.1 Differences

- The **depth** clause was introduced and added to the **fuse** directive to fuse all loops up to a specified depth (see [Section 6.4.7](#) and [Section 17.4](#)).
- The **attach** modifier was added to the **map** clause on **map-entering constructs** to provide more control over whether **pointer attachment** happens for a **list item** that has a **base pointer** or **base referring pointer** (see [Section 11.3.4](#)).
- To support the use of small, fast memory on devices for data with dynamic lifetimes, the **dyn_groupprivate** clause and the **groupprivate-information routines** were added (see [Section 13.9](#) and [Section 33.13](#)).
- Support of the **begin declare_variant** directive was extended to Fortran (see [Section 15.6](#)).
- The **sizes-selector** modifier was added to the **sizes** clause to specify the number of **tiles** (see [Section 17.2](#)).
- The **flatten** directive was added to coalesce multiple loops of a loop nest into a single loop (see [Section 17.3](#)).
- To support the use of scalable **SIMD instructions**, the **scaled-modifier** was added to the **simdlen** clause (see [Section 18.4.3](#)).
- The **omp_default_device** predefined identifier was added and is a **conforming device number** (see [Section 26.1](#)).
- The wording for the **omp_target_is_accessible** routine was updated to clarify when it returns zero (see [Section 31.2.2](#)).

B.3 Version 5.2 to 6.0 Differences

- All features **deprecated** in versions 5.0, 5.1 and 5.2 were removed.
- Full support for C23, C++23, and Fortran 2023 was added (see [Section 1.6](#)).
- Full support of Fortran 2018 was completed (see [Section 1.6](#)).
- The **environment variable** syntax was extended to support initializing **ICVs** for the **host device** and **non-host devices** with a single **environment variable** (see [Section 3.2](#) and [Chapter 4](#)).
- The handling of the **nthreads-var** ICV was updated (see [Section 3.4](#)) and the **nthreads** argument of the **num_threads** clause was changed to a list (see [Section 18.1.2](#)) to support context-specific reservation of inner parallelism.
- **Numeric abstract name** values are now allowed for the **OMP_NUM_THREADS**, **OMP_THREAD_LIMIT** and **OMP_TEAMS_THREAD_LIMIT** environment variables (see [Section 4.1.3](#), [Section 4.1.4](#) and [Section 4.2.2](#)).

1	• The environment variable <code>OMP_PLACES</code> was extended to support an increment between consecutive places when creating a place list from an abstract name (see Section 4.1.6).
2	
3	• The environment variable <code>OMP_AVAILABLE_DEVICES</code> was added and the environment variable <code>OMP_DEFAULT_DEVICE</code> was extended to support device selection by traits (see Section 4.3.7 and Section 4.3.8).
4	
5	
6	• The uid trait was added to the permissible traits in the environment variables <code>OMP_AVAILABLE_DEVICES</code> and <code>OMP_DEFAULT_DEVICE</code> and to the target device trait set (see Section 4.3.7 , Section 4.3.8 and Section 15.2).
7	
8	
9	• The environment variable <code>OMP_THREADS_RESERVE</code> was added to reserve a number of structured threads and free-agent threads (see Section 4.3.10).
10	
	<div>▼ C++ ▲</div>
11	• The <code>decl</code> attribute was added to improve the attribute syntax for declarative directives (see Section 5.1).
12	
	<div>▲ C++ ▼</div>
	<div>▼ C ▲</div>
13	• The OpenMP directive syntax was extended to include C attribute specifiers (see Section 5.1).
14	
	<div>▲ C ▼</div>
	<div>▼ Fortran ▲</div>
15	• Support for directives with the pure property in <code>DO CONCURRENT</code> constructs has been added (see Section 5.1).
16	
	<div>▲ Fortran ▼</div>
17	• To improve consistency in clause format, all inarguable clauses were extended to take an optional argument for which the default value yields equivalent semantics to the existing inarguable semantics (see Section 5.2).
18	
19	
20	• The adjust_args clause was extended to support positional specification of arguments (see Section 5.2.1 and Section 15.6.2).
21	
	<div>▼ Fortran ▲</div>
22	• The definitions of locator list items and assignable OpenMP types were extended to include function references that have data pointer results (see Section 5.2.1).
23	
	<div>▲ Fortran ▼</div>
	<div>▼ C / C++ ▲</div>
24	• The array section definition was extended to permit, where explicitly allowed, omission of the length when the size of the array dimension is not known (see Section 5.2.5).
25	
	<div>▲ C / C++ ▼</div>

- To support greater specificity on **compound constructs**, all **clauses** were extended to accept the *directive-name-modifier*, which identifies the **constituent directives** to which the **clause** applies (see [Section 5.5](#)).
- To allow specification of all **modifiers** of the **init** clause, extensions to the **interop** operation of the **append_args** clause were added (see [Section 5.7](#) and [Section 15.6.3](#)).
- The **init** clause was added to the **depobj** construct, and the **construct** now permits repeatable **init**, **update**, and **destroy** clauses (see [Section 5.7](#) and [Section 23.9.3](#)).
- The syntax that omits the argument to the **destroy** clause for the **depobj** construct was undeprecated (see [Section 5.8](#)).

Fortran

- **Atomic structured blocks** were extended to allow the **BLOCK** construct, pointer assignments and two intrinsic functions for enum and enumeration types (see [Section 6.3.3](#)).
- *conditional-update-statement* was extended to allow more forms and comparisons (see [Section 6.3.3](#)).

Fortran

- The concept of **canonical loop sequences** and the **looprange** clause were defined (see [Section 6.4.2](#) and [Section 6.4.8](#)).

Fortran

- For polymorphic types, restrictions were changed and behavior clarified for **data-sharing attribute clauses** and **data-mapping attribute clauses** (see [Chapter 7](#) and [Chapter 11](#)).

Fortran

- The **default** clause is now allowed on the **target** directive, and, similarly to the **defaultmap** clause, now accepts the *variable-category* modifier (see [Section 7.3.1](#)).
- The semantics of the **use_device_ptr** and **use_device_addr** clauses on a **target_data** construct were altered to imply a reference count update on entry and exit from the **region** for the corresponding objects that they reference in the **device data environment** (see [Section 7.3.7](#) and [Section 7.3.9](#)).
- The *saved* modifier, the **replayable** clause, and the **taskgraph** construct were added to support the recording and efficient **replay execution** of a sequence of **task-generating constructs** (see [Section 7.4](#), [Section 20.3](#), and [Section 20.6](#)).
- Support for **induction operations** was added (see [Chapter 8](#)) through the **induction** clause (see [Section 8.13](#)) and the **declare_induction** directive (see [Section 8.16](#)), which supports **user-defined induction**.
- Support for **reductions** over **private variables** with the **reduction** clause has been added (see [Chapter 8](#)).

C++

- The circumstances under which implicitly declared **reduction identifiers** are supported for **variables** of class type were clarified (see [Section 8.3](#) and [Section 8.6](#)).

C++

- The **scan directive** was extended to accept the **init_complete** clause to enable the identification of an **initialization phase** within the *final-loop-body* of an enclosing **simd construct** or **worksharing-loop construct** (or a **composite construct** that combines them) (see [Section 8.17](#) and [Section 8.17.3](#)).
- The **storage map-type** modifier was added as the preferred *map-type* when the **mapping operation** only allocates or releases storage on the **target device** (see [Section 11.3.1](#)).
- The **ref** modifier was added to the **map clause** to add more control over how the clause affects **list items** that are C++ references or Fortran pointer/allocatable **variables** (see [Section 11.3.3](#) and [Section 11.3](#)).
- The **property** of the *map-type* modifier was changed to *default* so that it can be freely placed and omitted even if other **modifiers** are used (see [Section 11.3](#)).
- The **self map-type-modifier** was added to the **map clause** and the **self implicit-behavior** was added to the **defaultmap clause** to request explicitly that the **corresponding list item** refers to the same object as the **original list item** (see [Section 11.3](#) and [Section 11.4](#)).
- The **map clause** was extended to permit mapping of **assumed-size arrays** (see [Section 11.3](#)).
- The **delete** keyword on the **map clause** was reformulated to be the *delete-modifier* (see [Section 11.3](#)).

Fortran

- The **automap** modifier was added to the **enter clause** to support automatic mapping and unmapping of Fortran allocatable **variables** when allocated and deallocated, respectively (see [Section 9.4](#)).

Fortran

- The **groupprivate** directive was added to specify that **variables** should be privatized with respect to a **contention group** (see [Section 9.2](#)).
- The **local** clause was added to the **declare_target** directive to specify that **variables** should be replicated locally for each **device** (see [Section 9.3](#)).
- The **allocator trait** **omp_atk_part_size** was added to specify the size of the **omp_atv_interleaved** allocator partitions (see [Section 13.2](#)).
- The **omp_atk_pin_device**, **omp_atk_preferred_device** and **omp_atk_target_access** memory allocator traits were defined to provide greater control of **memory** allocations that may be accessible from multiple **devices** (see [Section 13.2](#)).

- 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 13.2](#)).
- The **omp_atv_partitioner** value was added to the possible values of the **omp_atk_partition** allocator trait to allow ad-hoc user partitions (see [Section 13.2](#)).
- The **uses_allocators** clause was extended to permit more than one *clause-argument-specification* (see [Section 13.8](#)).
- The **need_device_addr** modifier was added to the **adjust_args** clause to support adjustment of arguments passed by reference (see [Section 15.6.2](#)).
- The **dispatch** construct was extended with the **interop** clause to support appending arguments specific to a call site (see [Section 15.7](#) and [Section 15.7.1](#)).

C / C++

- A **declare_target** directive that specifies **list items** must now be placed at the same scope as the declaration of those **list items**, and if the **directive** does not specify **list items** then it is treated as **declaration-associated** (see [Section 15.9.1](#)).

C / C++

- The **message** and **severity** clauses were added to the **parallel** directive to support customization of any **error termination** associated with the **directive** (see [Section 16.3](#), [Section 16.4](#), and [Section 18.1](#)).
- The **self_maps requirement** clause was added to require that all **mapping operations** are **self maps** (see [Section 16.5.1.6](#)).
- The **assumption** clause group was extended with the **no_openmp_constructs** clause to support identification of **regions** in which no **constructs** will be encountered (see [Section 16.6.1](#) and [Section 16.6.1.5](#)).
- A restriction for **loop-transforming constructs** was added that the **generated loop** must not be a **doacross-affected loop**, which implies that, in an **unroll** construct with an *unroll-factor* of one, a stand-alone **ordered** directive is now **non-conforming** (see [Chapter 17](#), [Section 17.10](#) and [Section 23.10.1](#)).
- The **apply** clause was added to enable more flexible composition of **loop-transforming constructs** (see [Section 17.1](#)).
- The **sizes** clause was updated to allow non-constant **list items** (see [Section 17.2](#)).
- The **fuse** construct was added to fuse two or more loops in a **canonical loop sequence** (see [Section 17.4](#)).
- The **interchange** construct was added to permute the order of loops in a loop nest (see [Section 17.5](#)).

- The **reverse** construct was added to reverse the iteration order of a loop (see [Section 17.6](#)).
- The **split** loop-transforming construct was added to apply **index-set splitting** to **canonical loop nests** (see [Section 17.7](#)).
- The **stripe** loop-transforming construct was added to apply **striping** to **canonical loop nests** (see [Section 17.8](#)).
- The **tile** construct was extended to allow **grid loops** and **tile loops** to be affected by the same **construct** (see [Section 17.9](#)).
- The *prescriptiveness* modifier was added to the **num_threads** clause and **strict** semantics were defined for the **clause** (see [Section 18.1.2](#)).
- To control which synchronizing **threads** are guaranteed to make progress eventually, the **safesync** clause on the **parallel** construct (see [Section 18.1.5](#)), the **omp_curr_progress_width** identifier (see [Section 26.1](#)) and the **omp_get_max_progress_width** routine were added (see [Section 30.6](#)).
- To make the **loop** construct and other **constructs** that specify the **order** clause with **concurrent** ordering more usable, calls to procedures in the **region** may now contain certain OpenMP **directives** (see [Section 18.3](#)).
- 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 18.3](#)).
- The **constructs** that may be encountered during the execution of a **region** that corresponds to a **construct** on which the **order** clause is specified with **concurrent** ordering, when the corresponding **regions** are not **strictly nested regions**, are no longer restricted (see [Section 18.3](#)).

Fortran

- The **workdistribute** directive was added to support Fortran array expressions in **teams** **constructs** (see [Section 19.5](#)).
- The **loop** construct was extended to allow a **DO CONCURRENT** loop as the **collapse-affected** loop (see [Section 19.8](#)).

Fortran

- The **taskloop** construct now includes the **task_iteration** directive as a **subsidiary directive** so that the **tasks** that it generates can include the semantics of the **affinity** and **depend** clauses (see [Section 20.2](#), [Section 20.2.3](#), [Section 20.10](#) and [Section 23.9.5](#)).
- The **threadset** clause was added to **task-generating constructs** to specify the **binding thread set** of the generated **task** (see [Section 20.8](#)).
- The **priority** clause was added to the **target_enter_data**, **target_exit_data**, **target_data**, **target** and **target_update** directives (see [Section 20.9](#), [Section 21.5](#), [Section 21.6](#), [Section 21.7](#), [Section 21.8](#) and [Section 21.9](#)).

- The **device_type** clause was added to the clauses that may appear on the **target** construct (see Section 21.1 and Section 21.8).
- When the **device** clause is specified with the **ancestor device-modifier** on the **target** construct, the **nowait** clause may now also be specified (see Section 21.2, Section 21.8 and Section 23.6).
- The **target_data** directive description was updated to make it a **composite construct**, to include a **taskgroup** region and to make the clauses that may appear on it reflect its constituent constructs and the **taskgroup** region (see Section 21.7).
- The *prefer-type* modifier of the **init** clause was updated to allow preferences other than foreign runtime identifiers (see Section 22.1.3).
- The *do_not_synchronize* argument for the **nowait** clause (see Section 23.6) and **nogroup** clause (see Section 23.7) 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 23.8.4, Section 23.8.5 and Section 23.8.6).
- The **transparent** clause was added to support multi-generational **task dependence** graphs (see Section 23.9.6).
- The **cancel** construct was extended to complete **tasks** that have not yet been fulfilled through an **event** variable and the **omp_fulfill_event** routine was restricted such that an **event** handle must be fulfilled before execution continues beyond a **barrier** (see Section 24.2 and Section 29.2.1).
- The rules for **compound-directive names** were simplified to be more intuitive and to allow more valid combinations of **immediately nested constructs** (see Section 25.1).
- The **omp_is_free_agent** and **omp_ancestor_is_free_agent** routines were added to test whether the **encountering thread**, or the **ancestor thread**, is a **free-agent thread** (see Section 29.1.4 and Section 29.1.5).
- The **omp_get_device_from_uid** and **omp_get_uid_from_device** routines were added to convert between unique identifiers and **device numbers** of devices (see Section 30.7 and Section 30.8).
- The **omp_get_device_num_teams**, **omp_set_device_num_teams**, **omp_get_device_teams_thread_limit**, and **omp_set_device_teams_thread_limit** routine were added to support getting and setting the *nteam*-var and *teams-thread-limit-var* ICVs for specific **devices** (see Section 30.11, Section 30.12, Section 30.13, and Section 30.14).
- The **omp_target_is_accessible** routine was updated to allow the *ptr* argument to accept any valid *ptr*, removing the restriction for it to be a valid **host device ptr**.
- The **omp_target_memset** and **omp_target_memset_async** routines were added to fill **memory** in a **device data environment** of a **device** (see Section 31.8.1 and Section 31.8.2).

Fortran

- Fortran versions of the runtime [routines](#) to operate on [interoperability objects](#) were added (see [Chapter 32](#)).

Fortran

- New [routines](#) were added to obtain [memory spaces](#) and [memory allocators](#) to allocate remote and shared [memory](#) (see [Chapter 33](#)).
- The [omp_get_memspace_num_resources](#) routine was added to support querying the number of available resources of a [memory space](#) (see [Section 33.2](#)).
- The [omp_get_memspace_pagesize](#) routine was added to obtain the page size supported by a given [memory space](#) (see [Section 33.3](#)).
- The [omp_get_submemspace](#) routine was added to obtain a [memory space](#) with a subset of the original storage resources (see [Section 33.4](#)).
- The [omp_init_mempartitioner](#), [omp_destroy_mempartitioner](#), [omp_init_mempartition](#), [omp_destroy_mempartition](#), [omp_mempartition_set_part](#), [omp_mempartition_get_user_data](#) routines were added to manipulate the [mempartitioner](#) and [mempartition](#) objects (see [Section 33.5](#)).
- The set of [callbacks](#) for which [set_callback](#) must return [ompt_set_always](#) no longer includes the [target_data_op](#), [target](#), [target_map](#) and [target_submit](#) callbacks, which were [deprecated](#) (see [Section 38.2.4](#), [Section 41.7](#), [Section 41.8](#), [Section 41.9](#) and [Section 41.10](#)).
- The more general values [ompt_target_data_transfer](#) and [ompt_target_data_transfer_async](#) were added to the [target_data_op](#) OMPT type and supersede the values [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](#) (see [Section 39.35](#)). The superseded values were [deprecated](#).
- The [get_buffer_limits](#) entry point was added to the [OMPT device](#) tracing interface so that a [first-party tool](#) can obtain an upper limit on the sizes of the trace buffers that it should make available to the implementation (see [Section 43.6](#)).

B.4 Version 5.1 to 5.2 Differences

- Major reorganization and numerous changes were made to improve the quality of the specification of OpenMP syntax and to increase consistency of restrictions and their wording. These changes frequently result in the possible perception of differences to preceding versions

of the OpenMP specification. However, those differences almost always resolve ambiguities, which may nonetheless have implications for existing implementations and programs.

- The *explicit-task-var ICV* replaced the *implicit-task-var ICV*, with the opposite meaning and semantics (see Chapter 3). The `omp_in_explicit_task` routine was added to query if a code region is executed from an explicit task region (see Section 29.1.2).

Fortran

- Expanded the directives that may be encountered in a pure procedure (see Chapter 5) by adding the `pure` property to metadirectives (see Section 15.4.3), assumption directives (see Section 16.6), the `nothing` directive (see Section 16.7), the `error` directive (see Section 16.1) and loop-transforming constructs (see Chapter 17).

Fortran

- For OpenMP directives, the `omp` sentinel and, for implementation defined directives that extend the OpenMP directives, the `ompx` sentinel for C/C++ and free source form Fortran and the `omx` sentinel for fixed source form Fortran (to accommodate character position requirements) were reserved (see Chapter 5). Reserved clause names that begin with the `ompx_` prefix for implementation defined clauses on OpenMP directives (see Chapter 5). Reserved names in the base language that start with the `omp_`, `ompt_`, `ompd_` and `ompx_` prefixes and reserved the `omp`, `ompx`, `ompt` and `ompd` namespaces for the OpenMP runtime API and for implementation defined extensions to that API (see Chapter 5).
- Allowed any clause that can be specified on a paired end directive to be specified on the directive (see Section 5.1), including, in Fortran, the `copyprivate` clause (see Section 10.2) and the `nowait` clause (see Section 23.6).
- Allowed the `if` clause on the `teams` construct (see Section 5.6 and Section 18.2).
- For consistency with the syntax of other definitions of the clause, the syntax of the `destroy` clause on the `depobj` construct with no argument was deprecated (see Section 5.8).
- For consistency with the syntax of other clauses, the syntax of the `linear` clause that specifies its argument and *linear-modifier* as *linear-modifier (list)* was deprecated and the *step* modifier was added for specifying the linear step (see Section 8.14).
- The *minus* (−) operator for reductions was deprecated (see Section 8.6).
- The syntax of modifiers without comma separators in the `map` clause was deprecated (see Section 11.3).
- To support the complete range of user-defined mappers and to improve consistency of `map` clause usage, the `declare_mapper` directive was extended to accept *iterator* modifiers and the `present` *map-type-modifier* (see Section 11.3 and Section 11.5).
- Mapping of a pointer list item was updated such that if a matched candidate is not found in the data environment, `firstprivate` semantics apply and the pointer retains its original value (see Section 11.3).

- The **enter** clause was added as a synonym for the **to** clause on **declare target** directives, and the corresponding **to** clause was **deprecated** to reduce parsing ambiguity (see [Section 9.4](#) and [Section 15.9](#)).

Fortran

- The **allocators** construct was added to support the use of OpenMP **allocators** for **variables** that are allocated by a Fortran **ALLOCATE** statement, and the application of **allocate** directives to an **ALLOCATE** statement was **deprecated** (see [Section 13.7](#)).

Fortran

- To support the full range of **allocators** and to improve consistency with the syntax of other **clauses**, the argument that specified the arguments of the **uses_allocators** clause as a comma-separated list in which each **list item** is a *clause-argument-specification* of the form *allocator[(traits)]* was **deprecated** (see [Section 13.8](#)).
- To improve code clarity and to reduce ambiguity in this specification, the **otherwise** clause was added as a synonym for the **default** clause on **metadirectives** and the corresponding **default** clause syntax was **deprecated** (see [Section 15.4.2](#)).

Fortran

- For consistency with other **constructs** with associated **base-language code**, the **dispatch** construct was extended to allow an optional paired **end directive** to be specified (see [Section 15.7](#)).

Fortran

C / C++

- To improve overall syntax consistency and to reduce redundancy, the delimited form of the **declare_target** directive was **deprecated** (see [Section 15.9.2](#)).

C / C++

- The behavior of the **order** clause with the *concurrent* argument was changed so that it only affects whether a **loop schedule** is **reproducible** if a **modifier** is explicitly specified (see [Section 18.3](#)).
- Support for the **allocate** and **firstprivate** clauses on the **scope** directive was added (see [Section 19.2](#)).
- The **work** OMPT type values for **worksharing-loop** constructs were added (see [Section 19.6](#)).
- To simplify usage, the **map** clause on a **target_enter_data** or **target_exit_data**, **construct** now has a default map type that provides the same behavior as the **to** or **from** map types, respectively (see [Section 21.5](#) and [Section 21.6](#)).
- The **interop** construct was updated to allow the **init** clause to accept an *interop_type* in any position of the **modifier** list (see [Section 22.1](#)).

- The **doacross** clause was added as a synonym for the **depend** clause with the keywords **source** and **sink** as *dependence-type modifiers* and the corresponding **depend** clause syntax was *deprecated* to improve code clarity and to reduce parsing ambiguity. Also, the **omp_cur_iteration** keyword was added to represent a *logical iteration vector* that refers to the current *logical iteration* (see [Section 23.9.7](#)).
- The **omp_pause_stop_tool** value was added to the **pause_resource** OpenMP type (see [Section 26.11.1](#)).

B.5 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 1.6](#)).
- Various changes throughout the specification were made to provide initial support of Fortran 2018 (see [Section 1.6](#)).
- To support *device-specific ICV* settings the *environment variable* syntax was extended to support *device-specific environment variables* (see [Section 3.2](#) and [Chapter 4](#)).
- The **OMP_PLACES** syntax was extended (see [Section 4.1.6](#)).
- The **OMP_NUM_TEAMS** and **OMP_TEAMS_THREAD_LIMIT** environment variables were added to control the number and size of *teams* on the *teams* construct (see [Section 4.2.1](#) and [Section 4.2.2](#)).
- The OpenMP *directive* syntax was extended to include C++ attribute specifiers (see [Section 5.1](#)).
- The **omp_all_memory** reserved locator was added (see [Section 5.2.2](#)), and the **depend** clause was extended to allow its use (see [Section 23.9.5](#)).
- Support for **private** and **firstprivate** as an argument to the **default** clause in C and C++ was added (see [Section 7.3.1](#)).
- The **has_device_addr** clause was added to the **target** construct to allow access to *variables* or *array sections* that already have a *device address* (see [Section 7.3.8](#) and [Section 21.8](#)).
- Support was added so that *iterators* may be defined and used in **map** clauses (see [Section 11.3](#)) or in *data-motion clauses* on a **target_update** directive (see [Section 21.9](#)).
- The **private** and **firstprivate** *data-sharing attribute* specifiers and the **present** argument was added to the **defaultmap** clause (see [Section 11.4](#)).
- Support for the **align** clause on the **allocate** directive and *allocator* and *align* modifiers on the **allocate** clause was added (see [Chapter 13](#)).

- The `target_device` trait set was added to the OpenMP context (see Section 15.1), and the `target_device` selector set was added to context selectors (see Section 15.2).
- For C/C++, the `declare_variant` directives were extended to support elision of preprocessed code and to allow enclosed function definitions to be interpreted as function variants (see Section 15.6).
- 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 function variants (see Section 15.6.4).
- 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 function variants (see Section 15.7).
- Support was added for indirect calls to the `device` version of a `procedure` in `target` regions (see Section 15.9).
- To allow users to control the compilation process and runtime error actions, the `error` directive was added (see Section 16.1).
- Assumption directives were added to allow users to specify invariants (see Section 16.6).
- To support clarity in metadirectives, the `nothing` directive was added (see Section 16.7).
- Loop-transforming constructs were added (see Chapter 17).
- The `masked` construct was added to support restricting execution to a specific thread to replace the deprecated `master` construct (see Section 18.5).
- The `scope` directive was added to support reductions without requiring a `parallel` or `worksharing` region (see Section 19.2).
- The `grainsize` and `num_tasks` clauses for the `taskloop` construct were extended with a `strict prescriptiveness` modifier to ensure a deterministic distribution of logical iterations to tasks (see Section 20.2).
- The `thread_limit` clause was added to the `target` construct to control the upper bound on the number of threads in the created contention group (see Section 21.8).
- The `interop` directive was added to enable portable interoperability with foreign execution contexts (see Section 22.1). Runtime routines that facilitate use of interoperability objects were also added (see Chapter 32).
- The `nowait` clause was added to the `taskwait` directive to support insertion of non-blocking join operations in a task dependence graph (see Section 23.5).
- Specification of the `seq_cst` clause on a `flush` construct was allowed, with the same meaning as a `flush` construct without a list and without a clause (see Section 23.8.1.5 and Section 23.8.6).

- Support was added for compare-and-swap and (for C and C++) minimum and maximum [atomic operations](#) through the [compare clause](#). Support was also added for the specification of the [memory order](#) to apply to a failed [atomic conditional update](#) with the [fail clause](#) (see [Section 23.8.3.2](#) and [Section 23.8.3.3](#)).
- To support inout sets, the [inoutset task-dependence-type modifier](#) was added to the [depend clause](#) (see [Section 23.9.5](#)).
- For the [alloctrait_key](#) OpenMP type, the [omp_atv_serialized](#) value was added and the [omp_atv_default](#) value was changed (see [Section 26.8](#)).
- The [omp_set_num_teams](#) and [omp_set_teams_thread_limit](#) routines were added to control the number of [teams](#) and the size of those [teams](#) on the [teams](#) construct (see [Section 28.2](#) and [Section 28.6](#)). Additionally, the [omp_get_max_teams](#) and [omp_get_teams_thread_limit](#) routines were added to retrieve the values that will be used in the next [teams](#) construct (see [Section 28.4](#) and [Section 28.5](#)).
- The [omp_target_is_accessible](#) routine was added to test whether a [host address](#) is accessible from a given [device](#) (see [Section 31.2.2](#)).
- The [omp_get_mapped_ptr](#) routine was added to support obtaining the [device pointer](#) that is associated with a [host pointer](#) for a given [device](#) (see [Section 31.2.3](#)).
- To support asynchronous [device memory](#) management, [omp_target_memcpy_async](#) and [omp_target_memcpy_rect_async](#) routines were added (see [Section 31.7.3](#) and [Section 31.7.4](#)).
- The [omp_calloc](#), [omp_realloc](#), [omp_aligned_alloc](#) and [omp_aligned_calloc](#) routines were added (see [Chapter 33](#)).
- The [omp_display_env](#) routine was added to provide information about [ICVs](#) and settings of [environment variables](#) (see [Section 36.4](#)).
- The [ompt_scope_beginend](#) value was added to the [scope_endpoint](#) OMPT type to indicate the coincident beginning and end of a scope (see [Section 39.27](#)).
- The [ompt_state_wait_barrier_implementation](#) and [ompt_state_wait_barrier_teams](#) values were added to the [state](#) OMPT type (see [Section 39.31](#)).
- The [ompt_sync_region_barrier_implicit_workshare](#), [ompt_sync_region_barrier_implicit_parallel](#), and [ompt_sync_region_barrier_teams](#) values were added to the [sync_region](#) OMPT type (see [Section 39.33](#)).
- Values for asynchronous data transfers were added to the [target_data_op](#) OMPT type (see [Section 39.35](#)).
- The [error callback](#) was added (see [Section 40.2](#)).

- The `target_data_op_emi`, `target_emi`, `target_map_emi`, and `target_submit_emi` callbacks were added to support external monitoring interfaces (see [Section 41.7](#), [Section 41.8](#), [Section 41.9](#) and [Section 41.10](#)).

B.6 Version 4.5 to 5.0 Differences

- The `memory` model was extended to distinguish different types of `flushes` according to specified `flush properties` (see [Section 1.3.4](#)) and to define a `happens-before order` based on synchronizing `flushes` (see [Section 1.3.5](#)).
- Various changes throughout the specification were made to provide initial support of C11, C++11, C++14, C++17 and Fortran 2008 (see [Section 1.6](#)).
- Full support of Fortran 2003 was completed (see [Section 1.6](#)).
- The `target-offload-var` ICV (see [Chapter 3](#)) and the `OMP_TARGET_OFFLOAD` environment variable (see [Section 4.3.9](#)) were added to support runtime control of the execution of `device constructs`.
- Control over whether `nested parallelism` is enabled or disabled was integrated into the `max-active-levels-var` ICV (see [Section 3.2](#)), the default value of which was made `implementation defined`, unless determined according to the values of the `OMP_NUM_THREADS` (see [Section 4.1.3](#)) or `OMP_PROC_BIND` (see [Section 4.1.7](#)) environment variables.
- The `OMP_DISPLAY_AFFINITY` (see [Section 4.3.4](#)) and `OMP_AFFINITY_FORMAT` (see [Section 4.3.5](#)) environment variables and the `omp_set_affinity_format` (see [Section 35.8](#)), `omp_get_affinity_format` (see [Section 35.9](#)), `omp_display_affinity` (see [Section 35.10](#)), and `omp_capture_affinity` (see [Section 35.11](#)) routines were added to provide OpenMP runtime `thread affinity` information.
- The `omp_set_nested` and `omp_get_nested` routines and the `OMP_NESTED` environment variable were `deprecated`.
- Support for `array shaping` (see [Section 5.2.4](#)) and for `array sections` with non-unit strides in C and C++ (see [Section 5.2.5](#)) was added to facilitate specification of discontinuous storage, and the `target_update` construct (see [Section 21.9](#)) and the `depend` clause (see [Section 23.9.5](#)) were extended to allow the use of `shape-operators` (see [Section 5.2.4](#)).
- The `iterator` modifier (see [Section 5.2.6](#)) was added to support expressions in a list that expand to multiple expressions.
- The `canonical loop nest` form was defined for Fortran and, for all `base languages`, extended to permit `non-rectangular loops` (see [Section 6.4.1](#)).
- The `relational-op` in a `canonical loop nest` for C/C++ was extended to include `!=` (see [Section 6.4.1](#)).

- To support conditional assignment to `lastprivate` variables, the *conditional* modifier was added to the `lastprivate` clause (see Section 7.3.5).
- The semantics of the `use_device_ptr` clause for pointer variables was clarified and the `use_device_addr` clause for using the device address of non-pointer variables inside the `target_data` construct was added (see Section 7.3.7, Section 7.3.9 and Section 21.7).
- The *inscan* modifier for the `reduction` clause (see Section 8.10) and the `scan` directive (see Section 8.17) were added to support inclusive scan and exclusive scan computations.
- To support task reductions, the *task* modifier was added to the `reduction` clause (see Section 8.10), the `task_reduction` clause (see Section 8.11) was added to the `taskgroup` construct (see Section 23.4), and the `in_reduction` clause (see Section 8.12) was added to the `task` (see Section 20.1) and `target` (see Section 21.8) constructs.
- To support `taskloop` reductions, the `reduction` (see Section 8.10) and `in_reduction` (see Section 8.12) clauses were added to the `taskloop` construct (see Section 20.2).
- The description of the `map` clause was modified to clarify the mapping order when multiple *map-type* modifiers are specified for a variable or structure elements of a variable on the same construct. The *close-modifier* was added as a hint for the runtime to allocate memory close to the target device (see Section 11.3).
- The capability to map C/C++ pointer variables and to assign the address of device memory that is mapped by an array section to them was added. Support for mapping of Fortran pointer and allocatable variables, including pointer and allocatable components of variables, was added (see Section 11.3).
- All uses of the `map` clause (see Section 11.3), as well as the `to` and `from` clauses on the `target_update` construct (see Section 21.9) and the `depend` clause on task-generating constructs (see Section 23.9.5) were extended to allow any lvalue expression as a list item for C/C++.
- The `defaultmap` clause (see Section 11.4) was extended to allow specification of the data-mapping attributes or data-sharing attributes for any of the scalar, aggregate, pointer, or allocatable classes on a per-region basis. Additionally, the `none` argument was added to support the requirement that all variables referenced in the construct must be explicitly mapped or privatized.
- The `declare_mapper` directive was added to support mapping of data types with direct and indirect members (see Section 11.5).
- Predefined memory spaces, predefined memory allocators and allocator traits and directives, clauses and routines (see Chapter 13 and Chapter 33) to use them were added to support different kinds of memories.

- **Metadirectives** (see [Section 15.4](#)) and **declare variant directives** (see [Section 15.6](#)) were added to support selection of **directive variants** and **function variants** at a call site, respectively, based on compile-time **traits** of the **enclosing context**.
- Support for nested **declare target directives** was added (see [Section 15.9](#)).
- To reduce programmer effort, implicit **declare target directives** for some **procedures** were added (see [Section 15.9](#) and [Section 21.8](#)).
- The **requires directive** (see [Section 16.5](#)) was added to support applications that require implementation-specific features.
- The **teams construct** (see [Section 18.2](#)) was extended to support execution on the **host device** without an enclosing **target construct** (see [Section 21.8](#)).
- The **loop construct** and the **order clause** with the **concurrent** argument were added to support compiler optimization and parallelization of loops for which **logical iterations** may execute in any order, including concurrently (see [Section 18.3](#) and [Section 19.8](#)).
- The collapse of **affected loops** that are **imperfectly nested loops** was defined for **simd constructs** (see [Section 18.4](#)), **worksharing-loop constructs** (see [Section 19.6](#)), **distribute constructs** (see [Section 19.7](#)) and **taskloop constructs** (see [Section 20.2](#)).
- The **simd construct** (see [Section 18.4](#)) was extended to accept the **if** and **nontemporal clauses** and, with the **concurrent** argument, **order clauses** and to allow the use of **atomic constructs** within it.
- The default *ordering-modifier* for the **schedule clause** on **worksharing-loop constructs** when the *kind* argument is not **static** and the **ordered clause** does not appear on the **construct** was changed to **nonmonotonic** (see [Section 19.6.3](#)).
- The **clauses** that can be specified on the **task construct** (see [Section 20.1](#)) were extended with the **affinity clause** (see [Section 20.10](#)) to support hints that indicate data affinity of **explicit tasks**.
- To support execution of **detachable tasks**, the **detach clause** for the **task construct** (see [Section 20.1](#)) and the **omp_fulfill_event routine** (see [Section 29.2.1](#)) were added.
- The **taskloop construct** (see [Section 20.2](#)) was added to the list of **constructs** that can be canceled by the **cancel constructs** (see [Section 24.2](#)).
- To support **reverse-offload regions**, the *ancestor* modifier was added to the **device clause** for the **target construct** (see [Section 21.2](#) and [Section 21.8](#)).
- The **target_update construct** (see [Section 21.9](#)) was modified to allow **array sections** that specify discontinuous storage.
- The **taskwait construct** was extended to accept the **depend clause** (see [Section 23.5](#) and [Section 23.9.5](#)).

- To support acquire and release semantics with weak memory ordering, the `acq_rel`, `acquire`, and `release` clauses (see [Section 23.8.1](#)) were added to the `atomic` construct (see [Section 23.8.5](#)) and `flush` construct (see [Section 23.8.6](#)), and the memory ordering semantics of `implicit flushes` on various `constructs` and `routines` were clarified (see [Section 23.8.7](#)).
- The `atomic` construct was extended with the `hint` clause (see [Section 23.8.5](#)).
- To support mutually exclusive inout sets, a `mutexinoutset` *task-dependence-type* was added to the `depend` clause (see [Section 23.9.1](#) and [Section 23.9.5](#)).
- The `depend` clause (see [Section 23.9.5](#)) was extended to support *iterator* modifiers and to support `depend` objects that can be created with the new `depobj` construct (see [Section 23.9.3](#)).
- New combined constructs (`master taskloop`, `parallel master`, `parallel master taskloop`, `master taskloop simd` and `parallel master taskloop simd`) (see [Section 25.1](#)) were added.
- Lock hints were renamed to `synchronization hints`, and the old names were `deprecated` (see [Section 26.9.5](#)).
- The `omp_get_supported_active_levels` routine was added to query the number of `active levels` of parallelism supported by the implementation (see [Section 27.11](#)).
- The `omp_get_device_num` routine (see [Section 30.4](#)) was added to support determination of the `device` on which a `thread` is executing.
- The `omp_pause_resource` and `omp_pause_resource_all` routines were added to allow the runtime to relinquish resources used by OpenMP (see [Section 36.2.1](#) and [Section 36.2.2](#)).
- Support for a `first-party tool` interface (see [Chapter 38](#)) was added.
- Support for a `third-party tool` interface (see [Chapter 44](#)) was added.
- Stubs for runtime library `routines` (previously Appendix A) were moved to a separate document.
- Interface declarations (previously Appendix B) were moved to a separate document.

B.7 Version 4.0 to 4.5 Differences

- Support for several features of Fortran 2003 was added (see [Section 1.6](#)).
- The `OMP_MAX_TASK_PRIORITY` environment variable was added to control the maximum `task priority` value allowed (see [Section 4.3.11](#)).

- The **if** clause was extended to accept a *directive-name-modifier* that allows it to apply to **combined constructs** (see [Section 5.5](#) and [Section 5.6](#)).
- An argument was added to the **ordered** clause of the **worksharing-loop construct** and the **ordered construct** was modified to support **doacross loop nests** (see [Section 6.4.6](#), [Section 19.6](#) and [Section 23.10.2](#)).
- The **implicitly determined data-sharing attribute** for **scalar variables** in **target regions** was changed to **firstprivate** (see [Section 7.1.1](#)).
- Use of some C++ reference types was allowed in some **data-sharing attribute clauses** (see [Section 7.3](#)).
- The **private**, **firstprivate** and **defaultmap** clauses were added to the **target construct** (see [Section 7.3.3](#), [Section 7.3.4](#), [Section 11.4](#) and [Section 21.8](#)).
- The *linear-modifier* was added to the **linear** clause (see [Section 8.14](#)).
- The **linear** clause was added to the **worksharing-loop construct** (see [Section 8.14](#) and [Section 19.6](#)).
- To support interaction with native **device** implementations, the **is_device_ptr** clause was added to the **target construct** and the **use_device_ptr** clause was added to the **target_data** construct (see [Section 7.3.6](#), [Section 7.3.7](#), [Section 21.7](#) and [Section 21.8](#)).
- Semantics for **reductions** on C/C++ **array sections** were added and restrictions on the use of arrays and pointers in reductions were removed (see [Section 8.10](#)).
- Support was added to the **map** clause to handle **structure** elements (see [Section 11.3](#)).
- To support unstructured data mapping for **devices**, the **map** clause (see [Section 11.3](#)) was updated and the **target_enter_data** (see [Section 21.5](#)) and **target_exit_data** (see [Section 21.6](#)) constructs were added.
- The **declare_target** directive was extended to allow mapping of **global variables** to be deferred to specific **device** executions and to allow an *extended-list* to be specified in C/C++ (see [Section 15.9](#)).
- The **simdlen** clause was added to the **simd construct** to support specification of the exact number of **logical iterations** desired per **SIMD chunk** (see [Section 18.4](#)).
- To support the use of the **simd construct** on loops with loop-carried backward dependences with or without a **worksharing-loop construct**, **clauses** were added to the **ordered construct** (see [Section 18.4](#), [Section 19.6](#) and [Section 23.10](#)).
- The **task construct** was extended to accept hints that the **priority** clause specifies (see [Section 20.1](#) and [Section 20.9](#)).
- The **taskloop construct** was added to support nestable parallel loops that create **explicit tasks** (see [Section 20.2](#)).

- To improve support for asynchronous execution of **target** regions, the **target** construct was extended to accept the **nowait** and **depend** clauses (see Section 21.8, Section 23.6 and Section 23.9.5).
- The **hint** clause was added to the **critical** construct (see Section 23.2).
- The **source** and **sink** dependence types were added to the **depend** clause to support **doacross** loop nests (see Section 23.9.5).
- To support a more complete set of **compound** constructs for **devices**, the **compound** constructs **target parallel**, **target parallel for** (C/C++), **target parallel for simd** (C/C++), **target parallel do** (Fortran) and **target parallel do simd** (Fortran) were added (see Section 25.1).
- The **omp_get_max_task_priority** routine was added to return the maximum supported **task priority** value (see Section 29.1.1).
- **Device memory routines** were added to allow explicit **memory** allocations, deallocations and transfers and **memory** associations (see Chapter 31).
- The **lock** API was extended with **lock routines** that support storing a hint with a **lock** to select a desired **lock** implementation for the intended usage of the **lock** by the application code (see Section 34.1.3 and Section 34.1.4).
- **Query routines** for **thread affinity** were added (see Section 35.2 to Section 35.7).
- C/C++ grammar (previously Appendix B) was moved to a separate document.

B.8 Version 3.1 to 4.0 Differences

- Various changes throughout the specification were made to provide initial support of Fortran 2003 (see Section 1.6).
- The **OMP_PLACES** environment variable (see Section 4.1.6), the **proc_bind** clause (see Section 18.1.3), and the **omp_get_proc_bind** routine (see Section 35.1) were added to support **thread affinity** policies.
- The **OMP_CANCELLATION** environment variable (see Section 4.3.6), the **cancel** construct (see Section 24.2), the **cancellation point** construct (see Section 24.3), and the **omp_get_cancellation** routine (see Section 36.1) were added to support the concept of **cancellation**.
- The **OMP_DEFAULT_DEVICE** environment variable (see Section 4.3.8), **device** constructs (see Chapter 21), and the **omp_get_num_teams**, **omp_get_team_num**, **omp_set_default_device**, **omp_get_default_device**, **omp_get_num_devices**, and **omp_is_initial_device** routines (see Chapter 28 and Chapter 30) were added to support execution on **devices**.

- The **OMP_DISPLAY_ENV** environment variable (see Section 4.7) was added to display the value of **ICVs** associated with the **OpenMP** environment variables.
- C/C++ array syntax was extended to support **array sections** (see Section 5.2.5).
- The **reduction** clause (see Section 8.10) was extended and the **declare_reduction** construct (see Section 8.15) was added to support **user-defined reductions**.
- **SIMD directives** were added to support **SIMD** parallelism (see Section 18.4).
- **Implementation defined task scheduling points** for **untied tasks** were removed (see Section 20.14).
- The **taskgroup** construct (see Section 23.4) was added to support deep **task** synchronization.
- The **atomic** construct was extended to support **atomic captured updates** with the **capture** clause, to allow new **atomic update** forms, and to support **sequentially consistent atomic operations** with the **seq_cst** clause (see Section 23.8.1.5, Section 23.8.3.1 and Section 23.8.5).
- The **depend** clause (see Section 23.9.5) was added to support **task dependences**.
- Examples (previously Appendix A) were moved to a separate document.

B.9 Version 3.0 to 3.1 Differences

- The **bind-var ICV** (see Section 3.1) and the **OMP_PROC_BIND** environment variable (see Section 4.1.7) were added to support control of whether **threads** are bound to **processors**.
- The **nthreads-var ICV** was modified to be a list of the number of **threads** to use at each nested **parallel region** level (see Section 3.1) and the algorithm for determining the number of **threads** used in a **parallel region** was modified to handle a list (see Section 18.1.1).
- **Data environment** restrictions were changed to allow **intent (in)** and **const**-qualified types for the **firstprivate** clause (see Section 7.3.4).
- **Data environment** restrictions were changed to allow Fortran pointers in **firstprivate** (see Section 7.3.4) and **lastprivate** (see Section 7.3.5) clauses.
- New **reduction** operators **min** and **max** were added for C/C++ (see Section 8.3).
- The **mergeable** and **final** clauses (see Section 20.5 and Section 20.7) were added to the **task** construct (see Section 20.1) to support optimization of **task data environments**.
- The **taskyield** construct was added to allow user-defined **task scheduling points** (see Section 20.12).

- The **atomic** construct was extended to include **read**, **write**, and **capture** forms, and an **update** clause was added to apply the already existing form of the **atomic** construct (see Section 23.8.2, Section 23.8.3.1 and Section 23.8.5).
- The nesting restrictions were clarified to disallow **closely nested regions** within an **atomic region** so that an **atomic region** can be consistently defined with other **regions** to include all code in the **atomic** construct (see Section 25.1).
- The **omp_in_final** routine was added to support specialization of **final task regions** (see Section 29.1.3).
- Descriptions of examples (previously Appendix A) were expanded and clarified.
- Incorrect use of **omp_integer_kind** in Fortran interfaces was replaced with **selected_int_kind(8)**.

B.10 Version 2.5 to 3.0 Differences

- The concept of **tasks** was added to the execution model (see Section 1.2 and Chapter 2).
- The OpenMP **memory** model was extended to cover atomicity of **memory** accesses (see Section 1.3.1). The description of the behavior of **volatile** in terms of **flushes** was removed.
- The definition of **active parallel region** was changed so that a **parallel region** is active if it is executed by a **team** to which more than one **thread** is assigned (see Chapter 2).
- The definition of the *nest-var*, *dyn-var*, *nthreads-var* and *run-sched-var* ICVs were modified to provide one copy of these ICVs per **task** instead of one copy for the whole OpenMP program (see Section 3.1). The **omp_set_num_threads** and **omp_set_dynamic** routines were specified to support their use from inside a **parallel region** (see Section 27.1 and Section 27.7).
- The *thread-limit-var* ICV, the **OMP_THREAD_LIMIT** environment variable and the **omp_get_thread_limit** routine were added to support control of the maximum number of **threads** (see Section 3.1, Section 4.1.4 and Section 27.5).
- The *max-active-levels-var* ICV, the **OMP_MAX_ACTIVE_LEVELS** environment variable and the **omp_set_max_active_levels** and **omp_get_max_active_levels** routines, and were added to support control of the number of nested **active parallel regions** (see Section 3.1, Section 4.1.5, Section 27.12 and Section 27.13).
- The *stacksize-var* ICV and the **OMP_STACKSIZE** environment variable were added to support control of **thread** stack sizes (see Section 3.1 and Section 4.3.2).
- The *wait-policy-var* ICV and the **OMP_WAIT_POLICY** environment variable were added to control the desired behavior of waiting **threads** (see Section 3.1 and Section 4.3.3).

- Predetermined data-sharing attributes were defined for Fortran assumed-size arrays (see Section 7.1.1).
- Static class member variables were allowed in **threadprivate** directives (see Section 9.1).
- Invocations of constructors and destructors for **private** and **threadprivate** class type variables were clarified (see Section 9.1, Section 7.3.3, Section 7.3.4, Section 10.1 and Section 10.2).
- The use of Fortran allocatable arrays was allowed in **private**, **firstprivate**, **lastprivate**, **reduction**, **copyin** and **copyprivate** clauses (see Section 9.1, Section 7.3.3, Section 7.3.4, Section 7.3.5, Section 8.10, Section 10.1 and Section 10.2).
- Support for **firstprivate** was added to the **default** clause in Fortran (see Section 7.3.1).
- Implementations were precluded from using the storage of the original list item to hold the new list item on the primary thread for list item in the **private** clause, and the value was made well defined on exit from the **parallel** region if no attempt is made to reference the original list item inside the **parallel** region (see Section 7.3.3).
- Determination of the number of threads in **parallel** regions was updated (see Section 18.1.1).
- The assignment of logical iterations to threads in a **worksharing-loop construct** with a **static** schedule kind was made deterministic (see Section 19.6).
- The **worksharing-loop construct** was extended to support association with more than one perfectly nested loop through the **collapse** clause (see Section 19.6).
- Loop-iteration variables for **worksharing-loop constructs** were allowed to be random access iterators or of unsigned integer type (see Section 19.6).
- The schedule kind **auto** was added to allow the implementation to choose any possible mapping of logical iterations in a **worksharing-loop constructs** to threads in the **team** (see Section 19.6).
- The **task** construct was added to support explicit tasks (see Section 20.1).
- The **taskwait** construct was added to support task synchronization (see Section 23.5).
- The **omp_set_schedule** and **omp_get_schedule** routines were added to set and to retrieve the value of the *run-sched-var* ICV (see Section 27.9 and Section 27.10).
- The **omp_get_level** routine was added to return the number of nested **parallel regions** that enclose the **task** that contains the call (see Section 27.14).
- The **omp_get_ancestor_thread_num** routine was added to return the thread number of the ancestor thread of the current thread (see Section 27.15).

- 1 • The `omp_get_team_size` routine was added to return the size of the `team` to which the
2 [ancestor thread](#) of the current `thread` belongs (see [Section 27.16](#)).
- 3 • The `omp_get_active_level` routine was added to return the number of `active parallel`
4 [regions](#) that enclose the `task` that contains the call (see [Section 27.17](#)).
- 5 • Lock ownership was defined in terms of `tasks` instead of `threads` (see [Chapter 34](#)).

C Nesting of Regions

This appendix describes a set of restrictions on the nesting of **regions**. The restrictions on nesting are as follows:

- A **teams region** must be strictly nested either within the **implicit parallel region** that surrounds the whole **OpenMP program** or within a **target region**. If a **teams construct** is nested within a **target construct**, that **target construct** must contain no statements, declarations or **directives** outside of the **teams construct** (see [Section 18.2](#)).
- Only **regions** that are generated by **teams-nestable constructs** or **teams-nestable routines** may be **strictly nested regions** of **teams regions** (see [Section 18.2](#)).
- The only **routines** for which a call may be nested inside a **region** that corresponds to a **construct** on which the **order clause** is specified with **concurrent** as the *ordering* argument are **order-concurrent-nestable routines** (see [Section 18.3](#)).
- Only **regions** that correspond to **order-concurrent-nestable constructs** or **order-concurrent-nestable routines** may be **strictly nested regions** of **regions** that correspond to **constructs** on which the **order clause** is specified with **concurrent** as the *ordering* argument (see [Section 18.3](#)).
- The only OpenMP **constructs** that can be encountered during execution of a **simd region** are **SIMDizable constructs** (see [Section 18.4](#)).
- 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 a **team-generating construct**. This follows from various restrictions requiring, in general, that **team-executed regions** (which include **worksharing regions** and **barrier regions**) are executed by all **threads** in a **team** or by none at all (see [Chapter 19](#) and [Section 23.3.1](#)).
- A **distribute region** must be strictly nested inside a **teams region** (see [Section 19.7](#)).
- A **loop region** that binds to a **teams region** must be strictly nested inside a **teams region** (see [Section 19.8.1](#)).
- During execution of a **target region**, other than **target constructs** for which a **device clause** on which the *ancestor device-modifier* appears, **device-affecting constructs** must not be encountered (see [Section 21.8](#)).
- A **critical region** must not be nested (closely or otherwise) inside a **critical region** with the same *name* (see [Section 23.2](#)).

- OpenMP **constructs** may not be encountered during execution of an **atomic region** (see Section 23.8.5).
- An **ordered region** that corresponds to an **ordered construct** with the **threads** or **doacross** clause may not be closely nested inside a **critical**, **ordered**, **loop**, **task**, or **taskloop** region (see Section 23.10).
- If the **simd parallelization-level** clause is specified on an **ordered construct**, the **ordered region** must bind to a **simd region** or one that corresponds to a **compound construct** for which the **simd construct** is a **leaf construct** (see Section 23.10.2).
- If the **threads parallelization-level** clause is specified on an **ordered construct**, the **ordered region** must bind to a **worksharing-loop region** or one that corresponds to a **compound construct** for which a **worksharing-loop construct** is a **leaf construct** (see Section 23.10.2).
- If the **threads parallelization-level** clause is specified on an **ordered construct** and the **binding region** corresponds to a **compound construct** then the **simd construct** must not be a **leaf construct** unless the **simd parallelization-level** clause is also specified (see Section 23.10.2).
- If *cancel-directive-name* is **taskgroup**, the **cancel construct** must be closely nested inside a **task construct** and the **cancel region** must be closely nested inside a **taskgroup region**. Otherwise, the **cancel construct** must be closely nested inside a **construct** for which *directive-name* is *cancel-directive-name* (see Section 24.2).
- A **cancellation point construct** for which *cancel-directive-name* is **taskgroup** must be closely nested inside a **task construct**, and the **cancellation point region** must be closely nested inside a **taskgroup region**. Otherwise, a **cancellation point construct** must be closely nested inside a **construct** for which *directive-name* is *cancel-directive-name* (see Section 24.3).

D Conforming Compound Directive Names

This appendix provides the grammar from which one may derive the full list of conforming [compound-directive names](#) (see [Section 25.1](#)) after excluding any productions for [compound-directive name](#) that would violate the following constraints:

- [Leaf-directive names](#) must be unique.
- The nesting of [constructs](#) indicated by the [compound construct](#) must be conforming.
- For Fortran, where spaces are optional, the resulting [compound-directive name](#) must have unambiguous [leaf-directive names](#) (e.g., plus signs should be used to separate [leaf-directive names](#) to disambiguate **taskloop** and **task loop** as [constituent-directive names](#)).

```
compound-dir-name :  
    composite-loop-dir-name  
    parallelism-generating-combined-dir-name  
    thread-selecting-combined-dir-name  
  
composite-loop-dir-name :  
    distribute-composite-dir-name  
    taskloop-composite-dir-name  
    worksharing-loop-composite-dir-name  
  
parallelism-generating-combined-dir-name :  
    parallel-combined-dir-name  
    target-combined-dir-name  
    target_data-combined-dir-name  
    task-combined-dir-name  
    teams-combined-dir-name  
  
thread-selecting-combined-dir-name :  
    masked-combined-dir-name  
    single-combined-dir-name  
  
distribute-composite-dir-name :  
    distribute parallel-worksharing-loop-dir-name  
    distribute simd-dir-name
```

```

1  taskloop-composite-dir-name:
2      taskloop simd-dir-name
3
4  worksharing-loop-composite-dir-name:
5      for simd-dir-name
6      do simd-dir-name
7
8  parallel-combined-dir-name:
9      parallel partitioned-worksharing-dir-name
10     parallel simd-dir-name
11     parallel target-task-generating-dir-name
12     parallel task-dir-name
13     parallel taskloop-dir-name
14     parallel thread-selecting-dir-name
15
16  target-combined-dir-name:
17     target loop-dir-name
18     target parallel-dir-name
19     target simd-dir-name
20     target task-dir-name
21     target taskloop-dir-name
22     target teams-dir-name
23
24  target_data-combined-dir-name:
25     target_data loop-dir-name
26     target_data parallel-dir-name
27     target_data simd-dir-name
28
29  task-combined-dir-name:
30     task loop-dir-name
31     task parallel-dir-name
32     task simd-dir-name
33
34  teams-combined-dir-name:
35     teams parallel-dir-name
36     teams partitioned-nonworksharing-workdist-dir-name
37     teams simd-dir-name
38     teams target-task-generating-dir-name
39     teams task-dir-name
40     teams taskloop-dir-name
41
42  masked-combined-dir-name:
43     masked loop-dir-name

```

```

1      masked parallel-dir-name
2      masked simd-dir-name
3      masked target-task-generating-dir-name
4      masked task-dir-name
5      masked taskloop-dir-name
6
7      single-combined-dir-name :
8          single loop-dir-name
9          single parallel-dir-name
10         single simd-dir-name
11         single target-task-generating-dir-name
12         single task-dir-name
13         single taskloop-dir-name
14
15     parallel-worksharing-loop-dir-name :
16         parallel worksharing-loop-dir-name
17
18     simd-dir-name :
19         simd
20
21     partitioned-worksharing-dir-name :
22         loop-dir-name
23         single-dir-name
24         worksharing-loop-dir-name
25         sections
26         workshare
27
28     target-task-generating-dir-name :
29         target_data-dir-name
30         target-dir-name
31         target_enter_data
32         target_exit_data
33         target_update
34
35     task-dir-name :
36         task-combined-dir-name
37         task
38
39     taskloop-dir-name :
40         taskloop-composite-dir-name
41         taskloop
42
43     thread-selecting-dir-name :

```

```

1      masked-dir-name
2      single-dir-name
3
4      loop-dir-name :
5          loop
6
7      parallel-dir-name :
8          parallel-combined-dir-name
9          parallel
10
11     teams-dir-name :
12         teams-combined-dir-name
13         teams
14
15     partitioned-nonworksharing-workdist-dir-name :
16         distribute-dir-name
17         loop-dir-name
18         workdistribute
19
20     worksharing-loop-dir-name :
21         worksharing-loop-composite-dir-name
22         for
23         do
24
25     single-dir-name :
26         single-combined-dir-name
27         single
28
29     target_data-dir-name :
30         target_data-combined-dir-name
31         target_data
32
33     target-dir-name :
34         target-combined-dir-name
35         target
36
37     masked-dir-name :
38         masked-combined-dir-name
39         masked
40
41     distribute-dir-name :
42         distribute-composite-dir-name
43         distribute

```

Index

Symbols

`_OPENMP` macro, [139](#), [140](#), [150](#), [175](#)

A

absent, [380](#)
access type, [564](#)
acq_rel, [504](#)
acquire, [505](#)
acquire flush, [11](#)
adjust_args, [348](#)
affinity, [407](#)
affinity, [464](#)
aggregate maps , [300](#)
align, [320](#)
aligned, [331](#)
alloc_memory, [862](#)
allocate, [321](#), [323](#)
allocator, [321](#)
allocator_handle type, [565](#)
allocators, [326](#)
alloctrait type, [566](#)
alloctrait_key type, [568](#)
alloctrait_val type, [573](#)
alloctrait_value type, [570](#)
append_args, [350](#)
apply Clause, [389](#)
array mapping, [300](#)
array section mapping, [300](#)
array sections, [169](#)
array shaping, [168](#)
assume, [385](#)
assumed maps, [300](#)
assumed-size array mapping, [300](#)
assumed-type variable mapping, [300](#)
assumes, [385](#)
assumption clauses, [380](#)

assumption directives, [379](#)
asynchronous device memory routines, [624](#)
at, [370](#)
atomic, [514](#)
atomic, [508](#)
atomic construct, [922](#)
atomic_default_mem_order, [373](#)
attach-modifier, [299](#)
attribute clauses, [223](#)
attributes, data-mapping, [289](#)
attributes, data-sharing, [215](#)
auto, [438](#)

B

barrier, [495](#)
barrier, implicit, [496](#)
base language format, [188](#)
begin declare_target, [366](#)
begin declare_variant, [353](#)
begin metadirective, [344](#)
begin assumes, [386](#)
bind, [444](#)
branch, [360](#)
buffer_complete, [802](#)
buffer_request, [801](#)

C

callback_device_host_fn, [871](#)
device_finalize, [799](#)
callbacks, [848](#)
cancel, [540](#), [785](#)
cancel-directive-name, [539](#)
cancellation constructs, [539](#)
 cancel, [540](#)
 cancellation_point, [544](#)
cancellation_point, [544](#)

- canonical loop nest form, 201
- canonical loop sequence form, 207
- capture**, 510
- capture, atomic**, 514
- clause format, 161
- clauses
 - absent**, 380
 - acq_rel**, 504
 - acquire**, 505
 - adjust_args**, 348
 - affinity**, 464
 - align**, 320
 - aligned**, 331
 - allocate**, 323
 - allocator**, 321
 - append_args**, 350
 - apply Clause**, 389
 - assumption*, 380
 - at**, 370
 - atomic*, 508
 - atomic_default_mem_order**, 373
 - attribute data-sharing, 223
 - bind**, 444
 - branch*, 360
 - cancel-directive-name*, 539
 - capture**, 510
 - collapse**, 210
 - collector**, 268
 - combiner**, 265
 - compare**, 511
 - contains**, 381
 - copyin**, 285
 - copyprivate**, 286
 - counts**, 396
 - data copying, 285
 - data motion, 310
 - data-sharing, 223
 - default**, 223
 - defaultmap**, 305
 - depend**, 527
 - destroy**, 186
 - detach**, 465
 - device**, 471
 - device_safesync**, 379
 - device_type**, 470
 - dist_schedule**, 441
 - doacross**, 532
 - dyn_groupprivate**, 328
 - dynamic_allocators**, 374
 - enter**, 281
 - exclusive**, 272
 - extended-atomic*, 510
 - fail**, 512
 - filter**, 422
 - final**, 461
 - firstprivate**, 227
 - from**, 313
 - full*, 400
 - grainsize**, 452
 - graph_id**, 458
 - has_device_addr**, 235
 - hint**, 492
 - holds**, 381
 - if Clause**, 184
 - in_reduction**, 255
 - inbranch**, 361
 - inclusive**, 272
 - indirect**, 368
 - induction**, 257
 - inductor**, 268
 - init**, 185
 - init_complete**, 273
 - initializer**, 265
 - interop**, 356
 - is_device_ptr**, 233
 - lastprivate**, 230
 - linear**, 259
 - link**, 283
 - local**, 280
 - map**, 292
 - match**, 347
 - memory-order*, 504
 - memscope**, 513
 - mergeable**, 460
 - message**, 370

- `no_openmp`, 382
- `no_openmp_constructs`, 383
- `no_openmp_routines`, 383
- `no_parallelism`, 384
- `nocontext`, 358
- `nogroup`, 503
- `nontemporal`, 418
- `notinbranch`, 361
- `novariants`, 357
- `nowait`, 501
- `num_tasks`, 453
- `num_teams`, 415
- `num_threads`, 406
- `order`, 416
- `ordered`, 211
- `otherwise`, 343
- parallelization-level*, 537
- partial*, 401
- `permutation`, 394
- `priority`, 463
- `private`, 226
- `proc_bind`, 410
- `read`, 508
- `reduction`, 251
- `relaxed`, 506
- `release`, 506
- `replayable`, 460
- requirement*, 373
- `reverse_offload`, 375
- `safelen`, 419
- `safesync`, 411
- `schedule`, 437
- `self_maps`, 378
- `seq_cst`, 507
- `severity`, 371
- `shared`, 225
- `simd`, 538
- `simdlen`, 420
- `sizes`, 391
- looprange*, 213
- `task_reduction`, 255
- `thread_limit`, 472
- `threads`, 537
- `threadset`, 462
- `to`, 311
- `transparent`, 531
- `unified_address`, 376
- `unified_shared_memory`, 377
- `uniform`, 331
- `untied`, 459
- `update`, 509, 526
- `use`, 489
- `use_device_addr`, 236
- `use_device_ptr`, 234
- `uses_allocators`, 326
- `weak`, 512
- `when`, 342
- `write`, 509
- `collapse`, 210
- combined and composite directive
 - names, 545
- `compare`, 511
- `compare, atomic`, 514
- compilation sentinels, 176, 177
- compliance, 15
- composition of constructs, 545
- compound construct semantics, 551
- compound directive names, 545
- conditional compilation, 175
- consistent loop schedules, 210
- construct syntax, 151
- constructs
 - `allocators`, 326
 - `atomic`, 514
 - `barrier`, 495
 - `cancel`, 540
 - cancellation constructs, 539
 - `cancellation_point`, 544
 - compound constructs, 551
 - `critical`, 493
 - `depobj`, 525
 - device constructs, 470
 - `dispatch`, 355
 - `distribute`, 439
 - `do`, 436
 - `flatten`, 392

- flush**, 518
- for**, 435
- fuse**, 392
- interchange**, 393
- interop**, 488
- loop**, 442
- masked**, 421
- ordered**, 533, 534, 536
- parallel**, 402
- reverse**, 395
- scope**, 425
- sections**, 426
- simd**, 417
- single**, 424
- split**, 395
- stripe**, 397
- target**, 481
- target_data**, 478
- target_enter_data**, 474
- target_exit_data**, 476
- target_update**, 486
- task**, 446
- task_iteration**, 454
- taskgraph**, 455
- taskgroup**, 498
- tasking constructs, 446
- taskloop**, 449
- taskwait**, 499
- taskyield**, 466
- teams**, 412
- tile**, 398
- unroll**, 400
- work-distribution, 423
- workdistribute**, 431
- workshare**, 428
- worksharing, 423
- worksharing-loop construct, 433
- contains**, 381
- control_tool**, 796
- control_tool** type, 585
- control_tool_result** type, 587
- controlling OpenMP thread affinity, 407
- copyin**, 285

- copyprivate**, 286
- counts**, 396
- critical**, 493

D

- data copying clauses, 285
- data mapping, 289
- data motion clauses, 310
- data sharing, 215
- data-sharing attribute clauses, 223
- data-sharing attribute rules, 215
- Declare Target, 362
- declare variant, 346
- declare_induction**, 266
- declare_mapper**, 307
- declare_reduction**, 262
- declare_simd**, 358
- declare_target**, 364
- declare_variant**, 351
- default**, 223
- defaultmap**, 305
- depend**, 527
- depend object, 525
- depend** type, 578
- dependences, 524
- dependences**, 786
- depobj**, 525
- deprecated features, 927
- destroy**, 186
- detach**, 465
- device**, 471
- device constructs
 - device constructs, 470
 - target**, 481
 - target_update**, 486
- target_enter_data**, 474
- target_exit_data**, 476
- device data environments overview, 8
- device directives, 470
- device information routines, 612
- device memory information routines, 624
- device memory routines, 623
- device_initialize**, 798
- device_load**, 800

- device_safesync**, 379
- device_to_host**, 872
- device_type**, 470
- device_unload**, 801
- directive format, 153
- directive syntax, 151
- directive-name-modifier**, 178
- directives, 953
 - allocate**, 321
 - assume**, 385
 - assumes**, 385
 - assumptions, 379
 - begin assumes**, 386
 - begin declare_target**, 366
 - begin declare_variant**, 353
 - begin metadirective**, 344
 - Declare Target, 362
 - declare variant, 346
 - declare_induction**, 266
 - declare_mapper**, 307
 - declare_reduction**, 262
 - declare_simd**, 358
 - declare_target**, 364
 - declare_variant**, 351
 - error**, 369
 - groupprivate**, 278
 - memory management directives, 315
 - metadirective**, 341, 344
 - nothing**, 386
 - requires**, 372
 - scan Directive**, 269
 - section**, 428
 - threadprivate**, 274
 - variant directives, 335
- dispatch**, 355, 780
- dist_schedule**, 441
- distribute**, 439
- do**, 436
- doacross**, 532
- dyn_groupprivate**, 328
- dynamic**, 438
- dynamic thread adjustment, 920
- dynamic_allocators**, 374

E

- enter**, 281
- environment variables, 130
 - OMP_AFFINITY_FORMAT**, 140
 - OMP_ALLOCATOR**, 147
 - OMP_AVAILABLE_DEVICES**, 142
 - OMP_CANCELLATION**, 142
 - OMP_DEBUG**, 150
 - OMP_DEFAULT_DEVICE**, 143
 - OMP_DISPLAY_AFFINITY**, 139
 - OMP_DISPLAY_ENV**, 150
 - OMP_DYNAMIC**, 132
 - OMP_MAX_ACTIVE_LEVELS**, 133
 - OMP_MAX_TASK_PRIORITY**, 146
 - OMP_NUM_TEAMS**, 137
 - OMP_NUM_THREADS**, 132
 - OMP_PLACES**, 134
 - OMP_PROC_BIND**, 135
 - OMP_SCHEDULE**, 137
 - OMP_STACKSIZE**, 138
 - OMP_TARGET_OFFLOAD**, 144
 - OMP_TEAMS_THREAD_LIMIT**, 137
 - OMP_THREAD_LIMIT**, 133
 - OMP_THREADS_RESERVE**, 145
 - OMP_TOOL**, 148
 - OMP_TOOL_LIBRARIES**, 148
 - OMP_TOOL_VERBOSE_INIT**, 149
 - OMP_WAIT_POLICY**, 139
- event, 609
- event callback registration, 727
- event routines, 609
- event_handle** type, 558
- exclusive**, 272
- execution control, 712
- execution model, 2
- extended-atomic*, 510

F

- fail**, 512
- features history, 927
- filter**, 422
- final**, 461
- firstprivate**, 227

fixed source form conditional compilation
sentinels, 177

fixed source form directives, 160

flatten, 392

flush, 518, 795

flush operation, 10

flush synchronization, 11

flush-set, 10

for, 435

frames, 744

free source form conditional compilation
sentinel, 176

free source form directives, 160

free_memory, 862

from, 313

full, 400

fuse, 392

G

general OpenMP types, 556

get_thread_context_for_thread_id,
868

glossary, 19

grainsize, 452

graph_id, 458

groupprivate, 278

groupprivate routines, 683

guided, 438

H

happens before, 11

has_device_addr, 235

header files, 553

hint, 492

history of features, 927

holds, 381

host_to_device, 872

I

ICVs (internal control variables), 118

if Clause, 184

impex type, 578

implementation, 915

implicit barrier, 496

implicit data-mapping, 305

implicit data-mapping attribute rules, 289

implicit flushes, 520

implicit_task, 783

in_reduction, 255

inbranch, 361

include files, 553

inclusive, 272

indirect, 368

induction, 257

inductor, 268

informational and utility directives, 369

init, 185

init_complete, 273

interchange, 393

internal control variables, 915

internal control variables (ICVs), 118

interop, 356

interop type, 558

interop_rc type, 559–561

interoperability, 488

Interoperability routines, 643

intptr type, 556

intrinsic identifiers, 178

introduction, 2

is_device_ptr, 233

iterators, 172

L

lastprivate, 230

linear, 259

link, 283

list item privatization, 220

local, 280

lock routines, 688

lock type, 579

lock_destroy, 793

lock_init, 792

loop, 442

loop concepts, 200

loop iteration spaces, 208

loop iteration vectors, 208

loop-transforming constructs, 388

M

- map**, 292
- map type decay, 297
- map-type*, 296
- mappable storage blocks, 302
- mapper**, 291
- mapper identifiers, 291
- masked**, 421, 777
- match**, 347
- memory allocator retrieving routines, 668
- memory allocators, 316
- memory copying routines, 633
- memory management, 315
- memory management directives
 - memory management directives, 315
- memory management routines, 651
- memory model, 7
- memory setting routines, 640
- memory space retrieving routines, 651, 675
- memory spaces, 315
- memory-order*, 504
- memory_read**, 865
- mempartition** type, 573
- mempartitioner** routines, 659
- mempartitioner** type, 573
- mempartitioner_compute_proc**
 - type, 575
- mempartitioner_lifetime** type, 574
- mempartitioner_release_proc**
 - type, 576
- memscope**, 513
- memspace_handle** type, 577
- mergeable**, 460
- message**, 370
- metadirective, 341
- metadirective**, 344
- modifier
 - attach-modifier*attach-modifier*, 299
 - directive-name-modifier*directive-name-modifier*, 178
 - map-type*map-type*, 296
 - reduction-identifier*reduction-identifier*, 250

- ref-modifier*ref-modifier*, 297
- task-dependence-type*task-dependence-type*, 525

- modifying and retrieving ICV values, 124
- modifying ICVs, 121
- mutex_acquire**, 790, 792
- mutex_acquired**, 792, 794
- mutex_released**, 794

N

- nest_lock**, 795
- nest_lock** type, 580
- nesting, 951
- no_openmp**, 382
- no_openmp_constructs**, 383
- no_openmp_routines**, 383
- no_parallelism**, 384
- nocontext**, 358
- nogroup**, 503
- nontemporal**, 418
- normative references, 16
- nothing**, 386
- notinbranch**, 361
- novariants**, 357
- nowait**, 501
- num_tasks**, 453
- num_teams**, 415
- num_threads**, 406

O

- OMP_AFFINITY_FORMAT**, 140
- omp_aligned_alloc**, 678
- omp_caligned_alloc**, 680
- omp_alloc**, 677
- OMP_ALLOCATOR**, 147
- omp_ancestor_is_free_agent**, 608
- OMP_AVAILABLE_DEVICES**, 142
- omp_calloc**, 679
- OMP_CANCELLATION**, 142
- omp_capture_affinity**, 710
- OMP_DEBUG**, 150
- OMP_DEFAULT_DEVICE**, 143
- omp_destroy_allocator**, 667
- omp_destroy_lock**, 693

- `omp_destroy_mempartition`, 663
- `omp_destroy_mempartitioner`, 661
- `omp_destroy_nest_lock`, 693
- `OMP_DISPLAY_AFFINITY`, 139
- `omp_display_affinity`, 709
- `OMP_DISPLAY_ENV`, 150
- `omp_display_env`, 716
- `OMP_DYNAMIC`, 132
- `omp_free`, 682
- `omp_fulfill_event`, 609
- `omp_get_active_level`, 600
- `omp_get_affinity_format`, 708
- `omp_get_ancestor_thread_num`, 598
- `omp_get_cancellation`, 712
- `omp_get_default_allocator`, 674
- `omp_get_default_device`, 613
- `omp_get_device_allocator`, 670
- `omp_get_device_and_host_allocator`, 671
- `omp_get_device_and_host_memspace`, 654
- `omp_get_device_from_uid`, 616
- `omp_get_device_memspace`, 653
- `omp_get_device_num`, 614
- `omp_get_device_num_teams`, 619
- `omp_get_device_teams_thread_limit`, 621
- `omp_get_devices_all_allocator`, 672
- `omp_get_devices_all_memspace`, 655
- `omp_get_devices_allocator`, 669
- `omp_get_devices_and_host_allocator`, 671
- `omp_get_devices_and_host_memspace`, 654
- `omp_get_devices_memspace`, 652
- `omp_get_dyn_groupprivate_ptr`, 684
- `omp_get_dyn_groupprivate_size`, 685
- `omp_get_dynamic`, 593
- `omp_get_groupprivate_limit`, 686
- `omp_get_initial_device`, 618
- `omp_get_interop_int`, 644
- `omp_get_interop_name`, 647
- `omp_get_interop_ptr`, 645
- `omp_get_interop_rc_desc`, 649
- `omp_get_interop_str`, 646
- `omp_get_interop_type_desc`, 648
- `omp_get_level`, 597
- `omp_get_mapped_ptr`, 626
- `omp_get_max_active_levels`, 597
- `omp_get_max_progress_width`, 615
- `omp_get_max_task_priority`, 606
- `omp_get_max_teams`, 603
- `omp_get_max_threads`, 590
- `omp_get_memspaces_num_resources`, 656
- `omp_get_num_devices`, 613
- `omp_get_num_interop_properties`, 644
- `omp_get_num_places`, 703
- `omp_get_num_procs`, 615
- `omp_get_num_teams`, 601
- `omp_get_num_threads`, 589
- `omp_get_partition_num_places`, 705
- `omp_get_partition_place_nums`, 706
- `omp_get_place_num`, 705
- `omp_get_place_num_procs`, 703
- `omp_get_place_proc_ids`, 704
- `omp_get_proc_bind`, 702
- `omp_get_schedule`, 594
- `omp_get_submemspace`, 658
- `omp_get_supported_active_levels`, 595
- `omp_get_team_num`, 602
- `omp_get_team_size`, 599
- `omp_get_teams_thread_limit`, 604
- `omp_get_thread_limit`, 591
- `omp_get_thread_num`, 589
- `omp_get_uid_from_device`, 617
- `omp_get_wtick`, 715
- `omp_get_wtime`, 715

- `omp_in_explicit_task`, 607
- `omp_in_final`, 607
- `omp_in_parallel`, 591
- `omp_init_allocator`, 666
- `omp_init_lock`, 689, 690
- `omp_init_mempartition`, 662
- `omp_init_mempartitioner`, 659
- `omp_init_nest_lock`, 690, 691
- `omp_is_free_agent`, 608
- `omp_is_initial_device`, 618
- `OMP_MAX_ACTIVE_LEVELS`, 133
- `OMP_MAX_TASK_PRIORITY`, 146
- `omp_mempartition_get_user_data`, 665
- `omp_mempartition_set_part`, 664
- `omp_memspace_get_pagesize`, 657
- `OMP_NUM_TEAMS`, 137
- `OMP_NUM_THREADS`, 132
- `omp_pause_resource`, 713
- `omp_pause_resource_all`, 714
- `OMP_PLACES`, 134
- `omp_pool`, 462
- `OMP_PROC_BIND`, 135
- `omp_realloc`, 681
- `OMP_SCHEDULE`, 137
- `omp_set_affinity_format`, 707
- `omp_set_default_allocator`, 673
- `omp_set_default_device`, 612
- `omp_set_device_num_teams`, 620
- `omp_set_device_teams_thread_limit`, 621
- `omp_set_dynamic`, 592
- `omp_set_lock`, 695
- `omp_set_max_active_levels`, 596
- `omp_set_nest_lock`, 696
- `omp_set_num_teams`, 601
- `omp_set_num_threads`, 588
- `omp_set_schedule`, 593
- `omp_set_teams_thread_limit`, 604
- `OMP_STACKSIZE`, 138
- `omp_target_alloc`, 627
- `omp_target_associate_ptr`, 630
- `omp_target_disassociate_ptr`, 631
- `omp_target_free`, 629
- `omp_target_is_accessible`, 625
- `omp_target_is_present`, 624
- `omp_target_memcpy`, 634
- `omp_target_memcpy_async`, 637
- `omp_target_memcpy_rect`, 635
- `omp_target_memcpy_rect_async`, 638
- `omp_target_memset`, 640
- `omp_target_memset_async`, 641
- `OMP_TARGET_OFFLOAD`, 144
- `omp_team`, 462
- `OMP_TEAMS_THREAD_LIMIT`, 137
- `omp_test_lock`, 699
- `omp_test_nest_lock`, 700
- `OMP_THREAD_LIMIT`, 133
- `OMP_THREADS_RESERVE`, 145
- `OMP_TOOL`, 148
- `OMP_TOOL_LIBRARIES`, 148
- `OMP_TOOL_VERBOSE_INIT`, 149
- `omp_unset_lock`, 697
- `omp_unset_nest_lock`, 698
- `OMP_WAIT_POLICY`, 139
- `ompd_bp_device_begin`, 909
- `ompd_bp_device_end`, 909
- `ompd_bp_parallel_begin`, 909
- `ompd_bp_parallel_end`, 910
- `ompd_bp_target_begin`, 912
- `ompd_bp_target_end`, 913
- `ompd_bp_task_begin`, 912
- `ompd_bp_task_end`, 912
- `ompd_bp_teams_begin`, 911
- `ompd_bp_teams_end`, 911
- `ompd_bp_thread_begin`, 908
- `ompd_bp_thread_end`, 908
- `ompd_dll_locations_valid`, 846
- `ompd_dll_locations`, 845
- `OMPT predefined identifiers`, 733
- `ompt_callback_error_t`, 774
- `OpenMP affinity support types`, 582
- `OpenMP allocator structured blocks`, 192
- `OpenMP argument lists`, 165
- `OpenMP atomic structured blocks`, 193

- OpenMP compliance, 15
- OpenMP context-specific structured blocks, 191
- OpenMP function dispatch structured blocks, 192
- OpenMP interoperability support types, 558
- OpenMP operations, 168
- OpenMP parallel region support types, 556
- OpenMP resource relinquishing types, 584
- OpenMP stylized expressions, 190
- OpenMP synchronization types, 578
- OpenMP tasking support types, 558
- OpenMP tool types, 585
- OpenMP types, 188
- order**, 416
- ordered**, 211, 533, 534, 536
- otherwise**, 343

P

- parallel**, 402
- parallel region support routines, 588
- parallel_begin**, 775
- parallel_end**, 776
- parallelism generating constructs, 402
- parallelization-level*, 537
- partial*, 401
- pause_resource** type, 584
- permutation**, 394
- predefined identifiers, 554
- prefer_type**, 490
- print_string**, 873
- priority**, 463
- private**, 226
- proc_bind**, 410
- proc_bind** type, 582

R

- rc**, 853
- read**, 508
- read**, **atomic**, 514
- read_memory**, 866
- read_string**, 866
- collector**, 268
- combiner**, 265

- initializer**, 265
- reduction**, 251, 790
- reduction and induction data control, 238
- task-dependence-type*, 250
- ref-modifier*, 297
- relaxed**, 506
- release**, 506
- release flush, 11
- replayable**, 460
- requirement*, 373
- requires**, 372
- reserved locators, 168
- resource relinquishing routines, 713
- reverse**, 395
- reverse_offload**, 375
- routine argument properties, 555
- routine bindings, 554
- runtime**, 438
- runtime library definitions, 553

S

- safelen**, 419
- safesync**, 411
- saved**, 237
- scan** Directive, 269
- sched** type, 556
- schedule**, 437
- scheduling, 467
- scope**, 425
- section**, 428
- sections**, 426
- self_maps**, 378
- seq_cst**, 507
- severity**, 371
- shared**, 225
- simd**, 417, 538
- simd data control, 331
- simdlen**, 420
- single**, 424
- sizeof_type**, 870
- sizes**, 391
- looprange*, 213
- split**, 395
- stand-alone directives, 159

- static**, 438
- static-lifetime data control, 274
- storage block mapping, 302
- strip**, 397
- strong flush, 10
- structure mapping, 300
- structured blocks, 191
- symbol_addr_lookup**, 863
- sync_region**, 788, 789
- sync_region_wait**, 789
- synchronization constructs, 492
- synchronization constructs and clauses, 492
- synchronization hint type, 580

T

- target**, 481, 806
- target asynchronous device memory
 - routines, 624
- target groupprivate routines, 683
- target memory copying routines, 633
- target memory information routines, 624
- target memory routines, 623
- target memory setting routines, 640
- target_data**, 478
- target_data_op**, 803
- target_data_op_emi**, 803
- target_emi**, 806
- target_map**, 808
- target_map_emi**, 808
- target_submit**, 810
- target_submit_emi**, 810
- target_update**, 486
- task**, 446
- task scheduling, 467
- task-dependence-type*, 525
- task_create**, 781
- task_dependence**, 787
- task_iteration**, 454
- task_reduction**, 255
- task_schedule**, 782
- taskgraph**, 455
- taskgroup**, 498
- tasking constructs, 446
- tasking routines, 606

- tasking support, 606
- taskloop**, 449
- taskwait**, 499
- taskyield**, 466
- teams**, 412
- teams region routines, 601
- thread affinity, 407
- thread affinity routines, 702
- thread_begin**, 772
- thread_end**, 773
- thread_limit**, 472
- threadprivate**, 274
- threads**, 537
- threadset**, 462
- tile**, 398
- timer, 715
- timing routines, 715
- to**, 311
- tool control, 718
- tool initialization, 724
- tool interfaces definitions, 721, 845
- tool support, 718
- tools header files, 721, 845
- tracing device activity, 728
- transparent**, 531
- types
 - access**, 564
 - allocator_handle**, 565
 - alloctrait**, 566
 - alloctrait_key**, 568
 - impex**, 578
 - alloctrait_val**, 573
 - alloctrait_value**, 570
 - control_tool**, 585
 - control_tool_result**, 587
 - depend**, 578
 - event_handle**, 558
 - interop_rc**, 559–561
 - interop**, 558
 - intptr**, 556
 - lock**, 579
 - mempartition**, 573
 - mempartitioner**, 573

- `mempartitioner_compute_proc`, 575
- `mempartitioner_lifetime`, 574
- `mempartitioner_release_proc`, 576
- `memspace_handle`, 577
- `nest_lock`, 580
- `pause_resource`, 584
- `proc_bind`, 582
- `sched`, 556
- `sync_hint`, 580
- `uintptr`, 556
- constructs, 423
- worksharing constructs, 423
- worksharing-loop construct, 433
- `write`, 509
- `write, atomic`, 514
- `write_memory`, 867

U

- `uintptr` type, 556
- `unified_address`, 376
- `unified_shared_memory`, 377
- `uniform`, 331
- `unroll`, 400
- `untied`, 459
- `update`, 509, 526
- `update, atomic`, 514
- `use`, 489
- `use_device_addr`, 236
- `use_device_ptr`, 234
- `uses_allocators`, 326

V

- variables, environment, 130
- variant directives, 335

W

- wait identifier, 768
- wall clock timer, 715
- `error`, 369
- `weak`, 512
- `when`, 342
- `work`, 778
- work-distribution
 - constructs, 423
- work-distribution constructs, 423
- `workdistribute`, 431
- `workshare`, 428
- worksharing