OpenMP Technical Report 15: Version 6.1 Public Comment Draft

This Technical Report is the final preview of the OpenMP Application Programming Interface Specification version 6.1. Version 6.1 has support for multi-dimensional teams and for new 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. In Fortran, the **`map`** clause now supports polymorphic variables, and invocation of type-bound procedures in **`target`** regions has been added. See Appendix B.2 for the complete list of changes relative to version 6.0.

EDITORS

Bronis R. de Supinski
Michael Klemm

July 9, 2026
Expires September 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 Public Comment Draft July 2026

Copyright ©1997-2026 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 draft version includes the following internal GitHub issues (corresponding Trac ticket numbers in parentheses when they exist) applied to the 6.0 LaTeX sources: 1877 (945), 2142, 2489, 2594, 2961, 3041, 3069, 3211, 3225, 3284, 3288, 3295, 3311, 3324, 3357, 3362, 3433, 3557, 3631, 3676, 3682, 3684, 3700, 3749, 3813, 3833, 3835, 3869, 3877, 3890, 3976, 4132-4133, 4143, 4264, 4320, 4326, 4332, 4371, 4397-4398, 4406-4408, 4414, 4425, 4432-4433, 4436, 4443-4444, 4448-4450, 4453-4455, 4457, 4459, 4467, 4470, 4472, 4476, 4483, 4485, 4487, 4491-4492, 4495, 4497-4498, 4502, 4504-4505, 4509, 4511-4512, 4514-4516, 4521, 4523, 4525, 4527, 4530-4535, 4540, 4542, 4551, 4553, 4555, 4557-4558, 4561-4562, 4565, 4567, 4569, 4571-4572, 4576, 4578-4580, 4584, 4586, 4588, 4591, 4594, 4597-4601, 4603, 4610, 4613, 4617-4621, 4623, 4627-4628, 4632, 4637, 4639-4641, 4655-4656, 4658, 4661, 4663-4668, 4670, 4672-4674, 4678-4679, 4682, 4684, 4690, 4692, 4701, 4714, 4724, 4730, 4744-4745, 4751, 4755-4757, 4759, 4772-4773, 4776-4778, 4784, 4787-4788, 4791, 4793, 4797, 4802-4803, 4806, 4815-4817, 4849, 4855, 4873, 4875, 4889, 4891-4892, 4896, 4900, 4917

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 . . . 9
1.3.3 Memory Management . . . 9
1.3.4 The Flush Operation . . . 10
1.3.5 Flush Synchronization and Happens-Before Order . . . 12
1.3.6 OpenMP Memory Consistency . . . 13
1.4 Tool Interfaces . . . 14
1.4.1 OMPT . . . 15
1.4.2 OMPD . . . 15
1.5 OpenMP Compliance . . . 16
1.6 Normative References . . . 16
1.7 Organization of this Document . . . 18

2 Glossary 20

3 Internal Control Variables 131

3.1 ICV Descriptions . . . 131
3.2 ICV Initialization . . . 134
3.3 Modifying and Retrieving ICV Values . . . 137
3.4 Spatial ICVs . . . 140
3.5 How the Per-Data Environment ICVs Work . . . 141

3.6 ICV Override Relationships . . . 143

4 Environment Variables 144

4.1 Parallel Region Environment Variables . . . 145
4.1.1 Abstract Name Values . . . 145
4.1.2 **OMP_DYNAMIC** . . . 146
4.1.3 **OMP_NUM_THREADS** . . . 146
4.1.4 **OMP_THREAD_LIMIT** . . . 147
4.1.5 **OMP_MAX_ACTIVE_LEVELS** . . . 147
4.1.6 **OMP_PLACES** . . . 148
4.1.7 **OMP_PROC_BIND** . . . 150
4.2 Teams Environment Variables . . . 152
4.2.1 **OMP_NUM_TEAMS** . . . 152
4.2.2 **OMP_TEAMS_THREAD_LIMIT** . . . 152
4.3 Program Execution Environment Variables . . . 153
4.3.1 **OMP_SCHEDULE** . . . 153
4.3.2 **OMP_STACKSIZE** . . . 153
4.3.3 **OMP_WAIT_POLICY** . . . 154
4.3.4 **OMP_DISPLAY_AFFINITY** . . . 155
4.3.5 **OMP_AFFINITY_FORMAT** . . . 155
4.3.6 **OMP_CANCELLATION** . . . 157
4.3.7 **OMP_AVAILABLE_DEVICES** . . . 157
4.3.8 **OMP_DEFAULT_DEVICE** . . . 159
4.3.9 **OMP_TARGET_OFFLOAD** . . . 159
4.3.10 **OMP_THREADS_RESERVE** . . . 160
4.3.11 **OMP_MAX_TASK_PRIORITY** . . . 161
4.4 Memory Allocation Environment Variables . . . 162
4.4.1 **OMP_ALLOCATOR** . . . 162
4.5 OMPT Environment Variables . . . 163
4.5.1 **OMP_TOOL** . . . 163
4.5.2 **OMP_TOOL_LIBRARIES** . . . 163
4.5.3 **OMP_TOOL_VERBOSE_INIT** . . . 164
4.6 OMPD Environment Variables . . . 165
4.6.1 **OMP_DEBUG** . . . 165

4.7 **OMP_DISPLAY_ENV** . . . 165

5 Directive and Construct Syntax **166**

5.1 Directive Format . . . 168
5.1.1 Free Source Form Directives . . . 175
5.1.2 Fixed Source Form Directives . . . 175
5.2 Clause Format . . . 176
5.2.1 OpenMP List Arguments . . . 181
5.2.2 Reserved Locators . . . 184
5.2.3 OpenMP Operations . . . 184
5.2.4 Array Shaping . . . 184
5.2.5 Array Sections . . . 185
5.2.6 **iterator** Modifier . . . 187
5.3 Conditional Compilation . . . 190
5.3.1 Free Source Form Conditional Compilation Sentinel . . . 190
5.3.2 Fixed Source Form Conditional Compilation Sentinels . . . 191
5.4 Intrinsic Identifiers . . . 191
5.5 *directive-name-modifier* Modifier . . . 192
5.6 **if** Clause . . . 198
5.7 **init** Clause . . . 199
5.7.1 *interop-type* Modifiers . . . 201
5.8 **destroy** Clause . . . 201
5.9 *variable-selector* Modifiers . . . 202

6 Base Language Formats and Restrictions **205**

6.1 OpenMP Types and Identifiers . . . 205
6.2 OpenMP Stylized Expressions . . . 207
6.3 Structured Blocks . . . 208
6.3.1 OpenMP Allocator Structured Blocks . . . 209
6.3.2 OpenMP Function Dispatch Structured Blocks . . . 209
6.3.3 OpenMP Atomic Structured Blocks . . . 210
6.4 Loop Concepts . . . 218
6.4.1 Canonical Loop Nest Form . . . 218
6.4.2 Canonical Loop Sequence Form . . . 225

6.4.3 OpenMP Loop-Iteration Spaces and Vectors . . . 226
6.4.4 Consistent Loop Schedules . . . 228
6.4.5 **collapse** Clause . . . 228
6.4.6 **ordered** Clause . . . 229
6.4.7 **depth** Clause . . . 230
6.4.8 **looprange** Clause . . . 231

II Data Control 232

7 Data-Sharing Control 233
7.1 Data-Sharing Attribute Rules . . . 233
7.1.1 Variables Referenced in a Construct . . . 233
7.1.2 Variables Referenced in a Region but not in a Construct . . . 237
7.2 List Item Privatization . . . 237
7.3 Data-Sharing Attribute Clauses . . . 240
7.3.1 **shared** Clause . . . 241
7.3.2 **private** Clause . . . 242
7.3.3 **firstprivate** Clause . . . 243
7.3.4 **lastprivate** Clause . . . 246
7.3.5 **is_device_ptr** Clause . . . 249
7.3.6 **use_device_ptr** Clause . . . 250
7.3.7 **has_device_addr** Clause . . . 251
7.3.8 **use_device_addr** Clause . . . 252
7.4 **default** Clause . . . 253
7.5 **saved** Modifier . . . 255

8 Reduction and Induction Data Control 256
8.1 OpenMP Reduction and Induction Identifiers . . . 256
8.2 OpenMP Reduction and Induction Expressions . . . 257
8.2.1 OpenMP Combiner Expressions . . . 257
8.2.2 OpenMP Initializer Expressions . . . 258
8.2.3 OpenMP Inductor Expressions . . . 260
8.2.4 OpenMP Collector Expressions . . . 261
8.3 Implicitly Declared OpenMP Reduction Identifiers . . . 261

8.4 Implicitly Declared OpenMP Induction Identifiers . . . 263
8.5 Properties Common to Reduction and **induction** Clauses . . . 263
8.6 Properties Common to All Reduction Clauses . . . 265
8.7 Reduction Scoping Clauses . . . 267
8.8 Reduction Participating Clauses . . . 268
8.9 *reduction-identifier* Modifier . . . 268
8.10 **reduction** Clause . . . 269
8.10.1 *reduction-type* Modifiers . . . 272
8.11 **task_reduction** Clause . . . 274
8.12 **in_reduction** Clause . . . 275
8.13 **induction** Clause . . . 276
8.14 **linear** Clause . . . 278
8.15 **declare_reduction** Directive . . . 281
8.15.1 **combiner** Clause . . . 284
8.15.2 **initializer** Clause . . . 284
8.15.3 **reduction_kind** Clause . . . 285
8.16 **declare_induction** Directive . . . 285
8.16.1 **inductor** Clause . . . 287
8.16.2 **collector** Clause . . . 288
8.17 **scan** Directive . . . 288
8.17.1 **inclusive** Clause . . . 291
8.17.2 **exclusive** Clause . . . 291
8.17.3 **init_complete** Clause . . . 292

9 Data-Localizing Control **293**
9.1 **threadprivate** Directive . . . 293
9.2 **groupprivate** Directive . . . 297
9.3 **local** Clause . . . 300

10 Data-Copying Control **301**
10.1 **copyin** Clause . . . 301
10.2 **copyprivate** Clause . . . 302

11 Data-Mapping Control **305**
11.1 Implicit Data-Mapping Attribute Rules . . . 305

11.2 Mapper Identifiers and **mapper** Modifiers . . . 307
11.3 **enter** Clause . . . 308
11.4 **link** Clause . . . 310
11.5 **map** Clause . . . 310
11.5.1 *map-type* Modifier . . . 315
11.5.2 Map Type Decay . . . 316
11.5.3 *map-type-modifying* Modifiers . . . 316
11.5.4 **ref-modifier** Modifier . . . 317
11.5.5 **attach-modifier** Modifier . . . 318
11.5.6 Assumed-Size Array Mapping . . . 320
11.5.7 Array and Structure Mapping . . . 320
11.5.8 Storage Block Mapping . . . 323
11.5.9 Implicit Data-Mapping . . . 325
11.6 **defaultmap** Clause . . . 326
11.7 **declare_mapper** Directive . . . 327

12 Data-Motion Control **330**
12.1 **to** Clause . . . 331
12.2 **from** Clause . . . 332

13 Memory Management **335**
13.1 Memory Spaces . . . 335
13.2 Memory Allocators . . . 336
13.3 **align** Clause . . . 340
13.4 **allocator** Clause . . . 341
13.5 **allocate** Directive . . . 341
13.6 **allocate** Clause . . . 343
13.7 **allocators** Construct . . . 346
13.8 **uses_allocators** Clause . . . 347
13.9 **dyn_groupprivate** Clause . . . 348

14 SIMD Data Control **351**
14.1 **uniform** Clause . . . 351
14.2 **aligned** Clause . . . 351

III General Directives and Clauses 353

15 Variant Directives 354

15.1 OpenMP Contexts . . . 354
15.2 Context Selectors . . . 356
15.3 Matching and Scoring Context Selectors . . . 359
15.4 Metadirectives . . . 360
15.4.1 **when** Clause . . . 361
15.4.2 **otherwise** Clause . . . 362
15.4.3 **metadirective** . . . 363
15.4.4 **begin metadirective** . . . 363
15.5 Declare Variant Directives . . . 364
15.5.1 **match** Clause . . . 366
15.5.2 **adjust_args** Clause . . . 366
15.5.3 **append_args** Clause . . . 369
15.5.4 **declare_variant** Directive . . . 370
15.5.5 **begin declare_variant** Directive . . . 372
15.6 **dispatch** Construct . . . 374
15.6.1 **interop** Clause . . . 375
15.6.2 **novariants** Clause . . . 376
15.6.3 **nocontext** Clause . . . 377
15.7 **declare_simd** Directive . . . 377
15.7.1 *branch* Clauses . . . 379
15.8 Declare Target Directives . . . 381
15.8.1 **declare_target** Directive . . . 382
15.8.2 **begin declare_target** Directive . . . 385
15.8.3 **indirect** Clause . . . 386

16 Informational and Utility Directives 388

16.1 **error** Directive . . . 388
16.2 **at** Clause . . . 389
16.3 **message** Clause . . . 389
16.4 **severity** Clause . . . 390

16.5 **requires** Directive . . . 391
16.5.1 *requirement* Clauses . . . 392
16.6 Assumption Directives . . . 398
16.6.1 *assumption* Clauses . . . 398
16.6.2 **assumes** Directive . . . 403
16.6.3 **assume** Directive . . . 404
16.6.4 **begin assumes** Directive . . . 404
16.7 **nothing** Directive . . . 405

17 Loop-Transforming Constructs 406
17.1 **apply** Clause . . . 407
17.2 **shift** Clause . . . 408
17.3 **sizes** Clause . . . 409
17.4 **flatten** Construct . . . 410
17.5 **fuse** Construct . . . 411
17.6 **interchange** Construct . . . 412
17.6.1 **permutation** Clause . . . 413
17.7 **reverse** Construct . . . 414
17.8 **split** Construct . . . 414
17.8.1 **counts** Clause . . . 415
17.9 **stripe** Construct . . . 416
17.10 **tile** Construct . . . 417
17.11 **unroll** Construct . . . 418
17.11.1 **full** Clause . . . 419
17.11.2 **partial** Clause . . . 420

18 Parallelism Generation and Control 421
18.1 **parallel** Construct . . . 421
18.1.1 Determining the Number of Threads for a **parallel** Region . . . 424
18.1.2 **num_threads** Clause . . . 425
18.1.3 Controlling OpenMP Thread Affinity . . . 427
18.1.4 **proc_bind** Clause . . . 429
18.1.5 **safesync** Clause . . . 430

18.2 **teams** Construct . . . 431
18.2.1 **num_teams** Clause . . . 434
18.3 *dims-modifier* Modifier . . . 436
18.4 **order** Clause . . . 437
18.5 **simd** Construct . . . 438
18.5.1 **nontemporal** Clause . . . 439
18.5.2 **safelen** Clause . . . 440
18.5.3 **simdlen** Clause . . . 440
18.6 **masked** Construct . . . 442
18.6.1 **filter** Clause . . . 443

19 Work-Distribution Constructs 444
19.1 **single** Construct . . . 445
19.2 **scope** Construct . . . 446
19.3 **sections** Construct . . . 447
19.3.1 **section** Directive . . . 449
19.4 **workshare** Construct . . . 449
19.5 **workdistribute** Construct . . . 452
19.6 Worksharing-Loop Constructs . . . 454
19.6.1 **for** Construct . . . 456
19.6.2 **do** Construct . . . 457
19.6.3 **schedule** Clause . . . 458
19.7 **distribute** Construct . . . 461
19.7.1 **dist_schedule** Clause . . . 463
19.8 **loop** Construct . . . 464
19.8.1 **bind** Clause . . . 465

20 Tasking Constructs 467
20.1 **task** Construct . . . 467
20.2 **taskloop** Construct . . . 470
20.2.1 **grainsize** Clause . . . 474
20.2.2 **num_tasks** Clause . . . 475
20.2.3 **task_iteration** Directive . . . 475

20.3 **taskgraph** Construct . . . 476
20.3.1 **graph_id** Clause . . . 480
20.3.2 **graph_reset** Clause . . . 480
20.4 **untied** Clause . . . 481
20.5 **mergeable** Clause . . . 482
20.6 **replayable** Clause . . . 482
20.7 **final** Clause . . . 483
20.8 **threadset** Clause . . . 484
20.9 **priority** Clause . . . 485
20.10 **affinity** Clause . . . 486
20.11 **detach** Clause . . . 487
20.12 Initial Task . . . 488
20.13 Task Scheduling . . . 489
20.13.1 Task Scheduling Points . . . 489
20.13.2 **taskyield** Construct . . . 489
20.13.3 Task Scheduling and its Constraints . . . 490
20.13.4 Synchronization Scheduling Constraints . . . 491

21 Device Directives and Clauses **493**
21.1 **device_type** Clause . . . 493
21.2 **device** Clause . . . 494
21.3 **thread_limit** Clause . . . 495
21.4 Device Initialization . . . 496
21.5 **target_enter_data** Construct . . . 497
21.6 **target_exit_data** Construct . . . 500
21.7 **target_data** Construct . . . 502
21.8 **target** Construct . . . 505
21.9 **target_update** Construct . . . 511

22 Interoperability **514**
22.1 **interop** Construct . . . 514
22.1.1 OpenMP Foreign Runtime Identifiers . . . 515
22.1.2 **use** Clause . . . 515
22.1.3 **prefer-type** Modifier . . . 516

23 Synchronization Constructs and Clauses 518
23.1 **hint** Clause . . . 518
23.2 **critical** Construct . . . 519
23.3 Barriers . . . 521
23.3.1 **barrier** Construct . . . 521
23.3.2 Implicit Barriers . . . 523
23.3.3 Implementation-Specific Barriers . . . 524
23.4 **taskgroup** Construct . . . 524
23.5 **taskwait** Construct . . . 526
23.6 **nowait** Clause . . . 528
23.7 **nogroup** Clause . . . 530
23.8 OpenMP Memory Ordering . . . 531
23.8.1 *memory-order* Clauses . . . 531
23.8.2 *atomic* Clauses . . . 535
23.8.3 *extended-atomic* Clauses . . . 537
23.8.4 **memscope** Clause . . . 540
23.8.5 **atomic** Construct . . . 541
23.8.6 **flush** Construct . . . 545
23.8.7 Implicit Flushes . . . 547
23.9 OpenMP Dependences . . . 548
23.9.1 *task-dependence-type* Modifier . . . 548
23.9.2 Depend Objects . . . 549
23.9.3 **depobj** Construct . . . 549
23.9.4 **update** Clause . . . 550
23.9.5 *depend-list-type* Modifiers . . . 550
23.9.6 **depend** Clause . . . 551
23.9.7 **transparent** Clause . . . 555
23.9.8 **doacross** Clause . . . 556
23.10 **ordered** Construct . . . 557
23.10.1 Stand-alone **ordered** Construct . . . 558
23.10.2 Block-associated **ordered** Construct . . . 560
23.10.3 *parallelization-level* Clauses . . . 561

24 Cancellation Constructs **563**
24.1 *cancel-directive-name* Clauses . . . 563
24.2 **cancel** Construct . . . 564
24.3 Cancellation Points . . . 566
24.4 **cancellation_point** Construct . . . 568

25 Composition of Constructs **570**
25.1 Compound Directive Names . . . 570
25.2 Clauses on Compound Constructs . . . 573
25.3 Compound Construct Semantics . . . 576

IV Runtime Library Routines **577**

26 Runtime Library Definitions **578**
26.1 Predefined Identifiers . . . 579
26.2 Routine Bindings . . . 579
26.3 Routine Argument Properties . . . 580
26.4 General OpenMP Types . . . 580
26.4.1 OpenMP **intptr** Type . . . 580
26.4.2 OpenMP **uintptr** Type . . . 581
26.5 OpenMP Parallel Region Support Types . . . 581
26.5.1 OpenMP **sched** Type . . . 581
26.6 OpenMP Tasking Support Types . . . 583
26.6.1 OpenMP **event_handle** Type . . . 583
26.7 OpenMP Interoperability Support Types . . . 583
26.7.1 OpenMP **interop** Type . . . 583
26.7.2 OpenMP **interop_fr** Type . . . 584
26.7.3 OpenMP **interop_property** Type . . . 585
26.7.4 OpenMP **interop_rc** Type . . . 587
26.8 OpenMP Memory Management Types . . . 589
26.8.1 OpenMP **access** Type . . . 589
26.8.2 OpenMP **allocator_handle** Type . . . 590
26.8.3 OpenMP **alloctrait** Type . . . 591
26.8.4 OpenMP **alloctrait_key** Type . . . 593

26.8.5 OpenMP **alloctrait_value** Type . . . 596
26.8.6 OpenMP **alloctrait_val** Type . . . 598
26.8.7 OpenMP **mempartition** Type . . . 599
26.8.8 OpenMP **mempartitioner** Type . . . 599
26.8.9 OpenMP **mempartitioner_lifetime** Type . . . 600
26.8.10 OpenMP **mempartitioner_compute_proc** Type . . . 600
26.8.11 OpenMP **mempartitioner_release_proc** Type . . . 602
26.8.12 OpenMP **memspace_handle** Type . . . 603
26.9 OpenMP Synchronization Types . . . 604
26.9.1 OpenMP **depend** Type . . . 604
26.9.2 OpenMP **impex** Type . . . 604
26.9.3 OpenMP **lock** Type . . . 605
26.9.4 OpenMP **nest_lock** Type . . . 606
26.9.5 OpenMP **sync_hint** Type . . . 606
26.10 OpenMP Affinity Support Types . . . 608
26.10.1 OpenMP **proc_bind** Type . . . 608
26.11 OpenMP Resource Relinquishing Types . . . 609
26.11.1 OpenMP **pause_resource** Type . . . 609
26.12 OpenMP Tool Types . . . 611
26.12.1 OpenMP **control_tool** Type . . . 611
26.12.2 OpenMP **control_tool_result** Type . . . 612

27 Parallel Region Support Routines 614

27.1 **omp_set_num_threads** Routine . . . 614
27.2 **omp_get_num_threads** Routine . . . 615
27.3 **omp_get_thread_num** Routine . . . 615
27.4 **omp_get_max_threads** Routine . . . 616
27.5 **omp_get_thread_limit** Routine . . . 616
27.6 **omp_in_parallel** Routine . . . 617
27.7 **omp_set_dynamic** Routine . . . 618
27.8 **omp_get_dynamic** Routine . . . 618
27.9 **omp_set_schedule** Routine . . . 619
27.10 **omp_get_schedule** Routine . . . 620
27.11 **omp_get_supported_active_levels** Routine . . . 621

27.12 **omp_set_max_active_levels** Routine . . . 621
27.13 **omp_get_max_active_levels** Routine . . . 622
27.14 **omp_get_level** Routine . . . 623
27.15 **omp_get_ancestor_thread_num** Routine . . . 623
27.16 **omp_get_team_size** Routine . . . 624
27.17 **omp_get_active_level** Routine . . . 625

28 Teams Region Routines **626**
28.1 **omp_get_num_teams** Routine . . . 626
28.2 **omp_set_num_teams** Routine . . . 626
28.3 **omp_get_team_num** Routine . . . 627
28.4 **omp_get_max_teams** Routine . . . 628
28.5 **omp_get_teams_thread_limit** Routine . . . 629
28.6 **omp_set_teams_thread_limit** Routine . . . 629

29 Multidimensional Support Routines **631**
29.1 **omp_get_num_threads_dim** Routine . . . 631
29.2 **omp_get_thread_num_dim** Routine . . . 631
29.3 **omp_get_num_teams_dim** Routine . . . 632
29.4 **omp_get_team_num_dim** Routine . . . 633
29.5 **omp_get_thread_limit_dim** Routine . . . 634
29.6 **omp_get_supported_active_team_dims** Routine . . . 635
29.7 **omp_get_supported_active_league_dims** Routine . . . 635
29.8 **omp_get_supported_threads_dim** Routine . . . 636
29.9 **omp_get_supported_teams_dim** Routine . . . 637

30 Tasking Support Routines **639**
30.1 Tasking Routines . . . 639
30.1.1 **omp_get_max_task_priority** Routine . . . 639
30.1.2 **omp_in_explicit_task** Routine . . . 640
30.1.3 **omp_in_final** Routine . . . 640
30.1.4 **omp_is_free_agent** Routine . . . 641
30.1.5 **omp_ancestor_is_free_agent** Routine . . . 641
30.2 Event Routine . . . 642
30.2.1 **omp_fulfill_event** Routine . . . 642

31 Device Information Routines **645**

31.1 **omp_set_default_device** Routine . . . 645
31.2 **omp_get_default_device** Routine . . . 646
31.3 **omp_get_num_devices** Routine . . . 646
31.4 **omp_get_device_num** Routine . . . 647
31.5 **omp_get_num_procs** Routine . . . 647
31.6 **omp_get_max_progress_width** Routine . . . 648
31.7 **omp_get_device_from_uid** Routine . . . 649
31.8 **omp_get_uid_from_device** Routine . . . 649
31.9 **omp_is_initial_device** Routine . . . 650
31.10 **omp_get_initial_device** Routine . . . 651
31.11 **omp_get_device_num_teams** Routine . . . 651
31.12 **omp_set_device_num_teams** Routine . . . 652
31.13 **omp_get_device_teams_thread_limit** Routine . . . 653
31.14 **omp_set_device_teams_thread_limit** Routine . . . 654
31.15 **omp_get_filtered_devices** Routine . . . 655
31.16 **omp_get_device_distances** Routine . . . 656

32 Device Memory Routines **657**

32.1 Asynchronous Device Memory Routines . . . 658
32.2 Device Memory Information Routines . . . 658
32.2.1 **omp_target_is_present** Routine . . . 658
32.2.2 **omp_target_is_accessible** Routine . . . 659
32.2.3 **omp_get_mapped_ptr** Routine . . . 660
32.3 **omp_target_alloc** Routine . . . 661
32.4 **omp_target_free** Routine . . . 662
32.5 **omp_target_associate_ptr** Routine . . . 663
32.6 **omp_target_disassociate_ptr** Routine . . . 665
32.7 Memory Copying Routines . . . 666
32.7.1 **omp_target_memcpy** Routine . . . 667
32.7.2 **omp_target_memcpy_rect** Routine . . . 669
32.7.3 **omp_target_memcpy_async** Routine . . . 670
32.7.4 **omp_target_memcpy_rect_async** Routine . . . 671

32.8 Memory Setting Routines . . . 673
32.8.1 **omp_target_memset** Routine . . . 674
32.8.2 **omp_target_memset_async** Routine . . . 675

33 Interoperability Routines **677**
33.1 **omp_get_num_interop_properties** Routine . . . 678
33.2 **omp_get_interop_int** Routine . . . 678
33.3 **omp_get_interop_ptr** Routine . . . 679
33.4 **omp_get_interop_str** Routine . . . 681
33.5 **omp_get_interop_name** Routine . . . 682
33.6 **omp_get_interop_type_desc** Routine . . . 683
33.7 **omp_get_interop_rc_desc** Routine . . . 683

34 Memory Management Routines **685**
34.1 Memory Space Retrieving Routines . . . 685
34.1.1 **omp_get_devices_memspace** Routine . . . 686
34.1.2 **omp_get_device_memspace** Routine . . . 687
34.1.3 **omp_get_devices_and_host_memspace** Routine . . . 687
34.1.4 **omp_get_device_and_host_memspace** Routine . . . 688
34.1.5 **omp_get_devices_all_memspace** Routine . . . 689
34.2 **omp_get_memspace_num_resources** Routine . . . 690
34.3 **omp_get_memspace_pagesize** Routine . . . 691
34.4 **omp_get_submemspace** Routine . . . 691
34.5 OpenMP Memory Partitioning Routines . . . 693
34.5.1 **omp_init_mempartitioner** Routine . . . 693
34.5.2 **omp_destroy_mempartitioner** Routine . . . 694
34.5.3 **omp_init_mempartition** Routine . . . 695
34.5.4 **omp_destroy_mempartition** Routine . . . 696
34.5.5 **omp_mempartition_set_part** Routine . . . 697
34.5.6 **omp_mempartition_get_user_data** Routine . . . 698
34.6 **omp_init_allocator** Routine . . . 699
34.7 **omp_destroy_allocator** Routine . . . 701
34.8 Memory Allocator Retrieving Routines . . . 702
34.8.1 **omp_get_devices_allocator** Routine . . . 703

34.8.2 **omp_get_device_allocator** Routine . . . 704
34.8.3 **omp_get_devices_and_host_allocator** Routine . . . 704
34.8.4 **omp_get_device_and_host_allocator** Routine . . . 705
34.8.5 **omp_get_devices_all_allocator** Routine . . . 706
34.9 **omp_set_default_allocator** Routine . . . 707
34.10 **omp_get_default_allocator** Routine . . . 708
34.11 Memory Allocating Routines . . . 709
34.11.1 **omp_alloc** Routine . . . 711
34.11.2 **omp_aligned_alloc** Routine . . . 712
34.11.3 **omp_calloc** Routine . . . 713
34.11.4 **omp_aligned_calloc** Routine . . . 714
34.11.5 **omp_realloc** Routine . . . 715
34.12 **omp_free** Routine . . . 716
34.13 Groupprivate Memory Routines . . . 717
34.13.1 **omp_get_dyn_gprivate_ptr** Routine . . . 718
34.13.2 **omp_get_dyn_gprivate_nofb_ptr** Routine . . . 719
34.13.3 **omp_get_dyn_gprivate_size** Routine . . . 720
34.13.4 **omp_get_dyn_gprivate_memspace** Routine . . . 721
34.13.5 **omp_get_gprivate_limit** Routine . . . 721

35 Lock Routines **723**

35.1 Lock Initializing Routines . . . 724
35.1.1 **omp_init_lock** Routine . . . 724
35.1.2 **omp_init_nest_lock** Routine . . . 725
35.1.3 **omp_init_lock_with_hint** Routine . . . 726
35.1.4 **omp_init_nest_lock_with_hint** Routine . . . 727
35.2 Lock Destroying Routines . . . 728
35.2.1 **omp_destroy_lock** Routine . . . 728
35.2.2 **omp_destroy_nest_lock** Routine . . . 729
35.3 Lock Acquiring Routines . . . 729
35.3.1 **omp_set_lock** Routine . . . 730
35.3.2 **omp_set_nest_lock** Routine . . . 731
35.4 Lock Releasing Routines . . . 732
35.4.1 **omp_unset_lock** Routine . . . 733

35.4.2 **omp_unset_nest_lock** Routine . . . 734
35.5 Lock Testing Routines . . . 735
35.5.1 **omp_test_lock** Routine . . . 735
35.5.2 **omp_test_nest_lock** Routine . . . 736

36 Thread Affinity Routines **738**
36.1 **omp_get_proc_bind** Routine . . . 738
36.2 **omp_get_num_places** Routine . . . 739
36.3 **omp_get_place_num_procs** Routine . . . 739
36.4 **omp_get_place_proc_ids** Routine . . . 740
36.5 **omp_get_place_num** Routine . . . 741
36.6 **omp_get_partition_num_places** Routine . . . 741
36.7 **omp_get_partition_place_nums** Routine . . . 742
36.8 **omp_set_affinity_format** Routine . . . 742
36.9 **omp_get_affinity_format** Routine . . . 743
36.10 **omp_display_affinity** Routine . . . 745
36.11 **omp_capture_affinity** Routine . . . 746

37 Execution Control Routines **748**
37.1 **omp_get_cancellation** Routine . . . 748
37.2 Resource Relinquishing Routines . . . 749
37.2.1 **omp_pause_resource** Routine . . . 749
37.2.2 **omp_pause_resource_all** Routine . . . 750
37.3 Timing Routines . . . 751
37.3.1 **omp_get_wtime** Routine . . . 751
37.3.2 **omp_get_wtick** Routine . . . 751
37.4 **omp_display_env** Routine . . . 752

38 Tool Support Routines **754**
38.1 **omp_control_tool** Routine . . . 754

V OMPT 756

39 OMPT Overview 757
39.1 OMPT Interfaces Definitions . . . 757
39.2 Activating a First-Party Tool . . . 757
39.2.1 **ompt_start_tool** Procedure . . . 757
39.2.2 Determining Whether to Initialize a First-Party Tool . . . 759
39.2.3 Initializing a First-Party Tool . . . 760
39.2.4 Monitoring Activity on the Host with OMPT . . . 763
39.2.5 Tracing Activity on Target Devices . . . 764
39.3 Finalizing a First-Party Tool . . . 768

40 OMPT Data Types 769
40.1 OMPT Predefined Identifiers . . . 769
40.2 OMPT **any_record_ompt** Type . . . 769
40.3 OMPT **buffer** Type . . . 771
40.4 OMPT **buffer_cursor** Type . . . 771
40.5 OMPT **callback** Type . . . 772
40.6 OMPT **callbacks** Type . . . 773
40.7 OMPT **cancel_flag** Type . . . 775
40.8 OMPT **data** Type . . . 775
40.9 OMPT **dependence** Type . . . 776
40.10 OMPT **dependence_type** Type . . . 777
40.11 OMPT **device** Type . . . 778
40.12 OMPT **device_time** Type . . . 778
40.13 OMPT **dispatch** Type . . . 779
40.14 OMPT **dispatch_chunk** Type . . . 779
40.15 OMPT **frame** Type . . . 780
40.16 OMPT **frame_flag** Type . . . 781
40.17 OMPT **hwid** Type . . . 782
40.18 OMPT **id** Type . . . 782
40.19 OMPT **interface_fn** Type . . . 783
40.20 OMPT **mutex** Type . . . 783
40.21 OMPT **native_mon_flag** Type . . . 784

40.22 OMPT **parallel_flag** Type . . . 785
40.23 OMPT **record** Type . . . 786
40.24 OMPT **record_abstract** Type . . . 786
40.25 OMPT **record_native** Type . . . 788
40.26 OMPT **record_ompt** Type . . . 788
40.27 OMPT **scope_endpoint** Type . . . 789
40.28 OMPT **set_result** Type . . . 790
40.29 OMPT **severity** Type . . . 791
40.30 OMPT **start_tool_result** Type . . . 792
40.31 OMPT **state** Type . . . 793
40.32 OMPT **subvolume** Type . . . 795
40.33 OMPT **sync_region** Type . . . 796
40.34 OMPT **target** Type . . . 797
40.35 OMPT **target_data_op** Type . . . 798
40.36 OMPT **target_map_flag** Type . . . 800
40.37 OMPT **task_flag** Type . . . 801
40.38 OMPT **task_status** Type . . . 802
40.39 OMPT **thread** Type . . . 803
40.40 OMPT **wait_id** Type . . . 804
40.41 OMPT **work** Type . . . 805

41 General Callbacks and Trace Records 806

41.1 Initialization and Finalization Callbacks . . . 807
41.1.1 **initialize** Callback . . . 807
41.1.2 **finalize** Callback . . . 808
41.1.3 **thread_begin** Callback . . . 808
41.1.4 **thread_end** Callback . . . 809
41.2 **error** Callback . . . 810
41.3 Parallelism Generation Callback Signatures . . . 811
41.3.1 **parallel_begin** Callback . . . 811
41.3.2 **parallel_end** Callback . . . 812
41.3.3 **masked** Callback . . . 813
41.4 Work Distribution Callback Signatures . . . 814
41.4.1 **work** Callback . . . 814

41.4.2 **dispatch** Callback . . . 815
41.5 Tasking Callback Signatures . . . 817
41.5.1 **task_create** Callback . . . 817
41.5.2 **task_schedule** Callback . . . 818
41.5.3 **implicit_task** Callback . . . 819
41.6 **cancel** Callback . . . 820
41.7 Synchronization Callback Signatures . . . 821
41.7.1 **dependences** Callback . . . 821
41.7.2 **task_dependence** Callback . . . 822
41.7.3 OMPT **sync_region** Type . . . 823
41.7.4 **sync_region** Callback . . . 824
41.7.5 **sync_region_wait** Callback . . . 825
41.7.6 **reduction** Callback . . . 826
41.7.7 OMPT **mutex_acquire** Type . . . 826
41.7.8 **mutex_acquire** Callback . . . 827
41.7.9 **lock_init** Callback . . . 828
41.7.10 OMPT **mutex** Type . . . 828
41.7.11 **lock_destroy** Callback . . . 829
41.7.12 **mutex_acquired** Callback . . . 829
41.7.13 **mutex_released** Callback . . . 830
41.7.14 **nest_lock** Callback . . . 830
41.7.15 **flush** Callback . . . 831
41.8 **control_tool** Callback . . . 832

42 Device Callbacks and Tracing **834**

42.1 **device_initialize** Callback . . . 834
42.2 **device_finalize** Callback . . . 835
42.3 **device_load** Callback . . . 836
42.4 **device_unload** Callback . . . 837
42.5 **buffer_request** Callback . . . 837
42.6 **buffer_complete** Callback . . . 838
42.7 **target_data_op_emi** Callback . . . 839
42.8 **target_emi** Callback . . . 842
42.9 **target_map_emi** Callback . . . 844

42.10 **target_submit_emi** Callback 846

43 General Entry Points **849**

43.1 **function_lookup** Entry Point 849
43.2 **enumerate_states** Entry Point 850
43.3 **enumerate_mutex_impls** Entry Point 851
43.4 **set_callback** Entry Point 852
43.5 **set_selective_callback** Entry Point 853
43.6 **get_callback** Entry Point 854
43.7 **get_thread_data** Entry Point 855
43.8 **get_num_procs** Entry Point 856
43.9 **get_num_places** Entry Point 856
43.10 **get_place_proc_ids** Entry Point 857
43.11 **get_place_num** Entry Point 857
43.12 **get_partition_place_nums** Entry Point 858
43.13 **get_proc_id** Entry Point 859
43.14 **get_state** Entry Point 859
43.15 **get_parallel_info** Entry Point 860
43.16 **get_task_info** Entry Point 861
43.17 **get_task_memory** Entry Point 863
43.18 **get_target_info** Entry Point 864
43.19 **get_num_devices** Entry Point 865
43.20 **get_unique_id** Entry Point 865
43.21 **finalize_tool** Entry Point 866

44 Device Tracing Entry Points **867**

44.1 **get_device_num_procs** Entry Point 867
44.2 **get_uid_from_device** Entry Point 867
44.3 **get_device_time** Entry Point 868
44.4 **translate_time** Entry Point 869
44.5 **set_trace_ompt** Entry Point 869
44.6 **set_trace_native** Entry Point 870
44.7 **get_buffer_limits** Entry Point 871
44.8 **start_trace** Entry Point 872

44.9 **pause_trace** Entry Point . . . 873
44.10 **flush_trace** Entry Point . . . 873
44.11 **stop_trace** Entry Point . . . 874
44.12 **advance_buffer_cursor** Entry Point . . . 875
44.13 **get_record_type** Entry Point . . . 875
44.14 **get_record_ompt** Entry Point . . . 876
44.15 **get_record_native** Entry Point . . . 877
44.16 **get_record_abstract** Entry Point . . . 878

VI OMPD **879**

45 OMPD Overview **880**
45.1 OMPD Interfaces Definitions . . . 881
45.2 Thread and Signal Safety . . . 881
45.3 Activating a Third-Party Tool . . . 881
45.3.1 Enabling Runtime Support for OMPD . . . 881
45.3.2 **ompd_dll_locations** . . . 881
45.3.3 **ompd_dll_locations_valid** Breakpoint . . . 882

46 OMPD Data Types **883**
46.1 OMPD **addr** Type . . . 883
46.2 OMPD **address** Type . . . 883
46.3 OMPD **address_space_context** Type . . . 884
46.4 OMPD **callbacks** Type . . . 884
46.5 OMPD **device** Type . . . 886
46.6 OMPD **device_type_sizes** Type . . . 887
46.7 OMPD **frame_info** Type . . . 887
46.8 OMPD **icv_id** Type . . . 888
46.9 OMPD **rc** Type . . . 889
46.10 OMPD **seg** Type . . . 890
46.11 OMPD **scope** Type . . . 891
46.12 OMPD **size** Type . . . 892
46.13 OMPD **team_generator** Type . . . 892
46.14 OMPD **thread_context** Type . . . 893

46.15 OMPD **thread_id** Type . . . 893
46.16 OMPD **wait_id** Type . . . 894
46.17 OMPD **word** Type . . . 894
46.18 OMPD Handle Types . . . 895
46.18.1 OMPD **address_space_handle** Type . . . 895
46.18.2 OMPD **parallel_handle** Type . . . 895
46.18.3 OMPD **task_handle** Type . . . 896
46.18.4 OMPD **thread_handle** Type . . . 896

47 OMPD Callback Interface **897**
47.1 Memory Management of OMPD Library . . . 897
47.1.1 **alloc_memory** Callback . . . 898
47.1.2 **free_memory** Callback . . . 898
47.2 Accessing Program or Runtime Memory . . . 899
47.2.1 **symbol_addr_lookup** Callback . . . 899
47.2.2 OMPD **memory_read** Type . . . 901
47.2.3 **write_memory** Callback . . . 903
47.3 Context Management and Navigation . . . 904
47.3.1 **get_thread_context_for_thread_id** Callback . . . 904
47.3.2 **sizeof_type** Callback . . . 906
47.4 Device Translating Callbacks . . . 907
47.4.1 OMPD **device_host** Type . . . 907
47.4.2 **device_to_host** Callback . . . 908
47.4.3 **host_to_device** Callback . . . 908
47.5 **print_string** Callback . . . 908

48 OMPD Routines **910**
48.1 OMPD Library Initialization and Finalization . . . 910
48.1.1 **ompd_initialize** Routine . . . 910
48.1.2 **ompd_get_api_version** Routine . . . 911
48.1.3 **ompd_get_version_string** Routine . . . 912
48.1.4 **ompd_finalize** Routine . . . 913
48.2 Process Initialization and Finalization . . . 913
48.2.1 **ompd_process_initialize** Routine . . . 913

48.2.2 **ompd_device_initialize** Routine 914
48.2.3 **ompd_get_device_thread_id_kinds** Routine 915
48.3 Address Space Information 916
48.3.1 **ompd_get_omp_version** Routine 916
48.3.2 **ompd_get_omp_version_string** Routine 917
48.4 Thread Handle Routines 918
48.4.1 **ompd_get_thread_in_parallel** Routine 918
48.4.2 **ompd_get_thread_handle** Routine 919
48.4.3 **ompd_get_thread_id** Routine 920
48.4.4 **ompd_get_device_from_thread** Routine 921
48.5 Parallel Region Handle Routines 922
48.5.1 **ompd_get_curr_parallel_handle** Routine 922
48.5.2 **ompd_get_enclosing_parallel_handle** Routine 922
48.5.3 **ompd_get_task_parallel_handle** Routine 923
48.6 Task Handle Routines 924
48.6.1 **ompd_get_curr_task_handle** Routine 924
48.6.2 **ompd_get_generating_task_handle** Routine 925
48.6.3 **ompd_get_scheduling_task_handle** Routine 926
48.6.4 **ompd_get_task_in_parallel** Routine 926
48.6.5 **ompd_get_task_function** Routine 927
48.6.6 **ompd_get_task_frame** Routine 928
48.7 Handle Comparing Routines 929
48.7.1 **ompd_parallel_handle_compare** Routine 929
48.7.2 **ompd_task_handle_compare** Routine 930
48.7.3 **ompd_thread_handle_compare** Routine 930
48.8 Handle Releasing Routines 931
48.8.1 **ompd_rel_address_space_handle** Routine 931
48.8.2 **ompd_rel_parallel_handle** Routine 932
48.8.3 **ompd_rel_task_handle** Routine 932
48.8.4 **ompd_rel_thread_handle** Routine 933
48.9 Querying Thread States 934
48.9.1 **ompd_enumerate_states** Routine 934
48.9.2 **ompd_get_state** Routine 935

48.10 Display Control Variables . . . 936
48.10.1 **ompd_get_display_control_vars** Routine . . . 936
48.10.2 **ompd_rel_display_control_vars** Routine . . . 937
48.11 Accessing Scope-Specific Information . . . 937
48.11.1 **ompd_enumerate_icvs** Routine . . . 937
48.11.2 **ompd_get_icv_from_scope** Routine . . . 938
48.11.3 **ompd_get_icv_string_from_scope** Routine . . . 939
48.11.4 **ompd_get_tool_data** Routine . . . 940

49 OMPD Breakpoint Symbol Names 942
49.1 **ompd_bp_thread_begin** Breakpoint . . . 942
49.2 **ompd_bp_thread_end** Breakpoint . . . 942
49.3 **ompd_bp_device_begin** Breakpoint . . . 943
49.4 **ompd_bp_device_end** Breakpoint . . . 943
49.5 **ompd_bp_parallel_begin** Breakpoint . . . 944
49.6 **ompd_bp_parallel_end** Breakpoint . . . 944
49.7 **ompd_bp_teams_begin** Breakpoint . . . 945
49.8 **ompd_bp_teams_end** Breakpoint . . . 945
49.9 **ompd_bp_task_begin** Breakpoint . . . 946
49.10 **ompd_bp_task_end** Breakpoint . . . 946
49.11 **ompd_bp_target_begin** Breakpoint . . . 947
49.12 **ompd_bp_target_end** Breakpoint . . . 947

VII Appendices 949

A OpenMP Implementation-Defined Behaviors 950

B Features History 962
B.1 Deprecated Features . . . 962
B.2 Version 6.0 to 6.1 Differences . . . 963
B.3 Version 5.2 to 6.0 Differences . . . 965
B.4 Version 5.1 to 5.2 Differences . . . 973
B.5 Version 5.0 to 5.1 Differences . . . 975
B.6 Version 4.5 to 5.0 Differences . . . 978

B.7 Version 4.0 to 4.5 Differences 982
B.8 Version 3.1 to 4.0 Differences 984
B.9 Version 3.0 to 3.1 Differences 985
B.10 Version 2.5 to 3.0 Differences 985

C Nesting of Regions **988**

D Conforming Compound Directive Names **990**

Index **994**

List of Figures

39.1 First-Party Tool Activation Flow Chart 759

List of Tables

3.1 ICV Scopes and Descriptions . . . 131
3.2 ICV Initial Values . . . 134
3.3 Ways to Modify and to Retrieve ICV Values . . . 137
3.4 ICV Override Relationships . . . 143

4.1 Predefined Place-group Abstract Names . . . 145
4.2 Available Field Types for Formatting OpenMP Thread Affinity Information . . . 156
4.3 Reservation Types for **OMP_THREADS_RESERVE** . . . 160

5.1 Syntactic Properties for Clauses, Arguments, Modifiers, Clause groups, and Modifier groups . . . 178

8.1 Implicitly Declared C/C++ Reduction Identifiers . . . 261
8.2 Implicitly Declared Fortran Reduction Identifiers . . . 262
8.3 Implicitly Declared C/C++ Induction Identifiers . . . 263
8.4 Implicitly Declared Fortran Induction Identifiers . . . 263

11.1 Map-Type Decay of Map Type Combinations . . . 316

13.1 Predefined Memory Spaces . . . 335
13.2 Allocator Traits . . . 336
13.3 Predefined Allocators . . . 339

19.1 **work** OMPT types for Worksharing-Loop . . . 455

20.1 **task_create** Callback Flags Evaluation . . . 469

26.1 Routine Argument Properties . . . 580
26.2 Required Values of the **interop_property** OpenMP Type . . . 586
26.3 Required Values for the **interop_rc** OpenMP Type . . . 588
26.4 Allowed Key-Values for **alloctrait** OpenMP Type . . . 592
26.5 Standard Tool Control Commands . . . 612

39.1 OMPT Callback Interface Runtime Entry Point Names and Their Type Signatures . . . 762
39.2 Callbacks for which **set_callback** Must Return **ompt_set_always** . . . 764
39.3 OMPT Tracing Interface Runtime Entry Point Names and Their Type Signatures . . . 765

42.1 Association of dev1 and dev2 arguments for target data operations 842

46.1 Mapping of Scope Type and OMPD Handles 892

Part I

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 that are defined implicitly or explicitly by OpenMP directives and their associated user code, if any. Each such directive has the task-generating property and, thus is a task-generating construct. Any task-generating construct that generates one or more explicit tasks has the explicit-task-generating property and, thus, is an explicit-task-generating construct. Alternatively, any task-generating construct that generates one or more implicit tasks has the implicit-task-generating property and, thus, is an implicit-task-generating construct. 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, which is an explicit 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 implicit task region (corresponding to an initial task) that is generated by an implicit parallel region. 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 implicit task region (corresponding to an initial task 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 an explicit-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. 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 7.3.2 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 11.5.

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 is 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 variables 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 35, 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.

Since OpenMP threads execute tasks, flushes are specified by where they occur in the execution of tasks. In order to ensure that a given task begins with a memory state that is consistent with the memory at the point at which it is generated, an implicit acquire-at-entry flush occurs in the task that synchronizes with an implicit release-at-entry flush that occurs in its generating task. Further, to ensure that its state is available to other tasks, an implicit release-at-exit flush occurs in each task. Additional implicit flushes that occur in synchronization constructs and routines ensure that the state of memory is consistent with their semantics. In addition, the OpenMP API includes the **`flush`** directive so that users may explicitly enforce memory consistency.

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*.

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

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

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

1.3.6 OpenMP Memory Consistency

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

- If two operations performed by different threads are sequentially consistent atomic operations or they are strong flushes that flush the same variable, then they must be completed as if in some sequential order, seen by all affected threads.
- If two operations performed by the same thread are sequentially consistent atomic operations or they access, modify, or, with a strong flush, flush the same variable, then they must be completed as if in the program order of that thread, as seen by all affected threads.
- If two operations are performed by different threads and one happens before the other, then they must be completed as if in that happens-before order, as seen by all affected threads, if:

- both operations access or modify the same variable;
- both operations are strong flushes that flush the same variable; or
- both operations are sequentially consistent atomic operations.

- Any two atomic operations from different **`atomic`** regions must be completed as if in the same order as the strong flushes implied in their regions, as seen by all affected threads.

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

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

OpenMP programs that:

- Do not use non-sequentially consistent atomic constructs;
- Do not rely on the accuracy of a *false* result from **`omp_test_lock`** and **`omp_test_nest_lock`**; and
- Correctly avoid data races as required in Section 1.3.1,

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 39.2.4 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 45, Chapter 46, Chapter 47, Chapter 48, and Chapter 49 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

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

Fortran / C++

Fortran/C++ specific text...

Fortran / C++

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

Fortran (cont.)

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

C/C++ (cont.)

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

Note – Non-normative text...

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. **145**, 36, 85, 145, 148, 149, 152, 951, 963, 965

access group

A group of tasks that has the same value for the grouping criterion defined by an access group type. 20, 349, 350, 717, 718, 954

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. 20, 59, 349, 589, 717, 718, 722

accessible device

The host device or any non-host device accessible for execution. 135, 157–159, 396

accessible storage

A storage block that may be accessed by a given thread. 325, 509, 660

acquire flush

A flush that has the acquire flush property. **11**, 11, 12, 62, 101, 113, 431, 546, 548, 643

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. 20, 56, 546

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. 21, 84, 113, 147, 151, 622, 951, 959, 981

active parallel region

A parallel region comprised of implicit tasks that are being executed by a team to which multiple threads are assigned. 21, 119, 131, 132, 151, 294, 295, 617, 622–625, 950, 953, 985–987

active spatial dimension

An execution space dimension that is configured with more than one execution space unit. 65, 113

active target region

A **`target`** region that is executed on a device other than the device that encountered the **`target`** construct. 142

address range

The addresses of a contiguous set of storage locations. 54, 77, 110, 660

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. 21, 55, 89, 104, 125, 163, 164, 395, 660, 760, 884, 895, 900, 902, 904–907, 911, 914, 915, 917–919, 934, 936, 938

address space context

A tool context that refers to an address space within an OpenMP process. 884

address space handle

A handle that refers to an address space within an OpenMP process. 892, 914–916, 921, 932

affected iteration

A logical iteration of the affected loops of a loop-nest-associated directive. 67, 104, 107, 419

affected loop

A loop from a semantic loop nest, or a **`DO CONCURRENT`** loop in Fortran, that is affected by a given loop-nest-associated directive. **226**, 5, 21, 69, 74, 75, 91, 105, 123, 129, 173, 221, 223, 226–228, 234, 235, 242, 247, 248, 272, 278–280, 289, 290, 406, 410–412, 415, 416, 418, 464, 955, 980

affected loop nest

The subset of semantic loop nests of an associated loop sequence that are selected by the **looprange** clause. **231**, 36, 102, 227, 228, 411, 412

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, 22, 42, 128, 183, 203, 295, 487, 950

aligned-memory-allocating routine

A memory-management routine that has the aligned-memory-allocating-routine property. **709**, 710, 712, 714

aligned-memory-allocating-routine property

The property that a memory-allocating routine ensures the allocated memory is aligned with respect to an *alignment* argument. **709**, 22, 712, 714

all-constituents property

The property that a clause applies to all leaf constructs that permit it when the clause appears on a compound directive. **178**, 179, 573

all-contention-group-tasks binding property

The binding property that the binding task set is all tasks in the contention group. **579**, 724–731, 733–736

all-data-environments clause

A clause that has the all-data-environments property. 81, 250, 253

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. 22, 250, 252, 275

all-device-tasks binding property

The binding property that the binding task set is all tasks on a specified device. **749**

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. **579**, 639, 647, 685–691, 693–699, 701, 703–706, 739, 740, 856, 857

allocator

A memory allocator. 23, 162, 336–348, 394, 508, 591, 592, 601, 604, 694, 695, 700, 702, 707–710, 716, 954, 968, 974

allocator structured block

A context-specific structured block that may be associated with an **`allocators`** directive. **209**, 346

allocator trait

A trait of an allocator. 162, 336, 338, 339, 342, 344, 593, 595, 598, 694, 700, 701, 954, 968, 969, 980

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. **178**, 343, 573

all tasks

All tasks participating in the OpenMP program or in a specified limiting context. 20, 22, 23, 29, 268, 298, 337, 349, 579, 750

all-tasks binding property

The binding property that the binding task set is all tasks. **750**, 749, 750

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. 9, 13, 22, 23, 29, 248, 540, 579, 685, 751, 856, 857

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. **579**

ancestor thread

For a given thread, its parent thread or one of the ancestor threads of its parent thread. 23, 624, 642, 971, 987

antecedent task

A task that must complete before its dependent tasks can be executed. **552**, 44, 54, 65, 95, 115, 479, 492, 552, 553, 823

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. 186, 187, 251–253, 264, 306, 320, 321

array element

A single member of an array as defined by the base language. 24, 77, 128, 252, 253, 257, 263, 264, 278, 291, 292, 305, 313, 320, 330, 331

array item

An array, an array section, or an array element. 574

array-map expansion

The process by which array elements of an array or array section that appears as a list item in a `map` clause become implicit list items with the same modifiers as the array or array section. **320**, 106, 313, 320, 325

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. 23–25, 27–29, 37, 38, 40, 82, 108, 112, 128, 130, 158, 182, 185–187, 240, 250–253, 257, 258, 260, 263, 264, 273, 278, 291, 292, 311, 312, 320–322, 328, 330, 433, 486, 553, 574, 966, 976, 979, 981, 982, 984

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*. 106, 979

assignable OpenMP type instance

An instance of an OpenMP type to which an assignment can be performed. **205**, 205

assigned list item

A list item to which assignment is performed as the result of a data-motion clause. 331–333

assigned thread

A thread that has been assigned an implicit task of a parallel region. 3, 4, 96, 118, 121, 454, 615, 632

assigning map type

A map type for which the mapping operations may include an assignment operation. **315**

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. 74, 75, 221, 226, 229, 230, 410

associated loop sequence

The associated canonical loop sequence of a loop-sequence-associated directive. 22, 231, 412

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, 25, 79, 336, 339

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. 25, 78, 112, 130, 186, 187, 221, 235, 236, 240, 250, 253, 311–313, 315, 320, 580, 968, 986

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. 25, 398, 400, 973, 976

assumption scope

The scope for which the invariants specified by an assumption directive must hold. 398–405

asynchronous device routine

A routine that has the asynchronous-device routine property. 549, 657, 658, 671, 673, 676

asynchronous-device routine property

The property of a device routine that it performs its operation asynchronously. 25, 658, 670, 671, 675

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. 25, 806, 839, 849

async-signal-safe entry point

An entry point that has the async-signal-safe property. 849

async-signal-safe property

The property of a routine or entry point that it is async signal safe. **7**, 25, 849, 855–861, 863–865

atomic captured update

An atomic update that captures the read value into a separate storage location. 36, 124, 126, 215, 538, 542, 984

atomic conditional update

An atomic update for which the write to the storage location is conditionally applied. 36, 124, 126, 213, 538, 539, 542, 544, 977

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, 12, 14, 26, 99, 101, 105, 323, 339, 518, 542–544, 977

atomic read

An atomic operation that reads from but does not write to a given storage location. 99, 212, 535, 542

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. 9, 12, 339, 541

atomic structured block

A context-specific structured block that may be associated with an **`atomic`** directive. **210**, 31, 99, 126, 130, 210, 216, 541–543, 963, 967

atomic update

An atomic operation that reads from and writes to a given storage location. 25, 26, 124, 126, 213, 536, 538, 542–544, 984

atomic write

An atomic operation that writes to but does not read from a given storage location. 130, 212, 537, 542

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. An attached pointer can be re-attached via pointer attachment to point to a different data entity, as explained in Section 11.5.5, but it remains an attached pointer for the remainder of its lifetime in the device data environment. 26, 94, 319, 324, 331, 508, 509

attach-ineligible

An attribute of a pointer for which pointer attachment may not be performed. 318, 319, 321

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. 234, 237, 238, 342

available device

An available non-host device or, where explicitly specified, the host device. 157, 159, 355, 655, 690, 707, 750

available non-host device

A non-host device that can be used for the current OpenMP program execution. 27, 157

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. **521**, 5, 6, 27, 53, 62, 122, 303, 422, 442, 444–447, 449, 454, 490, 491, 521–524, 529, 543, 566, 643, 749, 763, 794, 825, 971, 988

barrier scheduling constraint

The synchronization scheduling constraint that ensures that the synchronization semantics of a barrier, regardless of whether it is an explicit barrier or an implicit barrier, are enforced. **491**, 491, 521, 523

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. 54, 250, 253, 320, 664, 718, 720

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. 23, 574

base function

A procedure that is declared and defined in the base language. 43, 57, 102, 128, 358, 359, 364–369, 371–374, 955

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, 7, 8, 13, 16–19, 24, 25, 27–29, 37, 39, 40, 42, 48, 50, 53, 56, 57, 60, 89, 95, 96, 98, 102, 103, 110, 112, 113, 166, 170, 172, 174, 175, 181, 183, 184, 186, 188, 205–207, 212, 213, 219, 223, 224, 239, 252, 253, 256, 257, 259, 264, 265, 277, 282, 283, 286, 287, 293, 307, 320, 328, 329, 339, 340, 342, 346, 348, 366, 371, 373, 398, 451, 542, 560, 578, 580, 610, 950, 973, 979

base-language code

Code that is part of the base program, exclusive of any directives or implementation code. 30, 43, 50, 69, 124, 166, 171, 205, 218, 974

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. 7, 28

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. 23, 27, 39, 82, 234, 235, 253, 278, 305, 313, 318, 319, 321, 325, 326, 477, 479, 506, 508, 573, 574, 964

base program

A program written in a base language. 2, 28, 89

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 such that 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. 29, 235, 318, 506

base referring pointer

For a given data entity that has a base referencing variable, the referring pointer of the base referencing variable. 235, 319, 326, 964

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. 27, 241, 265, 266, 295, 305, 325, 477, 479, 508, 573, 574

binding implicit task

The implicit task of the current team assigned to the encountering thread. 29, 61, 141, 428, 707–709

binding-implicit-task binding property

The binding property that the binding task set is the binding implicit task. 707, 708

binding property

A property of a construct or a routine that determines the binding region, binding task set and/or binding thread set. 22, 23, 29, 51, 58, 579

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. 5, 29, 91, 228, 452, 464, 466, 522, 558, 561, 564, 568, 579, 743, 744, 944–948, 960, 989

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. 22, 23, 29, 58, 137, 374, 477, 498, 500, 502, 505, 511, 514, 525, 529, 579, 657, 707, 708, 750, 849, 944–948

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, 22, 23, 29, 51, 91, 93, 102, 121, 129, 228, 245, 248, 421, 431, 437, 438, 442, 444–447, 449, 452–454, 461, 464–468, 471, 472, 477, 481, 489, 519, 521, 526, 529, 540, 541, 543, 545, 549, 558, 560, 564, 567, 568, 579, 685, 743, 744, 849, 856, 857, 960, 970, 971

block-associated directive

A directive for which its associated base-language code is a structured block. **171**, 38, 91, 169, 171–173, 208, 346, 374, 404, 421, 431, 442, 445–447, 449, 452, 467, 476, 502, 505, 519, 524, 541, 560

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. 226

C

callback

A tool callback. xxix, 15, 30, 34, 47, 48, 80–82, 86, 87, 89, 92, 95, 101, 104, 105, 114, 126, 267, 312, 382, 388, 423, 433, 442, 445, 447–449, 451, 453, 455, 462, 468, 469, 472, 488, 491, 497, 499, 501, 504, 507, 512, 520, 522–525, 527, 528, 543, 544, 547, 553, 557, 559, 560, 565, 567, 568, 643, 657, 658, 662–664, 666–668, 670, 671, 673–676, 724–729, 731–734, 736, 737, 755, 757, 758, 760, 761, 763–768, 781, 786, 791, 792, 799, 806–841, 843, 845, 847, 849, 850, 853–855, 866, 867, 870–872, 874, 876, 880, 881, 885, 886, 890, 897–907, 909, 911, 913, 916, 918, 934, 936, 938, 940, 960–962, 965, 972, 978

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, 790, 791, 854, 871

callback registration

A process that makes a tool callback available to an OpenMP implementation to enable callback dispatch. 101, 105, 761, 763

canceled taskgroup set

A taskgroup set that has been canceled. **567**, 567

cancellable construct

A construct that has the cancellable property. **563**, 564, 569

cancellable property

The property that a construct may be subject to cancellation. **563**, 30, 421, 447, 456, 457, 524

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. **567**, 6, 27, 30, 31, 157, 444, 522, 563–567, 569, 748, 821, 984

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. **566**, 6, 31, 127, 132, 157, 455, 491, 522, 523, 565–567, 569, 803

cancellation-point property

The property that a construct includes a cancellation point. **566**, 421, 445–447, 449, 455–457, 521–523, 564, 568

candidate

A replacement candidate. 361, 364

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. 782

canonical loop nest

A loop nest that complies with the rules and requirements defined in Section 6.4.1. **218**, 24, 30, 31, 54, 74, 75, 105, 172, 218, 221, 223, 224, 226, 229, 231, 247, 289, 405, 459, 460, 576, 969, 970, 979

canonical loop sequence

A sequence of canonical loop nests that complies with the rules defined in Section 6.4.2. **225**, 25, 75, 105, 172, 231, 967, 969

canonical order

The order that can be imposed on the execution of OpenMP operations within a construct. 84, 90, 105, 119, 121, 272

capture structured block

An atomic structured block that may be associated with an **`atomic`** directive that expresses capture semantics. **214**, 36, 126, 214, 215, 544

C/C++-only property

The property that an OpenMP feature or a routine argument is only supported in C/C++. **580**, 581, 744, 746, 754, 769–773, 775–793, 795–798, 800–805, 807–815, 817–832, 834–839, 842, 844, 846, 849–861, 863–878, 883, 884, 886–896

C/C++ pointer property

The property that a routine argument has a pointer type in C/C++ but not the **`POINTER`** attribute in Fortran. **580**, 600, 602, 620, 693–697, 699, 724–731, 733–736, 754

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. 32, 44, 54, 57, 65, 106, 123, 492, 526, 527, 552, 555, 605, 964

chunk

A contiguous non-empty subset of the collapsed iterations of a loop-collapsing construct. 104, 153, 454, 455, 459, 460, 462, 463, 471, 576, 620, 780, 816, 960, 964

class type

For C++, the type of any variable declared with one of the **`class`**, **`struct`**, or **`union`** keywords. 238, 240, 244, 245, 247, 248, 261, 266, 274, 277, 295, 301–304, 322, 324, 509

clause

A mechanism to specify customized directive behavior. xxix, 4–6, 8, 9, 22–24, 32–36, 40–52, 55, 57–59, 61, 68, 69, 71, 76–78, 81, 84, 85, 87, 88, 90, 91, 95–97, 99–105, 109, 112–114, 116–118, 124, 125, 132, 135, 139, 142–144, 146, 151, 162, 166–168, 170, 171, 176–181, 183–190, 192, 193, 198, 200, 202–204, 225–227, 229–231, 233–245, 247–257, 260, 263–282, 284–286, 288–295, 298, 300–333, 335, 340–352, 355–403, 405, 407–413, 415–417, 419, 420, 422, 424–427, 430–433, 435–450, 454, 455, 459–468, 471, 472, 474–487, 491, 493–496, 498–501, 503, 506–509, 511, 512, 514–518, 520, 526–529, 531–563, 565, 566, 573–576, 580, 607, 614, 616, 628, 639, 643, 652, 658, 661, 662, 664, 701, 702, 707, 708, 710, 721, 738, 776, 778, 803, 810, 821, 822, 845, 953–957, 963–971, 973–977, 979–986, 988, 989

clause set

A set of clauses for which restrictions on their use or other properties of their use on a given directive are specified. **178**, 32, 34, 52, 57, 90, 102, 124, 125, 178–180, 233, 392, 399, 472

clause group

A clause set for which restrictions or properties related to their use on all directives are specified. xxix, **176**, 178, 379, 392, 399, 531, 535, 537, 561, 563, 969

clause-list trait

A trait that is defined with properties that match the clauses that may be specified for a given directive. 354–357

closely nested construct

A construct nested inside another construct with no other construct nested between them. 49, 451, 453, 466, 565, 569

closely nested directive

A directive nested inside another construct with no other construct nested between them. 173

closely nested region

A region nested inside another region with no parallel region nested between them. 93, 276, 444, 466, 565, 569, 985

code block

A contiguous region of memory that contains code of an OpenMP program to be executed on a device. 496

collapse-affected loop

For a loop-collapsing construct, a loop that is affected by the **`collapse`** clause. **5**, 33, 74, 117, 227, 228, 238, 286, 439, 454, 460, 463–465, 474–476, 561, 970

collapsed iteration space

The logical iteration space of the collapse-affected loops of a loop-collapsing construct. **227**, 33, 286, 289, 440, 455, 459, 462, 463, 964

collapsed iteration

A logical iteration of the collapse-affected loops of a loop-collapsing construct. 32, 33, 36, 66, 75, 104, 129, 228, 238, 261, 277, 279–281, 289, 290, 437, 438, 441, 444, 454, 455, 459–464, 471, 560, 561, 576, 815, 816, 964

collapsed logical iteration

A collapsed iteration. 228, 238

collapsed loop

A loop resulting from a loop-collapsing construct that iterates over the collapsed iteration space. 227, 279, 953

collective step expression

An expression in terms of a step expression and a collector that eliminates recursive calculation in an induction operation. **66**, 33, 261

collector

A binary operator used to eliminate recursion in an induction operation. **66**, 33, 288

collector expression

An OpenMP stylized expression that evaluates to the value of the collective step expression of a collapsed iteration. **261**, 66, 192, 261, 263, 286, 288

combined construct

A construct that is a shortcut for specifying one construct immediately nested inside a leaf construct. **575**, 23, 34, 35, 981, 982

combined directive

A compound directive that is used to form a combined construct. 34, 35, 570

combined-directive name

The name of a combined directive. **570**

combiner

A binary operator used by a reduction operation. **266**, 34, 99, 205, 256, 269, 270, 272

combiner expression

An OpenMP stylized expression that specifies how a reduction combines partial results into a single value. **257**, 99, 192, 257, 258, 264, 265, 268, 282, 284, 289, 962

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. 787, 788

common-type-callback property

The property that a callback has a type that at least one other callback has. 824–830, 902, 903, 908

commutative reduction

A reduction for which the combiner can be applied to the set of values in any order. 272

compatible context selector

A context selector that matches the OpenMP context in which a directive is encountered. **359**, 360–362, 364, 365

compatible map type

A map type that is consistent with the data-motion attribute of a given data-motion clause. 330, 332, 333

compatible property

The property that a clause, an argument, a modifier, a clause set, or a modifier set does not have the exclusive property. 178

compilation unit

For C/C++, a translation unit. For Fortran, a program unit. 9, 46, 173, 296, 297, 299, 308, 314, 342, 343, 345, 388, 391–393, 397, 403, 508, 662, 701, 702, 710, 953, 954

compile-time error termination

Error termination that is performed during compilation. **6**, 392, 427, 956

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. **417**, 93

complex modifier

A modifier that may take at least one argument when it is specified. **177**, 35, 177, 180, 187

complex property

The property that a modifier is a complex modifier. 199, 516

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. **16**, 2, 5, 16, 18, 35, 44, 61, 125, 154, 166, 460, 543, 578, 723, 757, 850, 880, 881, 957

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. 23, 289, 316, 576, 968, 971

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. 35, 108, 503, 571

composite-directive name

The directive name of a composite directive. **570**

compound construct

A construct that corresponds to a compound directive. 36, 68, 87, 91, 107, 193, 198, 273, 354, 561, 572–576, 966, 983, 989, 990

compound directive

A combined directive or a composite directive. 22, 34, 35, 37, 71, 179, 570, 571, 573

compound-directive name

The directive name of a compound directive. **570**, 48, 570, 572, 971, 990

compound target construct

A compound construct for which **target** is a constituent construct. 305, 306, 574

conceptual abstract name

An abstract name that refers to an implementation defined abstraction that is relevant to the execution model described by this specification. **145**, 20, 85, 94, 145

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. **214**, 215, 216, 544, 962

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. **213**, 213, 214, 544

conforming device number

A device number that may be used in a conforming program. **7**, 159, 336, 357, 359, 495, 593, 635–638, 645, 652–654, 656, 657, 686, 702, 722, 750, 964

conforming program

An OpenMP program that follows all rules and restrictions of the OpenMP specification. 2, 16, 35, 36, 84, 88, 126, 361, 460

C-only property

The property that an OpenMP feature is only supported in C. 757, 773, 884, 889, 891, 892, 898, 899, 901–904, 906–908, 910–940

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. **228**, 36, 228, 418, 444

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. **180**, 56, 103, 170, 179, 182, 200, 202, 205, 224, 227–231, 273, 292, 335, 340, 342, 344, 348, 352, 358, 359, 380, 386, 390, 393–398, 401–403, 408, 413, 419, 420, 425, 435, 436, 440, 441, 481, 482, 485, 495, 518, 532–539, 562, 563, 579, 969

constituent construct

For a given construct, a construct that corresponds to one of the constituent directives of the executable directive. 23, 36, 87, 107, 193, 198, 273, 399, 558, 560, 572–574, 971

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. 35, 37, 179, 193, 316, 503, 573, 576, 966

constituent-directive name

The directive name of a constituent directive. **570**, 570, 576, 990

construct

The lexical extent of an executable directive, that is, an executable directive, its paired end directive (if any), and the associated structured block (if any), not including the code in any called procedures; or, where explicitly referring to a base-language entity, a construct in the base language. The use of all capital letters when referring to the name of a construct is reserved for Fortran constructs. 2–8, 16, 20–23, 26, 29–33, 35, 37, 38, 41, 43–45, 47, 48, 51, 53–56, 58–65, 68–71, 75–77, 82–86, 88, 90–93, 96, 97, 99, 100, 102, 103, 105–107, 111–114, 116–122, 125, 127, 129, 130, 132, 133, 136, 139, 141–143, 151, 152, 157, 168, 171, 173, 174, 180, 188, 190, 193, 198, 200, 202, 203, 205, 206, 209, 214, 215, 223, 224, 227, 228, 230, 233–242, 244, 245, 247–254, 257, 265–268, 270–273, 275–277, 279, 280, 286, 289, 290, 294, 295, 297, 299, 303, 305, 306, 311–313, 316, 318–327, 330, 340, 341, 344–346, 348, 350, 354, 368, 370, 374–378, 393–395, 399, 400, 402, 407, 410–412, 414–419, 421–424, 431–433, 435, 437, 438, 441–446, 448–456, 458–469, 471–473, 475–479, 481–487, 489, 491, 493–501, 503–509, 511, 512, 514, 515, 518–520, 522, 523, 525–529, 531–549, 552, 553, 555–561, 563–570, 572–576, 583, 607, 628–630, 653, 654, 657, 752, 758, 766, 780, 786, 794, 795, 803, 807, 811, 815, 819–822, 832, 834, 845, 892, 893, 944–946, 955–957, 963, 964, 966–971, 973–977, 979–985, 987–990

`construct` selector set

A selector set that may match the construct trait set. **357**, 354, 357–360, 365

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. **354**, 37, 38, 354, 355, 357, 359, 360, 377

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. 27, 29, 37, 185, 313, 323

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 C/C++, a structure pointer `p` to which the `->` operator applies is equivalent to the application of a `.` (dot) operator to `(*p)` for the purposes of determining containing structures. For C++, an implicit member access `x` is equivalent to `(*this).x` for the purposes of determining containing structures. 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. 29, 38, 235, 313, 317, 318, 320, 323

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, 22, 29, 71, 89, 112, 132, 133, 147, 152, 160, 298, 337, 349, 395, 426, 430, 489, 496, 519, 540, 541, 579, 589, 617, 629, 630, 634, 653, 654, 717, 723, 957, 968, 977

context-matching construct

A construct that has the context-matching property. 357

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. 38, 374, 421, 431, 438, 456, 457, 505

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. **356**, 34, 49, 50, 110, 356, 358–362, 364–366, 372–374, 391, 954, 955, 976

context-specific structured block

Structured blocks that conform to specific syntactic forms and restrictions that are required for certain block-associated directives. **208**, 23, 26, 57, 209, 210

core

A physically indivisible hardware execution unit on a device onto which one or more hardware threads may be mapped via distinct execution contexts. 70, 84, 108, 145, 787

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. 76, 248, 253, 303, 304, 308, 309, 311, 313, 314, 319, 321–325, 330–332, 344, 347, 382, 397, 506, 508, 511, 664, 968

corresponding pointer

For a given pointer variable or a given referring pointer, the corresponding variable or handle that exists in a device data environment. 91, 319

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. 319, 507

corresponding storage

For a given storage block, its corresponding storage block. For a given mapped variable, the corresponding storage of its original storage block. 39, 78, 93, 101, 105, 251, 315, 320–326, 331, 508, 659, 801, 957

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. 9, 39, 76, 312, 323, 324

C pointer

For C/C++, a base language pointer variable. For Fortran, a scalar variable of type **`C_PTR`**. 47, 127, 251, 369

current device

The device on which the current task is executing. 9, 10, 22, 24, 47, 61, 79, 115, 163, 355, 477, 494, 547, 579, 623, 625, 628, 685, 702, 709, 710, 743, 744, 853, 854, 865

current task

For a given thread, the task corresponding to the task region that it is executing. 39, 51, 61, 311, 336, 526, 614, 616–619, 622, 623, 625, 642, 645, 646, 650, 738

current task region

The region that corresponds to the current task. 5, 117, 438, 468, 489, 526, 564, 567, 925

current team

All threads in the team executing the innermost enclosing parallel region. 29, 91, 119, 120, 133, 236, 438, 442, 443, 445–447, 449, 454, 477, 484, 489, 521, 525, 526, 558, 560, 564, 567, 568, 624, 643, 794

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, 337

D

data-copying property

The property that a clause copies a list item from one data environment to other data environments. 301, 302

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. 26, 27, 29, 39, 41, 42, 46, 58, 60, 63, 66, 96, 99, 100, 106, 121, 127, 319

data environment

The variables associated with the execution of a given region. 4, 6, 9, 22, 40–42, 45, 51, 61, 77, 78, 81, 84, 91, 115, 131–133, 137, 141–143, 233, 250, 275, 303, 311, 330, 467, 471, 478, 487, 498, 500, 506, 511, 657, 864, 974, 985

data-environment attribute

A data-sharing attribute or a data-mapping attribute. 40, 233, 327, 351

data-environment attribute clause

A clause that explicitly determines the data-environment attributes of the list items in its list argument. **233**, 40, 383, 384, 440, 478, 479, 487

data-environment attribute property

The property that a clause is a data-environment clause. 241–243, 246, 249–252, 269, 274–276, 278, 300, 308, 310, 347, 351

data-environment clause

A clause that is a data-environment attribute clause or otherwise affects the data environment. **233**, 40, 233, 486

data-localizing attribute

For a given data entity in a data environment, an attribute that determines the execution context with respect to which the entity is localized. 46, 58, 121, 236, 241, 265, 293

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. **233**, 40, 41, 53, 63, 236, 305, 327, 980

data-mapping attribute clause

A clause that explicitly determines the data-mapping attributes of the list items in its list argument. **233**, 9, 39, 41, 53, 84, 305, 308, 348, 498, 500, 506, 967

data-mapping attribute property

The property that a clause is a data-mapping clause. 308, 310

data-mapping clause

A clause that is a data-mapping attribute clause or otherwise affects the data environment of the target device. **233**, 41, 78, 233

data-mapping construct

A construct that has the data-mapping property. 50, 76, 235, 312, 316, 323, 324, 503

data-mapping property

The property of a construct on which a data-mapping attribute clause may be specified. 41, 497, 500, 502, 505

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. 34, 330

data-motion attribute property

The property that a clause is a data-motion clause. 331, 332

data-motion clause

A clause that specifies data movement between a device set that is specified by the construct on which it appears. 24, 34, 41, 307, 328, 330–333, 511, 512, 976

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, 14, 41, 242, 244, 248, 267, 277, 303, 324, 331, 339, 442, 461, 543

data-sharing attribute

For a given data entity in a data environment and an OpenMP region, an attribute that determines the lifetime and visibility of the entity with respect to execution of the region. 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. **233**, 40, 42, 53, 55, 60, 63, 66, 69–71, 96, 99, 106, 127, 233–237, 240, 254, 305, 327, 498, 500, 503, 506, 511, 573, 953, 976, 980

data-sharing attribute clause

A clause that explicitly determines the data-sharing attributes of the list items in its list argument. **233**, 8, 22, 23, 42, 53, 91, 179, 233, 236, 237, 239–242, 255, 256, 344, 345, 348, 465, 467, 471, 506, 508, 575, 967, 982

data-sharing attribute property

The property that a clause is a data-sharing clause. 241–243, 246, 249–252, 269, 274–276, 278, 347, 487

data-sharing clause

A clause that is a data-sharing attribute clause. **233**, 42, 233, 235, 236, 241

declaration-associated directive

A declarative directive for which its associated base language code is a procedure declaration. **172**, 171–174, 370, 377, 383, 384, 969

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. 373, 385, 405

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. 42, 43, 53, 67, 128, 170, 171, 173–175, 180, 281, 285, 293, 297, 327, 341, 370, 372, 377, 382, 385, 399, 493, 966

declare target directive

A declarative directive that has the declare-target property. 9, 77, 84, 235, 257, 298, 305, 313, 314, 322, 354, 381–383, 385, 386, 392, 396, 397, 506, 508, 610, 955, 974, 980

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, 382, 385

declare variant directive

A declarative directive that declares a function variant for a given base function. 50, 354, 364, 365, 373, 955, 976, 980

default-attribute clause

A clause that has the default-attribute property. 43, 203, 204, 305

default-attribute order

A partial order among default-attribute clauses on a given construct that determines precedence among the clauses for matched variables. **203**, 203, 204

default-attribute property

The property that a clause defines a default implicitly determined data-environment attribute for variables that do not have an explicitly determined data-environment attribute or a predetermined data-sharing attribute. 43, 253, 326

default mapper

The mapper that is used for a **`map`** clause for which the *mapper* modifier is not explicitly specified. 95, 307

deferrable task

A task that can be deferred with respect to its generating task region. 498, 500, 503, 506, 511, 529, 658

deferred task

A task that is deferred with respect to its generating task region. 526

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.

43, 124, 149, 164, 986

delimited directive

A directive for which the associated base-language code is explicitly delimited by the use of a required paired end directive. **173**, 171, 173, 363, 364, 372, 385, 404

dependence

An ordering relation between two instances of executable code that must be enforced by a compliant implementation. **548**, 44, 115, 200, 477, 548–550, 552, 553, 556, 559, 658, 776, 817, 822, 823

dependence-compatible task

Two tasks between which a task dependence may be established. **552**, 95, 115, 123, 548, 553–555, 605, 964

dependent task

A task that because of a task dependence cannot be executed until its antecedent tasks have completed. **552**, 23, 113, 115, 374, 492, 503, 527, 552, 553, 658, 803, 823

depend list

An list argument that consists of depend list items. **181**, 551

depend list item

A list item that represents a storage location for which a task dependence may be expressed. **181**, 44, 551

depend object

An OpenMP object that supplies user-computed dependences to **`depend`** clauses. **604**, 200, 477, 528, 549, 551–553, 658, 822, 964, 981

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. 44, 175, 215, 216, 282, 578, 657, 771, 774, 799, 840, 843, 845, 847, 950, 962, 965, 972–976, 979, 981

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. 38, 40, 44, 122, 472, 490, 567

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. **487**, 468, 479, 583, 643, 981

device

An implementation-defined logical execution engine.

COMMENT: A device could have one or more processors.

3, 4, 7–10, 21–24, 29, 33, 38, 39, 41, 42, 45–48, 51, 56, 60, 65, 66, 79, 83, 84, 88, 93, 94, 108, 113–116, 125, 131–133, 142, 144, 145, 157–159, 163, 201, 251, 300, 305, 308, 309, 311, 330, 331, 337–340, 349, 354, 355, 357, 359, 368, 381, 382, 395, 396, 478, 493, 496–498, 501, 506–509, 512, 540, 581, 610, 611, 617, 635–638, 643, 645, 647–650, 652–657, 659–661, 664–667, 673, 674, 685–690, 700, 702–707, 709, 710, 722, 743, 749, 750, 753, 764–767, 769, 771, 772, 778, 783, 787, 806, 834–841, 847, 849, 850, 857, 865, 867–869, 871–874, 876–878, 883, 886, 890, 897, 900, 907, 911, 915–918, 921, 943, 947, 948, 950, 955, 957, 959, 960, 965, 968, 971, 972, 975–977, 979, 981–984

device address

An address of an object that may be referenced on a target device. **8**, 8, 47, 69, 127, 249–252, 368, 395, 396, 661, 950, 976, 979

device-affecting construct

A construct that has the device-affecting property. 507, 653, 654, 988

device-affecting property

The property that a device construct can modify the state of the device data environment of a specified target device. 45, 497, 500, 502, 505, 511

device-associated property

The property of a clause that a device must be associated with the construct on which it appears. 249–252

device construct

A construct that has the device property. 2, 45–47, 60, 69, 114, 115, 127, 159, 312, 392, 494, 656, 798, 822, 843, 847, 978, 984

device data environment

The initial data environment associated with a device. **9**, 9, 26, 39, 41, 45, 60, 76–80, 93, 96, 127, 142, 233, 249, 250, 252, 253, 275, 305, 308, 309, 311, 313, 315, 319–326, 330, 331, 368, 382, 396, 498, 500, 506, 508, 509, 511, 652–654, 659, 661, 663–666, 673, 841, 950, 967, 971

device global requirement clause

A *requirement* clause that has the device global requirement property. 392

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. 45, 392, 394–397

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. **645**, 46, 645–656

device-information routine

A routine that has the device-information property. **645**, 645

device-local attribute

For a given device, a data-localizing attribute of a data entity that it is visible only to tasks that execute on that device and it has static storage duration or, for Fortran, is a named constant. **300**, 46, 235, 237

device-local variable

A variable that has the device-local attribute with respect to a given device. **300**, 9, 313, 381, 382, 396, 950

device-memory-information routine

A routine that has the device-memory-information routine property. **658**, 657

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. **658**, 46, 658–660

device memory routine

A device routine that has the device memory routine property. **657**, 46, 114, 115, 610, 657, 658, 722, 841, 953, 983

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. **657**, 46, 658–663, 665, 667, 669–671, 674, 675

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, 36, 131, 132, 135, 136, 144, 157–159, 340, 494, 586, 646, 647, 649, 651–656, 664, 666, 750, 753, 835–837, 841, 843, 865, 959, 971

device pointer

An implementation defined handle that refers to a device address and is represented by a C pointer. 8, 69, 127, 249, 250, 368, 395, 658, 661, 662, 664–667, 709, 950, 977

device procedure

A procedure that can be executed on a target device, as part of a **`target`** region. 115, 310, 381, 382, 392, 396

device property

The property of a construct that it accepts the **`device`** clause. 45, 382, 385, 497, 500, 502, 505, 511, 514

device region

A region that corresponds to a device construct. 776, 783, 806, 841, 843, 845, 847

device routine

An OpenMP API routine that may require access to one or more specified devices. 25, 46, 159

`device` selector set

A selector set that may match the device trait set. **357**, 357–359, 964

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. **135**, 136, 144, 145, 157, 975

device-tracing callback

A callback that has the device-tracing property. **834**

device-tracing entry point

An entry point that has the device-tracing property. 834, 835

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. **834**, 47, 834–839, 842, 844, 846

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. **355**, 47, 354–356, 373, 374, 964

device-translating callback

A callback that has the device-translating property. **906**, 908

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. **906**, 47, 908

directive

A base language mechanism to specify OpenMP program behavior. 2, 3, 6–9, 13, 14, 16, 18, 20, 23, 25, 26, 28, 30–32, 34–38, 42, 43, 48–53, 57, 61, 65, 67, 69, 71, 75, 77, 81, 87–89, 98–100, 105, 107–109, 112, 113, 117, 122, 124, 126, 127, 130, 132, 144, 162, 166–181, 183–186, 190, 193, 200, 202, 205, 209, 210, 212–214, 221, 223, 224, 227–229, 231, 233, 236, 238, 240, 241, 247, 256, 264, 265, 273, 275, 279–283, 285–287, 289–293, 295–299, 305, 307–311, 319, 323, 324, 327–329, 335, 337, 338, 341–347, 352, 354–356, 358–360, 362–364, 371–374, 376–379, 381, 383–393, 395, 398, 399, 403–409, 415, 419, 422, 424, 427, 433, 435, 439, 441, 449, 451, 464, 466, 467, 471, 473, 476, 487, 494–496, 498–503, 506, 508, 511, 515, 516, 520, 528, 529, 531, 535, 542, 543, 547–549, 563, 565, 566, 568, 572, 580, 607, 610, 662, 701, 702, 707, 708, 710, 723, 807, 810, 952, 953, 955, 956, 962–964, 966–971, 973–977, 979, 980, 982–984, 986, 988

directive name

The name of a directive or a corresponding construct. 35, 37, 48, 49, 71, 107, 108, 168, 181, 192, 193, 198, 199, 201, 229–231, 241, 242, 244, 246, 249–253, 269, 274–276, 279, 284, 285, 288, 291, 292, 300, 301, 303, 308, 310, 311, 326, 332, 333, 340, 341, 344, 347, 349, 351, 352, 361, 362, 366, 367, 369, 376, 377, 380, 386, 389, 390, 393–403, 407, 409, 413, 415, 419, 420, 425, 430, 435, 437, 440, 441, 443, 459, 463, 465, 474, 475, 480–487, 493–495, 516, 518, 529, 530, 532–540, 550, 551, 555, 556, 562, 563, 570, 572

directive-name list

An list argument that consists of directive-name list items. **181**

directive-name list item

A list item that is a directive name. **181**, 48

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. **570**, 48, 570–572

directive specification

The directive specifier and list of clauses that specify a given directive. **168**, 49, 168, 181

directive-specification list

An list argument that consists of directive-specification list items. **181**

directive-specification list item

A list item that is a directive specification. **181**, 48, 183

directive specifier

The directive name and, where permitted, the directive arguments that are specified for a given directive. **168**, 48

directive variant

A directive specification that can be used in a metadirective. **360**, 38, 102, 360–364, 980

dispatch-compliant clause

A clause that has the dispatch-compliant property. 355, 365

dispatch-compliant property

The property that a clause complies with a variant substitution requirement for a function dispatch. 49, 366, 369

dispatch-requirement clause

A clause that has the dispatch-requirement property. 355, 374

dispatch-requirement property

The property that a clause specifies a variant substitution requirement for a function dispatch. 49, 249, 251, 375, 528, 551

divergent threads

Two threads are divergent if one executes a diverging code path and the other does not due to a conditional statement. 7, 49, 398

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. 7, 49

doacross-affected loop

For a worksharing-loop construct, an ordered-affected loop in which a stand-alone **`ordered`** directive is closely nested. 230, 406, 558, 969

doacross dependence

The requirement that execution of a stand-alone **`ordered`** directive of a source iteration must complete before a matching stand-alone **`ordered`** directive of a sink iteration. **548**, 108, 109, 556, 559, 776

dynamic context selector

Any context selector that is not a static context selector. 373

dynamic groupprivate block

An unnamed, groupprivate storage block with dynamic storage duration that is instantiated from a groupprivate-instantiating construct. 20, 50, 58, 59, 349, 350, 717–722, 954

dynamic-groupprivate-information routine

A groupprivate-information routine that has the dynamic-groupprivate-information-routine property. 350, 717, 718

dynamic-groupprivate-information-routine property

The property of a groupprivate-information routine that it provides information about dynamic groupprivate blocks. 50, 717–721

dynamic replacement candidate

A replacement candidate that may be selected at runtime to replace a given metadirective. **361**, 361, 362, 364, 365

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. 50, 234, 237

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. **356**, 127, 354, 356, 358

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. **373**, 373

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. **322**, 323, 324

enclosing context

For C/C++, the innermost scope enclosing a directive. For Fortran, the innermost scoping unit enclosing a directive. 51, 81, 91, 106, 236, 270, 273, 277, 278, 282, 286, 304, 360, 376, 377, 448, 451, 453, 461, 462, 980

enclosing data environment

For a given directive, the data environment of its enclosing context. 41, 55, 60, 69, 70, 103, 127, 477–479

encountering device

For a given construct, the device on which the encountering task of the construct executes. 250, 324, 330, 332, 333, 498, 499, 501, 506–508, 512, 957

encountering task

For a given region, the current task of the encountering thread. 6, 51, 54, 57, 117, 270, 330, 388, 422, 423, 431, 432, 455, 468, 472, 478, 479, 484, 487, 489–492, 506, 507, 514, 521, 523, 525–527, 564, 565, 567, 579, 624, 640, 729, 732–737, 766, 806, 821, 822, 843, 862, 864, 944–946

encountering-task binding property

The binding property that the binding thread set is the encountering task. **579**

encountering thread

For a given region, the thread that encounters the corresponding construct, structured block sequence, or routine. 4, 5, 29, 51, 65, 100, 151, 270, 421, 422, 427, 428, 431, 432, 442, 443, 464, 466, 468, 492, 506, 515, 521, 546, 549, 579, 624, 642, 647, 741, 743–746, 749, 755, 786, 832, 849, 855, 858, 859, 862, 865, 963, 971

encountering-thread binding property

The binding property that the binding thread set is the encountering thread. **579**

end-clause property

The property that a clause may appear on an end directive. **168**, 302, 528

end directive

For a given directive, a paired directive that lexically delimits the code associated with that directive. **168**, 37, 43, 51, 168, 171–174, 179, 209, 210, 214, 224, 364, 373, 385, 386, 520, 973, 974

ending address

The address of the last storage location of a list item or, for a mapped variable, of its original list item. 54, 77, 320

entry point

A runtime entry point. 25, 47, 87, 760, 761, 763, 765–767, 772, 781, 783, 791, 807, 834, 835, 838, 849–878, 960, 961, 965, 972

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. 52, 581, 584, 585, 587, 589, 590, 593, 596, 600, 603–606, 608, 609, 611, 612, 773, 775, 777, 779, 781, 783–786, 788–791, 793, 796–798, 800–803, 805, 852, 889, 891, 892, 938

environment variable

Unless specifically stated otherwise, an OpenMP environment variable. 2, 6, 134, 135, 144–155, 157–165, 752, 753, 936, 951, 952, 963, 965, 966, 975, 977–979, 982, 984–986

error termination

A **`fatal`** action preformed in response to an error. **6**, 34, 103, 427, 435, 496, 969

event

A point of interest in the execution of a thread or a task. 11, 15, 30, 44, 101, 104, 115, 122, 123, 266, 267, 312, 382, 388, 422, 423, 432, 433, 442, 445–449, 451, 453, 455, 461, 462, 468, 469, 472, 479, 487, 488, 491, 496–499, 501, 504, 507, 512, 520, 522–525, 527, 528, 543, 544, 547, 553, 557, 559, 560, 565, 567, 568, 583, 639, 642, 643, 657, 658, 661–668, 670, 671, 673–676, 724–734, 736, 737, 755, 757, 760, 761, 763, 766, 771, 787, 789, 791, 803, 806, 808, 819–822, 825, 827, 829, 833, 834, 838–840, 843, 845, 847, 849, 853–855, 861, 870, 872, 876–878, 880, 942, 944, 946–948, 960, 961, 965, 971

exception-aborting directive

A directive that has the exception-aborting property. 401, 952

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. 52, 167, 505

exclusion operator

An operator that indicates the exclusion of a set of items from a list. 148, 150

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. When applied to a modifier set, the property applies only to modifiers within that set. **179**, 34, 178, 180, 269, 272, 279, 289, 344, 367, 379, 418, 435, 445, 467, 471, 531, 535, 550, 552, 563

exclusive scan computation

A scan computation for which the value read does not include the updates performed in the same logical iteration. **292**, 292, 979

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, 37, 67, 70, 74, 75, 110, 112, 127, 167, 170, 171, 173, 208, 221, 346, 361, 374, 388, 389, 410–412, 414, 416–418, 421, 431, 438, 442, 445–447, 449, 452, 456, 457, 461, 464, 467, 470, 476, 489, 497, 500, 502, 505, 511, 514, 519, 521, 524, 526, 541, 545, 549, 558, 560, 564, 568

execution space

The space formed by a league or team, consisting of one or more execution space units, and configured with a specific spatial dimensionality. 21, 53, 65, 71, 81, 83, 109, 110, 113, 119, 125, 140, 141, 436, 963

execution space unit

The base unit that participates in an execution space. The execution space units in a team space and a league space are the threads of the associated team and the teams of the associated league, respectively. 21, 53, 109, 110, 141

explicit barrier

A barrier that is specified by a **`barrier`** construct. **521**, 27, 491, 521

explicitly associated directive

A declarative directive for which its associated base language declarations are explicitly specified in a variable list or extended list argument. **171**, 171, 174, 293, 297, 341, 382

explicitly determined data-environment attribute

An explicitly determined data-sharing attribute or explicitly determined data-mapping attribute. 43, 327

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. 53, 305

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. **236**, 53, 233, 236, 254

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, 111, 167, 168, 453, 488, 749, 866

explicit task

A task that is not an implicit task. **5**, 3–7, 27, 30, 31, 44, 54, 56, 92, 116, 117, 122, 132, 271, 273, 422, 427, 428, 467, 468, 471, 472, 489–491, 503, 506, 521, 522, 526, 548, 569, 639, 749, 780, 817, 863, 928, 981, 983, 987

explicit-task-generating construct

A construct that has the explicit-task-generating property. **3**, 5, 55, 100, 102, 106, 117, 151, 236, 237, 468, 477, 479, 484, 489, 503, 506, 572, 967, 970, 980, 988

explicit-task-generating property

The property that a construct generates one or more explicit tasks that are child tasks of the encountering task. **3**, 54, 467, 470, 497, 500, 502, 505, 511, 526

explicit task region

A region that corresponds to an explicit task. 8, 100, 241, 468, 572, 640, 973

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. **555**, 123, 469, 479, 552, 555, 605, 964

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. **319**, 78, 320

extended list

An list argument that consists of extended list items. **181**, 53

extended list item

A variable list item or the name of a procedure. **181**, 54, 183

extension trait

A trait that is implementation defined. **356**, 354

F

finalized taskgraph record

A taskgraph record in which all information required for a replay execution has been saved. **478**, 78, 478

final loop body

For a canonical loop nest, the loop body of the innermost semantic loop. 224, 968

final task

A task that generates included final tasks when it encounters explicit-task-generating constructs on which the `final` clause may be specified. **483**, 55, 132, 469, 478, 479, 481, 484, 487, 640, 985

first-party tool

A tool that executes in the address space of the program that it is monitoring. **757**, 15, 30, 86, 87, 163, 755, 757, 759, 965, 972, 982

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. **244**, 55, 234–237, 306, 477, 506, 507, 974, 982

firstprivate variable

A private variable that has the firstprivate attribute with respect to a given construct. 472, 479, 957

flat-memory-copying property

The property that a memory-copying routine copies a unidimensional, contiguous storage block. **666**, 55, 667, 670

flat-memory-copying routine

A routine that has the flat-memory-copying property. **666**, 666, 668, 671

flatten-affected loop

A transformation-affected loop of a `flatten` construct. 411

flattened iteration space

The logical iteration space of a flattened loop. **411**, 411

flattened loop

The output loop of a `flatten` construct. **410**, 55, 410, 411, 955

flush

An operation that a thread performs to enforce consistency between its view of memory and the view of memory of any other threads. 7, 10–14, 20, 56, 63, 101, 111, 121, 444, 490, 518, 540, 546, 978, 985

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, 978

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. 56, 201, 514, 515, 586, 977

foreign runtime environment

A runtime environment that exists outside the OpenMP runtime with which the OpenMP implementation may interoperate. 56, 68, 96, 514, 517, 584, 586

foreign runtime identifier

A base language string literal or a constant expression of integer OpenMP type that represents a foreign runtime environment. **205**, 515, 517, 957, 971

foreign task

An instance of executable code that is executed in a foreign execution context. 201, 479, 515, 957

Fortran-only property

The property that an OpenMP feature is only supported in Fortran. 579

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. 31, 56, 780–782, 806, 862, 888, 928, 929

free-agent thread

An unassigned thread on which an explicit task is selected for execution or a primary thread for an explicit parallel region that was a free-agent thread when it encountered the **`parallel`** construct. 56, 112, 121, 132, 148, 151, 160, 161, 427–429, 490, 641, 642, 795, 956, 966, 971

free-agent-thread scheduling constraint

A task-scheduling constraint that ensures the number of free-agent threads does not exceed *free-agent-thread-limit-var*. **490**, 490

free property

The property that a clause, an argument, a modifier, a clause set, or a modifier set does not have the ultimate property. 178

full taskwait

A point in the execution of a program that ensures all child tasks of the encountering task complete execution before the encountering task may resume execution. 57, 118, 527

full taskwait region

The region that corresponds to a **`taskwait`** construct that specifies a full taskwait. 492, 527

function

A routine or procedure that returns a type that can be the right-hand side of a base language assignment operation. 174, 175, 342, 373, 454, 615–618, 621–629, 631–637, 639–641, 646–651, 653, 655, 658–661, 663, 665, 667, 669–671, 674, 675, 678, 679, 681–683, 686–691, 697–699, 703–706, 708, 711–715, 718–721, 735, 736, 738, 739, 741, 743, 746, 748–751, 754, 757, 807, 832, 849–861, 863–865, 867–870, 872–878, 898, 899, 901–904, 906–908, 910–940

function dispatch

A base function call for which variant substitution may be controlled. 49, 209, 354, 370, 374, 376

function-dispatch structured block

A context-specific structured block that may be associated with a **`dispatch`** directive. **209**, 209, 210, 354, 367, 369, 374

function variant

A definition of a procedure that may be used as an alternative to the base language definition. 38, 43, 50, 102, 128, 354, 364–371, 373, 374, 376, 514, 976, 980

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. 408, 410, 414, 418

generated task

The task that is generated as a result of the generating task encountering a task-generating construct. 5, 116, 142, 236, 468, 471, 472, 476, 481, 482, 484, 486, 489, 492, 506, 514, 515, 526, 527, 529, 548, 552, 553, 555, 817, 822

generating task

For a given region, the task for which execution by a thread generated the region. 13, 29, 57, 58, 124, 141, 142, 374, 468, 491, 498, 500, 502, 505, 511, 514, 548, 657, 925

generating-task binding property

The binding property that the binding task set is the generating task. **657**, 661–663, 665, 667, 669–671, 674, 675

generating task region

For a given region, the region that corresponds to its generating task. 32, 43, 65, 118, 124, 925

global

A program aspect such as a scope that covers the whole OpenMP program. 61, 131–133, 135, 145, 342, 983

grid loop

The output 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. 86, 416–418, 955, 970

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. **349**, 349, 722

groupprivate attribute

For a given group of tasks, a data-localizing attribute of a data entity that it is visible and/or accessible only to those tasks. **298**, 50, 58, 59, 266, 295, 298, 299, 349, 350

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. **349**, 349

groupprivate-information routine

A memory-management routine that has the groupprivate-information-routine property. 50, 717, 964

groupprivate-information-routine property

The property of a memory-management routine that it provides information about groupprivate variables and dynamic groupprivate blocks. 58, 717–721

groupprivate-instantiating construct

A construct that has the groupprivate-instantiating property. 50, 58, 349, 350, 722

groupprivate-instantiating property

The property that a construct can instantiate a dynamic groupprivate block via a **`dyn_groupprivate`** clause. 59, 431, 505

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. **349**, 58, 349, 720, 722, 954

groupprivate variable

A variable that has the groupprivate attribute with respect to a given group of tasks. **298**, 20, 58, 59, 298, 299, 313, 381, 383, 385, 454, 506, 722

H

handle

A reference that uniquely identifies an abstraction. 21, 39, 47, 59, 82, 83, 87, 92, 100, 104, 116, 129, 200–202, 313, 336, 337, 592, 593, 685, 690–692, 700–702, 707, 708, 710, 721, 771, 772, 782, 859, 884, 890, 891, 893, 895–898, 905, 914, 915, 917–919, 922, 923, 925–929, 931, 935, 939, 940, 950, 961

handle-comparing property

The property that a routine compares two handle arguments. **929**, 59, 929, 930

handle-comparing routine

A routine that has the handle-comparing property. **929**, 929, 961

handle property

The property that a type is used to represent handles. **895**, 59, 884, 893, 895, 896

handle-releasing property

The property that a routine releases a handle. **931**, 59, 931–933

handle-releasing routine

A routine that has the handle-releasing property. **931**, 931

handle type

An OpenMP type, OMPD type, or OMPT type that has the handle property. **895**

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, 60, 339, 396, 515, 978

hard pause

An instance of a resource-relinquishing routine that specifies that the OpenMP state is not required to persist. **610**, 610, 611

hardware thread

An indivisible hardware execution unit on which only one OpenMP thread can execute at a time. **6**, 6, 38, 98, 145, 148, 579, 648, 787

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. **251**

host address

An address of an object that may be referenced on the host device. 60, 396, 977

host device

The device on which the OpenMP program begins execution. 3–5, 7, 9, 20, 27, 60, 66, 84, 113, 136, 144, 145, 155, 156, 159, 331, 339, 355, 359, 395, 493, 509, 610, 627, 630, 647, 650–656, 659, 661, 664, 685, 688, 689, 702, 705, 706, 722, 750, 753, 757, 761, 763, 765, 766, 783, 856, 867–870, 878, 892, 914–916, 921, 965, 971, 980

host pointer

A pointer that refers to a host address. 395, 396, 659, 661, 664–666, 977

I

ICV

An internal control variable. **131**, 7, 61, 88, 109, 131–139, 141–148, 151–155, 157, 159–165, 294, 341, 357, 374, 394, 423, 424, 427, 428, 432, 435, 444, 455, 460, 467, 471, 478, 485, 494, 496, 498, 500, 506, 511, 564, 566, 582, 609, 614, 616–623, 625–634, 639–641, 645–648, 652–654, 707–709, 738–740, 742–746, 748, 752, 753, 759, 760, 856, 858, 881, 888, 893, 919, 927, 938–940, 956, 957, 959, 960, 965, 971, 973, 975, 977, 978, 984–987

ICV-defaulted clause

A clause that has the ICV-defaulted property. 478

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. 61, 341, 494

ICV modifying property

The property of a routine or clause that its effect includes modifying the value of an ICV. 495, 614, 618, 619, 621, 626, 629, 645, 652, 654, 742

ICV retrieving property

The property of a routine that its effect includes returning the value of an ICV. 616, 618, 620, 622, 623, 625–629, 632–634, 639–641, 646, 647, 651, 653, 738–743, 748

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. **131**, 61, 131, 135, 137, 141, 142, 145, 478, 498, 500, 506, 511

idle thread

An unassigned thread that is not currently executing any task. **488**, 795

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. 61, 114, 345, 433, 971

imperfectly nested loop

A nested loop that is not a perfectly nested loop. 980

implementation code

Implicit code that is introduced by the OpenMP implementation. 28, 42, 53, 100, 780

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, 36, 44, 47, 52, 54, 83, 97, 98, 100, 113, 126, 134–136, 142, 145–149, 151–157, 159, 161, 163–167, 176, 227, 237, 249, 251, 295, 308, 335–339, 349, 352, 355, 356, 358, 360, 362, 364–366, 371, 378, 381, 388, 390, 391, 409, 411, 416–420, 422, 426, 427, 429, 432, 435, 438, 441, 445, 448, 455, 459–461, 472, 475, 479, 496, 508, 515, 517, 543, 578–580, 584, 586, 587, 591, 604, 608, 620–622, 650, 656, 664, 666, 667, 678, 682, 723, 740, 743–746, 753, 755, 757, 758, 761, 763, 765, 780, 787, 791, 794, 825, 841, 851, 857–859, 868, 881, 909, 912, 929, 934, 938, 950–961, 973, 978, 984

`implementation` selector set

A selector set that may match the implementation trait set. **357**, 357–359

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. **355**, 62, 354, 356, 373, 374

implicit acquire-at-entry flush

An implicit acquire flush that occurs at entry to a region. 13, 422, 431, 489, 492, 548, 552, 560, 733, 734

implicit acquire-at-exit flush

An implicit acquire flush that occurs at exit from a region. 431, 489, 521, 525, 543, 548, 559

implicit acquire flush

An implicit flush that is also an acquire flush. 62, 490, 519, 521, 543, 547, 548

implicit-acquire-flush-at-entry property

The property of a construct or routine that an implicit acquire flush occurs on entry to its corresponding region. **547**, 421, 431, 467, 470, 497, 500, 502, 505, 511, 519, 526, 560

implicit-acquire-flush-at-exit property

The property of a construct or routine that an implicit acquire flush occurs at exit from its corresponding region. **547**, 431, 467, 470, 497, 500, 502, 505, 511, 521, 524, 526, 541, 558, 730, 731, 735, 736

implicit-acquire-flush property

The property of a construct or routine that it includes an implicit acquire flush. **547**, 421, 431, 467, 470, 497, 500, 502, 505, 511, 519, 521, 524, 526, 541, 558, 560, 730, 731, 735, 736

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. 27, 28, 313

implicit barrier

A barrier that is specified as part of the semantics of a construct other than the **`barrier`** construct. **523**, 4–6, 27, 63, 422, 446, 448, 450, 452, 455, 461, 491, 521, 523, 524, 529, 566, 794

implicit-barrier-flushes property

The property of a construct that its corresponding region includes implicit flushes that result from an implicit barrier. **523**, 421, 445–447, 449, 456, 457

implicit-barrier property

The property that a construct includes an implicit barrier. **523**, 421, 445–447, 449, 455–457, 523

implicit flush

A flush that is specified as part of the semantics of a construct or routine other than the **`flush`** construct. 12, 13, 62–64, 114, 117, 118, 422, 431, 506, 523, 526, 547, 548, 567, 981

implicit-flush property

The property that a construct or routine includes one or more implicit flushes. **547**, 421, 431, 445–447, 449, 456, 457, 467, 470, 497, 500, 502, 505, 511, 519, 521, 524, 526, 541, 548, 558, 560, 564, 568, 642, 730, 731, 733–736

implicitly determined data-environment attribute

An implicitly determined data-sharing attribute or implicitly determined data-mapping attribute. 43, 327, 963

implicitly determined data-mapping attribute

A data-mapping attribute that applies to a data entity for which no data-mapping attribute is otherwise determined. **305**, 63, 305, 325, 327, 800

implicitly determined data-sharing attribute

A data-sharing attribute that applies to a data entity for which no data-sharing attribute is otherwise determined. **236**, 63, 106, 233, 236, 237, 240, 254, 305, 306, 327, 982

implicit member access

For C++, an id-expression within a non-static member function that refers to a non-static data member of the object for which the function is invoked. 38, 183, 234, 306

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, 38, 63, 64, 67, 105, 151, 298, 427, 433, 466, 488, 627, 630, 653, 655, 749, 892, 988

implicit release-at-entry flush

An implicit release flush that occurs at entry to a region. 13, 64, 422, 431, 489, 521, 543, 548, 559

implicit release-at-exit flush

An implicit release flush that occurs at exit from a region. 13, 431, 489, 492, 521, 525, 548, 552, 560, 567, 643, 733, 734

implicit release flush

An implicit flush that is also a release flush. 63, 64, 490, 519, 521, 543, 547, 548

implicit-release-flush-at-entry property

The property of a construct or routine that an implicit release-at-entry flush occurs at entry to its corresponding region. **547**, 421, 431, 467, 470, 497, 500, 502, 505, 511, 521, 526, 541, 558

implicit-release-flush-at-exit property

The property of a construct or routine that an implicit release flush occurs at exit from its corresponding region. **547**, 421, 431, 467, 470, 497, 500, 502, 505, 511, 519, 526, 560, 642, 733, 734

implicit-release-flush property

The property that a construct or routine includes an implicit release flush. **547**, 421, 431, 467, 470, 497, 500, 502, 505, 511, 519, 521, 526, 541, 558, 560, 642, 733, 734

implicit-strong-flush property

The property of a construct or routine that its corresponding region includes an implicit flush that has the strong flush property. **547**, 541

implicit strong flush

An implicit flush that is also a strong flush. 542, 543, 547

implicit-strong-flush-at-entry property

The property of a construct or routine that its corresponding region includes at entry an implicit flush that has the strong flush property. **547**, 541

implicit-strong-flush-at-exit property

The property of a construct or routine that its corresponding region includes at exit an implicit flush that has the strong flush property. **547**, 541

implicit task

A task generated by an implicit parallel region or generated when a **`parallel`** construct is encountered during execution. 3, 4, 8, 21, 24, 29–31, 38, 40, 54, 65, 67, 68, 89, 92, 96, 112, 116, 118, 119, 131–133, 141–143, 236, 244, 269, 272, 273, 301, 303, 304, 422, 423, 428, 431, 444–455, 461, 462, 506, 521, 548, 567, 569, 742, 780, 806, 820, 858, 863, 892, 926–928

implicit-task-generating construct

A construct that has the implicit-task-generating property. **3**, 506

implicit-task-generating property

The property that a construct generates one or more implicit tasks. **3**, 65, 421, 431, 505

implicit taskgroup

A **`taskgroup`** directive that is specified as part of the semantics of a construct other than the **`taskgroup`** construct. 65, 116, 531

implicit-taskgroup property

The property that a construct includes an implicit taskgroup. 470, 471, 476, 479, 502, 503

implicit task region

A region that corresponds to an implicit task. 3, 4, 143, 820

implicit taskwait

A **`taskwait`** directive that is specified as part of the semantics of a construct other than the **`taskwait`** construct. 65, 374

implicit-taskwait property

The property that a construct includes an implicit taskwait. 374

importing task

A task that permits a preceding dependence-compatible task to be an antecedent task of one of its child tasks. **555**, 123, 469, 479, 552, 555, 605

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. 63, 623, 624

inactive spatial dimension

A dimension in an execution space that is not an active spatial dimension. 631–634

inactive target region

A **`target`** region that is executed on the same device that encountered the **`target`** construct. 142

included task

A task for which execution is sequentially included in the generating task region. That is, an included task is an undeferred task and executed by the encountering thread. 7, 32, 55, 65, 100, 468, 481, 484, 492, 498, 500, 503, 506, 511, 514, 526, 529, 657

inclusive scan computation

A scan computation for which the value read includes the updates performed in the same logical iteration. **291**, 291, 979

index-set splitting

The splitting of the logical iteration space into partitions that each are executed by an output loop. **414**, 969

indirect call

A procedure invocation for which the invoked procedure is determined at runtime through a function pointer (C/C++), a pointer to a member function (C++), a dummy procedure (Fortran), or a procedure pointer (Fortran). 66, 365, 508, 509, 976

indirect device invocation

An indirect call to the device version of a procedure on a device other than the host device, for which the call refers to the host version of the procedure. 386, 387

induction

A use of an induction operation. 66, 256

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. **277**, 71

induction expression

A collector expression or an inductor expression. **257**, 257

induction identifier

An OpenMP identifier that specifies an inductor OpenMP operation to use in an induction. **256**, 256, 257, 263–265, 277, 286, 287

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, 66, 67, 71, 110, 127, 192, 256, 260, 277, 288, 967

induction variable

A variable for which an induction operation determines its values. 66, 67, 260, 286

inductor

A binary operator used by an induction operation. 66, 260

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. **260**, 66, 192, 260, 261, 263–265, 277, 286, 288

informational directive

A directive that is neither declarative nor executable, but otherwise conveys user code properties to the compiler. **388**, 127, 170, 391, 399, 403, 404

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. **289**, 289, 290, 292, 968

initializer

An OpenMP operation that uses an initializer expression. **266**, 67, 99, 256, 261, 262, 266, 269, 270

initializer expression

An OpenMP stylized expression that determines the initializer for the private copies of list items in a reduction clause. **258**, 67, 99, 192, 259–261, 264, 265, 268, 282, 284, 289, 381

initial task

An implicit task associated with an implicit parallel region. 4, 5, 29, 67, 105, 142, 272, 428, 432, 433, 453, 461, 462, 488, 496, 506, 507, 739, 766, 780, 820, 847, 856, 863, 947, 948

initial task region

A region that corresponds to an initial task. 3, 131, 132, 617, 623, 625

initial team

The team that comprises an initial thread executing an implicit parallel region. 4, 7, 71, 119, 431, 432, 461, 463, 893

initial thread

The thread that executes an implicit parallel region. 3, 4, 67, 93, 96, 121, 151, 153, 294, 431, 432, 452, 461, 466, 488, 506, 804, 951, 953

innermost-leaf property

The property that a clause applies to the innermost leaf construct that permits it when it appears on a compound construct. **178**, 199, 242, 249, 278, 291, 292, 302, 487, 535–539, 550, 562, 573

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. **316**, 78, 92, 125, 316

input phase

The portion of a logical iteration that contains all computations that update a list item for which a scan computation is performed. **289**, 67, 126, 289, 291, 292

input place group

The place group that is used to determine the *place-group-var* and the place assignments of the implicit tasks of a **`parallel`** region. **427**, 68, 427, 428

input place partition

The place partition of the input place group. 118, 428, 430

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. **580**, 649, 655, 656, 659, 660, 663, 665, 668–670, 672, 678, 680–684, 686–697, 699–701, 703–707, 742, 745, 746, 752, 758, 787, 795, 810–814, 817, 820, 821, 823, 826, 828, 830–832, 834, 836, 839, 842, 844, 849, 899, 901, 903, 905, 907, 909, 910, 919

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. **580**, 655, 656, 678, 680, 681, 693, 695, 697, 744, 746, 850, 851, 912, 917, 934, 936, 937, 940

internal control variable

A conceptual variable that specifies runtime behavior of a set of threads or tasks in an OpenMP program. **131**, 60, 950

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. **583**, 69, 200, 370, 376, 514–517, 583, 588, 677, 684, 958, 972, 977

interoperability property

A property associated with an interoperability object. **514**, 69, 586, 677–680, 682, 683

interoperability-property-retrieving property

The property that a routine retrieves an interoperability property from an interoperability object. **677**, 69, 678, 679, 681

interoperability-property-retrieving routine

A routine that has the interoperability-property-retrieving property. **677**, 677, 679–681

interoperability routine

A routine that has the interoperability-routine property. **677**, 514, 586, 588, 677, 684

interoperability-routine property

The property that a routine provides a mechanism to inspect the properties associated with an interoperability object. **677**, 69, 678, 679, 681–683

interop-type set

For a given interoperability object, the set of *interop-type* modifiers specified in the **`init`** clause that initializes the object. **514**, 202, 370, 514–516

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. **220**, 93, 220, 221, 227, 228, 476

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. **205**, 192, 257–261, 415

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. **249**

ISO C binding property

The property of a routine that its Fortran version has the **`BIND(C)`** attribute. 70, 600, 602, 657–663, 665, 667, 669–671, 674, 675, 691, 695, 697, 698, 711–716

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. **580**, 70, 600, 659–663, 665, 668–670, 672, 674, 675, 695, 697, 711–716, 718, 719, 722, 832, 836, 839

iteration count

The number of times that the loop body of a given loop is executed. **226**, 225–228, 286, 411, 415, 419, 953, 955

iterator

A programming mechanism to specify a set of values. 187, 188, 219, 227, 313, 439, 976, 987

iterator specifier

A tuple that specifies an *iterator-identifier* and its associated iterator value set. **187**, 70, 181, 188

iterator-specifier list

An list argument that consists of iterator-specifier list items. **181**

iterator-specifier list item

A list item that is an iterator specifier. **181**, 70

iterator value set

The set of values that correspond to a given instance of an *iterator* modifier. **188**, 70, 187, 189

L

last-level cache

The last cache in a memory hierarchy that is used by a set of cores. 145

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. **246**, 70, 234

lastprivate variable

A private variable that has the lastprivate attribute with respect to a given construct. 979

leaf construct

For a given construct, a construct that corresponds to one of the leaf directives of the executable directive. 22, 33, 48, 68, 91, 193, 354, 561, 573–576, 989

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. 37, 70, 71, 108, 572

leaf-directive name

The directive name of a leaf directive. **570**, 570, 572, 990

league

The set of teams formed by a **`teams`** construct, each of which is associated with a different contention group. **4**, 53, 71, 119, 132, 272, 431, 432, 435, 462, 463, 626, 786, 820

league-generating property

The property that a construct generates a league of teams. 431

league size

The spatial size that indicates the number of initial teams in a league. 119, 132, 626, 633

league space

The execution space of a league. 53, 431, 432, 435, 633, 634, 636, 638

lexicographic order

The total order of two logical iteration vectors $\omega_a = (i_1, \ldots, i_n)$ and $\omega_b = (j_1, \ldots, j_n)$, denoted by $\omega_a \leq_{\text{lex}} \omega_b$, where either $\omega_a = \omega_b$ or $\exists m \in \{1, \ldots, n\}$ such that $i_m < j_m$ and $i_k = j_k$ for all $k \in \{1, \ldots, m-1\}$. 406, 416, 417

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. **279**, 71

linear variable

A private variable that has the linear attribute with respect to a given construct. **279**

list

A comma-separated set. 40–42, 52, 71, 72, 94, 109, 142, 147, 148, 150, 176, 177, 181, 272, 381, 386, 425, 486, 760, 951

list argument

A list that is used as the argument to a directive, clause, or modifier. **176**, 44, 48, 54, 70, 72, 89, 96, 97, 123, 128, 177, 178, 181, 183, 233, 237, 291, 292, 296, 436

list item

A member of a list. 23, 24, 35, 39–42, 44, 48, 49, 51, 53, 67, 68, 70, 72, 76–78, 81, 84, 88, 89, 91, 93, 95–97, 101, 106, 109, 110, 112, 123, 128, 142, 144, 147, 159, 177–179, 181, 183–185, 187, 189, 192, 233–235, 237–245, 247–253, 257, 258, 260–267, 269–281, 289–292, 295–300, 302–305, 309–315, 317–326, 328–333, 342–346, 352, 355, 368, 369, 376, 381–384, 386, 399, 400, 407–410, 415–417, 425, 427, 436, 440, 455, 462, 464, 471, 477, 479, 486, 487, 496, 498, 500, 503, 506–509, 511, 546, 547, 551–554, 565, 573–576, 939, 953, 964, 966, 968, 969, 974, 980, 986

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. 336

locator list

An list argument that consists of locator list items. **181**, 179, 330, 479

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. **182**, 72, 181–183, 200, 477, 479, 548–551, 553, 554, 964

lock

An OpenMP variable that is used in lock routines to enforce mutual exclusion. 72, 73, 82, 83, 88, 108, 125, 543, 604, 607, 723–737, 795, 804, 831, 852, 859, 959, 983

lock-acquiring property

The property that a routine may acquire a lock by putting it into the locked state. **729**, 72, 723, 730, 731

lock-acquiring routine

A routine that has the lock-acquiring property. **729**, 548, 723, 730, 735, 827, 829, 830

lock-destroying property

The property that a routine destroys a lock by putting it into the uninitialized state. **727**, 72, 728, 729

lock-destroying routine

A routine that has the lock-destroying property. **727**, 728, 729, 829

locked state

The lock state that indicates the lock has been set by some task. **723**, 72, 73, 732–734

lock-initializing property

The property that a routine initializes a lock by putting it into the unlocked state. **724**, 73, 724–727

lock-initializing routine

A routine that has the lock-initializing property. **724**, 724–727, 827, 828

lock property

The property that a routine operates on locks. **723**, 73

lock-releasing property

The property that a routine may unset a lock by returning it to the unlocked state. **732**, 73, 723, 733, 734

lock-releasing routine

A routine that has the lock-releasing property. **732**, 548, 723, 732, 829, 830

lock routine

A routine that has the lock property. **723**, 72, 579, 723, 959

lock state

The state of a lock that determines if it can be set. **723**, 72, 125, 723, 732–734

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. **735**, 73, 735, 736

lock-testing routine

A routine that has the lock-testing property. **735**, 735, 827, 829, 830

logical iteration

An instance of the executed loop body of a semantic 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. 5, 21, 33, 35, 53, 66, 68, 73, 74, 90, 102, 104, 111, 122, 126, 192, 225, 227, 228, 247, 272, 405, 406, 411, 412, 414–419, 440, 471, 472, 474, 475, 780, 816, 955, 957, 975, 977, 980, 983, 986, 987

logical iteration space

For a semantic loop nest, or a **DO CONCURRENT** loop in Fortran, the sequence $0,\ldots,N-1$ where N is the number of distinct logical iterations. **227**, **411**, 33, 55, 58, 66, 73, 90, 122, 192, 227, 410, 414–417

logical iteration vector

An n-tuple $(i_1, \ldots, i_n)$ that identifies a logical iteration of a semantic 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. **227**, 71, 74, 97, 227, 228, 416, 418, 975

logical iteration vector space

The set of logical iteration vectors that each correspond to a logical iteration of a semantic loop nest. 227, 228, 416, 417

loop body

A structured block that encompasses the executable statements that are iteratively executed by a loop statement. 69, 70, 73, 415, 476

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. 32, 33, 228, 237, 238, 279, 437

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. 74, 190, 219, 222–225, 227, 228, 234–236, 247, 279, 280, 406, 464, 476, 557, 574, 576, 987

loop-iteration vector

An n-tuple $(i_1, \ldots, i_n)$ 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. 74, 227, 556, 557

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. **227**, 227

loop-nest-associated construct

A loop-nest-associated directive and its associated loop nest. 66, 69, 71, 74, 91, 102, 104, 107, 129, 173, 224, 228, 278, 280, 406–408, 416, 418, 444, 556, 576

loop-nest-associated directive

An executable directive for which the associated user code must be a canonical loop nest. **172**, 21, 24, 74, 75, 171, 172, 221, 226, 234, 235, 247, 277, 280, 406, 410, 412, 414, 416–418, 438, 456, 457, 461, 464, 470, 561

loop nest depth

For a semantic loop nest, the maximal number of loops, including the outermost loop, that can be affected by a loop-nest-associated directive. 74, 224, 229, 230, 410

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. **228**, 36, 102, 104, 228, 437, 444, 454, 461, 464, 975

loop-sequence-associated construct

A loop-sequence-associated directive and its associated canonical loop sequence. 75, 91, 231, 406

loop-sequence-associated directive

An executable directive for which the associated user code must be a canonical loop sequence. **172**, 25, 75, 171, 411

loop sequence length

For a semantic loop sequence, the number of consecutive semantic loop nests regardless of their nesting into blocks. **226**, 231

loop-sequence-transforming construct

A loop-sequence-associated construct with the loop-transforming property. 226

loop-transforming construct

A loop-transforming directive and its associated loop nest or associated canonical loop sequence. **406**, 57, 85, 91, 92, 123, 220, 224, 225, 228, 405–408, 410, 414, 473, 969, 970, 973, 976

loop-transforming directive

A directive with the loop-transforming property. 57, 75, 406, 408–410, 415

loop-transforming property

The property that a construct may be used as *loop-nest-generating-construct* in a canonical loop nest or as *loop-sequence-generating-construct* in a canonical loop sequence. 75, 405, 410–412, 414, 416–418

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. 111, 172

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. **315**, 76, 325, 397, 499

map-entering construct

A construct that has the map-entering property. 76, 315, 319, 320, 322, 323, 572, 610, 964

map-entering map type

A map type that specifies the clause on which it is specified is a map-entering clause. **315**, 315, 316

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. 76, 316, 497, 502, 505

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. **315**, 76, 501

map-exiting construct

A construct that has the map-exiting property. 76, 315, 323, 572

map-exiting map type

A map type that specifies the clause on which it is specified is a map-exiting clause. **315**, 315, 316

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. 76, 316, 500, 502, 505

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. 319, 323–325, 331

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.

77, 309, 310, 313, 331

mapped address range

For a given original list item, the address range that starts from its starting address and ends with its ending address. **319**, 78, 111, 320

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. 39, 51, 77, 78, 91, 110, 509, 610

mapper

An operation that defines how variables of given type are to be mapped or updated with respect to a device data environment. 43, 50, 77, 127, 205, 305, 307, 316, 321–323, 328–333

mapper expansion

The process by which a list item with a *mapper* modifier in a clause is removed from the clause and replaced with the clauses specified in the declared mapper, which are then applied to the construct on which the clause appears. **320**, 106, 313, 321, 325, 330

mapper identifier

An OpenMP identifier that specifies the name of a user-defined mapper. **307**, 307, 330

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, 24, 26, 76, 78, 101, 105, 106, 111, 305, 312, 315, 317, 320, 321, 324, 325, 397, 611, 795, 800, 801, 968, 969

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. 24, 34, 68, 76, 92, 125, 312, 315, 316, 323, 324

map-type decay

A process applied to input map type, according to an underlying map type, that results in an output map type. **316**, 68, 92, 312, 316, 321, 503

matchable candidate

A mapped variable for which corresponding storage was created in a device data environment. **319**, 78, 320

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. **320**, 112, 250, 312, 320, 974

matching taskgraph record

A finalized taskgraph record that has a matching value for the scalar expression that identifies a **`taskgraph`** region. **478**, 102, 477, 478, 480, 481

memory

A storage resource for storing and retrieving variables that are accessible by threads. **7**, 6–11, 13, 20, 22, 33, 46, 55, 56, 58, 70, 72, 78–80, 84, 99, 101, 110, 111, 113, 120, 121, 130, 132, 162, 184, 248, 300, 335–339, 349, 395, 396, 506, 531–534, 541, 546, 548, 589, 601, 607, 657, 661, 663, 666, 673, 674, 685, 698, 702, 709, 710, 716, 781, 837, 840, 841, 864, 885, 890, 897–900, 902, 904, 911, 918, 936, 938, 940, 950, 968, 971, 972, 977–980, 983, 985

memory-allocating routine

A memory-management routine that has the memory-allocating-routine property. **709**, 22, 80, 98, 99, 130, 709, 710, 716

memory-allocating-routine property

The property that a memory-management routine allocates memory. **709**, 78, 711–715

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. **10**, 10, 22, 24, 25, 79, 132, 313, 336–345, 349, 394, 508, 591, 702, 707–710, 716, 954, 965, 968, 972, 980

memory-allocator-retrieving property

The property that a memory-management routine retrieves a memory allocator handle. **702**, 79, 703–706

memory-allocator-retrieving routine

A memory-management routine that has the memory-allocator-retrieving property. **702**, 702–707

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. **666**, 79, 667, 669–671

memory-copying routine

A routine that has the memory-copying property. **666**, 55, 99, 666, 667

memory fallback behavior

A behavior executed when a request for allocating or providing memory fails. 349

memory-management routine

A routine that has the memory-management-routine property. **685**, 22, 58, 78–80, 685, 690–692

memory-management-routine property

The property that a routine manages memory on the current device. **685**, 79, 686–691, 693–699, 701, 703–708, 711–716, 718–721

memory part

A storage block that resides on a single storage resource within a memory space. 79, 694, 696, 698

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. 79, 80, 338, 599, 601, 602, 694, 696–699

memory partitioner

An OpenMP object that represents mechanisms to create and to destroy memory partitions. 80, 337, 338, 593, 599–601, 693–699

memory-partitioning property

The property that a memory-management routine creates or destroys or otherwise affects memory partitions or memory partitioners. **693**, 80, 693–698

memory-partitioning routine

A memory-management routine that has the memory-partitioning property. **693**

memory-reading callback

A callback that has the memory-reading property. **901**, 902, 903

memory-reading property

The property that a callback reads memory from an OpenMP program. **901**, 80, 902, 903

memory-reallocating routine

A memory-management routine that has the memory-reallocating-routine property. **709**, 710, 715

memory-reallocating-routine property

The property that a memory-allocating routine deallocates memory in addition to allocating it. 80, 715

memory-setting property

The property that a routine fills memory in a device data environment with a specified value. **673**, 80, 674, 675

memory-setting routine

A routine that has the memory-setting property. **673**, 673–676

memory space

A representation of storage resources from which memory can be allocated or deallocated. More than one memory space may exist. **685**, 9, 10, 24, 25, 59, 79, 80, 115, 162, 313, 335, 338, 348, 349, 601, 685, 690–692, 698, 700, 702, 721, 954, 972, 980

memory-space-retrieving property

The property that a memory-management routine retrieves a memory space handle. **685**, 80, 686–689

memory-space-retrieving routine

A memory-management routine that has the memory-space-retrieving property. **685**, 685–690

mergeable task

A task that may be a merged task if it is an undeferred task. **482**, 115, 469, 482, 514, 526

merged task

A task with a minimal data environment. 80, 469, 482, 491, 503, 780, 843, 946, 947

metadirective

A directive that conditionally resolves to another directive. **360**, 49, 50, 102, 171, 360–364, 399, 954, 973, 974, 976, 980

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. 22, 81, 250, 253

modifier

A mechanism to customize clause behavior for its specified arguments. xxix, 24, 34–36, 43, 52, 57, 69–71, 77, 81, 84, 85, 88, 90, 95, 96, 101, 102, 108, 124, 125, 143, 151, 177–181, 183, 187–190, 193, 200, 201, 203, 204, 247, 248, 250, 255, 265, 268, 269, 273, 307, 311–313, 315–317, 320–322, 327, 330, 331, 348, 349, 352, 364, 367–369, 378, 379, 384, 435, 454, 460, 461, 477–479, 486, 503, 514, 516, 517, 548, 551–554, 573, 574, 801, 954, 956, 957, 963, 964, 966–971, 973–977, 979, 981

modifier group

A modifier set for which restrictions or properties related to their use on all clauses are specified. xxix, 178, 201, 203, 204, 317, 321, 367, 368, 551, 963

modifier set

A set of modifiers for which restrictions on their use or other properties of their use on a given clause are specified. **178**, 34, 52, 57, 81, 90, 102, 124, 125, 178, 180

multidimensional-information routine

A routine that has the multidimensional-information-routine property. 631

multidimensional-information-routine property

The property of a routine that it provides information about execution spaces. 81, 631–637

multidimensional value

A spatial value when interpreting the execution space as multidimensional. A multidimensional value is composed of one integer element for each dimension of the execution space. Section 3.4 defines the relationship between the unidimensional value and the multidimensional value of each spatial value. 81, 109, 110, 125, 139–141, 426, 631–634

mutex-acquiring callback

A callback that has the mutex-acquiring property. **827**

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. **827**, 81, 827, 828

mutex-execution callback

A callback that has the mutex-execution property. **828**

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. **828**, 82, 829, 830

mutual-exclusion construct

A construct that has the mutual-exclusion property. 82, 827, 829, 830

mutual-exclusion property

The property that a construct provides mutual-exclusion semantics. 82, 519, 541, 558, 560

mutually exclusive tasks

Tasks that may be executed in any order, but not at the same time. 82, 492, 553

mutually-exclusive-task scheduling constraint

The synchronization scheduling constraint that ensures at most one of a set of mutually exclusive tasks executes at a time. **492**, 491

N

named-handle property

The property that a handle is an integer kind in Fortran that is distinguished by the name of the handle. 583, 599, 604–606

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. 82, 185

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. 354–357, 359

natively-supported active spatial dimensions

The number of supported active spatial dimensions that are natively supported by the underlying device implementation for a type of execution space. The number of natively-supported active spatial dimensions is an implementation defined non-negative integer value. These dimensions form a consecutive sequence, starting with the first dimension. 83, 635, 636

native thread

An execution entity upon which an OpenMP thread may be implemented. 2, 3, 5, 6, 83, 89, 98, 121, 133, 153, 154, 422, 423, 432, 437, 780, 794, 795, 804, 807, 809, 810, 839, 849, 881, 893, 894, 900, 919–921, 931, 942

native thread context

A tool context that refers to a native thread. 886, 900, 902, 904, 905

native thread handle

A handle that refers to a native thread. 892, 918–921, 931, 933

native thread identifier

An identifier for a native thread defined by a native thread implementation. 156, 886, 894, 905, 916, 919, 920

native trace format

A format for implementation defined trace records that may be device-specific. 83, 765, 766, 876, 878

native trace record

A trace record in a native trace format. 766, 787, 788, 876, 878

nestable lock

A lock that can be acquired (i.e., set) multiple times by the same task before being released (i.e., unset). **723**, 83, 606, 723, 724, 731, 795, 831, 859

nestable lock property

The property that a routine operates on nestable locks. **723**, 83, 725, 727, 729, 731, 734, 736

nestable lock routine

A routine that has the nestable lock property. **723**, 606

nested construct

A construct (lexically) enclosed by another construct. 233

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, 978

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, 38, 84, 93, 404, 444

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. **237**, 39, 97, 126, 237–239, 242, 244, 245, 247, 249, 250, 277, 279, 289, 290, 312, 319, 323–325, 986

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. 145

noncommutative partial results list

A conceptually ordered list of partial results formed by a noncommutative reduction. 272

noncommutative reduction

A reduction for a given construct that must combine the set of values according to the canonical order of the construct. **272**, 84, 272

non-conforming program

An OpenMP program that is not a conforming program. 2, 35, 43, 126, 237, 295, 471, 490, 549, 723, 969

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. 358, 390, 955

non-host declare target directive

A declare target directive that does not specify a **`device_type`** clause with **`host`**. 381

non-host device

A device that is not the host device. 7, 20, 27, 113, 133, 136, 144, 158, 364, 395, 398, 422, 466, 493, 509, 647, 722, 750, 753, 915, 916, 921, 956, 965

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. **180**, 83, 121, 136, 147, 148, 152, 159–162, 179, 182, 191, 227, 336, 348, 359, 409, 415, 417, 421, 432, 485, 586, 621, 626, 631–634, 636, 637, 652, 655, 656, 692, 740, 755, 833, 857, 859, 958, 959

non-null pointer

A pointer that is not NULL. 677, 758, 760, 765, 807, 808, 878

non-null value

A value that is not NULL. 710, 792, 861, 863, 882, 900–902, 904, 935

non-property trait

A trait that is specified without additional properties. 354, 355, 359

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 semantic loop nest due to their presence. 406, 412

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. 85, 223, 225, 228, 230, 278, 280, 406, 411–413, 416, 418, 460, 463, 474, 475, 955, 979

non-sequentially consistent atomic construct

An **`atomic`** construct for which the **`seq_cst`** clause is not specified 14

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`**. 85, 90, 163, 250, 643, 650, 655, 659–661, 663, 665, 666, 673, 682, 683, 709, 710, 716, 718, 720, 744–746, 755, 758, 760, 765, 806, 818–820, 825, 833, 835, 836, 841, 843, 850, 853–855, 859–863, 882, 900, 902, 904, 909, 936, 960, 961

numeric abstract name

An abstract name that refers to a quantity associated with a conceptual abstract name. **145**, 20, 94, 146, 147, 152, 965

O

offsetting loop

The outer output loops of a **`stripe`** construct that determine the offsets within the grid cells used for each execution of the grid loops. **416**, 416, 955

OMPD

An interface that helps a third-party tool inspect the OpenMP state of a program that has begun execution. **880**, 2, 14–16, 86, 122, 132, 165, 207, 880–882, 884, 886, 888, 891, 892, 894, 897, 900, 905, 910–914, 920, 942

OMPD callback

A callback that has the OMPD property. 207, 887, 890, 891, 895, 897, 900, 902, 904, 905

OMPD library

A dynamically loadable library that implements the OMPD interface. **880**, 15, 48, 880–887, 890, 893–895, 897–900, 902–907, 909–916, 918, 931, 934, 936, 938, 940

OMPD property

The property that a callback, routine or type is included in OMPD and its namespace, which implies it has the **`ompd_`** prefix. 86, 883, 884, 886–896, 898, 899, 901–904, 906–908, 910–940

OMPD routine

A routine that has the OMPD property. 890, 891, 895, 910–915, 920, 922, 923, 925, 926, 939–941

OMPD type

A type that has the OMPD property. **206**, 34, 59, 90, 92, 207, 883–888, 890–900, 902, 904–907, 909

OMPT

An interface that helps a first-party tool monitor the execution of an OpenMP program. **757**, 2, 14, 15, 47, 87, 110, 163, 164, 207, 611, 750, 757–761, 763–766, 783, 786, 788, 794, 806–808, 834, 849, 850, 866, 867, 876, 877, 941, 960, 965, 972

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. 755, 760, 761, 768

OMPT callback

A callback that has the OMPT property. 207, 763, 772, 775, 806, 850, 866

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. 755, 759–761, 807

OMPT interface state

A state that indicates the permitted interactions between a first-party tool and the OpenMP implementation. 86, 87, 755, 759–761, 768, 807

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. 759, 760

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. 86, 87, 757, 758, 769, 771–773, 775–793, 795–798, 800–805, 807–815, 817–832, 834–839, 842, 844, 846, 849–878

OMPT-tool finalizer

An implementation of the **`finalize`** callback. **768**, 488, 758, 808

OMPT-tool initializer

An implementation of the **`initialize`** callback. **757**, 488, 758, 760, 761, 763, 807

OMPT type

A type that has the OMPT property. xxix, **206**, 34, 59, 90, 92, 207, 455, 757, 758, 760, 763, 766, 768, 769, 771, 772, 774–776, 778–792, 794, 796–800, 802–805, 807–813, 815–818, 820–825, 827–838, 840, 841, 843, 845, 847, 850–860, 862–878, 888, 894, 928, 934, 941, 960, 962, 972, 975, 978

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. **178**, 228, 229, 573

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. 88, 583, 599, 604–606, 678, 680–684, 771, 778, 834, 838, 875–877, 905, 913–915, 917, 918, 922–924, 927, 929–937, 939–941

opaque type

A type that has the opaque property. 68, 87, 88, 583, 599, 604–606

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/`. 88, 158, 355, 356, 515, 584, 586

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. 47, 89, 103, 131, 144, 257, 395, 396, 402, 578, 639, 685, 748, 754, 958

OpenMP architecture

The architecture on which a region executes. 89, 760

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. **354**, 34, 38, 110, 205, 354, 356–360, 362, 364–366, 368, 370, 371, 373, 374, 376, 377, 391, 587, 954, 955, 976

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. **144**, 47, 52, 131, 136, 144, 936, 984

OpenMP identifier

An identifier that has a specialized purpose for use in OpenMP programs, as defined by this specification. **205**, 66, 69, 77, 95, 99, 102, 177, 184, 205, 207, 257

OpenMP lock variable

A lock. 723

OpenMP object

Any object of an opaque type that allows programmers to save, to manipulate and to use state related to the OpenMP API. 44, 68, 79, 87, 549, 835, 839, 867, 875, 877, 878

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. **184**, 31, 66, 67, 89, 99, 181, 184, 205, 370, 446, 546

OpenMP operation list

An list argument that consists of OpenMP operation list items. **181**, 184

OpenMP operation list item

A list item that is an OpenMP operation. **181**, 89

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. 21, 89, 883, 884, 893, 900, 910, 911, 914, 915

OpenMP program

A program that consists of a base program that is annotated with OpenMP directives or that calls OpenMP API routines. 3, 5, 6, 8–10, 14, 15, 21, 23, 27, 33, 36, 37, 46–48, 50, 58, 60–63, 68, 80, 82, 84, 86, 88, 100, 102, 105, 123, 126, 131, 133, 144, 156, 166, 167, 184, 205, 237, 241, 268, 280, 295, 309, 328, 329, 335, 336, 354–356, 362, 396, 405, 433, 444, 485, 508, 509, 518, 520, 521, 544, 546, 549, 627, 630, 645, 653, 655, 666, 723, 738, 748, 750, 751, 754, 755, 757, 759, 760, 763, 780, 782, 806, 833, 853, 860, 865, 866, 872, 880–882, 885, 890, 893, 899, 902, 904, 907, 908, 942, 950, 986, 988

OpenMP property

The property that a routine, callback or type is in the OpenMP namespace, which implies it has the **`omp_`** prefix. 90, 580, 581, 583–585, 587, 589–591, 593, 594, 596, 598–600, 602–604, 606, 608, 609, 611–613, 619, 620, 678, 680–684, 686–697, 699–701, 703–707, 711–716, 718–722, 724–731, 733–736, 754

OpenMP stylized expression

A base language expression that is subject to restrictions that enable its use within an OpenMP implementation. 33, 34, 67, 177, 207, 257

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, 13, 23, 60, 83, 89, 93, 120, 121, 151, 152, 155, 614, 839, 916, 918–922, 925, 927, 935, 942

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, 23, 102, 121, 484, 489, 490

OpenMP type

A type that has the OpenMP property or a type that is an OMPD type or an OMPT type. **205**, 24, 34, 56, 59, 68, 87, 88, 92, 177, 181, 182, 184, 200, 202, 205–207, 227, 370, 413, 515, 551, 563, 578–581, 583, 584, 586, 588, 589, 591, 593, 595, 598–602, 604–613, 677, 833, 958, 975, 977

optional property

The property that a clause, an argument, a modifier, a clause set, or a modifier set does not have the required property. If any argument of a routine has the optional property, the argument can be omitted in C++ and Fortran and, for arguments of pointer type the value NULL can be passed in C and C++; additionally, the routine then has the overloaded property. 90, 176–178, 182, 229, 292, 361, 362, 367, 370, 377, 380, 383, 386, 393–398, 401–403, 407, 409, 419, 420, 430, 441, 458, 463, 480–482, 519, 529, 530, 532–539, 545, 555, 562, 580, 670, 672, 675, 678, 680, 681, 717–722

orderable construct

A construct that has the orderable property. 105, 119, 121, 273

orderable property

The property of a construct that it has a canonical order. 90, 273

`order`-`concurrent`-nestable construct

A construct that has the **`order`**-**`concurrent`**-nestable property. 437, 988

`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. 90, 410–412, 414, 416–418, 421, 438, 464, 541

`order`-`concurrent`-nestable routine

A routine that has the **`order`**-**`concurrent`**-nestable property. 437, 988

ordered-affected loop

For a worksharing-loop construct, a loop that is affected by the **`ordered`** clause. 49, 90, 230, 559

ordered iteration

A logical iteration of the ordered-affected loops of a worksharing-loop construct. 91, 108, 109, 548, 556, 559

ordered iteration space

The logical iteration space of the ordered-affected loops of a worksharing-loop construct. 556

ordered logical iteration

An ordered iteration. 556

original list item

The instance of a list item in the data environment of the enclosing context. 39, 51, 54, 76, 77, 91, 110, 111, 237–239, 241, 244, 245, 247–251, 255, 259, 264, 265, 267–275, 277–279, 289, 290, 302, 308, 311, 314, 323–325, 330, 331, 333, 382, 397, 459, 461, 463, 486, 511, 968, 986

original list-item updating clause

A clause that has the original list-item updating property 565

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. 91, 246, 269, 274, 276

original pointer

An original list item that corresponds to a corresponding pointer. 324

original storage

The storage of a given mapped variable. 9, 78, 101, 105, 313, 324–326, 801

original storage block

A storage block that contains the storage of one or more mapped variables in a data environment. 9, 39, 323

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. **8**, 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. 559

outermost-leaf property

The property that a clause applies to the outermost leaf construct that permits it when it appears on a compound construct. **178**, 251, 301, 528, 530, 573

output loop

A loop that represents the result of a loop-transforming construct and may be affected by a loop-nest-associated construct or loop-sequence-associated construct. 55, 58, 66, 86, 92, 105, 122, 224–226, 228, 406–408, 411, 412, 415–419, 473, 969

output loop nest

A semantic loop nest that consists of the output loops of a loop-transforming construct. 407, 412

output map type

The map type that results when map-type decay is applied to an input map type. **316**, 68, 78, 125, 312, 316, 321

overlapping type name

An OpenMP type for which its name has the overlapping type-name property. 816

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. 92, 779, 783, 796, 805, 814, 815, 823, 826, 828

overloaded property

The property that a routine has an overloaded C++ interface. 90, 92, 670, 671, 675, 678, 679, 681, 710–721

overloaded routine

A routine that has the overloaded property. 710, 716, 717

P

parallel handle

A handle that refers to a parallel region. **895**, 892, 922–924, 930, 932

parallelism-generating construct

A construct that has the parallelism-generating property. 248, 403, 406, 571

parallelism-generating property

The property that a construct enables parallel execution by generating one or more teams, explicit tasks, or SIMD instructions. 92, 421, 431, 438, 467, 470, 497, 500, 502, 505, 511

parallel region

A region that has a set of associated implicit tasks and an associated team of threads that execute those tasks. 4, 5, 21, 24, 33, 40, 56, 65, 92, 94, 112, 116, 118, 119, 130, 132, 143, 151, 155, 303, 304, 428, 429, 431, 442, 444–447, 449, 454, 464–467, 471, 521, 522, 524, 525, 572, 581, 614–616, 631, 632, 776, 783, 806, 820, 825, 860–863, 891, 895, 918, 922–924, 927, 930, 961, 985, 987

parameter list item

A list item that for C/C++ identifies one or more parameters of a procedure and for Fortran identifies one or more dummy arguments of a procedure. **182**, 97, 182, 183, 352

parent device

For a given **target** region, the device on which the corresponding **target** construct was encountered. 275, 395, 494, 506

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**, 23, 93, 428

partial tile

A tile that is not a complete tile. **417**, 418

partitioned construct

A construct that has the partitioned property. **444**, 93, 571

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. 93, 445, 447, 449, 452, 456, 457, 461, 464

partitioned worksharing construct

A construct that is both a partitioned construct and a worksharing construct. **4**, 93

partitioned worksharing region

A region that corresponds to a partitioned worksharing construct. 988

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. **221**, 61, 290, 411–413, 416–418, 558, 986

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. **396**, 9, 950

place

An unordered set of processors on a device. **147**, 4, 68, 94, 118, 120, 132, 148, 150, 151, 427–430, 656, 739–742, 856–858, 951, 956, 959, 963, 965

place group

An ordered list of place groups or places. **147**, 68, 94, 133, 145, 148–150, 428, 963

place-group abstract name

A conceptual abstract name that refers to a set of hardware abstractions of a given category that may be used to specify a place group. **145**, 94, 145, 148, 149

place-group child

A place group or place that is included in the list defined by a given place group. 145, 149, 428

place-group-count abstract name

A numeric abstract name that refers to a quantity associated with a place-group abstract name. **145**

place list

The ordered list that describes all OpenMP places available to the execution environment. 94, 432, 739, 856, 951, 963, 965

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. 741, 742, 858

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. 68, 120, 133, 428, 429

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. 508, 509

pointer attachment

The process of making a pointer variable an attached pointer. 26, 312, 319, 325, 964

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. **580**, 649, 655, 656, 669, 670, 672, 675, 681–684, 686, 688, 692, 700, 703, 705, 740, 742, 744–746, 758, 775, 787, 795, 807–815, 817–823, 826, 828, 830–832, 834, 836–839, 842, 844, 846, 849–851, 854, 855, 857–860, 862, 864, 867–878, 899, 901, 903, 905–907, 909–911, 913–915, 917–937, 939–941

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. **580**, 837, 844, 850, 851, 860, 862, 864, 898, 905, 912–915, 917–919, 921–927, 934, 937, 940

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. **180**, 113, 142, 146–148, 152, 153, 179, 182, 228–231, 336, 337, 340, 344, 352, 408, 409, 413, 420, 425, 430, 434–436, 440, 441, 458, 463, 474, 475, 495, 496, 579, 592, 593, 614, 628, 629, 654–656, 659, 669, 672, 686, 688, 700, 703, 705, 720, 795, 870, 951, 952, 956, 957, 959

post-modified property

The property of a clause that its modifiers must appear after its arguments. **176**, 180, 253, 278, 326, 351

preceding dependence-compatible task

For a given task, a dependence-compatible task that may be its antecedent task. **552**, 54, 65, 552, 553

predecessor task

For a given task, an antecedent task of that task, or any predecessor task of any of its antecedent tasks. **552**, 95, 115, 498, 501, 507, 512, 527, 552

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. **307**, 106, 252, 307, 314, 319–321, 330, 331

predefined identifier

Unless otherwise specified, an OpenMP identifier that is defined for use in arbitrary base language expressions. **205**, 7, 205, 578, 579, 645, 753, 769, 912, 958, 964

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. **234**, 43, 233–236, 240, 254, 305, 327, 506, 573, 986

preference specification

The specification of a set of preferences for interoperating with a foreign runtime environment. **516**, 96, 181, 517, 957

preference specification list

An list argument that consists of preference specification list items. **181**

preference specification list item

A list item that is a preference specification. **181**, 96, 516

pre-modified property

The property of a clause that its modifiers must appear before its arguments. **176**, 180

preprocessed code

For C/C++, a sequence of preprocessing tokens that result from the first six phases of translation, as defined by the base language. For Fortran, a sequence of source lines. 373, 976

present storage

A storage block that exists in a given device data environment. 313, 321–323, 325

primary thread

An assigned thread that has thread number 0. A primary thread may be an initial thread or the thread that encounters a **`parallel`** construct, forms a team, generates a set of implicit tasks, and then executes one of those tasks as thread number 0. **4**, 4, 5, 29, 56, 65, 93, 96, 119, 121, 294, 301, 302, 421, 422, 428, 429, 443, 445, 615, 861, 986

private attribute

For a given construct, a data-sharing attribute of a data entity that its lifetime is limited to that of the corresponding region and it is visible only to a single task generated by the construct or to a single SIMD lane used by the construct. **237**, 8, 23, 55, 66, 67, 69, 70, 96, 97, 99, 127, 179, 233–235, 237, 239, 248, 250, 253, 258, 259, 264, 269, 270, 272, 273, 275, 290, 304, 345, 406, 444, 564, 573, 986

private-only variable

A variable that has a private attribute and no other data-sharing attribute with respect to a given construct. **242**, 479

private variable

A variable that has a private attribute with respect to a given construct. **8**, 8, 55, 67, 70, 71, 97, 100, 238, 240, 289, 290, 301, 303, 450, 451, 453, 459, 462, 463, 967

privatization clause

The clause that may result in private variables that are new list items. **233**, 39, 84, 97, 240, 250

privatization property

The property that a clause privatizes list items. 242, 243, 246, 249, 250, 269, 274–276, 278, 487

privatized list item

A list item that appears in the list argument of a privatization clause, resulting in one or more private new list items. **237**, 237–240, 242, 249, 272, 273

procedure

A function (for C/C++ and Fortran) or subroutine (for Fortran). 16, 28, 31, 37, 42, 47, 54, 56, 57, 66, 72, 93, 100, 106, 112, 121, 122, 128, 164, 168, 173, 180–183, 192, 206, 210, 237, 242, 243, 245, 257, 279–282, 286, 306, 311, 318, 319, 329, 331, 352, 354, 355, 358, 359, 365, 367, 371, 372, 378–381, 383–385, 387, 441, 451–454, 493, 506, 507, 509, 578, 601, 602, 694, 696–699, 757, 758, 760, 768, 780, 782, 792, 862, 871, 885, 890, 891, 895, 900, 901, 905, 954, 955, 976, 980

procedure-argument list

An list argument that consists of procedure-argument list items. **181**

procedure-argument list item

A parameter list item, an actual argument position, or an actual argument position range. **182**, 97, 181, 192

procedure property

The property that a routine argument has a function pointer type in C/C++ and a procedure type in Fortran. **580**, 693, 872

processor

An implementation defined hardware unit on which one or more threads can execute. 45, 94, 133, 149, 154, 648, 740, 856, 857, 859, 867, 950, 951, 985

product order

The partial order of two logical iteration vectors $\omega_a = (i_1, \ldots, i_n)$ and $\omega_b = (j_1, \ldots, j_n)$, denoted by $\omega_a \leq_{\text{product}} \omega_b$, where $i_k \leq j_k$ for all $k \in \{1, \ldots, n\}$. 418

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, 98, 108

progress group

A group of consecutive threads in a team that may execute on the same progress unit. **430**, 430

progress unit

An implementation defined set of consecutive hardware threads on which native threads may execute a common stream of instructions. **6**, 7, 98, 430, 579, 648

property

A characteristic of an OpenMP feature. xxix, 22, 23, 25, 29–32, 34–36, 38, 40–43, 45–52, 54–59, 61–65, 68–73, 75, 76, 78–87, 89–93, 95–102, 104–110, 114, 116–121, 124–130, 178, 179, 187, 192, 198–201, 203, 228–231, 241–244, 246, 249–253, 255, 268, 269, 272, 274–276, 278, 279, 281, 284, 285, 287, 288, 291–293, 297, 298, 300, 301, 303, 307, 308, 310, 311, 315, 317, 318, 326–328, 331–333, 340, 341, 343, 344, 346–349, 351, 352, 355–357, 359–363, 366–370, 372, 374–377, 380, 382, 383, 385, 386, 388–391, 393–405, 407, 409–421, 425, 430, 431, 434–443, 445–447, 449, 452, 456–459, 461, 463–465, 467, 470, 474–476, 480–487, 489, 493–495, 497, 500, 502, 505, 511, 514–516, 518, 519, 521, 524, 526, 529, 530, 532–541, 545, 547–551, 555, 556, 558, 560, 562–564, 568, 573, 580, 581, 583–591, 593, 594, 596, 598–600, 602–606, 608, 609, 611–629, 631–637, 639–642, 645–656, 658–663, 665, 667–672, 674, 675, 677–684, 686–701, 703–708, 711–716, 718–722, 724–731, 733–736, 738–746, 748–752, 754, 757, 758, 769–773, 775–793, 795–798, 800–805, 807–815, 817–832, 834–839, 842, 844, 846, 849–878, 883, 884, 886–896, 898, 899, 901–941, 958, 968

pure property

The property that a directive has no observable side effects or state, yielding the same result every time it is encountered. 168, 170, 281, 285, 288, 293, 297, 327, 341, 361, 363, 370, 377, 382, 388, 403–405, 410–412, 414, 416–418, 438, 966, 973

R

raw-memory-allocating routine

A memory-allocating routine that has the raw-memory-allocating-routine property. **709**, 709–712

raw-memory-allocating-routine property

The property that a memory-allocating routine returns a pointer to uninitialized memory. **709**, 98, 711, 712

read-modify-write operation

An atomic operation that reads and then writes to a given storage location. 12

read structured block

An atomic structured block that may be associated with an **`atomic`** directive that expresses an atomic read operation. **212**, 212–214, 544

rectangular-memory-copying property

The property of a memory-copying routine that the memory that it copies forms a rectangular subvolume. **666**, 99, 669, 671

rectangular-memory-copying routine

A routine with the rectangular-memory-copying property. **666**, 666, 670, 673, 796, 841, 959

reduction

A use of a reduction operation. 34, 84, 99, 100, 117, 205, 256, 257, 259, 261, 262, 266–270, 273, 274, 471, 963, 967, 973, 976, 979, 982, 985

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. **266**, 100

reduction clause

A reduction-scoping clause or a reduction-participating clause. **256**, 67, 237, 240, 256–258, 263–265, 267, 268, 270, 274, 276, 282

reduction expression

A combiner expression or an initializer expression. **257**, 257

reduction identifier

An OpenMP identifier that specifies a combiner OpenMP operation to use in a reduction. **256**, 205, 256, 257, 261–266, 268, 282, 283, 285, 471, 967

reduction operation

An operation that applies a combiner and an associated initializer to a set of values. 34, 99, 104, 127, 192, 256

reduction-participating clause

A clause that defines the participants in a reduction. **256**, 99, 100, 256, 268, 269, 275

reduction-participating property

The property that a clause is a reduction-participating clause. 269, 275

reduction-scoping clause

A clause that defines the region in which a reduction is computed. **256**, 99, 100, 256, 267–269, 272, 274, 275, 471

reduction-scoping property

The property that a clause is a reduction-scoping clause. 269, 274

reduction variable

A private variable that has the reduction attribute with respect to a given construct. **266**, 266

referenced pointee

For a given referencing variable, the referenced data object to which the referring pointer points. 27, 28, 100, 251, 253, 317, 318, 323, 330

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. 27, 28, 100, 106, 128, 233, 235, 251, 253, 309, 312, 317, 318, 323, 330

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. 26, 27, 29, 39, 100, 233, 235, 253, 309, 317, 318, 323, 506

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 explicit-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. 3–6, 8, 9, 12, 14, 21–23, 27, 29, 30, 32, 33, 39, 40, 42, 44, 47, 49–54, 57, 58, 62–65, 67, 68, 77, 78, 84, 88, 90–93, 96, 100–103, 105–107, 111, 114–121, 124, 129–133, 139, 142, 145–147, 152, 155, 167, 174, 216, 217, 221, 228, 233, 237–239, 245, 247, 248, 251, 256,

257, 265–276, 294, 295, 301, 302, 312, 319, 320, 322–325, 331, 337, 339, 342, 345, 347, 349, 374–376, 381, 390, 394, 395, 401, 404, 421–425, 427, 431–433, 437–439, 442, 444–454, 461, 462, 464–468, 471, 472, 475, 477–479, 481, 487–490, 493, 494, 498, 500, 502–509, 511, 512, 514, 518–523, 525–527, 540, 541, 543–549, 557–561, 563–570, 579, 589, 611, 614, 615, 617, 622, 623, 625–628, 630–638, 640, 643, 645–655, 657, 673, 685, 701, 702, 707, 708, 710, 722, 724–730, 732–734, 736–738, 743–745, 747, 749, 750, 753–755, 763, 766, 776, 780, 786, 790, 794, 795, 797, 804, 806, 811–813, 815, 820, 824, 825, 827, 829, 830, 841, 843, 847, 852, 859, 861, 865, 915, 923, 942–948, 950, 952, 953, 955, 957, 960, 963, 964, 967, 969–971, 973, 976, 980, 982, 983, 985, 986, 988, 989

region endpoint

An event that indicates the beginning or end of a region that may be of interest to a tool. 763, 790

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. **180**, 179, 276, 279, 421, 432, 458, 463

registered callback

A callback for which callback registration has been performed. 15, 30, 86, 761, 763, 960

relaxed map

A mapping operation for which, if any part of the original storage of the list item has corresponding storage, correspondence is established or removed only for that portion; otherwise, it behaves like a strict map. 320, 321, 325

release flush

A flush that has the release flush property. **11**, 11, 12, 64, 101, 113, 546

release flush property

A flush with the release flush property orders memory operations that precede the flush before memory operations performed by a different thread with which it synchronizes. 56, 101, 546

release sequence

A set of modifying atomic operations that are associated with a release flush that may establish a synchronizes-with relation between the release flush and an acquire flush. 12, 543

repeatable property

The property that a clause or modifier may appear more than once in a given context with which it is associated. 178

replacement candidate

A directive variant or function variant that may be selected to replace a metadirective or base function. **360**, 31, 50, 361, 362, 364, 366, 371, 955

replayable construct

An explicit-task-generating construct that an implementation must record into a taskgraph record, if one is recorded. **477**, 102, 103, 116, 255, 477–479, 483

replay execution

An execution of a given **`taskgraph`** region that entails executing replayable constructs that are saved in a matching taskgraph record. **478**, 54, 103, 116, 255, 477–479, 957, 967

reproducible schedule

A loop schedule for the affected loop nest of a given loop-nest-associated construct that does not change between different executions of the construct that have the same binding thread set and have the same number of logical iterations. **444**, 228, 437, 454, 461, 464, 975

requested team space

In a **`parallel`** construct, the spatial list that consists of the requested number of threads for each dimension in the team space of the team that will execute the **`parallel`** region. **425**, 425, 427

required property

The property that a clause, a modifier, an argument, or at least one member of a clause set or modifier set is required and, thus, may not be omitted. **179**, 90, 176, 178–180, 200, 202, 268, 269, 274–276, 284, 287–289, 361, 366, 367, 392, 398, 409, 415, 502, 511, 514, 549, 556, 563, 580

reservation type

A thread-reservation type. 160

reserved locator

An OpenMP identifier that represents system storage that is not necessarily bound to any base language storage item. **184**, 182, 184, 550, 553, 554, 975

reserved thread

A thread in an OpenMP thread pool that must have a particular thread-reservation type when executing a task. 160

resource-relinquishing property

The property that a routine relinquishes some (or all) resources that the OpenMP program is currently using. **748**, 103, 749, 750

resource-relinquishing routine

A routine that has the resource-relinquishing property. **748**, 60, 108, 609–611, 748, 749

reverse-offload region

A region that is associated with a **target** construct that specifies a **device** clause with the **ancestor** *device-modifier*. 381, 981

right-hand side

The part of an assignment or expression that appears to the right of an equality or comparison operator. 213, 963

routine

Unless specifically stated otherwise, an OpenMP API routine. xxix, 2, 3, 6, 7, 13, 15, 16, 18, 20, 22, 23, 25, 29, 31, 36, 46, 51, 55, 57, 59, 61–64, 68–70, 72, 73, 79–81, 83–87, 89, 90, 92, 95, 97, 99–103, 108, 113, 117–119, 126, 127, 131, 136–139, 146, 157, 165, 200, 294, 337, 349, 350, 435, 437, 489, 491, 507, 508, 547, 548, 578–580, 582, 601, 602, 607, 609, 610, 612–643, 645–666, 668, 670, 671, 673, 675–681, 683–754, 758, 761, 806, 807, 816, 822, 831, 850, 862, 881, 890, 897, 910–931, 934–936, 938–941, 958–960, 965, 970–973, 977–988

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. **761**, 25, 51, 86, 87, 103, 757, 765, 807, 849

runtime error termination

An error termination that is performed during execution. **6**, 52, 167, 312, 325, 330, 349, 427, 435, 493, 494, 496, 652, 654, 657, 750, 952

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. 103, 431

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. **477**, 116, 255, 477, 478

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. 203, 207, 211, 218, 222, 234, 237, 248, 306, 840, 953, 982

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. **289**, 53, 66–68, 104, 126, 256, 272, 274, 289

scan phase

The portion of an affected iteration that includes all statements that read the result of a scan computation. **289**, 67, 289–292

schedulable task set

The set of tasks that may be selected for execution when a thread encounters a task scheduling point. **488**, 122, 489–492

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. **444**, 104, 228, 444

schedule-specification clause

A clause that has the schedule-specification property. **444**

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. 104, 437, 458, 463

schedule type

The part of a schedule specification that identifies the method by which the collapsed iterations are distributed to threads. 104, 133, 143, 153, 455, 460, 582, 619, 620, 958

scope handle

A handle that refers to an OpenMP scope. **891**, 939–941

segment

A portion of an address space associated with a set of address ranges. 21, 891

selectable callback

A callback that has the selectable property. 854

selectable property

The property that a callback can be registered selectively for events. 104, 811, 812, 814, 819, 824, 825, 827–830, 839, 842

selector argument

A callback argument that has the selector property. 854

selector property

The property of a callback argument that the argument is used for selective callback registration. 104, 811, 812, 814, 819, 823, 826, 828, 839, 842

selector set

Unless specifically stated otherwise, a trait selector set. 37, 47, 62, 114, 127, 359

self map

A mapping operation for which the corresponding storage is the same as its original storage. **324**, 93, 321, 323, 325, 397, 969

semantic loop

A syntactic loop or output loop of a canonical loop nest that may be an affected loop. 54, 105, 224–226, 406

semantic loop nest

The ordered set of semantic loops, from outermost to innermost, from a canonical loop nest. **224**, 21, 22, 73–75, 85, 92, 105, 220, 224–226, 406, 410–412, 416, 417, 419

semantic loop sequence

The sequence of semantic loop nests of a canonical loop sequence. **226**, 75, 220, 406, 407, 412, 414

separated construct

A construct for which its associated structured block is split into multiple structured block sequences by a separating directive. **173**, 105, 173, 174, 289, 290

separating directive

A directive that splits a structured block that is associated with a construct, the separated construct, into multiple structured block sequences. **173**, 105, 171, 173, 174, 288, 290, 449

sequentially consistent atomic operation

An atomic operation that is specified by an **`atomic`** construct for which the **`seq_cst`** clause is specified. 13, 14, 984

sequentially-orderable property

The property of an orderable construct that its canonical order is the original sequential order of the OpenMP program. 438, 447, 456, 457, 464

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.

106, 294, 743, 744

shallow map

A mapping operation for a list item in which only the storage directly occupied by the list item is mapped: the predefined default mapper applies to the list item and to any list item that results from any mapper expansion or array-map expansion due to those list items, except that any such list item that is a referencing variable is mapped as if with the **`ref_ptr`** *ref-modifier*. 326

shape-operator

For C/C++, an array shaping operator that reinterprets a pointer expression as an array with one or more specified dimensions. **184**, 184, 185, 330, 486, 553, 979

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. **241**, 7, 8, 106, 233–237, 242, 270, 277, 278, 468, 472, 498, 500, 506, 511, 953

shared variable

A variable that has the shared attribute with respect to a given construct. **8**, 8, 10–12, 14, 535–538

sharing task

A task for which the implicitly determined data-sharing attribute is shared unless explicitly specified otherwise. **236**, 106, 503

sharing-task property

The property that an explicit-task-generating construct generates sharing tasks. 502

sibling task

Two tasks are each a sibling task of the other if they are child tasks of the same task region. 106, 118, 552, 964

signal

A software interrupt delivered to a thread. 25, 106, 881

signal handler

A function called asynchronously when a signal is delivered to a thread. 7, 25, 781, 849, 881

SIMD

Single Instruction, Multiple Data, a lock-step parallelization paradigm. 279, 280, 351, 354, 378, 379, 441, 954, 955, 984

SIMD chunk

A set of iterations executed concurrently, each by a SIMD lane, by a single thread by means of SIMD instructions. **438**, 107, 378, 438, 440, 983

SIMD construct

A **`simd`** construct or a compound construct for which the **`simd`** construct is a constituent construct. 459, 460

SIMD instruction

A single machine instruction that can operate on multiple data elements. 2, 92, 107, 352, 438, 964

SIMDizable construct

A construct that has the SIMDizable property. 438, 988

SIMDizable property

The property that a construct may be encountered during execution of a **`simd`** region. 107, 410–412, 414, 416–418, 438, 464, 541, 560, 561

SIMD lane

A software or hardware mechanism capable of processing one data element from a SIMD instruction. 5, 8, 96, 107, 237–239, 242, 267–269, 272, 277, 279–281, 438, 441, 956

SIMD loop

A loop that includes at least one SIMD chunk. 351, 378

SIMD-partitionable construct

A construct that has the SIMD-partitionable property. 571

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. 107, 456, 457, 461, 470

simple-composite directive

A directive that has the simple-composite property. 107, 571, 572

simple-composite-directive name

A directive name of a simple-composite directive. 570, 572

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. 107, 502

simple lock

A lock that cannot be set if it is already owned by the task trying to set it. **723**, 108, 605, 723, 730

simple lock property

The property that a routine operates on simple locks. **723**, 108, 724, 726, 728, 730, 733, 735

simple lock routine

A routine that has the simple lock property. **723**, 605

simple modifier

A modifier that can never take an argument when it is specified. **177**, 177, 179, 180

simply contiguous array section

An array section that can be determined to have contiguous storage at compile time. In Fortran, this determination may result from the specification of the `CONTIGUOUS` attribute on the declaration of the array. 237, 953

simply happens before

For an event *A* to simply happen before an event *B*, *A* must precede *B* in simply happens-before order. **13**, 13

simply happens-before order

An ordering relation that is consistent with program order and the synchronizes-with relation. **13**, 60, 108

sink iteration

An ordered iteration with a stand-alone `ordered` directive that must not complete execution before a matching stand-alone `ordered` directive from a source iteration due to a doacross dependence. **556**, 49, 109

socket

The physical location to which a single chip of one or more cores of a device is attached. 145

soft pause

An instance of a resource-relinquishing routine that specifies that the OpenMP state is required to persist. **610**, 611

source iteration

An ordered iteration with a stand-alone **ordered** directive that must complete execution before a matching stand-alone **ordered** directive from a sink iteration due to a doacross dependence. **556**, 49, 108, 559

space-configuring clause

A clause that has the space-configuring property. 436

space-configuring property

The property of a clause that it specifies or limits the spatial size and spatial dimensionality of an execution space. 109, 425, 434, 495

spatial dimensionality

The number of dimensions, configurable to one or more, that define an execution space. 53, 109, 141, 425, 427, 431, 435, 436, 496

spatial ICV

An ICV that has the spatial property. 109, 110, 135, 140, 963

spatial index

A spatial value that indicates the index of an execution space unit within an execution space. 109, 119, 121, 140, 141

spatial index property

The property of a spatial ICV that it represents one or more spatial indices. 133

spatial limit

A spatial value that defines the maximum number of execution space units that any execution space may have within a specific context. 109, 140, 496

spatial limit property

The property of a spatial ICV that it represents one or more spatial limits. 133

spatial list

A list of integer values that represents a multidimensional spatial value that pertains to one or more execution spaces. Each list item corresponds to a specific dimension of the execution spaces, in the order of the ordinal sequence of the dimensions. 102, 135, 142, 424, 425, 427, 435, 436, 496

spatial property

The property of an ICV that it has a spatial value. 109

spatial size

A spatial value that indicates the number of execution space units within an execution space. 71, 109, 110, 119, 140, 141

spatial size property

The property of a spatial ICV that it represents one or more spatial sizes. 132, 133

spatial value

A value that specifies a spatial characteristic of an execution space or the execution space units within that space. A spatial value can be interpreted as a unidimensional value and as a multidimensional value. **140**, 81, 109, 110, 125, 135, 140

stand-alone directive

A unassociated directive that is also an executable directive. **171**, 49, 108, 109, 174, 175, 982, 983

standard trace format

A format for OMPT trace records. **765**, 771, 789, 876, 960

starting address

The address of the first storage location of a list item or, for a mapped variable of its original list item. 54, 77, 320

static context selector

The context selector for which traits in the OpenMP context can be fully determined at compile time. 49, 361, 363, 364

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. 9, 27, 46, 72, 121, 234, 237, 254, 259, 293, 296, 299, 305, 309, 310, 321, 322, 332, 336, 340, 342, 381, 382, 396, 477–479, 506, 950

step expression

A loop-invariant expression used by an induction operation. 33, 66, 67, 71, 190, 260, 261, 264, 286

storage block

The physical storage that corresponds to an address range in memory. 9, 20, 39, 50, 55, 76, 79, 91, 96, 111, 128, 312, 325, 350, 508, 957

storage location

A storage block in memory. 7–9, 21, 25–27, 44, 51, 72, 99, 110, 211, 216, 217, 250, 251, 275, 278, 279, 320, 339, 396, 440, 477, 541–544, 546, 547, 552–554, 661, 776, 953, 957

strictly nested region

A region nested inside another region with no other explicit region nested between them. 90, 119, 433, 437, 462, 466, 627, 630, 653, 655, 970, 988

strictly structured block

A single Fortran **`BLOCK`** construct, with a single entry at the top and a single exit at the bottom. 111, 172, 451

strict map

A mapping operation for which correspondence is established or removed for the full extent of the mapped address range of an original list item. **305**, 101, 325

string literal

For C/C++, a string literal. For Fortran, a character literal constant. 56, 159, 515, 517

striping

The reordering of logical iterations of a loop that follows a grid while skipping logical iterations in-between. **416**, 970

strong flush

A flush that has the strong flush property. **10**, 10, 11, 13, 14, 56, 64, 546

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. 56, 64, 111, 546

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. 38, 111, 112, 237, 252, 305, 307, 312–314, 321–323, 331–333, 508, 591, 758, 760, 768, 776, 779, 780, 786, 788, 789, 792, 795, 806–808, 816, 822, 863, 876, 883, 884, 887, 888, 895, 953, 982

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. **208**, 3, 8, 30, 37, 38, 44, 74, 91, 105, 124, 125, 151, 172–174, 208, 209, 218, 220, 221, 225, 250–253, 301–304, 378, 419, 422, 432, 442, 445, 446, 448,

450–453, 455, 461, 467, 468, 477, 481, 503, 520, 526, 560, 643, 766, 786, 803, 806, 816, 829, 830, 946, 947, 956

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. 30, 49, 51, 100, 105, 114, 173, 208, 220, 226, 247, 248, 289–292, 447–449, 956

structured operator

An operator that indicates a set of items to include and to reserve for structured threads. 148, 963

structured parallelism

Parallel execution through the implicit tasks of (possibly nested) parallel regions by the set of structured threads in a contention group. 160, 161

structured thread

A thread that is assigned to a team and is not a free-agent thread. 112, 118, 121, 133, 148, 160, 161, 426, 429, 489, 963, 966

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. 77, 128, 235, 305, 307, 311, 346, 508, 979

subroutine

A procedure for which a call cannot be used as the right-hand side of a base language assignment operation. 600, 602, 614, 618–621, 626, 629, 642, 645, 652, 654, 656, 662, 693–696, 701, 707, 716, 724–731, 733, 734, 740, 742, 745, 752, 772, 783, 808–815, 817–831, 834–839, 842, 844, 846, 866, 871

subsidiary directive

A directive that is not an executable directive and that appears only as part of a construct. 171, 174, 288–290, 448, 449, 471, 475–477, 970

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. **320**, 320

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, 490

successor task

For a given task, a dependent task of that task, or any successor task of a dependent task of that task. **552**, 113

supported active levels

An implementation defined maximum number of active levels of parallelism. 621, 622, 950

supported active spatial dimensions

The number of dimensions for a type of execution space that the implementation supports configuring as active spatial dimensions on a device. The number of supported active spatial dimensions is an implementation defined positive integer value. These dimensions form a consecutive sequence, starting with the first dimension. 83, 113, 635–638

supported device

The host device or any non-host device supported by the implementation, including any device-related requirements specified by the **`requires`** directive. 135, 158, 159, 493

synchronization construct

A construct that orders the completion of code executed by different threads. **518**, 2, 6, 13, 491, 548, 565, 821

synchronization hint

An indicator of the expected dynamic behavior or suggested implementation of a synchronization mechanism. **607**, 518, 607, 608, 723, 959, 981

synchronization scheduling constraint

A task-scheduling constraint that ensures the synchronization semantics of a construct (or its clauses) or a routine. **491**, 27, 82, 115, 116, 118, 124, 490, 491

synchronizes with

For an event *A* to synchronize with an event *B*, a synchronizes-with relation must exist from *A* to *B*. **12**, 11–13, 20, 422, 431, 490, 492, 519, 521, 525, 543, 548, 552, 559, 560, 643, 733, 734

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. **11**, 101, 108, 113, 548

synchronizing-region callback

A callback that has the synchronizing-region property. **824**, 826

synchronizing-region property

The property that a callback indicates the beginning or end of a synchronization-related region. **824**, 114, 824–826

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. 7, 114, 398, 430

syntactic loop

A loop that corresponds to *syntactic-loop-nest* in the grammar defined in Section 6.4.1. 105, 221, 223–225

T

`target`-consistent clause

A clause for which all expressions that are specified on it are **`target`**-consistent expressions. 433

`target`-consistent expression

An expression that has the **`target`**-consistent property. 114, 433

`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. 114, 198, 343, 348, 434, 495

target device

A device with respect to which the current device performs an operation, as specified by a device construct or device memory routine. **494**, 3, 4, 15, 41, 45, 47, 76, 115, 131, 132, 251, 253, 275, 312, 315, 323–325, 330, 332, 333, 344, 355, 396, 493, 494, 496–498, 501, 506, 507, 512, 645, 646, 657, 658, 661, 663–665, 757, 761, 764–766, 782, 783, 834, 835, 841, 843, 847, 865, 867, 868, 870–872, 878, 960, 968, 979

`target_device` selector set

A selector set that may match the target device trait set. **357**, 357–359, 964, 976

target device trait set

The trait set that consists of traits that define the characteristics of a device that the implementation supports. **355**, 114, 354–357, 359, 964, 965

target memory space

A memory space that is associated with at least one device that is not the current device when it is created. **685**, 338, 701, 702

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. 4, 275, 312, 498–501, 505–507, 511, 512, 657, 658, 667, 674, 780, 817, 822, 841, 843, 847, 863

target variant

A version of a device procedure that can only be executed as part of a `target` region. 354

task

A specific instance of executable code and its data environment that the OpenMP implementation can schedule for execution by a team. 3–6, 8, 9, 13, 20, 23, 29, 32, 39, 40, 43, 44, 46, 52, 54, 55, 57–59, 61, 64–66, 68, 72, 73, 80–83, 89, 92, 95, 96, 100, 102, 104, 106, 108, 113, 115–118, 121–126, 131–133, 141, 142, 151, 152, 201, 237–239, 241, 242, 244, 266–268, 270, 273–277, 294, 298, 313, 320, 322, 324–326, 336–338, 349, 350, 422, 423, 426, 430, 433, 442, 445, 446, 448, 449, 451, 453, 455, 461, 467–469, 471, 472, 474–479, 481, 482, 484–487, 489–492, 496, 503, 508, 514, 515, 519, 520, 522, 523, 525, 527, 529, 540, 541, 543, 544, 548, 552, 554, 557, 559, 560, 565, 567–569, 576, 579, 589, 605, 617, 620, 629, 630, 639, 643, 653, 654, 717, 718, 723–732, 780, 781, 783, 794, 802–804, 806, 817–823, 849, 862–864, 891, 895, 896, 925, 926, 928–930, 946, 947, 954, 957, 970, 971, 977, 984–987

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. 115, 117, 431, 468, 492, 548

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. **548**, 44, 115, 116, 123, 492, 548, 552, 555, 605, 639, 658, 776, 778, 964, 971, 977, 984

task-dependence scheduling constraint

The synchronization scheduling constraint that ensures that a dependent task does not execute before its predecessor tasks have reached task completion. **492**, 374, 491, 527, 552

task-generating construct

A construct that has the task-generating property. **3**, 3, 57, 81, 100, 117, 141, 142, 234, 236, 273, 529, 548, 553

task-generating property

The property that a construct generates one or more tasks; the generated tasks may be explicit tasks or implicit tasks. **3**, 116, 421, 431, 467, 470, 497, 500, 502, 505, 511, 526

taskgraph-altering clause

A clause that has the taskgraph-altering property. 477–479

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. 116, 474, 475, 551

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. **477**, 54, 102, 103, 116, 477–480, 957

taskgroup region

A region that enforces the taskgroup scheduling constraint, regardless of whether it is specified by an explicit **`taskgroup`** construct or an implicit taskgroup. 116, 491, 525

taskgroup scheduling constraint

The synchronization scheduling constraint that ensures that the synchronization semantics of a **`taskgroup`** region are enforced. **491**, **525**, 116, 471, 479, 491, 503

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, 116, 491, 525, 567

task handle

A handle that refers to a task region. 892, 925–927, 930, 933

task-inherited clause

A clause that has the task-inherited property. **476**

task-inherited property

The property of a clause that if it appears on a **`task_iteration`** directive, it will be inherited by the explicit tasks that are generated by an explicit-task-generating construct. 116, 486, 551

`taskloop`-affected loop

A collapse-affected loop of a **`taskloop`** construct. 190, 473, 476

task priority

A hint for the task execution order of tasks generated by a construct. **485**, 161, 485, 982, 983

task reduction

A reduction that is performed over a set of tasks that may include explicit tasks. **274**, 273, 275, 979

task region

A region consisting of all code encountered during the execution of a task. 4–6, 8, 39, 44, 93, 106, 113, 122, 126, 244, 294, 422, 432, 490, 491, 498, 500, 503, 511, 567, 641, 732, 776, 780, 783, 817, 862, 863, 924, 928, 946, 947

task-scheduling constraint

A limitation on the tasks that may be selected for execution when a task scheduling point is encountered. **490**, 5, 56, 113, 122, 490

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. 5, 104, 113, 117, 118, 266, 294, 422, 468, 479, 489–492, 521, 523, 525–527, 542, 666, 673, 803, 818, 819, 984, 985

task-scheduling-point-after-task-generation property

The property that a task scheduling point occurs in the encountering task of a task-generating construct after it generates a task. **489**, 467, 470, 497, 500, 502, 505, 511, 526

task-scheduling-point property

The property of a construct or routine that a task scheduling point occurs during its execution. **488**, 117, 118, 467, 470, 476, 489, 497, 521, 524, 526, 667, 669–671, 674, 675

task-scheduling-point-at-entry property

The property that a task scheduling point occurs in the encountering task at entry to a construct or a routine that has the task-scheduling-point property, after any implicit flushes that occur at entry to the construct or routine. **489**, 467, 470, 476, 497, 500, 502, 505, 511, 521, 526, 667, 669–671, 674, 675

task-scheduling-point-at-exit property

The property that a task scheduling point occurs at exit from the region in any task that executes a region associated with any construct or routine that has the task-scheduling-point property, after any implicit flushes that occur on exit from the construct or routine. **489**, 467, 470, 476, 489, 497, 500, 502, 505, 511, 524, 526, 667, 669–671, 674, 675

task switching

The act of a thread switching from the execution of one task to another task. 489, 490

task synchronization

The enforcement of a synchronization scheduling constraint through the use of task-synchronizing constructs, clauses or routines. 5, 468, 490

task-synchronizing property

The property of a construct, clause or routine that it implies a synchronization scheduling constraint. **491**, 118, 374, 421, 445–447, 449, 456, 457, 467, 470, 475, 476, 497, 500, 502, 503, 505, 511, 514, 521, 523, 524, 526, 528, 530, 551

taskwait order

An ordering relation that orders task *A* before its sibling task *B* if a full taskwait is encountered in the generating task region of the tasks after *A* is generated and before *B* is generated. 552

taskwait scheduling constraint

The synchronization scheduling constraint that ensures that the synchronization semantics of a full taskwait are enforced. **491**, 491, 527

team

A set of one or more assigned threads assigned to execute the set of implicit tasks of a parallel region. **4**, 3–5, 7, 21, 27, 40, 53, 65, 67, 71, 89, 92, 96, 98, 102, 112, 115, 119, 121, 124, 126, 129, 130, 133, 142, 143, 151, 152, 271, 272, 278, 280, 294, 301–303, 398, 421, 422, 425, 427, 428, 430–432, 435, 442, 444–450, 454, 455, 459, 460, 462, 463, 466, 491, 496, 519, 521, 522, 541, 560, 566, 615, 616, 626, 628, 631, 632, 634, 638, 652, 786, 794, 811, 820, 847, 861, 893, 918, 922–924, 927, 952, 956, 957, 975, 977, 985, 987, 988

team-eligible place

For a given parallel region, a place within its input place partition to which a structured thread may be assigned. 428, 429

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. 118, 445–447, 449, 456, 457, 464, 521

team-executed region

A region that is executed by all or none of the threads in the current team. **4**, 119, 988

team-generating construct

A construct that has the team-generating property. 988

team-generating property

The property that a construct generates a parallel region. 119, 421, 431, 505

team number

The spatial index 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 league size. 119, 133, 463, 628, 634, 820

team size

The spatial size that indicates the number of threads in a team. 133, 615, 631, 861

teams-nestable construct

A construct that has the **teams**-nestable property. 433, 988

teams-nestable property

The property that a construct or routine generates a region that may be a strictly nested region of a **teams** region. 119, 410–412, 414, 416–418, 421, 461, 464, 626, 627, 632, 633

teams-nestable routine

A routine that has the **teams**-nestable property. 433, 988

teams-orderable property

The property of an orderable construct that its canonical order is the order of the team numbers. 431

team space

The execution space of a team. 53, 102, 421, 424, 425, 427, 431, 496, 631–635, 637

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**, 129, 151, 429, 956

temporary view

The state of memory that is accessible to a particular thread. **7**, 7, 10, 11, 546, 548

third-party tool

A tool that executes as a separate process from the process that it is monitoring and potentially controlling. **880**, 15, 48, 86, 132, 880–882, 884–887, 890, 893–895, 897, 899, 900, 905, 907, 910, 911, 916, 942, 982

thread

Unless specifically stated otherwise, an OpenMP thread. 3–8, 10–16, 20, 21, 23, 24, 26, 27, 29, 36, 39–41, 49, 51–53, 55, 56, 58, 67, 68, 75, 78, 89, 92, 93, 96–98, 101, 102, 104, 106, 107, 111–115, 117–122, 124–126, 129, 131–133, 135, 146, 147, 151–153, 155, 156, 160, 161, 167, 228, 244, 245, 267, 268, 270, 271, 278, 280, 293–295, 301–304, 312, 336–339, 382, 388, 389, 395, 396, 401, 421–433, 442, 444–455, 459, 460, 462, 464–469, 471, 472, 477, 481, 484, 488–491, 496, 497, 499, 501, 507, 512, 518–525, 527–529, 540–544, 546, 547, 553, 557, 559–561, 565–567, 569, 579, 607, 614–619, 624, 629, 630, 632, 634, 637, 643, 653, 654, 661–667, 673, 674, 724–734, 736–738, 741, 752, 755, 757, 761, 766, 776, 786, 794, 795, 804, 809–811, 816, 820, 827, 831, 843, 849, 855, 858, 859, 861–864, 866, 877, 878, 885, 894–897, 900, 902, 904, 905, 910, 918, 920, 922–926, 928, 935, 942, 951–953, 956, 957, 963, 968, 970, 976, 977, 981, 985–988

thread affinity

A binding of threads to places within the current place partition. **427**, 131, 132, 151, 155–157, 294, 427–429, 738, 745, 746, 951, 956, 963, 978, 983, 984

thread-exclusive construct

A construct that has the thread-exclusive property. 988

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. 120, 519, 560

thread-limiting construct

A construct that has the thread-limiting property. 167

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. 120, 421, 431, 442, 445–447, 467, 505, 519, 560

thread number

For an assigned thread, the non-negative spatial index 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`**. **421**, 96, 121, 133, 294, 421, 430, 443, 459, 615, 624, 632, 820, 862, 918, 987

thread-pool-worker thread

A thread in an OpenMP thread pool that is not the initial thread. 804

threadprivate attribute

For a given OpenMP thread, a data-localizing 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. **293**, 121, 266, 293, 295, 297, 298, 302, 304, 986

threadprivate memory

The set of threadprivate variables associated with each thread. **7**, 295, 953

threadprivate variable

A variable that has the threadprivate attribute with respect to a given OpenMP thread. **293**, 121, 293–297, 301, 437, 454, 508

thread-reservation type

A categorization of a thread as either a structured thread or a free-agent thread. **160**, 102

thread-safe procedure

A procedure that performs the intended function even when executed concurrently (by multiple native threads). 16

thread-selecting construct

A construct that has the thread-selecting property. 571, 572

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. 121, 442, 445

thread-set

The set of threads for which a flush may enforce memory consistency. **10**, 10, 12, 13, 506, 540, 546, 547

threads-orderable property

The property of an orderable construct that its canonical order is the order of the thread numbers. 421, 446

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, 15, 757, 761, 794, 851, 859, 934, 935, 960

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, 122, 422, 481, 490

tied-task scheduling constraint

If any suspended tasks are tied to a thread and are not suspended at a barrier, a task-scheduling constraint that limits the explicit tied tasks in the schedulable task set that may be selected for execution by that thread to descendent tasks of all those suspended tasks. **490**, 490

tile

For a **`tile`** directive, the logical iteration space of the tile loops. 35, 93, 122, 410, 417, 420, 964

tile loop

The inner output loops of a **`tile`** construct that iterate over the logical iterations that correspond to a tile. **417**, 122, 417, 418, 420, 955, 970

tiling step

In a transformation defined as a sequence of steps, a step that performs a **`tile`** operation. 418, 420

tool

Code that can observe and/or modify the execution of an application. 2, 14–16, 18, 55, 101, 103, 120, 122, 123, 133, 163–165, 496, 504, 611–613, 668, 670, 671, 673, 675, 676, 749, 754, 755, 757–761, 763–767, 776, 781, 783, 787, 792, 794, 806–813, 815–818, 820–841, 843, 845, 847, 849–862, 864–878, 897–900, 902–907, 909–919, 922–924, 929, 931–934, 936–938, 940, 941, 960

tool callback

A procedure that a tool provides to an OpenMP implementation to invoke when an associated event occurs. 15, 30, 557, 576, 766, 807, 872, 960

tool context

An opaque reference provided by a tool to an OMPD library. A tool context uniquely identifies an abstraction. 21, 83, 122, 898, 904

tool defined

Behavior that must be documented by the tool implementation, and is allowed to vary among different compliant tools. 612, 755, 833

trace record

A data structure in which to store information associated with an occurrence of an event. 47, 83, 110, 126, 207, 765–767, 771, 786, 787, 789, 806, 822, 836, 838–841, 843, 845, 847, 867, 870–872, 874–878, 960, 962

trait

An aspect of an OpenMP implementation or the execution of an OpenMP program. 9, 10, 23, 32, 38, 47, 50, 54, 62, 82, 85, 88, 110, 115, 123, 127, 158, 159, 162, 335–340, 345, 347, 354–360, 368, 370, 373, 374, 376, 391, 592, 593, 601, 655, 694, 700, 709, 710, 954, 964, 965, 968, 976, 980

trait selector

A member of a trait selector set. **356**, 354, 357–360, 362, 365, 373, 964

trait selector set

A set of traits that are specified to match the trait set at a given point in an OpenMP program. **356**, 105, 123, 358

trait set

A grouping of related traits. **354**, 37, 47, 50, 62, 115, 123, 354, 357, 359

transformation-affected loop

An affected loop of a loop-transforming construct. 55, 228, 405, 406, 411–420

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. 308, 953, 954

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. **555**, 123, 479

type-name list

An list argument that consists of type-name list items. **181**, 188, 282, 283, 328

type-name list item

A list item that is the name of a type. **182**, 123, 181, 182, 184, 286, 287

U

ultimate property

The property that a clause, an argument, or a member of a clause set must be the lexically last clause or argument to appear on the directive. For a modifier or a modifier set, the property that the modifier or a member of the modifier set 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. **179**, 57, 178–180, 231, 268, 269, 274–276, 352, 362, 435, 458, 463

unassigned thread

A thread that is not currently assigned to any team. **3**, 3, 4, 56, 61, 121, 484, 489, 490, 615, 632, 795

unassociated directive

A directive that is not directly associated with any base-language code. **171**, 110, 171–173, 281, 285, 327, 363, 388, 391, 403, 405, 475, 489, 497, 500, 511, 514, 521, 526, 545, 549, 558, 564, 568

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. **213**, 213

undeferrable task

A task for which execution cannot be deferred with respect to its generating task region. That is, its generating task region must be suspended until execution of the structured block associated with the undeferrable task is completed. **468**, 124, 468, 469, 472, 479, 506, 548, 552, 825

undeferred 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 undeferred task is completed. 65, 80, 124, 469, 482, 491

undeferred-task scheduling constraint

The synchronization scheduling constraint that ensures that each undeferred task completes execution before its generating task is resumed. **491**, 491, 526

undefined

For variables, the property of not being defined; that is, the variable does not have a valid value. 9, 165, 565, 804, 855, 857, 858, 861, 862, 864, 865

underlying map type

The map type that determines which output map type results from an input map type. **316**, 78, 316

unidimensional value

A spatial value when interpreting the execution space as unidimensional. A unidimensional value is composed of a single integer element. Section 3.4 defines the relationship between the unidimensional value and the multidimensional value of each spatial value. 81, 110, 125, 135, 139–141, 160, 161, 617

unified address space

An address space that is used by all devices. **395**

uninitialized state

The lock state that indicates the lock must be initialized before it can be set. 72, 694, 696, 724, 728, 730, 735

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. 125, 769, 771, 775

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. For a modifier set, each member of the modifier set may appear at most once in the given context. **179**, 178–180, 187, 192, 198–201, 228–231, 241, 242, 244, 246, 249–253, 269, 274–276, 279, 284, 285, 287–289, 291, 292, 300, 301, 303, 307, 308, 310, 311, 317, 318, 326, 332, 333, 340, 341, 344, 347, 349, 351, 352, 361, 362, 366, 367, 369, 375–377, 380, 386, 389, 390, 393–403, 407–409, 413, 415, 419, 420, 425, 429, 430, 434–437, 440, 441, 443, 458, 459, 463, 465, 474, 475, 480–487, 493–495, 516, 518, 528–540, 550, 551, 555, 556, 562, 563

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. 125, 449, 450, 452, 453, 815

unlocked state

The lock state that indicates the lock can be set by any task. **723**, 73, 723, 724, 728, 730–734

unroll step

In a transformation defined as a sequence of steps, a step that performs an **unroll** operation. 420

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. 758, 811, 819, 826, 844, 846, 869

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.

8, 9, 35, 41, 61, 126, 144, 168, 202, 241, 245, 251, 259, 264, 296, 299, 322, 329, 337, 345, 395, 398, 485, 486, 507, 509, 524, 554, 565, 607, 635–638, 645–654, 657, 662, 664, 665, 677, 684, 700–702, 710, 716, 719, 720, 722, 723, 742–745, 747, 753, 866

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**, 115, 295, 469, 481, 490, 984

untraced-argument property

The property of an argument of a callback that it is omitted from the corresponding trace record of the callback. 809, 811, 817, 831, 832, 839, 842, 846

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. **214**, 215, 544

update structured block

An atomic structured block that may be associated with an **atomic** directive that expresses an atomic update operation. **213**, 36, 124, 214, 544

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. **289**, 126, 289

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. **252**

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. **250**

user-defined cancellation point

A cancellation point that is specified by a **`cancellation point`** construct. **568**, 569

user-defined induction

An induction operation that is defined by a **`declare_induction`** directive. **286**, 192, 286, 288, 967

user-defined mapper

A mapper that is defined by a **`declare_mapper`** directive. **328**, 77, 95, 205, 320, 328, 330, 974

user-defined reduction

A reduction operation that is defined by a **`declare_reduction`** directive. **282**, 192, 282, 284, 565, 566, 984

`user` selector set

A selector set that may match traits in the dynamic trait set. **358**, 358, 359

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. **388**, 127, 171, 388, 389, 405

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. **580**, 600, 659–663, 665, 668–670, 672, 674, 675, 711–716, 718–722, 795, 832, 836, 839

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.

7–14, 16, 22, 26–29, 32, 38–40, 42, 43, 46, 55, 56, 59, 66–72, 77–79, 82, 91, 96, 97, 100, 106, 110, 111, 114, 121, 124, 127, 128, 131, 171, 173, 174, 182, 183, 188, 192, 200, 202–207, 210, 211, 222, 223, 228, 233–242, 244, 245, 247, 248, 252, 254, 256–261, 265, 274, 277, 278, 281, 282, 286, 293–299, 301–311, 313–315, 321, 322, 327–329, 335, 336, 339–344, 346, 348, 358, 359, 362, 365, 366, 372, 376, 377, 381–386, 396, 406, 424, 432, 443, 444, 454, 459, 463, 465, 468, 472, 477–479, 484, 487, 493, 498, 500, 503, 506–509, 511, 546, 547, 555, 557, 573, 575, 723, 765, 804, 840, 841, 851, 852, 855, 860–863, 865, 881, 882, 894, 898, 900, 916, 950, 953, 954, 963, 967, 968, 974, 976, 979, 980, 983, 986

variable list

An list argument that consists of variable list items. **181**, 53, 265, 344

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. **182**, 54, 128, 181–183, 311, 477, 479

variant-generating directive

A declarative directive that has the variant-generating property. 361

variant-generating property

The property that a declarative directive generates a variant of a procedure. 128, 377, 382, 385

variant substitution

The replacement of a call to a base function by a call to a function variant. 49, 57, 355, 365, 374, 976

version string

A string that is formatted as a concatenation of key-value pairs that unambiguously represents the version of some aspect of the OpenMP implementation, where each key-value pair is formatted as **`[`***key*`=`*value*`]`. The *key* must not contain **`=`**, and *key* and *value* must not contain **`]`**. 128, 758, 912

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. **804**, 804, 859

white space

A non-empty sequence of space and/or horizontal tab characters. 48, 144, 153, 155, 168, 169, 174, 175, 190, 191, 571, 962

work distribution

The manner in which execution of a region that corresponds to a work-distribution construct is assigned to threads. 229

work-distribution construct

A construct that has the work-distribution property. **444**, 2, 93, 129, 130, 244, 245, 248, 266, 270, 444, 464, 814

work-distribution property

The property that a construct is cooperatively executed by threads in the binding thread set of the corresponding region. 129, 445–447, 449, 452, 456, 457, 461, 464

work-distribution region

A region that corresponds to a work-distribution construct. 245, 248, 444

worker thread

Unless specifically stated otherwise, a team-worker thread. 121, 423

worksharing construct

A construct that has the worksharing property. **444**, 4, 93, 129, 130, 245, 266, 269–273, 278, 280, 448, 454, 464, 523, 565, 566, 572, 573, 794

worksharing-loop construct

A construct that has the worksharing-loop property. **454**, 49, 90, 130, 153, 273, 277, 455, 456, 458, 459, 558–561, 565, 571, 815, 956, 968, 975, 980, 982, 983, 986, 987, 989

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. 129, 456, 457, 574

worksharing-loop region

A region that corresponds to a worksharing-loop construct. **454**, 133, 143, 454, 455, 558, 560, 567, 989

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. 129, 445–447, 449, 456, 457, 464

worksharing region

A region that corresponds to a worksharing construct. **444**, 4, 245, 270, 444, 522, 815, 976, 988

write structured block

An atomic structured block that may be associated with an **`atomic`** directive that expresses an atomic write operation. **212**, 212–214, 544

Z

zeroed-memory-allocating routine

A memory-allocating routine that has the zeroed-memory-allocating-routine property. **709**, 709, 713, 714

zeroed-memory-allocating-routine property

The property that a memory-allocating routine returns a pointer to memory that has been set to zero. **709**, 130, 713, 714

zero-length array section

An array section that does not include any elements of the array. 264, 312, 554

zero-offset assumed-size array

An assumed-size array for which the lower bound is zero. 250, 306, 321

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

Table 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
bind-var	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
cancel-var	global	Controls the desired behavior of the **`cancel`** construct and cancellation points
debug-var	global	Controls whether an OpenMP implementation will collect information that an OMPD library can access to satisfy requests from a third-party tool
def-allocator-var	implicit task	Controls the memory allocator used by memory allocation routines, directives and clauses that do not specify one explicitly
default-device-var	data environment	Controls the default target device
device-num-var	device	Device number of a given device
display-affinity-var	global	Controls the display of thread affinity
dyn-var	data environment	Enables dynamic adjustment of the number of threads used for encountered **`parallel`** regions
explicit-task-var	data environment	Boolean that is *true* if a given task is an explicit task, otherwise *false*
final-task-var	data environment	Boolean that is *true* if a given task is a final task, otherwise *false*
free-agent-thread-limit-var	data environment	Controls the maximum number of free-agent threads that may execute tasks in the contention group in parallel
free-agent-var	data environment	Boolean that is *true* if a free-agent thread is currently executing a given task, otherwise *false*
league-size-var	data environment	League size of a league; this ICV has the spatial size property
levels-var	data environment	Number of nested parallel regions such that all parallel regions are enclosed by the outermost initial task region on the device
max-active-levels-var	data environment	Controls the maximum number of nested active parallel regions when the innermost active parallel region is generated by a given task
max-task-priority-var	global	Controls the maximum value that can be specified in the **`priority`** clause

ICV	Scope	Description
nteams-var	device	Controls the number of teams requested for encountered **`teams`** regions
nthreads-var	data environment	Controls the number of threads requested for encountered **`parallel`** regions
num-devices-var	global	Number of available non-host devices
num-procs-var	device	The number of processors available on the device
place-group-var	implicit task	Controls the place group and associated place partition available for encountered **`parallel`** regions
run-sched-var	data environment	Controls the schedule used for worksharing-loop regions that specify the **`runtime`** schedule type
stacksize-var	device	Controls the stack size for threads that the OpenMP implementation creates
structured-thread-limit-var	data environment	Controls the maximum number of structured threads that may execute tasks in the contention group in parallel
target-offload-var	global	Controls the offloading behavior
team-generator-var	data environment	Generator type of current team that refers to a construct name or the OpenMP program
team-num-var	data environment	Team number of a given thread; this ICV has the spatial index property
team-size-var	data environment	Team size of the current team; this ICV has the spatial size property
teams-thread-limit-var	device	Controls the maximum number of threads that may execute tasks in parallel in each contention group that a **`teams`** construct creates
thread-limit-var	data environment	List that controls the maximum number of threads, in each dimension, that may execute tasks in the contention group in parallel; this ICV has the spatial limit property
thread-num-var	data environment	Thread number of an implicit task within its current team; this ICV has the spatial index property
tool-libraries-var	global	List of absolute paths to tool libraries
tool-var	global	Indicates that a tool will be registered
tool-verbose-init-var	global	Controls whether an OpenMP implementation will verbosely log the registration of a tool
wait-policy-var	device	Controls the desired behavior of waiting native threads

3.2 ICV Initialization

Table 3.2 shows the ICVs, associated environment variables, and initial values.

TABLE 3.2: ICV Initial Values

ICV	Environment Variable	Initial Value
active-levels-var	(none)	*0 (zero)*
affinity-format-var	**OMP_AFFINITY_FORMAT**	implementation defined
available-devices-var	**OMP_AVAILABLE_DEVICES**	*See below*
bind-var	**OMP_PROC_BIND**	implementation defined
cancel-var	**OMP_CANCELLATION**	*false*
debug-var	**OMP_DEBUG**	*disabled*
def-allocator-var	**OMP_ALLOCATOR**	implementation defined
default-device-var	**OMP_DEFAULT_DEVICE**	*See below*
device-num-var	(none)	*0 (zero)*
display-affinity-var	**OMP_DISPLAY_AFFINITY**	*false*
dyn-var	**OMP_DYNAMIC**	implementation defined
explicit-task-var	(none)	*false*
final-task-var	(none)	*false*
free-agent-thread-limit-var	**OMP_THREAD_LIMIT**, **OMP_THREADS_RESERVE**	*See below*
free-agent-var	(none)	*false*
league-size-var	(none)	*1 (one)*
levels-var	(none)	*0 (zero)*
max-active-levels-var	**OMP_MAX_ACTIVE_LEVELS**, **OMP_NUM_THREADS**, **OMP_PROC_BIND**	implementation defined
max-task-priority-var	**OMP_MAX_TASK_PRIORITY**	*0 (zero)*
nteams-var	**OMP_NUM_TEAMS**	*0 (zero)*
nthreads-var	**OMP_NUM_THREADS**	implementation defined
num-devices-var	(none)	implementation defined
num-procs-var	(none)	implementation defined
place-group-var	**OMP_PLACES**	implementation defined
run-sched-var	**OMP_SCHEDULE**	implementation defined
stacksize-var	**OMP_STACKSIZE**	implementation defined
structured-thread-limit-var	**OMP_THREAD_LIMIT**, **OMP_THREADS_RESERVE**	*See below*

ICV	Environment Variable	Initial Value
target-offload-var	**OMP_TARGET_OFFLOAD**	*default*
team-generator-var	(none)	*0 (zero)*
team-num-var	(none)	*0 (zero)*
team-size-var	(none)	*1 (one)*
teams-thread-limit-var	**OMP_TEAMS_THREAD_LIMIT**	*0 (zero)*
thread-limit-var	**OMP_THREAD_LIMIT**	implementation defined
thread-num-var	(none)	*0 (zero)*
tool-libraries-var	**OMP_TOOL_LIBRARIES**	*empty string*
tool-var	**OMP_TOOL**	*enabled*
tool-verbose-init-var	**OMP_TOOL_VERBOSE_INIT**	*disabled*
wait-policy-var	**OMP_WAIT_POLICY**	implementation defined

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

<ENVIRONMENT VARIABLE>`_ALL`

or

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

where *<ENVIRONMENT VARIABLE>* is the associated environment variable and *<device>* is the device number as specified in the **device** clause (see Section 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*.
- Each spatial ICV maintains a spatial value.
- The value of the *thread-limit-var* ICV is a spatial list.
- The initial value of *free-agent-thread-limit-var* is one less than the unidimensional value of the initial value of *thread-limit-var*.
- The initial value of *structured-thread-limit-var* is the unidimensional value of 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

Table 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
active-levels-var	(none)	**omp_get_active_level**
affinity-format-var	**omp_set_affinity_format**	**omp_get_affinity_format**
available-devices-var	(none)	(none)
bind-var	(none)	**omp_get_proc_bind**
cancel-var	(none)	**omp_get_cancellation**
debug-var	(none)	(none)
def-allocator-var	**omp_set_default_allocator**	**omp_get_default_allocator**
default-device-var	**omp_set_default_device**	**omp_get_default_device**
device-num-var	(none)	**omp_get_device_num**
display-affinity-var	(none)	(none)
dyn-var	**omp_set_dynamic**	**omp_get_dynamic**
explicit-task-var	(none)	**omp_in_explicit_task**
final-task-var	(none)	**omp_in_final**
free-agent-thread-limit-var	(none)	(none)
free-agent-var	(none)	**omp_is_free_agent**
league-size-var	(none)	**omp_get_num_teams**, **omp_get_num_teams_dim**
levels-var	(none)	**omp_get_level**
max-active-levels-var	**omp_set_max_active_levels**	**omp_get_max_active_levels**
max-task-priority-var	(none)	**omp_get_max_task_priority**
nteams-var	**omp_set_device_num_teams**	

ICV	Ways to Modify Value	Ways to Retrieve Value
		omp_get_device_num_teams
	omp_set_num_teams	**omp_get_max_teams**
nthreads-var	**omp_set_num_threads**	**omp_get_max_threads**
num-devices-var	(none)	**omp_get_num_devices**
num-procs-var	(none)	**omp_get_num_procs**
place-group-var	(none)	**omp_get_partition_num_places**, **omp_get_partition_place_nums**, **omp_get_place_num_procs**, **omp_get_place_proc_ids**
run-sched-var	**omp_set_schedule**	**omp_get_schedule**
stacksize-var	(none)	(none)
structured-thread-limit-var	(none)	(none)
target-offload-var	(none)	(none)
team-generator-var	(none)	(none)
team-num-var	(none)	**omp_get_team_num**, **omp_get_team_num_dim**
team-size-var	(none)	**omp_get_num_threads**, **omp_get_num_threads_dim**
teams-thread-limit-var	**omp_set_device_teams_thread_limit**	
		omp_get_device_teams_thread_limit
	omp_set_teams_thread_limit	
		omp_get_teams_thread_limit
thread-limit-var	**thread_limit**	**omp_get_thread_limit**, **omp_get_thread_limit_dim**
thread-num-var	(none)	**omp_get_thread_num**, **omp_get_thread_num_dim**
tool-libraries-var	(none)	(none)
tool-var	(none)	(none)
tool-verbose-init-var	(none)	(none)
wait-policy-var	(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-group-var* ICV are retrieved using the listed routines.

- The value of the *thread-limit-var* ICV is a list. The **omp_get_thread_limit** routine retrieves the unidimensional value of this ICV. The **omp_get_thread_limit_dim** routine retrieves the multidimensional value of this ICV.
- 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 36.9
- **omp_get_cancellation** Routine, see Section 37.1
- **omp_get_default_allocator** Routine, see Section 34.10
- **omp_get_default_device** Routine, see Section 31.2
- **omp_get_device_num** Routine, see Section 31.4
- **omp_get_device_num_teams** Routine, see Section 31.11
- **omp_get_device_teams_thread_limit** Routine, see Section 31.13
- **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 30.1.1
- **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 31.3
- **omp_get_num_procs** Routine, see Section 31.5
- **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 36.6
- **omp_get_partition_place_nums** Routine, see Section 36.7
- **omp_get_place_num_procs** Routine, see Section 36.3
- **omp_get_place_proc_ids** Routine, see Section 36.4
- **omp_get_proc_bind** Routine, see Section 36.1
- **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 30.1.2
- **omp_in_final** Routine, see Section 30.1.3
- **omp_set_affinity_format** Routine, see Section 36.8
- **omp_set_default_allocator** Routine, see Section 34.9
- **omp_set_default_device** Routine, see Section 31.1
- **omp_set_device_num_teams** Routine, see Section 31.12
- **omp_set_device_teams_thread_limit** Routine, see Section 31.14
- **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 Spatial ICVs

Spatial ICVs hold spatial values related to execution spaces. A spatial value can be interpreted as a unidimensional value and as a multidimensional value. A unidimensional value consists of a single integer element. In contrast, a multidimensional value consists of an integer element for each dimension. Unless otherwise specified, any reference to a spatial ICV or spatial value throughout this document refers to its unidimensional value.

A spatial value can be a spatial index, a spatial size, or a spatial limit. The invariant in Equation 3.1 holds between the unidimensional value and the multidimensional value of each spatial size and spatial limit.

$$S_{\text{unidim}} = S_{\text{dim}_0} \cdot S_{\text{dim}_1} \cdot S_{\text{dim}_2} \cdot \ \ldots \ \cdot S_{\text{dim}_{N-1}} \tag{3.1}$$

S_{unidim} is the integer unidimensional value, which indicates the total number of execution space units within the execution space. S_{dim_d} is the integer element of the multidimensional value in dimension d, which indicates the number of execution space units within the execution space in that dimension. N is the value that corresponds to the spatial dimensionality of the execution space. Equivalently, Equation 3.1 can be rewritten as Equation 3.2.

$$S_{\text{unidim}} = \prod_{n=0}^{N-1} S_{\text{dim}_n} \tag{3.2}$$

The invariant in Equation 3.3 holds between the unidimensional value and the multidimensional value of each spatial index.

$$I_{\text{unidim}} = I_{\text{dim}_0} + I_{\text{dim}_1} \cdot S_{\text{dim}_0} + \ \ldots \ + I_{\text{dim}_{N-1}} \cdot \prod_{n=0}^{N-2} S_{\text{dim}_n} \tag{3.3}$$

I_{unidim} is the integer unidimensional value, which indicates the linear index of an execution space unit within the execution space. I_{dim_d} is the integer element of the multidimensional value in dimension d, which indicates the index of the execution space unit within the execution space in that dimension. S_{dim_d} is the integer element of the multidimensional spatial size of the execution space in dimension d. N is the value that corresponds to the spatial dimensionality of the execution space. Equivalently, Equation 3.3 can be rewritten as in Equation 3.4.

$$I_{\text{unidim}} = \sum_{n=0}^{N-1} (I_{\text{dim}_n} \cdot \prod_{m=0}^{n-1} S_{\text{dim}_m}) \tag{3.4}$$

Equation 3.5 defines the inverse mapping of Equation 3.4, determining I_{dim_n} for a given I_{unidim} and dimension n.

$$I_{\text{dim}_n} = \left\lfloor \frac{I_{\text{unidim}}}{\prod_{m=0}^{n-1} S_{\text{dim}_m}} \right\rfloor \bmod S_{\text{dim}_n} \tag{3.5}$$

3.5 How the Per-Data Environment ICVs Work

When a task-generating 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 without the *dims-modifier* and 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 list 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 of the data environment of the resulting initial task is set to a spatial list of values with length equal to the number of dimensions specified in the **thread_limit** clause. The value of the *thread-limit-var* ICV for each dimension is set to the value specified in the clause for that dimension when *prescriptiveness-modifier* is **strict** and an implementation defined value between one and the value specified in the clause for that dimension when *prescriptiveness-modifier* is **relaxed**.

If a **target** construct with no **thread_limit** clause is encountered, the *thread-limit-var* ICV of the data environment of the resulting initial task is set to a spatial list consisting of a single implementation defined positive value.

If a **teams** construct with a **thread_limit** clause is encountered, the *thread-limit-var* ICV of the data environment of the initial task for each team is set to a spatial list of values with length equal to the number of dimensions specified in the **thread_limit** clause. The value of the *thread-limit-var* ICV for each dimension is set to the value specified in the clause for that dimension when *prescriptiveness-modifier* is **strict** and an implementation defined value between one and the value specified in the clause for that dimension when *prescriptiveness-modifier* is **relaxed**.

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 of the data environment of the initial task of each team is set to a spatial list consisting of a single implementation defined positive value that does not exceed the value of *teams-thread-limit-var*. If a **teams** construct with no **thread_limit** clause is encountered and the value of *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 a spatial list consisting of a single implementation defined positive value.

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 46.13.

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 46.13

3.6 ICV Override Relationships

Table 3.4 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
bind-var	**proc_bind**
def-allocator-var	**allocate**, **allocator**
nteams-var	**num_teams**
nthreads-var	**num_threads**
run-sched-var	**schedule**
teams-thread-limit-var	**thread_limit**

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

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

Cross References

- **allocate** Clause, see Section 13.6
- **allocator** Clause, see Section 13.4
- **num_teams** Clause, see Section 18.2.1
- **num_threads** Clause, see Section 18.1.2
- **proc_bind** Clause, see Section 18.1.4
- **schedule** Clause, see Section 19.6.3
- **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; if the value assigned to an environment variable is a list, then each list item may have leading and/or 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-group abstract names that are the strings defined in Table 4.1. If an environment variable is set to a value that includes a place-group abstract name, the behavior is as if the place-group abstract name were replaced with the list of place-group children associated with that abstract name on each device where the environment variable is applied.

TABLE 4.1: Predefined Place-group Abstract Names

Abstract Name	Meaning
`cores`	A place group where each place-group child corresponds to a single core of the device.
`ll_caches`	A place group where each place-group child corresponds to the set of cores for a single last-level cache of the device.
`numa_domains`	A place group where each place-group child corresponds to the set of cores for a single NUMA domain of the device.
`sockets`	A place group where each place-group child corresponds to the set of cores for a single socket of the device.
`threads`	A place group where each place-group child corresponds to a single hardware thread of the device.

For each place-group abstract name specified in Table 4.1, a corresponding place-group-count abstract name prefixed with **`n_`** also exists for which the associated value is the number of place-group children in the place group specified by the place-group 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 to a list with a single list item. 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-group-var* ICV. A list of places comprising a place group can be specified in the **OMP_PLACES** environment variable, optionally organized into a hierarchy of place groups. The value of **OMP_PLACES** can be a list of one or more types of values: an explicit list of place or place group specifications described by non-negative numbers, or a list of one or more place-group abstract name specifications that each implicitly refers to a list of places or place groups.

The **OMP_PLACES** environment variable can be defined with a comma-separated list of place group or place specifications. Each place group specification defines a place group with a comma-separated list of place group or place specifications enclosed by angle brackets. A place specification defines a place with 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 and place groups. Intervals can be specified using the *lower-bound* : *length* : *stride* notation to represent the following list of numbers: *lower-bound*, *lower-bound* + *stride*, ..., *lower-bound* + (*length* - 1)**stride*. When *stride* is omitted, a unit stride is assumed. Intervals can specify numbers within a place as well as sequences of places or place groups. When specifying a sequence of places or place groups, the stride defines an increment with respect to the set of numbers that defines the first place, or the first place of the first place group, in the sequence.

A structured operator "**structured**" can be used to indicate that the places or place groups enclosed in the immediately following parenthesized list are reserved for structured threads and free-agent threads cannot be assigned to them. The enclosed list may contain numbers, places, or place groups.

An exclusion operator "**!**" can also be used to exclude the number, place, or place group immediately following the operator from the innermost enclosing list of numbers, places, or place groups in which it appears. If the **OMP_PLACES** environment variable is defined using the exclusion operator only, the initial value of the *place-group-var* ICV is implementation defined but will exclude the specified place or place group.

Alternatively, the **OMP_PLACES** environment variable can be defined with a comma-separated list of place-group abstract name specifications. The specifications based on the place-group 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 interval, structured operator, and exclusion operator syntax described above for explicit place and place group specifications are also allowed for abstract name specifications.

An abstract name specification may be an abstract name that is appended with one or two positive numbers in parentheses, that is, *abstract_name* **(***len***)** or *abstract_name* **(***len***:***pstride***)** where *abstract_name* is a place-group abstract name listed in Table 4.1, *len* denotes the number of its

place-group children and *pstride* denotes the increment between consecutive place-group children. Alternatively, an abstract name specification may be an abstract name that is appended with a list of intervals (as described above) enclosed by square brackets, where each interval represents a list of indices to its place-group children that select the corresponding resources.

An abstract name specification may be followed by a dot (i.e., '**.**') that is then followed by a nested abstract name specification. This syntax denotes that the place-group children of the combined abstract name specification are the place-group children of the place group given by the nested abstract name specification in the context of the place group given by the enclosing abstract name specification. An abstract name specification may also be followed by a nested abstract name specification enclosed by angle brackets. This syntax denotes that the place-group children of the combined abstract name specification are the place groups given by the nested abstract name specification in the context of the place group given by the enclosing abstract name specification.

An abstract name specification may not be nested inside a **`threads`** abstract name specification, and only a **`threads`** abstract name specification may be nested inside a **`cores`** abstract name specification. Otherwise, both a **`threads`** and **`cores`** abstract name specification may be nested inside a specification for any of the other place-group abstract names listed in Table 4.1, and whether any other nesting not excluded by the above is permitted is implementation defined.

When requesting fewer place-group children than available for the place-group abstract name within the current place group context, the determination of which resources of type *abstract_name* are to be included as place-group children is implementation defined. When requesting more resources than available, the number of included place-group children 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-group abstract name.

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

$$
\begin{aligned}
\langle \text{list} \rangle &\models \langle \text{item} \rangle \mid \langle \text{list} \rangle,\langle \text{item} \rangle \\
\langle \text{item} \rangle &\models \langle \text{pg-item} \rangle \mid \langle \text{ag-item} \rangle \\
\langle \text{pg-list} \rangle &\models \langle \text{pg-item} \rangle \mid \langle \text{pg-list} \rangle,\langle \text{pg-item} \rangle \\
\langle \text{pg-item} \rangle &\models \langle \text{pg-interval} \rangle \mid \langle \text{p-interval} \rangle \mid \text{structured}(\langle \text{pg-list} \rangle) \\
\langle \text{pg-interval} \rangle &\models \langle \text{pgroup} \rangle:\langle \text{len} \rangle:\langle \text{stride} \rangle \mid \langle \text{pgroup} \rangle:\langle \text{len} \rangle \mid \langle \text{pgroup} \rangle \mid !\langle \text{pgroup} \rangle \\
\langle \text{pgroup} \rangle &\models <\langle \text{pg-list} \rangle> \\
\langle \text{p-interval} \rangle &\models \langle \text{place} \rangle:\langle \text{len} \rangle:\langle \text{stride} \rangle \mid \langle \text{place} \rangle:\langle \text{len} \rangle \mid \langle \text{place} \rangle \mid !\langle \text{place} \rangle \\
\langle \text{place} \rangle &\models \{\langle \text{res-list} \rangle\} \mid \langle \text{res} \rangle \\
\langle \text{ag-list} \rangle &\models \langle \text{ag-item} \rangle \mid \langle \text{ag-list} \rangle,\langle \text{ag-item} \rangle
\end{aligned}
$$

⟨ag-item⟩ ⊨ ⟨ag-interval⟩ | ⟨a-interval⟩ | structured(⟨ag-list⟩)
⟨ag-interval⟩ ⊨ ⟨agroup⟩:⟨len⟩:⟨stride⟩ | ⟨agroup⟩:⟨len⟩ | ⟨agroup⟩ | !⟨agroup⟩
⟨agroup⟩ ⊨ <⟨ag-list⟩>
⟨a-interval⟩ ⊨ ⟨a-spec⟩:⟨len⟩:⟨stride⟩ | ⟨a-spec⟩:⟨len⟩ | ⟨a-spec⟩ | !⟨a-spec⟩
⟨a-spec⟩ ⊨ ⟨a-spec-outer⟩ | ⟨a-spec-outer⟩.⟨a-spec⟩ | ⟨a-spec-outer⟩<⟨a-spec⟩>
⟨a-spec-outer⟩ ⊨ ⟨aname⟩ | ⟨aname⟩(⟨len⟩:⟨pstride⟩) | ⟨aname⟩(⟨len⟩) | ⟨aname⟩[⟨res-list⟩]
⟨aname⟩ ⊨ sockets | cores | ll_caches | numa_domains | threads
| *implementation-defined abstract name*
⟨res-list⟩ ⊨ ⟨res-interval⟩ | ⟨res-list⟩,⟨res-interval⟩
⟨res-interval⟩ ⊨ ⟨res⟩:⟨len⟩:⟨stride⟩ | ⟨res⟩:⟨len⟩ | ⟨res⟩ | !⟨res⟩
⟨res⟩ ⊨ *non-negative integer*
⟨len⟩ ⊨ *positive integer*
⟨stride⟩ ⊨ *integer*
⟨pstride⟩ ⊨ *positive integer*

Additionally, a list that is immediately enclosed by angle brackets, curly braces, or square brackets must specify at least one place group, place, or resource that is not excluded by the use of an exclusion operator.

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"
```

The fourth, fifth, and sixth definitions corresponds to the same place group of 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-group-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. An optional modifier of either **reset** or **noreset** can be specified first followed by a colon and one of **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 of the **place-list** 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 encountering thread is bound to the first place in the *place-group-var* ICV when binding is performed. An initial thread that is created by a **teams** construct is bound to the first place in its *place-group-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. Before binding the encountering thread at the beginning of an active parallel region its thread binding is stored. If the **reset** modifier is specified, then binding is applied to the encountering thread at the beginning of an active parallel region and reset to the stored value on leaving. If the **noreset** modifier is specified, or if neither is specified, then the encountering thread is bound prior to the first active parallel region, or immediately after encountering the first explicit-task-generating construct and that binding persists.

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-group-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-group-var* ICV, see Table 3.1
- **omp_get_proc_bind** Routine, see Section 36.1

- **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 *nteams-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

- *nteams-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 **omp_get_team_num**
T	**num_teams**	The value returned by **omp_get_num_teams**
L	**nesting_level**	The value returned by **omp_get_level**
n	**thread_num**	The value returned by **omp_get_thread_num**
N	**num_threads**	The value returned by **omp_get_num_threads**
a	**ancestor_tnum**	The value returned by **omp_get_ancestor_thread_num** with an argument of one less than the value returned by **omp_get_level**
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 32
- *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 the unidimensional value of *thread-limit-var* if the reservation type is **structured**, or the unidimensional value of *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
structured	Threads reserved for structured threads	1
free_agent	Threads reserved for free-agent threads	0

The **OMP_THREADS_RESERVE** environment variable sets the initial value of the *structured-thread-limit-var* and the *free-agent-thread-limit-var* ICVs according to Algorithm 4.1.

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

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

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*;

let *thread-limit-total* be the unidimensional value of *thread-limit-var*

if (*structured-reserve* $<$ 1) **then** *structured-reserve* = 1;

if (*free-agent-reserve* = *thread-limit-total*) **then** *free-agent-reserve* = *free-agent-reserve* - 1;

if (*threads-reserve* $\leq$ *thread-limit-total*) **then**

structured-thread-limit-var = *thread-limit-total* - *free-agent-reserve*;

free-agent-thread-limit-var = *thread-limit-total* - *structured-reserve*;

else behavior is implementation defined

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.

⟨allocator⟩ |= ⟨predef-allocator⟩ | ⟨predef-mem-space⟩ | ⟨predef-mem-space⟩:⟨traits⟩
⟨traits⟩ |= ⟨trait⟩=⟨value⟩ | ⟨trait⟩=⟨value⟩,⟨traits⟩
⟨predef-allocator⟩ |= *one of the predefined allocators from Table 13.3*
⟨predef-mem-space⟩ |= *one of the predefined memory spaces from Table 13.1*
⟨trait⟩ |= *one of the allocator trait names from Table 13.2*
⟨value⟩ |= *one of the allowed values from Table 13.2* | *non-negative integer*
| ⟨predef-allocator⟩

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 39
- *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 39
- *tool-libraries-var* ICV, see Table 3.1
- **ompt_start_tool** Procedure, see Section 39.2.1

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 39.2.2 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 39
- *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 45.3.1
- OMPD Overview, see Chapter 45
- *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 37.4). 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 37.4.

Cross References

- **omp_display_env** Routine, see Section 37.4

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 **ompx** 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 **ompx** 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 **!$omp** sentinel must be used for OpenMP directives that are defined in this specification; OpenMP reserves these uses for such directives.
- In fixed form source files, sentinels that end with **omp** must be used for OpenMP directives that are defined in this specification; OpenMP reserves these uses for such directives.
- In free form source files, the **!$ompx** 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 **omx** 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 **ompx_**, 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 **ompx_** 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**, **ompx**, **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++

Fortran

- 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.

Fortran

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.

White space is not significant except for distinguishing lexical tokens that comprise a directive and any of its clauses.

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, unless they are pure.

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:

C / C++

directive
structured-block

C / C++

Fortran

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

If *structured-block* is a loosely structured block, *end-directive* is required, 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 *loop-sequence-block*, 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 closely nested directives that 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.

C / C++

Fortran

5.1.1 Free Source Form Directives

The following sentinels are recognized in free form source files:

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

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.

Fortran

Fortran

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.

Fortran

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. The format of each clause in a *directive-specification* is given by *clause-specification*, which is:

clause-name`[(`*intra-clause-specification* `[;` *intra-clause-specification* `[;`*...*`]]``)`]

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 *intra-clause-specification* is required unless otherwise 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

An *intra-clause-specification* may be unmodified or modified. When unmodified, the format of *intra-clause-specification* is:

clause-argument-list

Unless otherwise specified, clauses have the pre-modified property, in which case the format of *intra-clause-specification* is:

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

Some clauses are explicitly specified to have the post-modified property, in which case the format of *intra-clause-specification* is:

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

A *clause-argument-list* is a list of at least one *argument-name* with the format:

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

In most clauses, the list only has a single *argument-name*. For many such clauses, *argument-name* refers to a list argument, which is a comma-separated list of a specific kind of list items (see Section 5.2.1).

If an argument is specified to have a list item type, then it is treated as having a list argument type that consists of only a single list item.

A *modifier-specification-list* is a comma-separated list of modifiers, each specified by a *modifier-specification* for which the format is:

> *modifier-specification [, modifier-specification [,...]]*

Modifiers may be simple modifiers or complex modifiers. For simple modifiers, the format of *modifier-specification* is:

> *modifier-name*

For complex modifiers, the format is:

> *modifier-name[* **(***modifier-argument-list***)** *]*

where *modifier-argument-list* is a comma-separated list as defined above for *clause-argument-list*. The *modifier-argument-list* and parentheses are required unless every *argument-name* is optional and omitted, in which case the format of the complex modifier is:

> *modifier-name*

Each *clause-name* and each *modifier-name* for a complex modifier represents an OpenMP keyword that identifies the clause or modifier. Each *argument-name* and each *modifier-name* for a simple modifier is an argument name or modifier name that is specified for the modifier, clause, or directive in which it may appear. Syntactically, each such *argument-name* or *modifier-name* represents one of the following:

- An OpenMP keyword;
- An OpenMP identifier;
- An OpenMP list argument;
- An expression of some OpenMP type; or
- An OpenMP stylized expression.

A particular lexical instantiation of an *argument-name* in an *intra-clause-specification* is an argument of the clause, while a lexical instantiation of a *modifier-specification* is a modifier that affects how or when the clause behavior is applied to the arguments.

The order of arguments must match the order in the *clause-specification* or *modifier-specification*. The order of modifiers in an *intra-clause-specification* is not significant unless otherwise specified.

C++

In an *intra-clause-specification*, if `:::` is present then the first `:` is treated as the separator that indicates the end of the *modifier-specification-list* or *clause-argument-list*, and the trailing `::` is treated as a scope resolution operator for the global namespace that applies to the first *argument-name* in *clause-argument-list* or the first *modifier-name* in *modifier-specification-list*.

C++

The clauses that a directive accepts may form clause sets, and the modifiers that a clause accepts may form modifier sets. For a given directive or clause, a clause set or modifier set may imply restrictions on the use of its members or otherwise capture properties for those members.

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. Similarly, all modifiers that are specified as a modifier group form a modifier set for which properties are specified with the specification of the modifier group. Additionally, some directives accept a clause group for which each member is a *directive-name* of a directive that has a specific property.

While specific properties may be defined for clauses, clause or directive arguments, modifiers, clause sets, or modifier sets, certain synactic properties that govern their use in a directive have general meanings and implications as indicated by the restrictions at the end of this section. These properties are listed in Table 5.1, along with the respective default properties for clauses, arguments, modifiers, clause groups, and modifier groups.

TABLE 5.1: Syntactic Properties for Clauses, Arguments, Modifiers, Clause groups, and Modifier groups

Property	Inverse Property	Clause defaults	Argument defaults	Modifier defaults	Clause group defaults	Modifier group defaults
required	optional	optional	required	optional	optional	optional
unique	repeatable	repeatable	unique	unique	unique	unique
exclusive	compatible	compatible	compatible	compatible	compatible	compatible
ultimate	free	free	free	free	free	free

A clause, argument, modifier, clause group, or modifier group 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, list arguments 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 *intra-clause-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 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 an *intra-clause-specification*.
- An argument with the exclusive property must not appear if an argument with a different *argument-name* appears in the *intra-clause-specification*.
- A modifier with the required property must appear in the *intra-clause-specification*.
- A modifier with the unique property may appear at most once in a *intra-clause-specification*.
- A modifier with the exclusive property must not appear if a modifier with a different *modifier-name* also appears in the *intra-clause-specification*.
- If a modifier set has the required property, at least one modifier in the set must be present in the *intra-clause-specification* for which the modifier set is specified.
- If a modifier is a member of a modifier set that has the unique property for a clause then the modifier has the unique property for that clause regardless of whether it has the unique property when it is not part of such a modifier set.
- If one modifier of a modifier set with the exclusive property appears in a *intra-clause-specification*, no other modifiers with a different *modifier-name* in that modifier set may appear in the *intra-clause-specification*.
- If a clause has the pre-modified property, a modifier with the ultimate property must be the last modifier in any *intra-clause-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 *intra-clause-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 with the constant property must be a compile-time constant.
- An argument with the positive property must be greater than zero.
- An argument with the non-negative property must be greater than or equal to zero.
- An argument 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.

C / C++

- If an expression is an argument to a directive, clause, or modifier and is the result of a comma operator, its result must be immediately enclosed in parentheses within the argument.

C / C++

Cross References

- Directive Format, see Section 5.1
- OpenMP List Arguments, see Section 5.2.1
- OpenMP Stylized Expressions, see Section 6.2
- OpenMP Types and Identifiers, see Section 6.1

5.2.1 OpenMP List Arguments

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 list arguments allow the corresponding semantics to apply to multiple list items. In any list argument 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 list argument 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 procedure-argument list consists of a comma-separated collection of one or more procedure-argument 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 depend list is a comma-separated collection of one or more depend list items, which are either all locator list items or all expressions of **depend** OpenMP type.

Each procedure-argument list item can be one of the following:

- A parameter list item;
- The position of an actual argument passed to a procedure call and specified by a positive integer, where 1 represents the first argument; or
- An argument position range specified by $lb : ub$ where both lb and ub must be a constant and positive expression of integer OpenMP type.

In both lb and ub, an expression using **`omp_num_args`**, that enables identification of arguments relative to the last argument of the call, can be used with the form:

`omp_num_args` *[$\pm$ logical_offset]*

where *logical_offset* is a constant and non-negative expression of integer OpenMP type. 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 argument positions $lb, lb + 1, \ldots, ub$ had been specified individually.

C / C++

Each parameter list item is the name of a procedure parameter. Each variable list item is a variable or an array section. Each locator list item is a reserved locator, an array section, or any lvalue expression including variables. Each type-name list item is a type name.

C / C++

Fortran

Each parameter list item is a dummy argument of a procedure. 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).

Each locator list item is a variable list item, a function reference with data pointer result, or a reserved locator. Each type-name list item is a type specifier in the base language.

When a named common block acquires an attribute because of its presence in a list argument, all members of that common block acquire the same attribute. 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 list arguments 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.

C / C++

- Each parameter list item starting with a colon must be separated by whitespace from a preceeding colon.

C / C++

- 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 unless it is referenced via an implicit member access.

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.

- Each 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 **:=** **([***s*$_1$**][***s*$_2$**]...[***s*$_n$**])** *cast-expression*

The result of applying the shape-operator to an expression is an lvalue expression with an n-dimensional array type with dimensions $s_1 \times s_2 \ldots \times s_n$ 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 : ]
[  : length ]
[ : : stride]
[ : : ]
[ : ]
```

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

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

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

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

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

The *stride* must evaluate to a positive integer.

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

When the *length* is absent and the size of the dimension is known, it defaults to $\lceil (size - lower\text{-}bound)/stride \rceil$, 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.

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
iterator	*all arguments*	Complex, name: **`iterator`** Arguments: ***iterator-specifier-list*** list of iterator specifier list item type (*default*)	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-specifier-list* argument is an iterator specifier with this format:

C / C++

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

C / C++

Fortran

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

Fortran

where:

- *iterator-identifier* is a base language identifier.
- *iterator-type* is a type that is permitted in a type-name list.
- *range-specification* is of the form *begin* : *end*[: *step*], where *begin* and *end* are expressions for which their types can be converted to *iterator-type* and *step* is an integral expression.

C / C++

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

C / C++

Fortran

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

Fortran

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

An iterator only exists in the context of the clause 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 $i_0, \ldots, i_{N-1}$ where:

- $i_0 = (\textit{iterator-type})\ \textit{begin}$;
- $i_j = (\textit{iterator-type})\ (i_{j-1} + \textit{step})$, where $j \geq 1$; and
- if $\textit{step} > 0$,
  - $i_0 < (\textit{iterator-type})\ \textit{end}$;
  - $i_{N-1} < (\textit{iterator-type})\ \textit{end}$; and
  - $(\textit{iterator-type})\ (i_{N-1} + \textit{step}) \geq (\textit{iterator-type})\ \textit{end}$;
- if $\textit{step} < 0$,
  - $i_0 > (\textit{iterator-type})\ \textit{end}$;
  - $i_{N-1} > (\textit{iterator-type})\ \textit{end}$; and
  - $(\textit{iterator-type})\ (i_{N-1} + \textit{step}) \leq (\textit{iterator-type})\ \textit{end}$.

C / C++

Fortran

The iterator value set of the *iterator* are the set of values $i_1, \ldots, i_N$ where:

- $i_1 = begin$;
- $i_j = i_{j-1} + step$, where $j \geq 2$; and
- if $step > 0$,
  - $i_1 \leq end$;
  - $i_N \leq end$; and
  - $i_N + step > end$;
- if $step < 0$,
  - $i_1 \geq end$;
  - $i_N \geq end$; and
  - $i_N + step < end$.

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.

Restrictions

Restrictions to *iterator* modifiers are as follows:

- 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*.

C / C++

- The *iterator-type* must be an integral or pointer type.
- The *iterator-type* must not be **`const`** qualified.

C / C++

Fortran

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

Fortran

- If the step expression of a *range-specification* equals zero, the behavior is unspecified.
- Each *iterator-identifier* can only be defined once in 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*.

Cross References

- **affinity** Clause, see Section 20.10
- **depend** Clause, see Section 23.9.6
- **from** Clause, see Section 12.2
- **map** Clause, see Section 11.5
- **to** Clause, see Section 12.1

5.3 Conditional Compilation

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

Fortran

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

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.

Fortran

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.

Fortran

5.4 Intrinsic Identifiers

Intrinsic Identifiers

Name	Type	Properties
omp_fill	integer	non-negative
omp_idx	see below	*default*
omp_in	see below	*default*
omp_num_args	integer	non-negative
omp_orig	see below	*default*
omp_out	see below	*default*
omp_priv	see below	*default*
omp_step	see below	*default*
omp_var	see below	*default*

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 procedure-argument list items and is a context-specific value that evaluates to the number of actual arguments that are passed 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 their types are described in Section 8.2.1. The **omp_priv** and **omp_orig** intrinsic identifiers can only be used in initializer expressions, and their types are described in Section 8.2.2.

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, and their types are described in Section 8.16. 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.

Cross References

- **declare_induction** Directive, see Section 8.16
- OpenMP Combiner Expressions, see Section 8.2.1
- OpenMP Initializer Expressions, see Section 8.2.2

5.5 *directive-name-modifier* Modifier

Modifiers

Name	Modifies	Type	Properties
directive-name-modifier	*all arguments*	Keyword: *directive-name* (a directive name)	unique

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**, **counts**, **default**, **defaultmap**, **depend**, **depth**, **destroy**, **detach**, **device**, **device_safesync**, **device_type**, **dist_schedule**, **doacross**, **dyn_groupprivate**, **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**, **looprange**, **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**, **reduction_kind**, **relaxed**, **release**, **replayable**, **reverse_offload**, **safelen**, **safesync**, **schedule**, **self_maps**, **seq_cst**, **severity**, **shared**, **shift**, **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 15.5.2
- **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 15.5.3
- **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 17.11.1
- **looprange** Clause, see Section 6.4.8
- **partial** Clause, see Section 17.11.2
- **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
- **counts** Clause, see Section 17.8.1
- **default** Clause, see Section 7.4
- **defaultmap** Clause, see Section 11.6
- **depend** Clause, see Section 23.9.6
- **depth** Clause, see Section 6.4.7
- **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 23.9.8
- **dyn_groupprivate** Clause, see Section 13.9
- **dynamic_allocators** Clause, see Section 16.5.1.2
- **enter** Clause, see Section 11.3
- **exclusive** Clause, see Section 8.17.2
- **fail** Clause, see Section 23.8.3.3

- **filter** Clause, see Section 18.6.1
- **final** Clause, see Section 20.7
- **firstprivate** Clause, see Section 7.3.3
- **from** Clause, see Section 12.2
- **grainsize** Clause, see Section 20.2.1
- **graph_id** Clause, see Section 20.3.1
- **graph_reset** Clause, see Section 20.3.2
- **has_device_addr** Clause, see Section 7.3.7
- **hint** Clause, see Section 23.1
- **holds** Clause, see Section 16.6.1.3
- **if** Clause, see Section 5.6
- **in_reduction** Clause, see Section 8.12
- **inbranch** Clause, see Section 15.7.1.1
- **inclusive** Clause, see Section 8.17.1
- **indirect** Clause, see Section 15.8.3
- **induction** Clause, see Section 8.13
- **inductor** Clause, see Section 8.16.1
- **init** Clause, see Section 5.7
- **init_complete** Clause, see Section 8.17.3
- **initializer** Clause, see Section 8.15.2
- **interop** Clause, see Section 15.6.1
- **is_device_ptr** Clause, see Section 7.3.5
- **lastprivate** Clause, see Section 7.3.4
- **linear** Clause, see Section 8.14
- **link** Clause, see Section 11.4
- **local** Clause, see Section 9.3
- **map** Clause, see Section 11.5
- **match** Clause, see Section 15.5.1
- **memscope** Clause, see Section 23.8.4

- **mergeable** Clause, see Section 20.5
- **message** Clause, see Section 16.3
- **no_openmp** Clause, see Section 16.6.1.4
- **no_openmp_constructs** Clause, see Section 16.6.1.5
- **no_openmp_routines** Clause, see Section 16.6.1.6
- **no_parallelism** Clause, see Section 16.6.1.7
- **nocontext** Clause, see Section 15.6.3
- **nogroup** Clause, see Section 23.7
- **nontemporal** Clause, see Section 18.5.1
- **notinbranch** Clause, see Section 15.7.1.2
- **novariants** Clause, see Section 15.6.2
- **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 18.4
- **ordered** Clause, see Section 6.4.6
- **otherwise** Clause, see Section 15.4.2
- **permutation** Clause, see Section 17.6.1
- **priority** Clause, see Section 20.9
- **private** Clause, see Section 7.3.2
- **proc_bind** Clause, see Section 18.1.4
- **read** Clause, see Section 23.8.2.1
- **reduction** Clause, see Section 8.10
- **reduction_kind** Clause, see Section 8.15.3
- **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

- **safelen** Clause, see Section 18.5.2
- **safesync** Clause, see Section 18.1.5
- **schedule** Clause, see Section 19.6.3
- **self_maps** Clause, see Section 16.5.1.6
- **seq_cst** Clause, see Section 23.8.1.5
- **severity** Clause, see Section 16.4
- **shared** Clause, see Section 7.3.1
- **shift** Clause, see Section 17.2
- **simd** Clause, see Section 23.10.3.2
- **simdlen** Clause, see Section 18.5.3
- **sizes** Clause, see Section 17.3
- **task_reduction** Clause, see Section 8.11
- **thread_limit** Clause, see Section 21.3
- **threads** Clause, see Section 23.10.3.1
- **threadset** Clause, see Section 20.8
- **to** Clause, see Section 12.1
- **transparent** Clause, see Section 23.9.7
- **unified_address** Clause, see Section 16.5.1.4
- **unified_shared_memory** Clause, see Section 16.5.1.5
- **uniform** Clause, see Section 14.1
- **untied** Clause, see Section 20.4
- **update** Clause, see Section 23.8.2.2
- **update** Clause, see Section 23.9.4
- **use** Clause, see Section 22.1.2
- **use_device_addr** Clause, see Section 7.3.8
- **use_device_ptr** Clause, see Section 7.3.6
- **uses_allocators** Clause, see Section 13.8
- **weak** Clause, see Section 23.8.3.4
- **when** Clause, see Section 15.4.1

- **`write`** Clause, see Section 23.8.2.3

5.6 `if` Clause

Name: `if`	**Properties:** target-consistent

Arguments

Name	Type	Properties
if-expression	expression of OpenMP logical type	*default*

Modifiers

Name	Modifies	Type	Properties
directive-name-modifier	*all arguments*	Keyword: *directive-name* (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 18.5
- **`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: `init`	**Properties:** innermost-leaf

Arguments

Name	Type	Properties
init-var	variable of OpenMP type	*default*

Modifier groups

interop-type

Modifiers

Name	Modifies	Type	Properties
prefer-type	*init-var*	Complex, name: **`prefer_type`** Arguments: ***prefer-type-specification*** list of preference specification list item type (*default*)	complex, unique
depinfo-modifier	*init-var*	Complex, Keyword: **`in`**, **`inout`**, **`inoutset`**, **`mutexinoutset`**, **`out`** Arguments: ***locator-list-item*** locator list item (*default*)	complex, unique
directive-name-modifier	*all arguments*	Keyword: *directive-name* (a directive name)	unique

Directives

`depobj`, **`interop`**

Semantics

The **init** clause initializes the *init-var* argument, an *uninitialized* handle of OpenMP type, by changing its state to *initialized* if the initialization requirements can be satisfied. This initialization creates an underyling object that can be referenced by the handle in various directives or routines.

When the **init** clause appears on a **depobj** construct, it specifies that *init-var* is initialized to refer to a depend object. The effect of this initialization is the creation of a depend object that encapsulates 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* should be initialized to refer to an interoperability object. This initialization creates an interoperability object that encapsulates a list of properties associated with any *interop-type* modifier. For any *interop-type* modifier, the properties **type**, **type_name**, **vendor**, **vendor_name** and **device_num** will be available. If the implementation cannot initialize *interop-var*, it is instead set to **omp_interop_none** and no interoperability object is created.

Restrictions

- *init-var* must not be constant.

C / C++

- If the **init** clause appears on a **depobj** construct, *init-var* must be an lvalue expression of **depend** OpenMP type that is *uninitialized*.

C / C++

Fortran

- If the **init** clause appears on a **depobj** construct, *init-var* must be a variable of **depend** OpenMP type that is *uninitialized*.

Fortran

- 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 be an *uninitialized* 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* modifier keyword has the unique property. Otherwise, *interop-type* modifiers 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

5.7.1 *interop-type* Modifiers

Modifier set

Properties: *default*	**Members:** *target-type*, *targetsync-type*

Modifiers

Name	Modifies	Type	Properties
target-type	*init-var*	Keyword: **`target`**	*default*
targetsync-type	*init-var*	Keyword: **`targetsync`**	*default*

Clauses

`init`

Semantics

The *interop-type* modifier group includes modifiers that each define a set of properties available for supporting interoperability. For any such modifier, the properties **`type`**, **`type_name`**, **`vendor`**, **`vendor_name`** and **`device_num`** will be available.

The *targetsync-type* modifier 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-type* modifier 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.

Cross References

- **`init`** Clause, see Section 5.7

5.8 `destroy` Clause

Name: **`destroy`**	**Properties:** *default*

Arguments

Name	Type	Properties
destroy-var	variable of OpenMP variable type	*default*

Modifiers

Name	Modifies	Type	Properties
directive-name-modifier	*all arguments*	Keyword: *directive-name* (a directive name)	unique

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

The **destroy** clause takes the *destroy-var* argument, an *initialized* handle of OpenMP type, releases any resources for the underyling object that was referenced by the handle and leaves it in an *uninitialized* state. The object to which *destroy-var* refers is unusable after destruction and the effect of using values associated with it is unspecified until it is initialized again

When the **destroy** clause appears on an **interop** construct, the interop-type set is inferred based on the interop-type set 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.

C / C++

- If the **destroy** clause appears on a **depobj** construct, *destroy-var* must be an lvalue expression of **depend** OpenMP type that is *initialized*.

C / C++

Fortran

- If the **destroy** clause appears on a **depobj** construct, *destroy-var* must be a variable of **depend** OpenMP type that is *initialized*.

Fortran

- If the **destroy** clause appears on an **interop** construct, *destroy-var* must be a variable of **interop** OpenMP type that is *initialized*.

Cross References

- **depobj** Construct, see Section 23.9.3
- **interop** Construct, see Section 22.1

5.9 *variable-selector* Modifiers

Modifier set

Properties: required	**Members:** *optional-modifier*, *variable-category*

Modifiers

Name	Modifies	Type	Properties
variable-category	*default-behavior*	Keyword: **aggregate**, **all**, **allocatable**, **pointer**, **scalar**	*default*
optional-modifier	*default-behavior*	Keyword: **optional**	*default*

Clauses

default, **defaultmap**

Semantics

The *variable-selector* modifier group defines the modifiers that collectively select the variables that are affected by a default-attribute clause.

The *variable-category* modifier specifies the category of variables for which the behavior given by the clause argument may be set. 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. The **aggregate** *variable-category* specifies non-pointer aggregate variables. The **pointer** *variable-category* specifies variables of pointer type.

C / C++

Fortran

The **scalar** *variable-category* specifies non-pointer, non-allocatable scalar variables. The **aggregate** *variable-category* specifies non-pointer, non-allocatable aggregate variables. The **allocatable** *variable-category* specifies variables with the **ALLOCATABLE** attribute. The **pointer** *variable-category* specifies variables with the **POINTER** attribute.

The *optional-modifier* specifies variables with the **OPTIONAL** attribute.

Fortran

If *variable-category* is specified as **all** then the behavior is as if for each other *variable-category* keyword the clause instead appears on the construct with that keyword in place of **all**.

The *variable-selector* modifiers define a default-attribute order among the default-attribute clauses that appear on a given construct. All **defaultmap** clauses precede all **default** clauses in default-attribute order for a **target** construct. Otherwise, clause *A* precedes clause *B* if the set of *variable-selector* modifiers as specified in *B* is a strict subset of the set of *variable-selector* modifiers as specified in *A*. Otherwise, clause *A* precedes clause *B* if the set of *variable-selector* modifiers are identical except clause *B*, but not clause *A*, is derived from a clause that specified the **all** *variable-category* as described above.

If a variable matches the *variable-selector* modifiers of multiple default-attribute clauses, it is only affected by a clause that is not preceded by any of the others in default-attribute order. If more than one such clause is specified, they will necessarily specify the same *default-behavior* per the semantics and restrictions defined in this section.

Restrictions

Restrictions to the *variable-selector* modifier group are as follows:

- A variable must not match the *variable-selector* modifiers of multiple default-attribute clauses if the clauses specify more than one *default-behavior* argument and one of the clauses does not precede all others in default-attribute order.

C / C++

- The specified *variable-category* must not be **`allocatable`**.
- The *optional-modifier* must not be specified.

C / C++

Cross References

- **`default`** Clause, see Section 7.4
- **`defaultmap`** Clause, see Section 11.6

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

An assignable OpenMP type instance is a variable or a function reference with data pointer result of that OpenMP type.

Fortran

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

C / C++

An expression of logical OpenMP type is of Boolean type, possibly after implicit conversion to a Boolean value. 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**.

C / C++

Fortran

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**.

Fortran

Cross References

- Intrinsic Identifiers, see Section 5.4
- OpenMP Foreign Runtime Identifiers, see Section 22.1.1
- OpenMP Reduction and Induction Identifiers, see Section 8.1
- Predefined Identifiers, see Section 26.1
- 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, given by *function-dispatch*.

C / C++

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

```
lvalue-expression = function-dispatch ( [expression-list] );
```

or

```
function-dispatch ( [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 = function-dispatch ( [arguments] )
```

or

```
CALL function-dispatch [ ( [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 *function-dispatch* expression can only be a direct call.

C++

Fortran

- *function-dispatch* must be a procedure name.
- *function-dispatch* must not be a procedure pointer.

Fortran

Cross References

- **dispatch** Construct, see Section 15.6

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

The right-hand side of the assignments of a read structured block or write structured block may have implicit or explicit type conversions as required or permitted by the base language.

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 operator expr* must be, depending on its type, either mathematically or logically equivalent to *x operator (expr)*. This requirement is satisfied if the operators in *expr* have precedence greater than *operator*, or by using parentheses around *expr* or subexpressions of *expr*.
- The expression *expr operator x* must be, depending on its type, either mathematically or logically equivalent to *(expr) operator x*. This requirement is satisfied if the operators in *expr* have precedence equal to or greater than *operator*, or by using parentheses around *expr* or subexpressions of *expr*.
- The expression *x equalop e* must be, depending on its type, either mathematically or logically equivalent to *x equalop (e)*. This requirement is satisfied if the operators in *e* have precedence equal to or greater than *equalop*, or by using parentheses around *e* or subexpressions of *e*.

- *intrinsic-procedure-name* must refer to the intrinsic procedure name and not to other program entities.
- *operator* must refer to the intrinsic operator and not to a user-defined operator.
- 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 and it satisfies the additional requirements for canonical loop nests that follow the grammar.

loop-nest A structured block that is one of the following:

syntactic-loop-nest

or

loop-nest-generating-construct

or

C / C++

```
{
  loop-nest
}
```

C / C++

or

Fortran

```
BLOCK
  loop-nest
END BLOCK
```

Fortran

syntactic-loop-nest One of the following:

C / C++

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

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.

Fortran

loop-body One of the following:

```
loop-nest
```

or

C / C++

```
{
  [intervening-code]
  loop-body
  [intervening-code]
}
```

C / C++

or

Fortran

```
BLOCK
  [block-specification-part]
  [intervening-code]
  loop-body
  [intervening-code]
END BLOCK
```

Fortran

or if none of the previous productions match

terminal-loop-body

loop-nest-generating-construct

A loop-transforming construct for which the associated structured block matches *loop-nest* or *loop-sequence-block* and that generates either a semantic loop nest or a semantic loop sequence that contains exactly one semantic loop nest.

intervening-code

C / C++

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

C / C++

Fortran

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

- loops;
- **`CYCLE`** statements;
- **`EXIT`** statements;

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

Fortran

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

terminal-loop-body

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

C / C++

init-expr One of the following:

var = *lb*

integer-type var = *lb*

C

pointer-type var = *lb*

C

C++

random-access-iterator-type var = *lb*

C++

test-expr One of the following:

var relational-op ub

ub relational-op var

relational-op One of the following:

<

<=

>

>=

!=

incr-expr One of the following:

++*var*

var++

- - *var*

var - -

var += *incr*

var - = *incr*

var = *var* + *incr*

var = *incr* + *var*

var = *var* - *incr*

The value of *incr*, respectively 1 and -1 for the increment and decrement operators, is the increment of the loop.

C / C++

var One of the following:

C / C++

A variable of a signed or unsigned integer type.

C / C++

C

A variable of a pointer type.

C

C++

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.

a1, *a2*, *incr* — Integer expressions that are loop invariant with respect to the outermost syntactic loop of the canonical loop nest.

If the loop is affected by a directive, the expressions are evaluated before the construct formed from that directive.

var-outer — The loop-iteration variable of a surrounding syntactic loop in the canonical loop nest.

C++

range-decl — A declaration of a variable or a declaration specifier sequence as defined by the base language for range-based **`for`** loops.

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

C++

Additionally, a canonical loop nest satisfies the following requirements:

C / C++

- *incr-expr* must cause *var* to increase on each iteration of the loop when *test-expr* is of the form *var relational-op b* and *relational-op* is < or <=. *incr-expr* must cause *var* to decrease on each iteration of the loop when *test-expr* is of the form *var relational-op b* and *relational-op* is > or >=. Increase and decrease are using the order induced by *relational-op*.
- *incr-expr* must cause *var* to decrease on each iteration of the loop when *test-expr* is of the form *var relational-op b* and *relational-op* is < or <=. *incr-expr* must cause *var* to increase on each iteration of the loop when *test-expr* is of the form *var relational-op b* and *relational-op* is > or >=. Increase and decrease are using the order induced by *relational-op*.
- *incr-expr* must increase *var* by one each iteration or must decrease *var* by one each iteration and the increment must be constant when *relational-op* is **!=**.
- *terminal-loop-body* must not contain any **break** statement that would cause the termination of the innermost syntactic loop.

C / C++

Fortran

- *terminal-loop-body* must not contain any **EXIT** statement that would cause the termination of the innermost syntactic loop.
- The directive of a loop-nest-associated construct must precede the outermost **DO** statement that shares the termination of a nonblock **DO** construct when its paired end directive immediately follows the nonblock **DO** construct as defined by the base language.

Fortran

- The loop-iteration variable must not appear in a **threadprivate** directive.

The semantic loop nest for a given canonical loop nest is derived as follows. Each syntactic loop that is not affected by a loop-transforming construct that matches *loop-nest-generating-construct* and each output loop of any such loop-transforming construct is considered for inclusion as a semantic loop of the semantic loop nest. The semantic loop nest is then the maximal number of semantic loops that satisfy the conditions below. The number of semantic loops in the semantic loop nest is the loop nest depth.

C / C++

- The final loop body must not contain any **break** statement that would cause the termination of the innermost semantic loop.

C / C++

Fortran

- The final loop body must not contain any **EXIT** statement that would cause the termination of the innermost semantic loop.

Fortran

- Each semantic loop must enclose or be nested in all other semantic loops in the semantic loop nest;
- A semantic loop must not be an output loop of a loop-transforming construct for which a clause argument refers to a loop-iteration variable of any enclosing syntactic loop in the semantic loop nest;
- The iteration count of a non-rectangular semantic loop must be loop invariant with respect to all enclosing semantic loops in the semantic loop nest except one; and
- The difference in the iteration count of a non-rectangular semantic loop between two logical iterations of an enclosing semantic loop must be loop invariant with respect to the outermost semantic loop.

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 *loop-sequence-block* in the following grammar:

loop-sequence-block

One of the following:

C / C++

```
{
  loop-sequence
}
```

C / C++

Fortran

```
loop-sequence
```

or

```
BLOCK
  loop-sequence
END BLOCK
```

Fortran

or

```
loop-sequence-generating-construct
```

loop-sequence A structured block sequence with executable statements that match *loop-sequence-block* or *loop-nest* (a canonical loop nest as defined in Section 6.4.1). The loops are bounds-independent loops with respect to *loop-sequence-block*.

loop-sequence-generating-construct
A loop-sequence-transforming construct. The output loops are bounds-independent loops with respect to *loop-sequence-block*.

The semantic loop sequence consists of the sequence of adjacent outermost semantic loop nests in *loop-sequence-block* and the output loops from *loop-sequence-generating-construct*. The number of semantic loop nests in that sequence is the loop sequence length.

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 semantic loops of an associated loop nest, called the affected loops, in accordance with its specified clauses. 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 / C++

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

C / C++

C

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

C

C++

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

- For range-based **for** loops, the loop iteration count is computed from *range-expr* using the type **std::iterator_traits<***random-access-iterator-type***>::difference_type** where *random-access-iterator-type* is the iterator type derived from *range-expr*.

C++

Fortran

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

Fortran

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

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

In general, an iteration of the innermost affected loop, in the context of the iterations of its enclosing affected loops, can be identified by a logical iteration vector, and all such logical iteration vectors form a logical iteration vector space. In particular, let the affected loops of a construct be n, where all affected loops have a loop-iteration variable. For such an affected loop nest, 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:

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

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 *[± 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, the declaration specifier sequence 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 output 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
n	expression of integer type	constant, positive

Modifiers

Name	Modifies	Type	Properties
directive-name-modifier	*all arguments*	Keyword: *directive-name* (a directive name)	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 18.5
- **`taskloop`** Construct, see Section 20.2

6.4.6 `ordered` Clause

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

Arguments

Name	Type	Properties
n	expression of integer type	optional, constant, positive

Modifiers

Name	Modifies	Type	Properties
directive-name-modifier	*all arguments*	Keyword: *directive-name* (a directive name)	unique

Directives

`do`, **`for`**

Semantics

The **ordered** clause is used to specify the ordered-affected loops for the purpose of identifying cross-iteration dependences. The argument *n* specifies the number of ordered-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
depth-expr	expression of integer type	positive, constant

Modifiers

Name	Modifies	Type	Properties
directive-name-modifier	*all arguments*	Keyword: *directive-name* (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 17.4
- **`fuse`** Construct, see Section 17.5

6.4.8 `looprange` Clause

Name: `looprange`	**Properties:** unique

Arguments

Name	Type	Properties
first	expression of integer type	constant, positive
count	expression of integer type	constant, positive, ultimate

Modifiers

Name	Modifies	Type	Properties
directive-name-modifier	*all arguments*	Keyword: *directive-name* (a directive name)	unique

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 length 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 17.5
- Canonical Loop Sequence Form, see Section 6.4.2

Part II

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 list argument. 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 with static storage duration that are declared in a scope inside the construct are shared.
- 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 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, it is firstprivate unless:
  - It is a list item in a **map** clause on the construct; or

C++

  - It is an implicit member access.

C++

- 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

- Named constants are device-local.
- Assumed-size arrays are shared in constructs that are not data-mapping 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 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 and do not have data-localizing attributes. In a task-generating construct, the data-sharing attributes of these variables are determined by the **default** clause, if specified without the **implicit** *default-behavior* (see Section 7.4). Otherwise, rules for variables with implicitly determined data-sharing attributes are as follows:

- In a **parallel** construct these variables are shared.
- On constructs that are not explicit-task-generating constructs, these variables reference the variables with the same names that exist in the enclosing context. If 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 explicit-task-generating construct, formal arguments passed by reference are firstprivate.

C++

Fortran

- In an orphaned explicit-task-generating construct, dummy arguments are firstprivate.

Fortran

- In an explicit-task-generating construct, 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 an explicit-task-generating construct, 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 an explicit-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 7.3.3.

Cross References

- **default** Clause, see Section 7.4

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 with static storage duration are shared.
- Objects with dynamic storage duration are shared.

Fortran

- Named constants are device-local.
- 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 list argument may be privatized for the construct on which they appear. Each task that references a privatized list item in any statement in the construct receives at least one new list item if the construct is a loop-collapsing construct, and otherwise each such task receives one new list item. Each SIMD lane used in a **simd** construct that references a privatized list item in any statement in the construct receives at least one new list item. Language-specific attributes for new list items are derived from the corresponding original list items. Inside the construct, all references to the original list items are replaced by references to the new list items received by the task or SIMD lane, 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++

Each 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 each 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.

Each 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.

Each 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.

Each 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

Restrictions

Restrictions to data-sharing clauses are as follows:

- Each list item that has a base variable with a data-localizing attribute must not appear in a data-sharing clause except where explicitly permitted by the clause.

7.3.1 `shared` Clause

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

Arguments

Name	Type	Properties
list	list of variable list item type	*default*

Modifiers

Name	Modifies	Type	Properties
directive-name-modifier	*all arguments*	Keyword: *directive-name* (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, named constants, 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

7.3.2 `private` Clause

Name: `private`	**Properties:** data-environment attribute, data-sharing attribute, innermost-leaf, privatization

Arguments

Name	Type	Properties
list	list of variable list item type	*default*

Modifiers

Name	Modifies	Type	Properties
directive-name-modifier	*all arguments*	Keyword: *directive-name* (a directive name)	unique

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 18.5
- **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

7.3.3 `firstprivate` Clause

Name: **firstprivate**	**Properties:** data-environment attribute, data-sharing attribute, privatization

Arguments

Name	Type	Properties
list	list of variable list item type	*default*

Modifiers

Name	Modifies	Type	Properties
saved	*list*	Keyword: **saved**	*default*
directive-name-modifier	*all arguments*	Keyword: *directive-name* (a directive name)	unique

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 7.3.2, 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.

C / C++

For variables of non-array type, the initialization occurs by copy assignment. For an array of elements of non-array type, each element is initialized as if by assignment from an element of the original array to the corresponding element of the new array.

C / C++

C++

For each variable of class type:

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

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

C++

Fortran

If the **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:

- Each 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.
- Each 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 7.3.2
- **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

7.3.4 `lastprivate` Clause

Name: `lastprivate`	**Properties:** data-environment attribute, data-sharing attribute, original list-item updating, privatization

Arguments

Name	Type	Properties
list	list of variable list item type	*default*

Modifiers

Name	Modifies	Type	Properties
lastprivate-modifier	*list*	Keyword: **conditional**	*default*
directive-name-modifier	*all arguments*	Keyword: *directive-name* (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 7.3.2. In addition, each new list item has the lastprivate attribute.

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, then:

- If the clause is specified on a loop-nest-associated directive and the number of logical iterations is zero, the original list item will have an unspecified value after the construct;
- Otherwise, if the new list item or its subcomponent from the sequentially last iteration of affected loop, or the lexically last structured block sequence associated with a **sections** construct, is assigned a value in that iteration or structured block sequence, then that value is assigned to the original list item or its subcomponent;
- Otherwise, the original list item or its subcomponent will have an unspecified value after the construct.

If the **conditional** modifier appears on the clause and an assignment to a list item is encountered in the construct then the original list item is assigned the value that is assigned to the new list item in the sequentially last iteration or lexically last structured block sequence in which such an assignment is encountered. However, if the list item is modified in the region by an assignment outside the construct or by an assignment that does not lexically assign to the list item then the value assigned to the original list item is unspecified.

If the list item is a loop-iteration variable of one of the affected loops then 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. If such an execution of the canonical loop nest would not assign a value to the list item and the **conditional** modifier is not present then the original list item will have an unspecified value after the construct.

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 **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 races, concurrent reads or updates of the original list item must be synchronized with any update of the original list item that occurs as a result of the **lastprivate** clause.

Restrictions

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

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

C++

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

C++

Fortran

- A variable that appears in a **lastprivate** clause must be definable.
- If the original list item has the **ALLOCATABLE** attribute, the corresponding list item of which the value is assigned to the original list item must have an allocation status of allocated upon exit from the sequentially last iteration of the affected loops or lexically last structured block sequence associated with a **sections** construct.
- If the list item is a polymorphic variable with the **ALLOCATABLE** attribute, the behavior is 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 7.3.2
- **sections** Construct, see Section 19.3
- **simd** Construct, see Section 18.5
- **taskloop** Construct, see Section 20.2

7.3.5 is_device_ptr Clause

Name: is_device_ptr	Properties: data-environment attribute, data-sharing attribute, device-associated, dispatch-requirement, innermost-leaf, privatization

Arguments

Name	Type	Properties
list	list of variable list item type	*default*

Modifiers

Name	Modifies	Type	Properties
directive-name-modifier	*all arguments*	Keyword: *directive-name* (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.

Restrictions

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

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

Cross References

- **dispatch** Construct, see Section 15.6
- **has_device_addr** Clause, see Section 7.3.7
- **target** Construct, see Section 21.8

7.3.6 use_device_ptr Clause

Name: use_device_ptr	Properties: all-data-environments, data-environment attribute, data-sharing attribute, device-associated, privatization

Arguments

Name	Type	Properties
list	list of variable list item type	*default*

Modifiers

Name	Modifies	Type	Properties
fallback	*list*	Keyword: **fb_nullify**, **fb_preserve**	*default*
directive-name-modifier	*all arguments*	Keyword: *directive-name* (a directive name)	unique

Directives

target_data

Semantics

Each list item in the **use_device_ptr** clause results in a new list item that has the use-device-ptr attribute and is a device pointer that refers to a device address. Since the **use_device_ptr** clause is an all-data-environments clause, it has this effect even for minimal data environments. The device address is determined as follows. Each list item is treated as if a zero-offset assumed-size array at the storage location to which the list item points is mapped by a **map** clause on the construct with a *map-type* of **storage**. If a matched candidate is found for the assumed-size array (see Section 11.5), the new list item refers to the device address that is the base address of the array section that corresponds to the assumed-size array in the device data environment. Otherwise, the behavior is determined by the *fallback* modifier:

- If the *fallback* modifier is **fb_preserve**, the new list item refers to the address stored in the original list item.
- If the *fallback* modifier is **fb_nullify**, the new list item is set to NULL.

If the *fallback* modifier is not specified, the behavior is as if the **fb_preserve** *fallback* modifier is present. All references to the list item inside the structured block associated with the construct are replaced with the new list item that is a private copy in the associated data environment on the encountering 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

7.3.7 has_device_addr Clause

Name: **has_device_addr**	**Properties:** data-environment attribute, data-sharing attribute, device-associated, dispatch-requirement, outermost-leaf

Arguments

Name	Type	Properties
list	list of variable list item type	*default*

Modifiers

Name	Modifies	Type	Properties
directive-name-modifier	*all arguments*	Keyword: *directive-name* (a directive name)	unique

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 each 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 each 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

Restrictions

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

- If a list item is an array section or an array element, the array base must be a base language identifier.

C / C++

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

C / C++

Fortran

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

Fortran

Cross References

- **dispatch** Construct, see Section 15.6
- **target** Construct, see Section 21.8

7.3.8 use_device_addr Clause

Name: use_device_addr	**Properties:** all-data-environments, data-environment attribute, data-sharing attribute, device-associated

Arguments

Name	Type	Properties
list	list of variable list item type	*default*

Modifiers

Name	Modifies	Type	Properties
directive-name-modifier	*all arguments*	Keyword: *directive-name* (a directive name)	unique

Directives

target_data

Semantics

Each list item in a **use_device_addr** clause has the use-device-addr attribute inside the construct. If the list item is present in the device data environment on entry to the construct, the list item is treated as if it is implicitly mapped by a **map** clause on the construct with a *map-type* of **storage** and all references to the list item inside the structured block associated with the

construct are to the corresponding list item in the device data environment. The list items in a **use_device_addr** clause may include array sections and assumed-size arrays. Since the **use_device_addr** clause is an all-data-environments clause, it has this effect even for minimal data environments.

If the list item is a referencing variable, the semantics of the **use_device_addr** clause apply to its referenced pointee. A private copy of the referring pointer that refers to the corresponding referenced pointee is used in place of the original referring pointer in the structured block.

C / C++

If a list item is an array section that has a base pointer, all references to the base pointer inside the structured block are replaced with a new pointer that contains the base address of the corresponding list item. This conversion may be elided if no corresponding list item is present.

C / C++

Restrictions

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

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

Cross References

- **target_data** Construct, see Section 21.7

7.4 default Clause

Name: default	**Properties:** default-attribute, post-modified

Arguments

Name	Type	Properties
default-behavior	Keyword: **firstprivate**, **implicit**, **none**, **private**, **shared**	*default*

Modifier groups

variable-selector

Modifiers

Name	Modifies	Type	Properties
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.

If *default-behavior* is **implicit** then the data-sharing attribute is implicitly determined by the rules given in Section 7.1.1.

If *default-behavior* is **none**, the data-sharing attribute is not implicitly determined. If *default-behavior* is **shared** then the clause has no effect on a **target** construct.

Otherwise, the data-sharing attributes of the selected variables will be the data-sharing attribute that corresponds to *default-behavior*.

Restrictions

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

- If *default-behavior* 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 *default-behavior* 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 11.6
- Variables Referenced in a Construct, see Section 7.1.1
- **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

7.5 `saved` Modifier

Modifiers

Name	Modifies	Type	Properties
saved	*list*	Keyword: **`saved`**	*default*

Clauses

`firstprivate`

Semantics

If the *saved* modifier is present in a data-sharing attribute clause that is specified on a replayable construct then its original list items of a replay execution are defined by the saved data environment of the replayable construct. The *saved* modifier has no effect if specified in a clause that does not appear on a replayable construct.

Cross References

- **`firstprivate`** Clause, see Section 7.3.3

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**. If the intrinsic procedure is renamed to a local name in a **USE** statement, the reduction identifier is as if that intrinsic procedure is referenced and the implicit initializer and combiner for the intrinsic procedure are applied.

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`
`&&`	`omp_priv = 1`	`omp_out = omp_in && omp_out`
`||`	`omp_priv = 0`	`omp_out = omp_in || omp_out`

table continued on next page

table continued from previous page

Identifier	Initializer	Combiner
`max`	`omp_priv =` *Minimal representable number in the reduction list item type*	`omp_out = omp_in > omp_out ? omp_in : omp_out`
`min`	`omp_priv =` *Maximal representable number in the reduction list item type*	`omp_out = omp_in < omp_out ? omp_in : omp_out`

C / C++

Fortran

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
`+`	`omp_priv = 0`	`omp_out = omp_in + omp_out`
`*`	`omp_priv = 1`	`omp_out = omp_in * omp_out`
`.and.`	`omp_priv = .true.`	`omp_out = omp_in .and. omp_out`
`.or.`	`omp_priv = .false.`	`omp_out = omp_in .or. omp_out`
`.eqv.`	`omp_priv = .true.`	`omp_out = omp_in .eqv. omp_out`
`.neqv.`	`omp_priv = .false.`	`omp_out = omp_in .neqv. omp_out`
`max`	`omp_priv =` *Minimal representable number in the reduction list item type*	`omp_out = max(omp_in, omp_out)`
`min`	`omp_priv =` *Maximal representable number in the reduction list item type*	`omp_out = min(omp_in, omp_out)`
`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 each 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.
- Each 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 each 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 each list item and of the step expression that appear in an **induction** clause must be valid for the inductor expression.

- Each 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.

A reduction clause may include list items for which the base variable has a data-localizing attribute unless:

- The clause appears on a **target** construct; or

- The base variable is groupprivate or threadprivate and the clause appears on a **teams** construct or a work-distribution construct for which the corresponding region binds to a **teams** region; or
- The base variable is threadprivate and the clause appears on a **parallel** construct or a worksharing construct.

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.

Execution Model Events

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

Tool Callbacks

A thread dispatches a registered **reduction** callback with **ompt_sync_region_reduction** in its *kind* argument and **ompt_scope_begin** as its *endpoint* argument for each occurrence of a *reduction-begin* event in that thread. Similarly, a thread dispatches a registered **reduction** callback with **ompt_sync_region_reduction** in its *kind* argument and **ompt_scope_end** as its *endpoint* argument for each occurrence of a *reduction-end* event in that thread. These callbacks occur in the context of the task that performs the reduction.

Restrictions

Restrictions common to reduction clauses are as follows:

C

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

C

C++

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

C++

Cross References

- **reduction** Callback, see Section 41.7.6
- OMPT **scope_endpoint** Type, see Section 40.27
- OMPT **sync_region** Type, see Section 40.33

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: `reduction`	**Properties:** data-environment attribute, data-sharing attribute, original list-item updating, privatization, reduction scoping, reduction participating

Arguments

Name	Type	Properties
list	list of variable list item type	*default*

Modifier groups

reduction-type

Modifiers

Name	Modifies	Type	Properties
reduction-identifier	*all arguments*	An OpenMP reduction identifier	required, ultimate
original-sharing-modifier	*list*	Complex, name: **`original`** Arguments: ***sharing*** Keyword: **`default`**, **`private`**, **`shared`** (*default*)	*default*
per-thread-modifier	*list*	Keyword: **`per_thread`**	*default*
directive-name-modifier	*all arguments*	Keyword: *directive-name* (a directive name)	unique

Modifier set

Properties: exclusive	**Members:** *inscan-modifier*, *noncommutative-modifier*, *per-thread-modifier*

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. The manner in which the reduction is performed for each list item is determined according to the *reduction-type* modifier that is specified. If a *reduction-type* modifier is not present, the behavior is as if the *default-modifier* is specified in the clause.

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 without the

per-thread-modifier 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`** and the clause appears on a work-distribution construct, original list items are assumed to be private in the enclosing context if the *per-thread-modifier* is present. Otherwise, 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 the *per-thread-modifier* is present, the behavior is the same as described for the specified *reduction-type*, except the original list item for a given encountering task is updated by combining its original value with only the final values of the private copies created for that task using the combiner of the specified *reduction-identifier*.

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.
- The *per-thread-modifier* must not be present for a work-distribution construct if the *original-sharing-modifier* is specified with the **`shared`** *sharing* argument.
- 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 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.

- Each 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.

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 it *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.

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 7.3.2
- **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 18.5
- **taskloop** Construct, see Section 20.2
- **teams** Construct, see Section 18.2

8.10.1 *reduction-type* Modifiers

Modifier set

Properties: exclusive	**Members:** *default-modifier*, *inscan-modifier*, *noncommutative-modifier*, *task-modifier*

Modifiers

Name	Modifies	Type	Properties
inscan-modifier	*list*	Keyword: **`inscan`**	*default*
noncommutative-modifier	*list*	Keyword: **`noncommutative`**	*default*
task-modifier	*list*	Keyword: **`task`**	*default*
default-modifier	*list*	Keyword: **`default`**	*default*

Clauses

`reduction`

Semantics

If the *default-modifier* or *noncommutative-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.

For commutative reductions, at the end of a region that corresponds to a construct for which the **`reduction`** clause was specified, if the *default-modifier* is present, 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 *noncommutative-modifier* is present, the reduction is noncommutative. For noncommutative reductions, the partial results in the private copies conceptually form an ordered list called the noncommutative partial results list, which is formed as if the partial results are ordered according to the canonical order of the construct. The combination of these partial results is executed as if two adjacent elements are removed from the list, combined using the combiner of the specified *reduction-identifier* such that the first element corresponds to **`omp_out`**, and the combined result is inserted back into the list at the same position as the removed elements. The removal, combining, and insertion of the result occur as if performed as a single atomic operation. The operation is repeated until only one element remains in the list, which is combined into the original list item.

If the *inscan-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-modifier* is present for a **parallel** or worksharing construct, 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-type* modifiers are as follows:

- Each list item specified with the *inscan-modifier* must appear as a list item in an **inclusive** or **exclusive** clause on a **scan** directive enclosed by the construct.
- If the *inscan-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 *inscan-modifier* is specified, a **reduction** clause without the *inscan-modifier* must not appear on the same construct.
- A **reduction** clause with the *inscan-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-modifier* must not be specified on a construct for which the **ordered** or **schedule** clause is specified.
- The *task-modifier* must not appear in a **reduction** clause if the **nowait** clause is specified on the same construct.
- A **reduction** clause with the *task-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.
- The *noncommutative-modifier* may only be specified on an orderable construct, which is any construct that has the orderable property.
- The *noncommutative-modifier* must not be specified on a task-generating construct.
- If a list item that appears in a **reduction** clause with the *noncommutative-modifier* is an array then the size of the array must be constant.
- If a list item that appears in a **reduction** clause with the *noncommutative-modifier* is an array section then its *lower-bound*, *length*, and *stride* expressions must be constant.

C / C++

- A variable of class type (or array thereof) that appears in a **reduction** clause with the *inscan-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

- **reduction** Clause, see Section 8.10

8.11 task_reduction Clause

Name: task_reduction	**Properties:** data-environment attribute, data-sharing attribute, original list-item updating, privatization, reduction scoping

Arguments

Name	Type	Properties
list	list of variable list item type	*default*

Modifiers

Name	Modifies	Type	Properties
reduction-identifier	*all arguments*	An OpenMP reduction identifier	required, ultimate
directive-name-modifier	*all arguments*	Keyword: *directive-name* (a directive name)	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: `in_reduction`	**Properties:** data-environment attribute, data-sharing attribute, privatization, reduction participating

Arguments

Name	Type	Properties
list	list of variable list item type	*default*

Modifiers

Name	Modifies	Type	Properties
reduction-identifier	*all arguments*	An OpenMP reduction identifier	required, ultimate
directive-name-modifier	*all arguments*	Keyword: *directive-name* (a directive name)	unique

Directives

`target`, **`target_data`**, **`task`**, **`taskloop`**

Semantics

The **`in_reduction`** clause is a reduction-participating clause, as described in Section 8.8, that specifies that a task participates in a reduction. For a given list item, the **`in_reduction`** clause defines a task to be a participant in a task reduction that is defined by an enclosing region for a matching list item that appears in a **`task_reduction`** clause or a **`reduction`** clause with the **`task`** *reduction-modifier*, where either:

1. The matching list item has the same storage location as the list item in the **`in_reduction`** clause; or
2. A private copy, derived from the matching list item, that is used to perform the task reduction has the same storage location as the list item in the **`in_reduction`** clause.

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

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

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

When specified on the **`target_data`** directive, the **`in_reduction`** clause has the all-data-environments property.

Restrictions

Restrictions to the **in_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 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: induction	**Properties:** data-environment attribute, data-sharing attribute, original list-item updating, privatization

Arguments

Name	Type	Properties
list	list of variable list item type	*default*

Modifiers

Name	Modifies	Type	Properties
induction-identifier	*list*	OpenMP induction identifier	required, ultimate
step-modifier	*list*	Complex, name: **step** Arguments: ***induction-step*** expression of induction-step type (region-invariant)	required
induction-modifier	*list*	Keyword: **relaxed**, **strict**	*default*
directive-name-modifier	*all arguments*	Keyword: *directive-name* (a directive name)	unique

Directives

distribute, **do**, **for**, **simd**, **taskloop**

Semantics

The **induction** clause provides a superset of the functionality provided by the **private** clause. Each list item that appears in an **induction** clause is subject to the **private** clause semantics described in Section 7.3.2, 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.

Restrictions

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

- All restrictions listed in Section 8.5 apply to this clause.
- The *induction-step* must not be an array or array section.
- If an array section or array element appears as a list item in an **induction** clause on a worksharing construct, all threads of the team must specify the same storage location.
- None of the affected loops of a loop-nest-associated construct that has an **induction** clause may be a non-rectangular loop.

C / C++

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

C / C++

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
- **private** Clause, see Section 7.3.2
- **simd** Construct, see Section 18.5
- **taskloop** Construct, see Section 20.2

8.14 linear Clause

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

Arguments

Name	Type	Properties
list	list of variable list item type	*default*

Modifiers

Name	Modifies	Type	Properties
step-simple-modifier	*list*	expression of integer type	exclusive, region-invariant, unique
step-complex-modifier	*list*	Complex, name: **`step`** Arguments: ***linear-step*** expression of integer type (region-invariant)	unique
linear-modifier	*list*	Keyword: **`ref`**, **`uval`**, **`val`**	unique
directive-name-modifier	*all arguments*	Keyword: *directive-name* (a directive name)	unique

Directives

`declare_simd`, **`do`**, **`for`**, **`simd`**

Semantics

The **`linear`** clause provides a superset of the functionality provided by the **`private`** clause. A list item that appears in a **`linear`** clause is subject to the **`private`** clause semantics described in Section 7.3.2, 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 not 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 15.7
- **do** Construct, see Section 19.6.2
- **for** Construct, see Section 19.6.1
- **private** Clause, see Section 7.3.2
- **simd** Construct, see Section 18.5
- **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
reduction-specifier	OpenMP reduction specifier	*default*

Clauses

combiner, **initializer**, **reduction_kind**

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
- **reduction_kind** Clause, see Section 8.15.3

8.15.1 `combiner` Clause

Name: `combiner`	**Properties:** unique, required

Arguments

Name	Type	Properties
combiner-expr	OpenMP combiner expression	*default*

Modifiers

Name	Modifies	Type	Properties
directive-name-modifier	*all arguments*	Keyword: *directive-name* (a directive name)	unique

Directives

`declare_reduction`

Semantics

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

Cross References

- **`declare_reduction`** Directive, see Section 8.15
- OpenMP Combiner Expressions, see Section 8.2.1

8.15.2 `initializer` Clause

Name: `initializer`	**Properties:** unique

Arguments

Name	Type	Properties
initializer-expr	OpenMP initializer expression	*default*

Modifiers

Name	Modifies	Type	Properties
directive-name-modifier	*all arguments*	Keyword: *directive-name* (a directive name)	unique

Directives

`declare_reduction`

Semantics

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

Cross References

- **`declare_reduction`** Directive, see Section 8.15
- OpenMP Initializer Expressions, see Section 8.2.2

8.15.3 `reduction_kind` Clause

Name: `reduction_kind`	**Properties:** unique

Arguments

Name	Type	Properties
reduction-kind	Keyword: **`noncommutative`**	*default*

Modifiers

Name	Modifies	Type	Properties
directive-name-modifier	*all arguments*	Keyword: *directive-name* (a directive name)	unique

Directives

`declare_reduction`

Semantics

If a **`reduction_kind`** clause that specifies **`noncommutative`** is present on a **`declare_reduction`** directive, then the effect is as if the *noncommutative-modifier* is specified on any **`reduction`** clause that uses the reduction identifier.

Restrictions

Restrictions to the **`reduction_kind`** clause are as follows:

- If a **`reduction_kind`** clause that specifies **`noncommutative`** is present on a **`declare_reduction`** directive, then the reduction identifier can only be used in **`reduction`** clauses in which the *default-modifier* or *noncommutative-modifier* is specified.

Cross References

- **`declare_reduction`** Directive, see Section 8.15
- **`reduction`** Clause, see Section 8.10

8.16 `declare_induction` Directive

Name: `declare_induction` **Category:** declarative	**Association:** unassociated **Properties:** pure

Arguments

`declare_induction``(`*induction-specifier*`)`

Name	Type	Properties
induction-specifier	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.
- Each 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: inductor	**Properties:** unique, required

Arguments

Name	Type	Properties
inductor-expr	OpenMP inductor expression	*default*

Modifiers

Name	Modifies	Type	Properties
directive-name-modifier	*all arguments*	Keyword: *directive-name* (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: `collector`	**Properties:** unique, required

Arguments

Name	Type	Properties
collector-expr	OpenMP collector expression	*default*

Modifiers

Name	Modifies	Type	Properties
directive-name-modifier	*all arguments*	Keyword: *directive-name* (a directive name)	unique

Directives

`declare_induction`

Semantics

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

Cross References

- **`declare_induction`** Directive, see Section 8.16
- OpenMP Collector Expressions, see Section 8.2.4

8.17 `scan` Directive

Name: `scan` **Category:** subsidiary	**Association:** separating **Properties:** pure

Separated directives

do, **for**, **simd**

Clauses

exclusive, **inclusive**, **init_complete**

Clause set

Properties: unique, required, exclusive	**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 ($\texttt{PRESUM}_{\texttt{last}}$) 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 $\texttt{PRESUM}_{\texttt{last}}(op, a_1, \ldots, a_N)$ is defined for a given binary operator *op* and a sequence of N values $a_1, \ldots, a_N$ as follows:

- if $N = 1$, a_1
- if $N > 1$, $op($ $\texttt{PRESUM}_{\texttt{last}}(op, a_1, \ldots, a_j)$, $\texttt{PRESUM}_{\texttt{last}}(op, a_k, \ldots, a_N)$), $1 \leq j + 1 = k \leq N$.

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 each 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 PRESUM$_{\texttt{last}}$(*combiner*, *orig-val*, $u_0, \ldots, u_i$). 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 PRESUM$_{\texttt{last}}$(*combiner*, *orig-val*, $u_0, \ldots, u_{i-1}$).

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 PRESUM$_{\texttt{last}}$(*combiner*, *orig-val*, $u_0, \ldots, u_k$).

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 each 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 18.5

8.17.1 `inclusive` Clause

Name: `inclusive`	**Properties:** innermost-leaf, unique

Arguments

Name	Type	Properties
list	list of variable list item type	*default*

Modifiers

Name	Modifies	Type	Properties
directive-name-modifier	*all arguments*	Keyword: *directive-name* (a directive name)	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 list argument. 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
list	list of variable list item type	*default*

Modifiers

Name	Modifies	Type	Properties
directive-name-modifier	*all arguments*	Keyword: *directive-name* (a directive name)	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 list argument. 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: `init_complete`	**Properties:** innermost-leaf, unique

Arguments

Name	Type	Properties
create_init_phase	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

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 Data-Localizing Control

This chapter presents various directives and clauses for controlling the data-localizing attributes of variables.

C / C++

The following base language attributes for variables imply a data-localizing attribute:

- Variables with the **`_Thread_local`** (in C) or **`thread_local`** (in C/C++) storage-class specifier are threadprivate.

C / C++

9.1 `threadprivate` Directive

Name: **`threadprivate`**	**Association:** explicit
Category: declarative	**Properties:** pure

Arguments

`threadprivate`(*list*)

Name	Type	Properties
list	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.
- None of the list items may have the groupprivate attribute.

C / C++

- Each list item must be a file-scope, namespace-scope, or static block-scope variable.
- No list item may have an incomplete type.
- The address of a threadprivate variable must not be an address constant.
- If the value of a variable referenced in an explicit initializer of a threadprivate variable is modified prior to the first reference to any instance of the threadprivate variable, the behavior is unspecified.
- A **threadprivate** directive for file-scope variables must appear outside any definition or declaration, and must lexically precede all references to any of the variables in its list argument.
- A **threadprivate** directive for namespace-scope variables must appear outside any definition or declaration other than the namespace definition itself and must lexically precede all references to any of the variables in its list argument.
- Each variable in the list argument 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 list argument.
- Each variable in the list argument 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 list argument.
- 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 18.4
- Determining the Number of Threads for a **parallel** Region, see Section 18.1.1

9.2 grouppriv ate Directive

Name: **grouppriv ate** **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.
- None of the list items may have the threadprivate attribute.

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: `local`	**Properties:** data-environment attribute

Arguments

Name	Type	Properties
list	list of variable list item type	*default*

Modifiers

Name	Modifies	Type	Properties
directive-name-modifier	*all arguments*	Keyword: *directive-name* (a directive name)	unique

Directives

`declare_target`

Semantics

The **`local`** clause specifies that each list item has the device-local attribute. Each reference to a list item on a given device will refer to a copy of the list item that is a device-local variable and is in memory associated with the device.

Cross References

- **`declare_target`** Directive, see Section 15.8.1

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: copyin	**Properties:** outermost-leaf, data copying

Arguments

Name	Type	Properties
list	list of variable list item type	*default*

Modifiers

Name	Modifies	Type	Properties
directive-name-modifier	*all arguments*	Keyword: *directive-name* (a directive name)	unique

Directives

parallel

Semantics

The **copyin** clause provides a mechanism to copy the value of a threadprivate variable of the primary thread to the threadprivate variable of each other member of the team that is executing the **parallel** region.

C / C++

The copy is performed after the team is formed and prior to the execution of the associated structured block. For variables of non-array type, the copy is by copy assignment. For an array of elements of non-array type, each element is copied as if by assignment from an element of the array of the primary thread to the corresponding element of the array of all other threads.

C / C++

C++

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

C++

Fortran

The copy is performed, as if by assignment, after the team is formed and prior to the execution of the associated structured block.

Named variables that appear in a threadprivate common block may be specified. The whole common block does not need to be specified.

On entry to any **parallel** region, the copy of each thread of a variable that is affected by a **copyin** clause for the **parallel** region will acquire the type parameters, allocation, association, and definition status of the copy of the primary thread, according to the following rules:

- If the original list item has the **POINTER** attribute, each copy receives the same association status as that of the copy of the primary thread as if by pointer assignment.
- If the original list item does not have the **POINTER** attribute, each copy becomes defined with the value of the copy of the primary thread as if by 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:

- Each 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

Arguments

Name	Type	Properties
list	list of variable list item type	*default*

Modifiers

Name	Modifies	Type	Properties
directive-name-modifier	*all arguments*	Keyword: *directive-name* (a directive name)	unique

Directives

`single`

Semantics

The **`copyprivate`** clause provides a mechanism to use a private variable to broadcast a value from the data environment of one implicit task to the data environments of the other implicit tasks that belong to the innermost enclosing parallel region. The effect of the **`copyprivate`** clause on the specified list items occurs after the execution of the structured block associated with the construct on which the clause is specified, and before any of the threads in the team have left the barrier at the end of the construct. To avoid data races, concurrent reads or updates of the list item must be synchronized with the update of the list item that occurs as a result of the **`copyprivate`** clause if, for example, the **`nowait`** clause is used to remove the barrier.

C / C++

In all other implicit tasks that belong to the parallel region, each specified list item becomes defined with the value of the corresponding list item in the implicit task associated with the thread that executed the structured block. For variables of non-array type, the definition occurs by copy assignment. For an array of elements of non-array type, each element is copied by copy assignment from an element of the array in the data environment of the implicit task that is associated with the thread that executed the structured block to the corresponding element of the array in the data environment of the other implicit tasks.

C / C++

C++

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

C++

Fortran

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. All mapping operations are strict maps unless otherwise specified. 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 or base pointer 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 or base pointer 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 any default-attribute clauses are present for the category of the variable and they specify a default behavior other than **implicit**, the data-mapping attribute or data-sharing attribute is determined by those clauses.

C++

- If a variable is referenced inside a **target** construct via an implicit member access, the effect is 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
mapper	*locator-list*	Complex, name: **`mapper`** Arguments: ***mapper-identifier*** OpenMP identifier (*default*)	unique

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 11.7
- **`from`** Clause, see Section 12.2
- Data-Motion Control, see Chapter 12
- **`map`** Clause, see Section 11.5
- **`to`** Clause, see Section 12.1

11.3 `enter` Clause

Name: `enter`	**Properties:** data-environment attribute, data-mapping attribute

Arguments

Name	Type	Properties
list	list of extended list item type	*default*

Modifiers

Name	Modifies	Type	Properties
automap-modifier	*list*	Keyword: **`automap`**	*default*
directive-name-modifier	*all arguments*	Keyword: *directive-name* (a directive name)	unique

Directives

`declare_target`

Semantics

The **`enter`** clause is a data-mapping attribute clause.

If a procedure name appears in an **`enter`** clause in the same compilation unit in which the definition of the procedure occurs then a device-specific version of the procedure is created for all 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 *automap-modifier* must not be present.

C / C++

Fortran

- If the *automap-modifier* is present, each list item must be an allocatable variable.

Fortran

Cross References

- **declare_target** Directive, see Section 15.8.1

11.4 `link` Clause

Name: `link`	**Properties:** data-environment attribute

Arguments

Name	Type	Properties
list	list of variable list item type	*default*

Modifiers

Name	Modifies	Type	Properties
directive-name-modifier	*all arguments*	Keyword: *directive-name* (a directive name)	unique

Directives

`declare_target`

Semantics

The **`link`** clause supports compilation of device procedures that refer to variables with static storage duration that appear as list items in the clause. The **`declare_target`** directive on which the clause appears does not map the list items. Instead, they are mapped according to the data-mapping rules described in Section 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 15.8.1
- Implicit Data-Mapping Attribute Rules, see Section 11.1

11.5 `map` Clause

Name: `map`	**Properties:** data-environment attribute, data-mapping attribute

Arguments

Name	Type	Properties
locator-list	list of locator list item type	*default*

Modifier groups

map-type-modifying

Modifiers

Name	Modifies	Type	Properties
mapper	*locator-list*	Complex, name: **mapper** Arguments: ***mapper-identifier*** OpenMP identifier (*default*)	unique
iterator	*all arguments*	Complex, name: **iterator** Arguments: ***iterator-specifier-list*** list of iterator specifier list item type (*default*)	unique
ref-modifier	*all arguments*	Keyword: **ref_ptee**, **ref_ptr**, **ref_ptr_ptee**	unique
attach-modifier	*all arguments*	Complex, name: **attach** Arguments: ***attachment-mode*** Keyword: **always**, **auto**, **never** (*default*)	unique
map-type	*all arguments*	Keyword: **from**, **storage**, **to**, **tofrom**	*default*
directive-name-modifier	*all arguments*	Keyword: *directive-name* (a directive name)	unique

Directives

declare_mapper, **target**, **target_data**, **target_enter_data**, **target_exit_data**

Semantics

The **map** clause specifies how an original list item is mapped from the data environment of the current task to a corresponding list item in the device data environment of the device identified by the construct. The list items that appear on a **map** clause may include array sections, assumed-size arrays, and structure elements. Each list item in a **map** clause may reference any *iterator-identifier* defined in its *iterator* modifier. Each 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 11.5.1 describes the various map types determined by the *map-type* modifier in a **map** clause, and Section 11.5.2 describes the process of map-type decay for determining an output map type.

Section 11.5.4 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 11.5.5 describes the effect of the *attach-modifier*, when specified, and otherwise describes when pointer attachment occurs for list items in a **map** clause.

Section 11.5.6 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 11.5.7 describes the effect of the **map** clause when an array, array section or structure is specified as a list item.

Section 11.5.8 describes the mapping of storage blocks based on the **map** clauses and their *map-type* and *map-type-modifying* modifiers on a given data-mapping construct, as well as the creation of any new list items within the corresponding storage blocks.

Section 11.5.9 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 before any array-map expansion or mapper expansion, 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.
- A variable that appears in an **enter** clause on a declare target directive that also specifies a **device_type** clause with the **nohost** *device-type-description* 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++

- Each 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.
- If a list item in a **map** clause is a member of a common block that does not appear in an **enter** clause on a declare target directive in the same compilation unit, that common block must not have the **BIND(C)** attribute.
- If two compilation units map storage from the same common block and the common block does not appear in an **enter** clause on a declare target directive in any compilation unit, the declarations of that common block in those compilation units must have the same number of members at the same positions with the same type, kind type parameter, and rank.

Fortran

Cross References

- **attach-modifier** Modifier, see Section 11.5.5
- **declare_mapper** Directive, see Section 11.7
- Array and Structure Mapping, see Section 11.5.7
- Array Sections, see Section 5.2.5
- Assumed-Size Array Mapping, see Section 11.5.6
- Implicit Data-Mapping, see Section 11.5.9
- Map Type Decay, see Section 11.5.2
- Storage Block Mapping, see Section 11.5.8
- **iterator** Modifier, see Section 5.2.6
- Mapper Identifiers and **mapper** Modifiers, see Section 11.2
- *map-type* Modifier, see Section 11.5.1
- **ref-modifier** Modifier, see Section 11.5.4
- OMPT **scope_endpoint** Type, see Section 40.27
- **target** Construct, see Section 21.8
- **target_data** Construct, see Section 21.7
- **target_data_op_emi** Callback, see Section 42.7

- **target_enter_data** Construct, see Section 21.5
- **target_exit_data** Construct, see Section 21.6
- **target_map_emi** Callback, see Section 42.9
- **target_update** Construct, see Section 21.9

11.5.1 *map-type* Modifier

Modifiers

Name	Modifies	Type	Properties
map-type	*all arguments*	Keyword: **from**, **storage**, **to**, **tofrom**	*default*

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 map type 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 *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 *map-type* is **to**, **tofrom**, or **storage**. The map type is a map-exiting map type if *map-type* 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 *map-type* is not specified for a clause on which it may be specified, it defaults to **storage** if the *delete-modifier* is present on the clause or if the list item for which *map-type* is not specified is an assumed-size array. Otherwise, *map-type* defaults to **tofrom** if a *map-type* modifier is not specified for a clause on which it may be specified, unless otherwise specified.

Fortran

When a *map-type* modifier is not specified for a clause on which it may be specified, it defaults to **storage** if the list item for which *map-type* is not specified is an assumed-type variable.

Fortran

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

Restrictions

Restrictions to the *map-type* modifier are as follows:

- If the clause on which the *map-type* modifier 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* modifier 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 11.5

11.5.2 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.

11.5.3 *map-type-modifying* Modifiers

Modifier set

Properties: *default*	**Members:** *always-modifier*, *close-modifier*, *delete-modifier*, *present-modifier*, *self-modifier*

Modifiers

Name	Modifies	Type	Properties
present-modifier	*locator-list*	Keyword: **`present`**	*default*
always-modifier	*locator-list*	Keyword: **`always`**	*default*
delete-modifier	*locator-list*	Keyword: **`delete`**	*default*
close-modifier	*locator-list*	Keyword: **`close`**	*default*
self-modifier	*locator-list*	Keyword: **`self`**	*default*

Clauses

`map`

The *map-type-modifying* modifier group defines the modifiers that combine with a *map-type* to determine details of a mapping operation. The effect of the modifiers in this group is described in Section 11.5.8.

Cross References

- Storage Block Mapping, see Section 11.5.8
- **`map`** Clause, see Section 11.5

11.5.4 `ref-modifier` Modifier

Modifiers

Name	Modifies	Type	Properties
ref-modifier	*all arguments*	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`** *ref-modifier* is specified, the semantics of the clause apply to the referring pointer of the referencing variable. If the **`ref_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_ptr_ptee`** *ref-modifier* is specified, the effect is as if the list item appeared instead in two identical **`map`** clauses except the modifier is replaced with the **`ref_ptr`** *ref-modifier* in one and with the **`ref_ptee`** *ref-modifier* in the other.

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

When the semantics of the clause apply to the referring pointer of the list item, the base referencing variable of the list item is the referencing variable for which the referenced pointee contains the storage of that referring pointer.

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:

- Each list item that appears in a clause with the *ref-modifier* must be a referencing variable.

C / C++

- Each list item that appears in a clause for which the *ref-modifier* is specified must have a containing structure.

C / C++

Cross References

- **map** Clause, see Section 11.5

11.5.5 `attach-modifier` Modifier

Modifiers

Name	Modifies	Type	Properties
attach-modifier	*all arguments*	Complex, name: **attach** Arguments: ***attachment-mode*** Keyword: **always**, **auto**, **never** (*default*)	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 data entity that the corresponding list item represents 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 11.5.7.

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.
- Each 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 11.5.7
- **map** Clause, see Section 11.5

11.5.6 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.

Each 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* modifier must be **storage**.

11.5.7 Array and Structure Mapping

If a list item is an array or array section, array-map expansion occurs: the array elements become implicit list items with the same modifiers (including *map-type*) as specified in the clause. If the mapping operation for the array or array section is relaxed 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. Additionally, if the mapping operation for the array or array section is relaxed, the implicit maps resulting from its array-map expansion are also relaxed.

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 11.7) 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 that mapper

expansion occurs: the list item is removed from the **map** clause and the clauses specified in the declared mapper are applied 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 from the *map-type-modifying* modifier group appears in the **map** clause then the effect is as if that modifier appears in each **map** clause specified in the declared mapper. For a **map** clause applied as a result of mapper expansion, the mapping operation is relaxed if it is not a self map, the mapping operation for the original list item is relaxed, and the clause does not have the *present-modifier*.

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 that is explicitly mapped, 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 is polymorphic and its dynamic type differs from its declared 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.
- Each list item must not be unlimited polymorphic unless the **map** clause is specified with the *self-modifier*.
- 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 11.7

11.5.8 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:

- The original list item is not accessible and cannot be made accessible from the target device;
- The corresponding list item is present prior to a task encountering the construct on which the clause appears, and the corresponding storage differs from the original storage; or
- The list item is a pointer that would be assigned a different value as a result of pointer attachment.

Restrictions

Additional restrictions to the **map** clause for mapping of storage blocks are as follows:

- If the mapping operation for a list item is strict and any part of the original storage of the list item has corresponding storage in the device data environment prior to a task encountering the construct associated with the **map** clause, all of the original storage must have corresponding storage in the device data environment prior to the task encountering the construct.
- If a list item in a **map** 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 list item appears in a **map** clause with the *self-modifier*, any other list item in a **map** clause on the same construct that has the same base variable or base pointer must also be specified with the *self-modifier*.

11.5.9 Implicit Data-Mapping

Before any array-map expansion or mapper expansion occurs, the effect of an implicitly determined data-mapping attribute on each list item is determined as follows:

- The mapping operation for the list item is relaxed unless the mapping operation is self, in which case it is strict.

- If a single contiguous part of the original storage of the list item has corresponding storage in the device data environment prior to a task encountering the construct on which the **map** clause appears, only that part of the original storage will have corresponding storage in the device data environment as a result of the **map** clause;
- If any part of the original storage of the list item is a base referring pointer or a base pointer to another list item that is explicitly mapped on the construct, the effect is as if that base referring pointer was also explicitly mapped on the construct with the **storage** *map-type* and the **ref_ptr** *ref-modifier*, and as if that base pointer was also explicitly mapped on the construct with the **storage** *map-type*, respectively;
- If any part of the original storage of the list item is explicitly mapped on the construct, the list item is implicitly mapped as a shallow map with a *map-type* of **storage**; and
- If the list item does not have any corresponding storage in the device data environment prior to a task encountering the construct associated with the **map** clause, and one or more contiguous parts of the original storage are list items that are explicitly mapped on the construct, only those parts of the original storage will have corresponding storage in the device data environment as a result of the **map** clauses on the construct.

11.6 `defaultmap` Clause

Name: `defaultmap`	**Properties:** default-attribute, post-modified

Arguments

Name	Type	Properties
default-behavior	Keyword: **firstprivate**, **from**, **implicit**, **none**, **present**, **private**, **self**, **storage**, **to**, **tofrom**	*default*

Modifier groups

variable-selector

Modifiers

Name	Modifies	Type	Properties
directive-name-modifier	*all arguments*	Keyword: *directive-name* (a directive name)	unique

Directives

target

Additional information

The value **alloc** may also be specified as *default-behavior* with identical meaning to the value **storage**. The value **default** may also be specified as *default-behavior* with identical meaning to the value **implicit**.

Semantics

The **defaultmap** clause controls the implicitly determined data-mapping attributes or data-sharing attributes of certain variables that are referenced in a **target** construct, in accordance with the rules given in Section 11.1.

If *default-behavior* is **implicit** then the data-environment attribute is implicitly determined by the rules given in Section 11.1.

If *default-behavior* corresponds to a data-sharing attribute, the behavior is as if a **default** clause with the same argument and *variable-selector* modifiers were instead applied to the **target** construct.

If *default-behavior* corresponds to a *map-type* or *map-type-modifying* modifier, the attribute is a data-mapping attribute determined by an implicit **map** clause with the corresponding modifier.

If *default-behavior* is **none** then no implicitly determined data-mapping attributes or implicitly determined data-sharing attributes are defined for variables selected by the specified *variable-selector* modifiers, except for variables that appear in the **enter** or **link** clause of a **declare_target** directive.

Restrictions

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

- If *default-behavior* is **none**, each variable that is specified by the *variable-selector* modifiers and is referenced in the construct but does not have a predetermined data-sharing attribute and does not appear in an **enter** or **link** clause on a **declare_target** directive must have an explicitly determined data-environment attribute on the construct.

Cross References

- Implicit Data-Mapping Attribute Rules, see Section 11.1
- **target** Construct, see Section 21.8

11.7 declare_mapper Directive

Name: **declare_mapper** **Category:** declarative	**Association:** unassociated **Properties:** pure

Arguments

`declare_mapper`(*mapper-specifier*)

Name	Type	Properties
mapper-specifier	OpenMP mapper specifier	*default*

Clauses

`map`

Additional information

The **`declare_mapper`** directive may alternatively be specified with **`declare mapper`** as the *directive-name*.

Semantics

User-defined mappers can be defined using the **`declare_mapper`** directive. The *mapper-specifier* argument declares the mapper using the following syntax:

C / C++

```
[ mapper-identifier : ] type var
```

C / C++

Fortran

```
[ mapper-identifier : ] type :: var
```

Fortran

where *mapper-identifier* is a mapper identifier, *type* is a type that is permitted in a type-name list, and *var* is a base language identifier.

The *type* and an optional *mapper-identifier* uniquely identify the mapper for use in a **`map`** clause or data-motion clause later in the OpenMP program.

If *mapper-identifier* is not specified, the behavior is as if *mapper-identifier* is **`default`**.

The variable declared by *var* is available for use in all **`map`** clauses on the directive, and no part of the variable to be mapped is mapped by default.

The effect that a user-defined mapper has on either a **`map`** clause that maps a list item of the given base language type or a data-motion clause that invokes the mapper and updates a list item of the given base language type is to replace the map or update with a set of **`map`** clauses or updates derived from the **`map`** clauses specified by the mapper, as described in Section 11.5 and Chapter 12.

Each list item in a **`map`** clause that appears on a **`declare_mapper`** directive may include array sections.

All **`map`** clauses that are introduced by a mapper are further subject to mappers that are in scope, except a **`map`** clause with list item *var* maps *var* without invoking a mapper.

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.
- 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 11.5

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. mapper expansion occurs: 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
locator-list	list of locator list item type	*default*

Modifiers

Name	Modifies	Type	Properties
present-modifier	*locator-list*	Keyword: **present**	*default*
mapper	*locator-list*	Complex, name: **mapper** Arguments: ***mapper-identifier*** OpenMP identifier (*default*)	unique
iterator	*all arguments*	Complex, name: **iterator** Arguments: ***iterator-specifier-list*** list of iterator specifier list item type (*default*)	unique
directive-name-modifier	*all arguments*	Keyword: *directive-name* (a directive name)	unique

Directives

target_update

Semantics

The **to** clause is a data-motion clause that specifies data movement to the target devices from the encountering device so the corresponding list items are the assigned list items and the compatible map types are **to** and **tofrom**.

C

Each list item for which a mapper does not exist is ignored if it has static storage duration and has the **constexpr** specifier.

C

C++

Each list item for which a mapper does not exist is ignored if it has static storage duration and either it has the **constexpr** specifier or it is a non-mutable member of a structure that has the **constexpr** specifier.

C++

Cross References

- **iterator** Modifier, see Section 5.2.6
- **target_update** Construct, see Section 21.9

12.2 from Clause

Name: from	**Properties:** data-motion attribute

Arguments

Name	Type	Properties
locator-list	list of locator list item type	*default*

Modifiers

Name	Modifies	Type	Properties
present-modifier	*locator-list*	Keyword: **`present`**	*default*
mapper	*locator-list*	Complex, name: **`mapper`** Arguments: ***mapper-identifier*** OpenMP identifier (*default*)	unique
iterator	*all arguments*	Complex, name: **`iterator`** Arguments: ***iterator-specifier-list*** list of iterator specifier list item type (*default*)	unique
directive-name-modifier	*all arguments*	Keyword: *directive-name* (a directive name)	unique

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

Each 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++

Each 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

Each 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
`omp_default_mem_space`	Represents the system default storage
`omp_large_cap_mem_space`	Represents storage with large capacity
`omp_const_mem_space`	Represents storage optimized for variables with constant values
`omp_high_bw_mem_space`	Represents storage with high bandwidth
`omp_low_lat_mem_space`	Represents storage with low latency
`omp_cgroup_mem_space`	Represents storage optimized for shared access within the contention group

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
sync_hint	**contended**, **uncontended**, **serialized**, **private**	**contended**
alignment	Non-negative integer powers of 2	1 byte
access	**all**, **memspace**, **device**, **cgroup**, **pteam**, **thread**	**memspace**
pool_size	Any positive integer	Implementation defined
fallback	**default_mem_fb**, **null_fb**, **abort_fb**, **allocator_fb**	See below
fb_data	An allocator handle	(none)
pinned	*true*, *false*	*false*
partition	**environment**, **nearest**, **blocked**, **interleaved**, **partitioner**	**environment**
pin_device	Conforming device number	(none)
preferred_device	Conforming device number	(none)
target_access	**single**, **multiple**	**single**

table continued on next page

table continued from previous page

Allocator Trait	Allowed Values	Default Value
atomic_scope	**all**, **device**	**device**
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
omp_default_mem_alloc	**omp_default_mem_space**	**fallback**:**null_fb**
omp_large_cap_mem_alloc	**omp_large_cap_mem_space**	(none)
omp_const_mem_alloc	**omp_const_mem_space**	(none)
omp_high_bw_mem_alloc	**omp_high_bw_mem_space**	(none)
omp_low_lat_mem_alloc	**omp_low_lat_mem_space**	(none)
omp_cgroup_mem_alloc	**omp_cgroup_mem_space**	**access**:**cgroup**
omp_pteam_mem_alloc	Implementation defined	**access**:**pteam**
omp_thread_mem_alloc	Implementation defined	**access**:**thread**

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.

13.3 align Clause

Name: align	**Properties:** unique

Arguments

Name	Type	Properties
alignment	expression of integer type	constant, positive

Modifiers

Name	Modifies	Type	Properties
directive-name-modifier	*all arguments*	Keyword: *directive-name* (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: `allocator`	**Properties:** ICV-defaulted, unique

Arguments

Name	Type	Properties
allocator	expression of OpenMP allocator_handle type	*default*

Modifiers

Name	Modifies	Type	Properties
directive-name-modifier	*all arguments*	Keyword: *directive-name* (a directive name)	unique

Directives

`allocate`

Semantics

The **`allocator`** clause specifies the memory allocator to be used for allocations associated with the construct on which the clause appears. Specifically, the allocator to which *allocator* evaluates is used for the allocations. On 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

- **`allocate`** Directive, see Section 13.5
- Memory Allocators, see Section 13.2
- *def-allocator-var* ICV, see Table 3.1

13.5 `allocate` Directive

Name: `allocate` **Category:** declarative	**Association:** explicit **Properties:** pure

Arguments

Name	Type	Properties
list	list of variable list item type	*default*

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 allocation happens at the same point at which the list item would have been allocated in the base language and its scope is that of the list item in the base language. At the end of the scope for a given list item the memory allocator that is used to allocate that list item is used to deallocate 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.
- Each list item must not be a function parameter.
- If a list item has automatic storage duration, the value of the variable is undefined unless the execution of the program passes through its declaration.

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

- Each 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.
- Each 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
list	list of variable list item type	*default*

Modifiers

Name	Modifies	Type	Properties
allocator-simple-modifier	*list*	expression of OpenMP allocator_handle type	exclusive, unique
allocator-complex-modifier	*list*	Complex, name: **allocator** Arguments: ***allocator*** expression of allocator_handle type (*default*)	unique
align-modifier	*list*	Complex, name: **align** Arguments: ***alignment*** expression of integer type (constant, positive)	unique
directive-name-modifier	*all arguments*	Keyword: *directive-name* (a directive name)	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; however, if a directiverefteams construct is an immediately nested construct of the **target** construct, the specified allocator must be a predefined memory allocator.
- **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

Fortran

13.7 allocators Construct

Name: allocators **Category:** executable	**Association:** block : allocator **Properties:** *default*

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:

- Each 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.
- Each 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

Fortran

13.8 `uses_allocators` Clause

Name: `uses_allocators`	**Properties:** data-environment attribute, data-sharing attribute

Arguments

Name	Type	Properties
allocator	expression of OpenMP allocator_handle type	*default*

Modifiers

Name	Modifies	Type	Properties
mem-space	*allocator*	Complex, name: **`memspace`** Arguments: ***memspace-handle*** expression of OpenMP memspace_handle type (*default*)	*default*
traits-array	*allocator*	Complex, name: **`traits`** Arguments: ***traits*** variable of alloctrait array type (*default*)	*default*
directive-name-modifier	*all arguments*	Keyword: *directive-name* (a directive name)	unique

Directives

`target`

Semantics

The **`uses_allocators`** clause enables the use of the specified *allocator* in the region associated with the directive on which the clause appears. 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.

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 34.7
- **omp_init_allocator** Routine, see Section 34.6
- **target** Construct, see Section 21.8

13.9 dyn_groupprivate Clause

Name: dyn_groupprivate	**Properties:** target-consistent

Arguments

Name	Type	Properties
size	expression of integer type	non-negative

Modifiers

Name	Modifies	Type	Properties
access-group	*size*	Keyword: **cgroup**	unique
fallback-modifier	*size*	Complex, name: **fallback** Arguments: ***fallback-mode*** Keyword: **abort**, **default_mem**, **null** (*default*)	unique
directive-name-modifier	*all arguments*	Keyword: *directive-name* (a directive name)	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_gprivate_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 memory fallback 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 34.13.

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 34.13
- **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: `uniform`	**Properties:** data-environment attribute

Arguments

Name	Type	Properties
parameter-list	list of parameter list item type	*default*

Modifiers

Name	Modifies	Type	Properties
directive-name-modifier	*all arguments*	Keyword: *directive-name* (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.

Cross References

- **`declare_simd`** Directive, see Section 15.7

14.2 `aligned` Clause

Name: `aligned`	**Properties:** data-environment attribute, post-modified

Arguments

Name	Type	Properties
list	list of variable list item type	*default*

Modifiers

Name	Modifies	Type	Properties
alignment	*list*	expression of integer type	constant, positive, ultimate, unique
directive-name-modifier	*all arguments*	Keyword: *directive-name* (a directive name)	unique

Directives

declare_simd, **simd**

Semantics

C / C++

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

C / C++

Fortran

The **aligned** clause declares that the target of each list item is aligned to the number of bytes expressed in *alignment*.

Fortran

The *alignment* modifier specifies the alignment that the program ensures related to the list items. If the *alignment* modifier is not specified, implementation defined default alignments for SIMD instructions on the target platforms are assumed.

Restrictions

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

- If the clause appears on a **declare_simd** directive, each list item must be a parameter list item of the associated procedure.

C

- The type of each list item must be an array or pointer type.

C

C++

- The type of each list item must be an array, pointer, reference to array, or reference to pointer type.

C++

Fortran

- Each list item must be an array.

Fortran

Cross References

- **declare_simd** Directive, see Section 15.7
- **simd** Construct, see Section 18.5

Part III

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 function dispatch 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 $c_1, \ldots, c_N$, 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 $c_1, \ldots, c_M$ of the **`construct`** selector set are added in the same order to the beginning of the construct trait set as $c_1, \ldots, c_M$ 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.

Traits in the construct trait set are non-property traits, unless otherwise specified.

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 *simdlen* property and one of the *inbranch* or *notinbranch* properties.

The *dispatch* trait is a clause-list trait that is defined with properties that match the dispatch-requirement clauses that can be specified on the **`dispatch`** directive with the same names and semantics, except that the **`interop`** clause results in a distinct *interop* property for each of its list item. If variant substitution results in a dispatch-compliant clause consuming one of the *dispatch* trait properties then the property is not otherwise considered active for that call.

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 *vendor(vendor-list)* name-list trait specifies the vendor identifiers for the device manufacturers. Each member of *vendor-list* is a *vendor-name*, for which the defined values are implementation defined but must include at least the values in the OpenMP Additional Definitions document.

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 implementation defined implicit requirements.

Implementations can define additional traits in the device trait set, target device trait set and implementation trait set; these traits are extension traits.

The dynamic trait set includes traits that define the dynamic properties of an OpenMP program at a point in its execution. The *data state* trait in the dynamic trait set refers to the complete data state of the OpenMP program that may be accessed at runtime.

15.2 Context Selectors

Context selectors are used to define the properties that can match an OpenMP context. OpenMP defines different trait selector sets, each of which contains different trait selectors.

The syntax for a context selector is *context-selector-specification* as described in the following grammar:

```
context-selector-specification:
    trait-set-selector[,trait-set-selector[,...]]

trait-set-selector:
    trait-set-selector-name={trait-selector[, trait-selector[, ...]]}

trait-selector:
    trait-selector-name[([trait-score: ] trait-property[, trait-property[, ...]])]

trait-property:
    trait-property-name
    trait-property-clause
    trait-property-expression
    trait-property-extension

trait-property-clause:
    clause
```

trait-property-name:
    identifier
    string-literal

trait-property-expression
    scalar-expression (for C/C++)
    scalar-logical-expression (for Fortran)
    scalar-integer-expression (for Fortran)

trait-score:
    score(*score-expression***)**

trait-property-extension:
    trait-property-name
    identifier **(***trait-property-extension*[**,** *trait-property-extension*[**,** *...*]]**)**
    constant integer expression

For trait selectors that correspond to name-list traits, each *trait-property* should be *trait-property-name* and, for any value that is a valid identifier, both the identifier and the 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 and restrictions is the same as for the matching clause, unless otherwise specified.

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. The **construct** selector set is an ordered list $c_1, \ldots, c_N$.

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.
- A *trait-property* must not be specified in a **dispatch** trait selector.
- 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.
- A trait of the **target_device** selector set may only be specified if the clause that contains the context selector is evaluated on the host device.
- 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:** *default*

Arguments

Name	Type	Properties
directive-variant	directive-specification list item	optional, unique

Modifiers

Name	Modifies	Type	Properties
context-selector	*directive-variant*	An OpenMP context-selector-specification	required, unique
directive-name-modifier	*all arguments*	Keyword: *directive-name* (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.
- *directive-variant* must not specify a **begin declare_variant** directive.

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
directive-variant	directive-specification list item	optional, unique

Modifiers

Name	Modifies	Type	Properties
directive-name-modifier	*all arguments*	Keyword: *directive-name* (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.
- *directive-variant* must not specify a **begin declare_variant** directive.

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 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* modifier.

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 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.
- A dispatch-compliant 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

15.5.1 `match` Clause

Name: `match`	**Properties:** unique, required

Arguments

Name	Type	Properties
context-selector	An OpenMP context-selector-specification	*default*

Modifiers

Name	Modifies	Type	Properties
directive-name-modifier	*all arguments*	Keyword: *directive-name* (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 15.5.5
- **`declare_variant`** Directive, see Section 15.5.4
- Context Selectors, see Section 15.2

15.5.2 `adjust_args` Clause

Name: `adjust_args`	**Properties:** dispatch-compliant

Arguments

Name	Type	Properties
argument-list	list of procedure-argument list item type	*default*

Modifier groups

adjust-op

Modifiers

Name	Modifies	Type	Properties
directive-name-modifier	*all arguments*	Keyword: *directive-name* (a directive name)	unique

Directives

`declare_variant`

Semantics

The **`adjust_args`** clause specifies how to adjust the arguments of the base function when a specified function variant is selected for replacement 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 modifier from the *adjust-op* modifier group 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 procedure call site is ignored.

Cross References

- **`declare_variant`** Directive, see Section 15.5.4

15.5.2.1 *adjust-op* Modifiers

Modifier set

Properties: required, exclusive	**Members:** *need-device-addr*, *need-device-ptr*, *nothing*

Modifiers

Name	Modifies	Type	Properties
nothing	*argument-list*	Keyword: **`nothing`**	*default*
need-device-addr	*argument-list*	Keyword: **`need_device_addr`**	*default*
need-device-ptr	*argument-list*	Complex, name: **`need_device_ptr`** Arguments: ***fallback*** Keyword: **`fb_nullify`**, **`fb_preserve`** (optional)	*default*

Clauses

adjust_args

Semantics

The *adjust-op* modifier group defines the adjustment operation to apply to arguments:

If the *nothing* modifier is present, the argument is passed to the selected function variant without modification.

If the *need-device-ptr* modifier is present, the arguments are converted to corresponding device pointers of the default device if they are not already device pointers. If the *dispatch* trait and *is_device_ptr* property for a given argument is active in the OpenMP context at 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 with the same *fallback* modifier as the *fallback* argument of the *need-device-ptr* modifier converts its pointer list items into device pointers on a **target_data** construct. If the *fallback* argument is not specified, the behavior is as if the **fb_nullify** *fallback* argument is present.

If the *need-device-addr* modifier is present, 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 *dispatch* trait and *has_device_addr* property for a given argument is active in the OpenMP context at 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

Restrictions to the *adjust-op* modifier group are as follows:

- If the *need-device-addr* modifier is present in a clause and the *dispatch* trait and *has_device_addr* property does not exist for a specified argument in the OpenMP context at the call to the base function, 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* modifier is present in a clause, 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* modifier must not be specified in a clause.

C

C++

- If the *need-device-ptr* modifier is present in a clause, 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* modifier is present in a clause, 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* modifier is present in a clause, 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* modifier is present in a clause, 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* modifier is present in a clause, the corresponding actual argument for each specified argument must be contiguous.

Fortran

Cross References

- **adjust_args** Clause, see Section 15.5.2
- **use_device_addr** Clause, see Section 7.3.8
- **use_device_ptr** Clause, see Section 7.3.6

15.5.3 append_args Clause

Name: append_args	**Properties:** unique, dispatch-compliant

Arguments

Name	Type	Properties
append-op-list	list of OpenMP operation list item type	*default*

Modifiers

Name	Modifies	Type	Properties
directive-name-modifier	*all arguments*	Keyword: *directive-name* (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 OpenMP context for the call contains the *dispatch* trait with an ordered list of *interop* properties that have not yet been consumed, the first such property is consumed for the function dispatch; if the property is not associated with **omp_interop_none**, its associated interoperability object is appended as an argument. Otherwise, if no such property exists or it is associated with **omp_interop_none**, the **interop** operation constructs an argument of **interop** OpenMP type using the *dispatch* trait of the OpenMP context. 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 OpenMP context contains the *dispatch* trait with one or more properties that correspond to clauses for an **interop** construct with a given interop-type set, those properties are consumed for the function dispatch and the behavior is as if the corresponding clauses are specified on the **interop** construct.

Any appended arguments that were not obtained from the *interop* property of the *dispatch* trait in the OpenMP context 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 15.5.4
- **destroy** Clause, see Section 5.8
- OpenMP Operations, see Section 5.2.3
- **init** Clause, see Section 5.7
- **interop** Construct, see Section 22.1

15.5.4 declare_variant Directive

Name: **declare_variant**	**Association:** declaration
Category: declarative	**Properties:** pure

Arguments

declare_variant (*[base–name:]variant-name*)

Name	Type	Properties
base-name	identifier of function type	optional
variant-name	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 expression in the **match** clause is interpreted as if it appeared in the scope of arguments of the base function.

C

C++

If `:::` appears in the directive argument, it is intepreted in the same manner as specified for an *intra-clause-specification* (see Section 5.2), as if *base-name* is a *modifier-name* in *modifier-specification-list* and *variant-name* is an *argument-name* in *clause-argument-list*.

variant-name and any expression 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

Any expression in the **match** clause is interpreted as if it appeared at the end of specification part that contains the **declare_variant** directive.

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, the name of a procedure pointer, a dummy procedure or a statement function.
- *variant-name* must not be the name of a procedure pointer, a dummy procedure or a statement function.
- If the base function contains an *entry-stmt*, the **declare_variant** directive must specify *base-name*.
- The procedure *base-name* must have an accessible explicit interface at the location of the directive.
- Any expression in the **match** clause must not contain any reference to an internal procedure nor to any local variable.

Fortran

Cross References

- **adjust_args** Clause, see Section 15.5.2
- **append_args** Clause, see Section 15.5.3
- Declare Variant Directives, see Section 15.5
- **match** Clause, see Section 15.5.1

15.5.5 begin declare_variant Directive

Name: **begin declare_variant**	Association: delimited
Category: declarative	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

Cross References

- Declare Variant Directives, see Section 15.5
- **match** Clause, see Section 15.5.1

15.6 `dispatch` Construct

Name: **dispatch** **Category:** executable	**Association:** block : function-dispatch **Properties:** context-matching, implicit-taskwait, task-synchronizing

Clauses

depend, **device**, **has_device_addr**, **interop**, **is_device_ptr**, **nocontext**, **novariants**, **nowait**

Binding

The binding task set for a **dispatch** region is the generating task. The **dispatch** region binds to the region of the generating task.

Semantics

The **dispatch** construct controls whether variant substitution occurs for *function-dispatch* in the associated function-dispatch structured block. The **dispatch** construct may also affect the arguments of the function variant if variant substitution occurs (see Section 15.5.2 and Section 15.5.3).

The **dispatch** construct has the implicit-taskwait property. If the **dispatch** directive specifies one or more **depend** clauses, and after any variant substitution the *depend* properties of the *dispatch* trait remain active in the OpenMP context of the function dispatch, the behavior is as if those **depend** clauses are applied to an implicit taskwait that is executed immediately before the call is executed. The construct therefore implies the task-dependence scheduling constraint for the dependent task that results from the implicit taskwait. The other dispatch-requirement clauses have no effect on the execution of the **dispatch** construct apart from the effect they may have on the 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* that is not **omp_interop_none**, 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*.

If the **interop** clause is present and **depend** clauses are specified, the behavior is as if an **interop** construct with a **use** clause with the same *interop-var-list* was encountered and with any **depend** clauses of the **dispatch** construct applied to it before entering the **dispatch** region.

If the **interop** clause is present and the **nowait** is not present, the behavior is as if an **interop** construct with a **use** clause with the same *interop-var-list* was encountered and with any **depend** clauses of the **dispatch** construct applied to it after the end of the **dispatch** region.

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 23.9.6
- **device** Clause, see Section 21.2
- OpenMP Function Dispatch Structured Blocks, see Section 6.3.2
- **has_device_addr** Clause, see Section 7.3.7
- **interop** Clause, see Section 15.6.1
- **is_device_ptr** Clause, see Section 7.3.5
- **nocontext** Clause, see Section 15.6.3
- **novariants** Clause, see Section 15.6.2
- **nowait** Clause, see Section 23.6
- **taskwait** Construct, see Section 23.5

15.6.1 interop Clause

Name: interop	**Properties:** dispatch-requirement, unique

Arguments

Name	Type	Properties
interop-var-list	list of variable of OpenMP interop type	*default*

Modifiers

Name	Modifies	Type	Properties
directive-name-modifier	*all arguments*	Keyword: *directive-name* (a directive name)	unique

Directives

dispatch

Semantics

The **interop** clause specifies an ordered list of one or more *interop* properties of the *dispatch* trait in the OpenMP context, one for each interoperability objects as specified by *interop-var-list*.

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 *function-dispatch* 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 15.6

15.6.2 novariants Clause

Name: novariants	**Properties:** unique

Arguments

Name	Type	Properties
do-not-use-variant	expression of OpenMP logical type	*default*

Modifiers

Name	Modifies	Type	Properties
directive-name-modifier	*all arguments*	Keyword: *directive-name* (a directive name)	unique

Directives

dispatch

Semantics

If *do-not-use-variant* evaluates to *true*, no function variant is selected for the function dispatch 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 15.6

15.6.3 `nocontext` Clause

Name: `nocontext`	**Properties:** unique

Arguments

Name	Type	Properties
do-not-update-context	expression of OpenMP logical type	*default*

Modifiers

Name	Modifies	Type	Properties
directive-name-modifier	*all arguments*	Keyword: *directive-name* (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 15.6

15.7 `declare_simd` Directive

Name: `declare_simd` **Category:** declarative	**Association:** declaration **Properties:** pure, variant-generating

Arguments

`declare_simd`*[* **(***proc-name***)** *]*

Name	Type	Properties
proc-name	identifier of function type	optional

Clause groups

branch

Clauses

`aligned`, **`linear`**, **`simdlen`**, **`uniform`**

Additional information

The **`declare_simd`** directive may alternatively be specified with **`declare simd`** as the *directive-name*.

Semantics

The association of one or more **declare_simd** directives with a procedure declaration or definition enables the creation of corresponding SIMD versions of the associated procedure that can be used to process multiple arguments from a single invocation in a SIMD loop concurrently.

If a SIMD version is created and the **simdlen** clause is not specified, the number of concurrent arguments for the function is implementation defined.

For purposes of the **linear** clause, any integer-typed parameter that is specified in a **uniform** clause on the directive is considered to be constant and so may be used in a *step-complex-modifier* as *linear-step*.

C / C++

The expressions that appear in the clauses of each directive are evaluated in the scope of the arguments of the procedure declaration or definition.

C / C++

C++

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

C++

Restrictions

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

- The procedure body must be a structured block.
- The execution of the procedure, when called from a SIMD loop, must not result in the execution of any constructs except for **atomic** constructs and **ordered** constructs on which the **simd** clause is specified.
- The execution of the procedure must not have any side effects that would alter its execution for concurrent iterations of a SIMD chunk.

C / C++

- If a **declare_simd** directive is specified for a declaration of a procedure then the definition of the procedure must have a **declare_simd** directive with identical clauses with identical arguments and modifiers.
- 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 18.5.3
- **uniform** Clause, see Section 14.1

15.7.1 *branch* Clauses

Clause set

Properties: exclusive	**Members:** **Clauses** **inbranch**, **notinbranch**

Directives

declare_simd

Semantics

The *branch* clause group defines a set of clauses that indicate if a procedure can be assumed to be or not to be encountered in a branch. If neither clause is specified, then the procedure may or may not be called from inside a conditional statement of the calling context.

Cross References

- **declare_simd** Directive, see Section 15.7

15.7.1.1 `inbranch` Clause

Name: `inbranch`	**Properties:** unique

Arguments

Name	Type	Properties
inbranch	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

`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 15.7

15.7.1.2 `notinbranch` Clause

Name: `notinbranch`	**Properties:** unique

Arguments

Name	Type	Properties
notinbranch	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

`declare_simd`

Semantics

If *notinbranch* evaluates to *true*, the **`notinbranch`** clause specifies that the procedure will never be called from inside a conditional statement of the calling context. If *notinbranch* evaluates to *false*, the procedure may be called from inside a conditional statement. If *notinbranch* is not specified, the effect is as if *notinbranch* evaluates to *true*.

Cross References

- **declare_simd** Directive, see Section 15.7

15.8 Declare Target Directives

Declare target directives apply to procedures and/or variables to ensure that they can be executed or accessed on a device. Variables are either replicated as device-local variables for each device through a **local** clause, are mapped for all device executions through an **enter** clause, or are mapped for specific device executions through a **link** clause. An implementation may generate different versions of a procedure to be used for **target** regions that execute on different devices. Whether it generates different versions, and whether it calls a different version in a **target** region from the version that it calls outside a **target** region, are implementation defined.

To facilitate device usage, OpenMP defines rules that implicitly specify declare target directives for procedures and variables. The remainder of this section defines those rules as well as restrictions that apply to all declare target directives.

C++

If a variable with static storage duration has the **constexpr** specifier and is not a groupprivate variable then the variable is treated as if it had appeared as a list item in an **enter** clause on a declare target directive.

C++

If a variable with static storage duration that is not a device-local variable (including that it is not a groupprivate variable) is declared in a device procedure then the variable is treated as if it had appeared as a list item in an **enter** clause on a declare target directive.

If a procedure is referenced outside of any reverse-offload region in a procedure that appears as a list item in an **enter** clause on a non-host declare target directive then the name of the referenced procedure is treated as if it had appeared in an **enter** clause on a declare target directive.

C / C++

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 15.8.2
- **`declare_target`** Directive, see Section 15.8.1
- **`enter`** Clause, see Section 11.3
- **`link`** Clause, see Section 11.4
- OMPT **`scope_endpoint`** Type, see Section 40.27
- **`target`** Construct, see Section 21.8
- **`target_data_op_emi`** Callback, see Section 42.7

15.8.1 `declare_target` Directive

Name: **`declare_target`** **Category:** declarative	**Association:** explicit **Properties:** declare-target, device, pure, variant-generating

Arguments

`declare_target` (***extended-list***)

Name	Type	Properties
extended-list	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 11.3
- Declare Target Directives, see Section 15.8
- **`indirect`** Clause, see Section 15.8.3
- **`link`** Clause, see Section 11.4
- **`local`** Clause, see Section 9.3

C / C++

15.8.2 `begin declare_target` Directive

Name: `begin declare_target` **Category:** declarative	**Association:** delimited **Properties:** declare-target, device, variant-generating

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 an explicit **device_type** or **indirect** clause is present on the contained declare target directive, then its argument determines, respectively, which versions are made available or whether calling with an indirect device invocation is permitted. Otherwise, the effect is as if the respective clause of the innermost enclosing **begin declare_target** is specified on the contained declare target directive. If a list item appears both in an implicit and explicit list, the explicit list determines the versions that 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 11.3
- Declare Target Directives, see Section 15.8
- **indirect** Clause, see Section 15.8.3

C / C++

15.8.3 indirect Clause

Name: indirect	**Properties:** unique

Arguments

Name	Type	Properties
invoked-by-fptr	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

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*.

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 15.8.2
- **declare_target** Directive, see Section 15.8.1

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: `error` **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 41.2
- **`message`** Clause, see Section 16.3
- **`severity`** Clause, see Section 16.4

16.2 `at` Clause

Name: `at`	**Properties:** unique

Arguments

Name	Type	Properties
action-time	Keyword: **`compilation`**, **`execution`**	*default*

Modifiers

Name	Modifies	Type	Properties
directive-name-modifier	*all arguments*	Keyword: *directive-name* (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
msg-string	expression of string type	*default*

Modifiers

Name	Modifies	Type	Properties
directive-name-modifier	*all arguments*	Keyword: *directive-name* (a directive name)	unique

Directives

error, **parallel**, **target**, **teams**

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. When the *action-time* is **execution**, the support for a non-constant expression in *msg-string* within a **target** region is implementation defined.

Restrictions

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

- 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
- **target** Construct, see Section 21.8
- **teams** Construct, see Section 18.2

16.4 severity Clause

Name: **severity**	**Properties:** unique

Arguments

Name	Type	Properties
sev-level	Keyword: **fatal**, **warning**	*default*

Modifiers

Name	Modifies	Type	Properties
directive-name-modifier	*all arguments*	Keyword: *directive-name* (a directive name)	unique

Directives

error, **parallel**, **target**, **teams**

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
- **target** Construct, see Section 21.8
- **teams** Construct, see Section 18.2

16.5 requires Directive

Name: **requires** **Category:** informational	**Association:** unassociated **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.

C

- The **requires** directive must only appear at file scope.

C

C++

- The **requires** directive must only appear at file or namespace scope.

C++

C / C++

- Any **requires** directive that specifies a device global requirement clause must appear lexically before any device constructs or device procedures.

C / C++

Fortran

- 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.

Fortran

16.5.1 *requirement* Clauses

Clause set

Properties: required	**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

Arguments

Name	Type	Properties
memory-order	Keyword: **`acq_rel`**, **`acquire`**, **`relaxed`**, **`release`**, **`seq_cst`**	*default*

Modifiers

Name	Modifies	Type	Properties
directive-name-modifier	*all arguments*	Keyword: *directive-name* (a directive name)	unique

Directives

`requires`

Semantics

The **`atomic_default_mem_order`** clause specifies the default memory ordering behavior for **`atomic`** constructs that an implementation must provide. The effect is as if its argument appears as a clause on any **`atomic`** construct that does not specify a *memory-order* clause and that permits the clause.

Restrictions

Restrictions to the **`atomic_default_mem_order`** clause are as follows:

- All **`requires`** directives in the same compilation unit that specify the **`atomic_default_mem_order`** requirement must specify the same argument.
- Any directive that specifies the **`atomic_default_mem_order`** clause must not appear lexically after any **`atomic`** construct on which a *memory-order* clause is not specified.

Cross References

- **`atomic`** Construct, see Section 23.8.5
- *memory-order* Clauses, see Section 23.8.1
- **`requires`** Directive, see Section 16.5

16.5.1.2 `dynamic_allocators` Clause

Name: `dynamic_allocators`	**Properties:** unique

Arguments

Name	Type	Properties
required	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

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 34.7
- **omp_init_allocator** Routine, see Section 34.6
- **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
required	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

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: unified_address	**Properties:** unique, device global requirement

Arguments

Name	Type	Properties
required	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

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 7.3.5
- **`requires`** Directive, see Section 16.5
- **`target`** Construct, see Section 21.8

16.5.1.5 `unified_shared_memory` Clause

Name: `unified_shared_memory`	**Properties:** unique, device global requirement

Arguments

Name	Type	Properties
required	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

`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 11.3
- **`requires`** Directive, see Section 16.5
- **`unified_address`** Clause, see Section 16.5.1.4

16.5.1.6 `self_maps` Clause

Name: `self_maps`	**Properties:** unique, device global requirement

Arguments

Name	Type	Properties
required	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

`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 11.3
- **`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
required	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

`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 set

Properties: required	**Members:** **Clauses** **`absent`**, **`contains`**, **`holds`**, **`no_openmp`**, **`no_openmp_constructs`**, **`no_openmp_routines`**, **`no_parallelism`**

Directives

`assume`, **`assumes`**, **`begin assumes`**

Semantics

The *assumption* clause group defines a clause set that 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:

- Each *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: `contains`	**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 **`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: `holds`	**Properties:** unique

Arguments

Name	Type	Properties
hold-expr	expression of OpenMP logical 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

When the **`holds`** clause appears on an assumption directive, the program guarantees that the listed expression evaluates to *true* in the assumption scope. The effect of the clause does not include any evaluation of the expression that affects the behavior of the program.

Cross References

- **assume** Directive, see Section 16.6.3
- **assumes** Directive, see Section 16.6.2
- **begin assumes** Directive, see Section 16.6.4

16.6.1.4 no_openmp Clause

Name: no_openmp	**Properties:** unique

Arguments

Name	Type	Properties
can_assume	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

assume, **assumes**, **begin assumes**

Semantics

If *can_assume* evaluates to *true*, the **no_openmp** clause implies the **no_openmp_constructs** clause and the **no_openmp_routines** clause. If *can_assume* is not specified, the effect is as if *can_assume* evaluates to *true*.

C++

The **no_openmp** clause also guarantees that no thread will throw an exception in the assumption scope if it is contained in a region that arises from an exception-aborting directive.

C++

Cross References

- **assume** Directive, see Section 16.6.3
- **assumes** Directive, see Section 16.6.2
- **begin assumes** Directive, see Section 16.6.4

16.6.1.5 no_openmp_constructs Clause

Name: no_openmp_constructs	**Properties:** unique

Arguments

Name	Type	Properties
can_assume	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

`assume`, **`assumes`**, **`begin assumes`**

Semantics

If *can_assume* evaluates to *true*, the **`no_openmp_constructs`** clause guarantees that no constructs are encountered in the assumption scope. 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_openmp_routines` Clause

Name: `no_openmp_routines`	**Properties:** unique

Arguments

Name	Type	Properties
can_assume	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

`assume`, **`assumes`**, **`begin assumes`**

Semantics

If *can_assume* evaluates to *true*, the **`no_openmp_routines`** clause guarantees that no OpenMP API routines are executed in the assumption scope. 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

16.6.1.7 no_parallelism Clause

Name: no_parallelism	Properties: unique

Arguments

Name	Type	Properties
can_assume	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

assume, **assumes**, **begin assumes**

Semantics

If *can_assume* evaluates to *true*, the **no_parallelism** clause guarantees that no parallelism-generating constructs will be 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

16.6.2 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

Clause groups

assumption

Semantics

The assumption scope of the **assume** directive is the corresponding region and any nested region of that region.

C / C++

16.6.4 `begin assumes` Directive

Name: `begin assumes` **Category:** informational	**Association:** delimited **Properties:** *default*

Clause groups

assumption

Semantics

The assumption scope of the **`begin assumes`** directive is the code that is executed and reached from any of the declared functions in the delimited code region. The delimited code region is a declaration sequence.

C / C++

16.7 `nothing` Directive

Name: `nothing` **Category:** utility	**Association:** unassociated **Properties:** pure, loop-transforming

Clauses

`apply`

Loop Modifiers for the `apply` Clause

loop-modifier	Number of Generated Loops	Description
`identity` *(default)*	1	the copy of the transformation-affected loop

Semantics

The **`nothing`** directive has no effect on the execution of the OpenMP program unless otherwise specified by the **`apply`** clause.

If the **`nothing`** directive immediately precedes a canonical loop nest then it forms a loop-transforming construct. It is associated with the outermost loop and generates one loop that 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 generates a loop nest or loop sequence for which each output loop can be a semantic loop of a semantic loop nest or semantic loop sequence (see Section 6.4.1 and Section 6.4.2).

Logical iterations of output loops that are not affected loops of loop-nest-associated constructs or loop-sequence-associated constructs execute sequentially in lexicographic order.

Loop-iteration variables of output 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:

- An output loop of a loop-transforming construct must not be a doacross-affected loop.
- 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.
- An output 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
- Canonical Loop Sequence Form, see Section 6.4.2

17.1 `apply` Clause

Name: `apply`	**Properties:** *default*

Arguments

Name	Type	Properties
applied-directives	list of directive-specification list item type	*default*

Modifiers

Name	Modifies	Type	Properties
loop-modifier	*applied-directives*	Complex, Keyword: **`flattened`**, **`fused`**, **`grid`**, **`identity`**, **`interchanged`**, **`intratile`**, **`offsets`**, **`reversed`**, **`split`**, **`unrolled`** Arguments: ***indices*** list of expression of integer type (optional)	optional
directive-name-modifier	*all arguments*	Keyword: *directive-name* (a directive name)	unique

Directives

`flatten`, **`fuse`**, **`interchange`**, **`nothing`**, **`reverse`**, **`split`**, **`stripe`**, **`tile`**, **`unroll`**

Semantics

The **`apply`** clause applies loop-nest-associated constructs, specified by the *applied-directives* list, to output loops of a loop-transforming construct. The *loop-modifier* specifies to which output loops the directives are applied. If the loop-transforming construct results in a semantic loop sequence, the output loops to which the directives are applied are the outermost loops of each output 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 output 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 $1, 2, \ldots, m$ is specified, where m is the number of output 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 $1, 2, \ldots, m$ 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 output 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

- **flatten** Construct, see Section 17.4
- **fuse** Construct, see Section 17.5
- **interchange** Construct, see Section 17.6
- **metadirective**, see Section 15.4.3
- **nothing** Directive, see Section 16.7
- **reverse** Construct, see Section 17.7
- **split** Construct, see Section 17.8
- **stripe** Construct, see Section 17.9
- **tile** Construct, see Section 17.10
- **unroll** Construct, see Section 17.11

17.2 shift Clause

Name: shift	Properties: unique

Arguments

Name	Type	Properties
shift-list	list of expression of integer type	non-negative, optional

Modifiers

Name	Modifies	Type	Properties
directive-name-modifier	*all arguments*	Keyword: *directive-name* (a directive name)	unique

Directives

`tile`

Semantics

For a given loop-transforming directive on which the clause appears, the **`shift`** clause modifies the manner in which the **`sizes`** clause subdivides the m-dimensional grid cells. Specifically, each list item in *shift-list* specifies an offset to the alignment of the cell along the corresponding dimension.

If the **`shift`** clause is used without a *shift-list* argument then the values of *shift-list* are implementation defined.

Restrictions

Restrictions to the **`shift`** clause are as follows:

- The number of list items in *shift-list* must be the same as the number of list items in the *size-list* argument of the **`sizes`** clause that appears on the same directive.
- The *sizes-selector* of the **`sizes`** clause must be **`tile`**.

Cross References

- **`tile`** Construct, see Section 17.10

17.3 `sizes` Clause

Name: `sizes`	**Properties:** unique, required

Arguments

Name	Type	Properties
size-list	list of expression of integer type	positive

Modifiers

Name	Modifies	Type	Properties
sizes-selector	*size-list*	Keyword: **`grid`**, **`tile`**	*default*
directive-name-modifier	*all arguments*	Keyword: *directive-name* (a directive name)	unique

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 semantic 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 17.9
- **tile** Construct, see Section 17.10

17.4 `flatten` Construct

Name: `flatten` **Category:** executable	**Association:** loop nest **Properties:** generally-composable, loop-transforming, order-concurrent-nestable, pure, simdizable, teams-nestable

Clauses

apply, **depth**

Loop Modifiers for the `apply` Clause

loop-modifier	Number of Generated Loops	Description
flattened *(default)*	1	the flattened loop

Semantics

The **flatten** construct combines the loops of a semantic loop nest into a single loop, the flattened loop. The number of flatten-affected loops is specified by the **depth** clause. If the **depth** clause is not present, the effect is as if a **depth** clause with a *depth-expr* equal to two was specified.

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

17.5 `fuse` Construct

Name: `fuse` **Category:** executable	**Association:** loop sequence **Properties:** loop-transforming, order-concurrent-nestable, pure, simdizable, teams-nestable

Clauses

apply, **depth**, **looprange**

Loop Modifiers for the `apply` Clause

loop-modifier	Number of Generated Loops	Description
fused *(default)*	1	the fused loop

Semantics

The **fuse** construct merges the affected loop nests specified by the **looprange** clause into a single semantic loop nest where execution of each logical iteration of the output 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 semantic loop sequence that consists of a semantic loop nest with at least *depth-expr* output loops. The affected loops of the affected loop nests in the semantic loop sequence are fused from outermost to innermost. At each level of fusion the output loop has the maximum number of logical iterations of any of the affected loops at that level. The i^{th} logical iteration of the output 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

17.6 `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` *(default)*	n	the output loops, in the new order

Semantics

The **`interchange`** construct has n transformation-affected loops, where $s_1, \ldots, s_n$ are the n items in the *permutation-list* argument of the **`permutation`** clause. Let $\ell_1, \ldots, \ell_n$ be the transformation-affected loops, from outermost to innermost. The output loop nest consists of the loops in the order $\ell_{s_1}, \ldots, \ell_{s_n}$. 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 17.6.1

17.6.1 permutation Clause

Name: permutation	**Properties:** unique

Arguments

Name	Type	Properties
permutation-list	list of expression of integer type	constant, positive

Modifiers

Name	Modifies	Type	Properties
directive-name-modifier	*all arguments*	Keyword: *directive-name* (a directive name)	unique

Directives

interchange

Semantics

The **permutation** clause specifies a list of n positive constant expressions of integer OpenMP type.

Restrictions

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

- Every integer from 1 to n must appear exactly once in *permutation-list*.
- n must be at least 2.

Cross References

- **interchange** Construct, see Section 17.6

17.7 `reverse` Construct

Name: `reverse` **Category:** executable	**Association:** loop nest **Properties:** generally-composable, loop-transforming, order-concurrent-nestable, pure, simdizable, teams-nestable

Clauses

`apply`

Loop Modifiers for the `apply` Clause

loop-modifier	Number of Generated Loops	Description
`reversed` *(default)*	1	the reversed loop

Semantics

The `reverse` construct has one transformation-affected loop, the outermost loop, where $0, 1, \ldots, n-2, n-1$ are the logical iteration numbers of that loop. The construct transforms that loop into a loop in which iterations occur in the order $n-1, n-2, \ldots, 1, 0$.

Cross References

- `apply` Clause, see Section 17.1

17.8 `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

loop-modifier	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 semantic 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 $c_1, \ldots, c_m$ 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 output loop in the sequence contains a copy of the loop body of the affected loop. The i^{th} output 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 17.8.1

17.8.1 counts Clause

Name: counts	**Properties:** unique, required

Arguments

Name	Type	Properties
count-list	list of expression of integer type	non-negative

Modifiers

Name	Modifies	Type	Properties
directive-name-modifier	*all arguments*	Keyword: *directive-name* (a directive name)	unique

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 an output loop of the directive. Specifically, each list item in *count-list* specifies the iteration count of one of the output loops. List items in *count-list* are not required to be unique.

Cross References

- **split** Construct, see Section 17.8

17.9 `stripe` Construct

Name: `stripe`	**Association:** loop nest
Category: executable	**Properties:** loop-transforming, order-concurrent-nestable, pure, simdizable, teams-nestable

Clauses

`apply`, **`sizes`**

Loop Modifiers for the `apply` Clause

loop-modifier	Number of Generated Loops	Description
`offsets`	m	the offsetting loops $o_1, \ldots, o_m$
`grid`	m	the grid loops $g_1, \ldots, g_m$

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 $s_1, \ldots, s_m$. 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 $\ell_1, \ldots, \ell_m$ be the transformation-affected loops, from outermost to innermost. The **`stripe`** construct results in a semantic loop nest that consists of $2m$ perfectly nested loops. Let $o_1, \ldots, o_m, g_1, \ldots, g_m$ be the output loops, from outermost to innermost. The loops $o_1, \ldots, o_m$ are the offsetting loops and the loops $g_1, \ldots, g_m$ are the grid loops.

Let $n_1, \ldots, n_m$ be number of logical iterations of each affected loop and $O = \{G_{\alpha_1,\ldots,\alpha_m} \mid \forall k \in \{1, \ldots, m\} : 0 \leq \alpha_1 < s_k\}$ the logical iteration vector space of the offsetting loops. The logical iteration $(i_1, \ldots, i_m)$ is executed in the logical iteration space of $G_{i_1 \mod s_1,\ldots,i_m \mod s_m}$.

The offsetting loops iterate over all $G_{\alpha_1,\ldots,\alpha_m}$ 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 output loops 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 17.3

17.10 `tile` Construct

Name: `tile` **Category:** executable	**Association:** loop nest **Properties:** loop-transforming, order-concurrent-nestable, pure, simdizable, teams-nestable

Clauses

`apply`, **`shift`**, **`sizes`**

Loop Modifiers for the `apply` Clause

loop-modifier	Number of Generated Loops	Description
`grid`	m	the grid loops $g_1, \ldots, g_m$
`intratile`	m	the tile loops $t_1, \ldots, t_m$

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 $s_1, \ldots, s_m$. Let $\ell_1, \ldots, \ell_m$ be the transformation-affected loops, from outermost to innermost. The **`tile`** construct generates a semantic loop nest that consists of $2m$ perfectly nested loops. Let $f_1, \ldots, f_m$ be the values of the list items of the **`shift`** clause. Let $g_1, \ldots, g_m, t_1, \ldots, t_m$ be the output loops, from outermost to innermost. The loops $g_1, \ldots, g_m$ are the grid loops and the loops $t_1, \ldots, t_m$ are the tile loops.

Let $\mathbb{N}$ be the set of non-negative integers, Ω be the logical iteration vector space of the transformation-affected loops and $G = \{T_{\alpha_1,\ldots,\alpha_m} \mid T_{\alpha_1,\ldots,\alpha_m} \neq \emptyset, (\alpha_1, \ldots, \alpha_m) \in \mathbb{N}^m\}$ 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 $T_{\alpha_1,\ldots,\alpha_m} = \{(i_1, \ldots, i_m) \in \Omega \mid \forall k \in \{1, \ldots, m\} : s_k\alpha_k \leq i_k + (f_k \mod s_k) < s_k\alpha_k + s_k\}$. 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 $T_{\alpha_1,\ldots,\alpha_m} = \{(i_1, \ldots, i_m) \in \Omega \mid \forall k \in \{1, \ldots, m\} : u_k^{\alpha_k} \leq i_k < u_k^{\alpha_k+1}\}$ where for each $k \in \{1, \ldots, m\}$, $0 = u_k^0, u_k^1, \ldots, u_k^{s_k} = n_k$ are implementation defined sequences where $\lfloor n_k/s_k \rfloor \leq u_k^{\alpha_k+1} - u_k^{\alpha_k} \leq \lceil n_k/s_k \rceil$ for any α_k.

The grid loops iterate over all tiles $\{T_{\alpha_1,\ldots,\alpha_m} \in G\}$ in lexicographic order with respect to their indices $(\alpha_1, \ldots, \alpha_m)$ and the tile loops iterate over the iterations in $T_{\alpha_1,\ldots,\alpha_m}$ 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 output loops 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 $u_k^0, u_k^1, \ldots, u_k^{s_k}$ must conform to the consistent schedules requirement.

Cross References

- **`apply`** Clause, see Section 17.1
- Consistent Loop Schedules, see Section 6.4.4
- **`shift`** Clause, see Section 17.2
- **`sizes`** Clause, see Section 17.3

17.11 `unroll` Construct

Name: `unroll` **Category:** executable	**Association:** loop nest **Properties:** generally-composable, loop-transforming, order-concurrent-nestable, pure, simdizable, teams-nestable

Clauses

`apply`, **`full`**, **`partial`**

Clause set

Properties: exclusive	**Members:** **`full`**, **`partial`**

Loop Modifiers for the `apply` Clause

loop-modifier	Number of Generated Loops	Description
`unrolled` *(default)*	1	the grid loop g_1 of the tiling step in the definition of the **`partial`** clause

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 semantic loop nest with one output loop 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 17.11.1
- **partial** Clause, see Section 17.11.2

17.11.1 full Clause

Name: full	**Properties:** unique

Arguments

Name	Type	Properties
fully_unroll	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

unroll

Semantics

If *fully_unroll* evaluates to *true*, the **full** clause specifies that the transformation-affected loop is *fully unrolled*. The construct generates 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 17.11

17.11.2 partial Clause

Name: partial	**Properties:** unique

Arguments

Name	Type	Properties
unroll-factor	expression of integer type	optional, constant, positive

Modifiers

Name	Modifies	Type	Properties
directive-name-modifier	*all arguments*	Keyword: *directive-name* (a directive name)	unique

Directives

unroll

Semantics

The **partial** clause specifies a two-step transformation. First, in the tiling step, the transformation-affected loop is tiled with a tile size of *unroll-factor*. Then, in the unroll step, the resulting 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 17.11

18 Parallelism Generation and Control

This chapter defines constructs for generating and controlling parallelism.

18.1 `parallel` Construct

Name: `parallel` **Category:** executable	**Association:** block **Properties:** cancellable, cancellation-point, context-matching, implicit-acquire-flush-at-entry, implicit-acquire-flush, implicit-barrier, implicit-barrier-flushes, implicit-flush, implicit-release-flush, implicit-release-flush-at-entry, implicit-release-flush-at-exit, implicit-task-generating, order-concurrent-nestable, parallelism-generating, task-generating, task-synchronizing, team-generating, teams-nestable, thread-limiting, threads-orderable

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, configured with a specific team space. 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. An implicit release-at-entry flush occurs in the encountering task before it generates any implicit tasks. This implicit flush synchronizes with the implicit acquire-at-entry flush that occurs in each of the implicit tasks. 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 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* **& ompt_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.

Restrictions

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

- A **parallel** construct may not be encountered if the *levels-var* ICV is greater than zero and the *thread-limit-var* ICV specifies more than one dimension.

Cross References

- **allocate** Clause, see Section 13.6
- **copyin** Clause, see Section 10.1
- **default** Clause, see Section 7.4
- **firstprivate** Clause, see Section 7.3.3
- **if** Clause, see Section 5.6
- **implicit_task** Callback, see Section 41.5.3
- **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 41.3.1
- **parallel_end** Callback, see Section 41.3.2
- OMPT **parallel_flag** Type, see Section 40.22
- **private** Clause, see Section 7.3.2
- **proc_bind** Clause, see Section 18.1.4
- **reduction** Clause, see Section 8.10
- **safesync** Clause, see Section 18.1.5
- OMPT **scope_endpoint** Type, see Section 40.27
- **severity** Clause, see Section 16.4
- **shared** Clause, see Section 7.3.1
- OMPT **task_flag** Type, see Section 40.37
- **thread_begin** Callback, see Section 41.1.3
- **thread_end** Callback, see Section 41.1.4

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 in the team space is determined according to Algorithm 18.1. The result of Algorithm 18.1 is a spatial list that specifies the number of threads for each dimension.

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:** space-configuring, unique

Arguments

Name	Type	Properties
nthreads-list	list of expression of integer type	positive

Modifiers

Name	Modifies	Type	Properties
dims-modifier	*all arguments*	Complex, name: **`dims`** Arguments: ***dimensionality*** expression of integer type (constant, positive)	unique
prescriptiveness-modifier	*nthreads-list*	Keyword: **`relaxed`**, **`strict`**	*default*
directive-name-modifier	*all arguments*	Keyword: *directive-name* (a directive name)	unique

Directives

`parallel`

Semantics

The **`num_threads`** clause specifies the requested team space of the team formed to execute a **`parallel`** region. The *dims-modifier* specifies the dimensions of the team space. If the *dims-modifier* is not specified, the spatial dimensionality of the team space is configured to one and the *nthreads-list* specifies a list with the requested number of threads for the encountered **`parallel`** region and for any nested **`parallel`** regions that are encountered during its execution. In this case, the requested team space is the spatial list that consists of the first list item

Algorithm 18.1 Determine Number of Threads

let *ThreadsBusy* = number of threads currently executing tasks in this contention group;

let *StructuredThreadsBusy* = number of structured threads currently executing tasks in this contention group;

if an **`if`** clause is specified **then let** *IfClauseValue* = value of *if-expression*;

else let *IfClauseValue* = *true*;

if a **`num_threads`** clause with *dims-modifier* is specified **then**

  let *ThreadsRequested* = *nthreads-list*;

else if a **`num_threads`** clause is specified **then**

  let *ThreadsRequested* = single-item list containing the first item of *nthreads-list*;

else let *ThreadsRequested* = single-item list containing the first element of *nthreads-var*;

let *ThreadsAvailable* = $\min($*thread-limit-var* $-$ *ThreadsBusy*,
structured-thread-limit-var $-$ *StructuredThreadsBusy*$) + 1$;

let *TotalThreadsRequested* = product of all items in *ThreadsRequested*;

let *ThreadLimit* = multidimensional value of *thread-limit-var*;

let *N* = number of items in *ThreadsRequested*;

if (*IfClauseValue* = *false*) **or** (*active-levels-var* $\geq$ *max-active-levels-var*) **then**

  return list of *N* items all equal to 1;

else if (*dyn-var* = *true*) **then**

  return an implementation defined list L with N items that satisfies:

  (a) for each item position D, $1 \leq L[D] \leq \min($*ThreadsRequested*$[D]$, *ThreadLimit*$[D])$; and

  (b) the product of all items in L is $\leq \min($*TotalThreadsRequested*, *ThreadsAvailable*$)$;

else if (*dyn-var* = *false*) **and** (*TotalThreadsRequested* $\leq$ *ThreadsAvailable*) **then**

  return list of N items, where each item, in position D, is
$\min($*ThreadsRequested*$[D]$, *ThreadLimit*$[D])$;

else if (*dyn-var* = *false*) **and** (*TotalThreadsRequested* $>$ *ThreadsAvailable*) **then**

  behavior is implementation defined;

of the *nthreads-list*. Otherwise, if the *dims-modifier* is specified, *nthreads-list* specifies the requested number of threads in each dimension for the encountered **parallel** region, and thus, the requested team space is the spatial list that consists of all of its list items. If the **num_threads** clause is not specified, the spatial dimensionality of the team space is configured to the number of dimensions specified by the *thread-limit-var* ICV.

Algorithm 18.1 determines the spatial list that specifies the number of threads for each dimension in the team space of the team that executes the **parallel** region. If the *prescriptiveness-modifier* is specified as **strict**, an **if** clause is either not present or *if-expression* evaluates to *true*, and an implementation determines that Algorithm 18.1 would result in a spatial list other than the requested team space then error termination is performed. If this is always the case, compile-time error termination may be performed and the effect of any **message** clause associated with the directive is implementation defined; otherwise, 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. Otherwise, if the *prescriptiveness-modifier* is specified as **relaxed**, the implementation uses the number of threads that Algorithm 18.1 determines. If the *prescriptiveness-modifier* is not specified, the effect is as if **strict** is specified when the *dims-modifier* is present, and as if **relaxed** is specified when the *dims-modifier* is not present.

Restrictions

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

- If the *dims-modifier* is specified, its *dimensionsionality* argument must not exceed the number of dimensions specified by the *thread-limit-var* ICV.

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 group, 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 group for all thread affinity policies is given by the

place-group-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 group for all thread affinity policies is the input place group of that parallel region. If the encountering thread is not a free-agent thread, the input place group for all thread affinity policies is given by the *place-group-var* ICV of its binding implicit task.

For the description of the thread affinity policies that follow, let T be the number of threads in the team to be assigned to a place, and let P be the number of places in the input place partition.

The **primary** thread affinity policy results in the execution environment assigning every thread in the team to the same place as the primary thread. The **close** thread affinity policy results in the execution environment assigning the threads in the team to places close to the place of the parent thread. Under the **primary** and **close** thread affinity policies, the *place-group-var* ICV of each implicit task is assigned the input place group.

The **spread** thread affinity policy results in the execution environment assigning the threads in the team in a sparse manner across the input place partition. The input place partition is subdivided into subpartitions, according to a traversal of the input place group (as described in the following paragraph), and each thread in the team and its implicit task is then assigned to one of these resulting subpartitions. The *place-group-var* ICV of each implicit task is then assigned the place group for its subpartition.

For the **spread** policy, let G be the number of place-group children of the input place group. If $T \leq G$, the list of place-group children are subdivided into T sets of $\lceil G/T \rceil$ or $\lfloor G/T \rfloor$ consecutive place-group children. The corresponding place group of each set is the left-to-right concatenation of members of the set that are places or the place-group children of members of the set that are place groups. Each thread is uniquely assigned to one of these place groups and a subpartition that is the place partition of that place group. If $G < T$, the threads are distributed over the place-group children such that each place-group child receives $\lceil T/G \rceil$ or $\lfloor T/G \rfloor$ consecutive threads in the team. For each place-group child that is a place, its threads are assigned to a subpartition that contains only that place. Otherwise, for each place-group child that receives only one thread, its thread is assigned to the place-group child and a subpartition that is the place partition of that place-group child. Otherwise, for each other place-group child, its assigned threads are distributed over the place-group children of that place group recursively, following the distribution procedure described above for the **spread** policy as if the input place group is replaced with the place group.

The set of team-eligible places in the input place partition for the **primary** thread affinity policy is the place to which the primary thread is or will be assigned and for the **close** thread affinity policy it is all places in the input place partition. For the **spread** thread affinity policy, a place is team-eligible in a given subpartition if the primary thread is already assigned to it upon encountering the parallel region or if it is the first place in its subpartition and the primary thread is not already assigned to a different place in the subpartition. Let Q be the number of team-eligible places. If $T > Q$, the threads that comprise the team are divided into Q sets of $\lceil T/Q \rceil$ or $\lfloor T/Q \rfloor$ consecutive threads, and otherwise they are divided into T sets of one thread each. The first set of threads is assigned to the place of the primary thread if that thread is already assigned, and

otherwise the set is assigned to one of the Q team-eligible places. The next set of threads is then assigned to the next team-eligible place, and so on with wrap around.

For the **primary** policy, free-agent threads that begin execution in the parallel region will be assigned the place to which the primary thread is assigned, unless that place is reserved for structured threads. For the **spread** policy when $T > P$ or for the **close** policy, after the T threads are assigned there will be $K = P - T \bmod P$ team-eligible places with one less thread than the rest. If $K < P$, each of the first K free-agent threads that begin execution in the parallel region will be assigned to one of these places that is not reserved for structured threads, in some implementation defined order. Any additional free-agent threads are handled as follows. For the **spread** policy when $T \leq P$, the first free-agent thread goes to the subpartition to which the primary thread is assigned, the next free-agent thread goes to the next subpartition, and so on with wrap around. Within each subpartition, free-agent threads are round-robin assigned to the allotted places that are not reserved for structured threads, with wrap around, starting with the place that follows its team-eligible place if it has more than one place.

If a free-agent thread cannot be assigned a place because all eligible places are reserved for structured threads, a free-agent thread may be assigned a place available in the next enclosing parallel region until an available place is found.

The determination of whether the thread affinity request can be fulfilled is implementation defined. If it cannot be fulfilled, then the affinity of team-worker threads and free-agent threads 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-group-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
affinity-policy	Keyword: **close**, **primary**, **spread**	*default*

Modifiers

Name	Modifies	Type	Properties
directive-name-modifier	*all arguments*	Keyword: *directive-name* (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: safesync	**Properties:** unique

Arguments

Name	Type	Properties
width	expression of integer type	positive, optional

Modifiers

Name	Modifies	Type	Properties
directive-name-modifier	*all arguments*	Keyword: *directive-name* (a directive name)	unique

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.
- If the team space of the team that executes the parallel region associated with the construct in which the clause appears has a spatial dimensionality greater than one, then the *width* argument must be one.

Cross References

- **parallel** Construct, see Section 18.1

18.2 teams Construct

Name: teams **Category:** executable	**Association:** block **Properties:** context-matching, groupprivate-instantiating, implicit-acquire-flush-at-entry, implicit-acquire-flush-at-exit, implicit-acquire-flush, implicit-flush, implicit-release-flush, implicit-release-flush-at-entry, implicit-release-flush-at-exit, implicit-task-generating, league-generating, parallelism-generating, task-generating, team-generating, teams-orderable, thread-limiting

Clauses

allocate, **default**, **dyn_groupprivate**, **firstprivate**, **if**, **message**, **num_teams**, **private**, **reduction**, **severity**, **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, configured with a specific league space. Each team is an initial team, and the initial thread in each team executes the **teams** region as an implicit task. An implicit release-at-entry flush occurs in the encountering task before it generates any implicit tasks. This implicit flush synchronizes with the implicit acquire-at-entry flush that occurs in each of the implicit tasks. Similarly, after task completion and the associated implicit release-at-exit flush occurs in each of the implicit tasks, an implicit acquire-at-exit flush occurs in the encountering task; each of the implicit release-at-exit flushes in the implicit tasks synchronizes with that acquire flush.

The number of teams created in the league space is determined by evaluating the **if** and **num_teams** clauses. Once the teams are created, the number of initial teams is region-invariant and, thus, does not change for the duration of the **teams** region. Within a **teams** region, initial team numbers uniquely identify each initial team. Initial team 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-group-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-group-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 7.4
- **distribute** Construct, see Section 19.7
- **dyn_groupprivate** Clause, see Section 13.9
- **firstprivate** Clause, see Section 7.3.3

- *default-device-var* ICV, see Table 3.1
- *place-group-var* ICV, see Table 3.1
- **`if`** Clause, see Section 5.6
- **`implicit_task`** Callback, see Section 41.5.3
- **`message`** Clause, see Section 16.3
- **`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 41.3.1
- **`parallel_end`** Callback, see Section 41.3.2
- OMPT **`parallel_flag`** Type, see Section 40.22
- **`private`** Clause, see Section 7.3.2
- **`reduction`** Clause, see Section 8.10
- OMPT **`scope_endpoint`** Type, see Section 40.27
- **`severity`** Clause, see Section 16.4
- **`shared`** Clause, see Section 7.3.1
- **`target`** Construct, see Section 21.8
- OMPT **`task_flag`** Type, see Section 40.37
- **`thread_begin`** Callback, see Section 41.1.3
- **`thread_end`** Callback, see Section 41.1.4
- **`thread_limit`** Clause, see Section 21.3

18.2.1 `num_teams` Clause

Name: `num_teams`	**Properties:** target-consistent, space-configuring, unique

Arguments

Name	Type	Properties
upper-bound-list	list of expression of integer type	positive

Modifiers

Name	Modifies	Type	Properties
dims-modifier	*all arguments*	Complex, name: **dims** Arguments: ***dimensionality*** expression of integer type (constant, positive)	unique
lower-bound	*upper-bound-list*	expression of integer type	positive, ultimate, unique
directive-name-modifier	*all arguments*	Keyword: *directive-name* (a directive name)	unique

Modifier set

Properties: exclusive	**Members:** *dims-modifier*, *lower-bound*

Directives

teams

Semantics

The **num_teams** clause specifies the bounds on the league space of the league formed by the construct on which it appears. The *dims-modifier* specifies the spatial dimensionality of the league space. The *upper-bound-list* argument is a spatial list that specifies the upper bound on the number of teams requested in each dimension. If the *dims-modifier* is not specified, the spatial dimensionality of the league space is configured to one, and the *lower-bound* modifier may be used to specify the lower 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 the upper bound in that dimension. 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 in each dimension.

If the *dims-modifier* is specified and the implementation cannot satisfy an upper bound in any of the dimensions specified in *upper-bound-list* then runtime error termination is performed. In this error termination scenario, the effect is as if an **error** directive has been encountered on which any provided **message** and **severity** clauses and an **at** clause with **execution** as *action-time* are specified. The maximum number of teams that may be specified in a given dimension can be obtained using the **omp_get_supported_teams_dim** routine.

If the **num_teams** clause is not specified on a construct then the effect is as if the clause was specified without the *dims-modifier* and the upper bound was specified as follows. If the value of the *nteams-var* ICV is greater than zero, the effect is as if the upper bound was specified as an implementation defined value greater than zero but less than or equal to the value of the *nteams-var* ICV. Otherwise, the effect is as if the 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 the value that corresponds to the list item in the *upper-bound-list*.

Cross References

- *nteams-var* ICV, see Table 3.1
- **teams** Construct, see Section 18.2

18.3 *dims-modifier* Modifier

Modifiers

Name	Modifies	Type	Properties
dims-modifier	*all arguments*	Complex, name: **dims** Arguments: ***dimensionality*** expression of integer type (constant, positive)	unique

Clauses

num_teams, **num_threads**, **thread_limit**

Semantics

The *dims-modifier* specifies the number of dimensions for the execution space that the space-configuring clause configures or limits. The *dimensionality* argument in the *dims-modifier* specifies the number of dimensions and determines the expected length of the spatial list argument in the clause in which it appears. The execution space is configured with a spatial dimensionality equal to the value of the *dimensionality* argument. If the *dims-modifier* is not specified then the execution space is configured with a spatial dimensionality of one.

Restrictions

- The number of list items in the list argument of the space-configuring clause in which the *dims-modifier* appears must be equal to the value specified in the *dimensionality* argument.

Cross References

- **num_teams** Clause, see Section 18.2.1
- **num_threads** Clause, see Section 18.1.2
- **thread_limit** Clause, see Section 21.3

18.4 `order` Clause

Name: `order`	**Properties:** schedule-specification, unique

Arguments

Name	Type	Properties
ordering	Keyword: **`concurrent`**	*default*

Modifiers

Name	Modifies	Type	Properties
order-modifier	*ordering*	Keyword: **`reproducible`**, **`unconstrained`**	*default*
directive-name-modifier	*all arguments*	Keyword: *directive-name* (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.

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 18.5

18.5 simd Construct

Name: **simd** **Category:** executable	**Association:** loop nest **Properties:** context-matching, sequentially-orderable, simdizable, order-concurrent-nestable, parallelism-generating, pure

Separating directives

scan

Clauses

aligned, **collapse**, **if**, **induction**, **lastprivate**, **linear**, **nontemporal**, **order**, **private**, **reduction**, **safelen**, **simdlen**

Binding

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

Semantics

The **simd** construct enables the execution of multiple collapsed iterations concurrently by using SIMD instructions. The number of collapsed iterations that are executed concurrently at any given time is implementation defined. Each concurrent iteration will be executed by a different SIMD lane. Each set of concurrent iterations is a SIMD chunk. Lexical forward dependences in the iterations of the original loop must be preserved within each SIMD chunk, unless an **order** clause that specifies **concurrent** is present.

When an **if** clause is present with an *if-expression* that evaluates to *false*, the preferred number of iterations to be executed concurrently is one, regardless of whether a **simdlen** clause is specified.

Restrictions

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

- If both **simdlen** and **safelen** clauses are specified, the value of the **simdlen** *length* must be less than or equal to the value of the **safelen** *length*.
- Only SIMDizable constructs may be encountered during execution of a **simd** region.

- If an **order** clause that specifies **concurrent** appears on a **simd** directive, the **safelen** clause must not also appear.

C / C++

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

C / C++

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 7.3.4
- **linear** Clause, see Section 8.14
- **nontemporal** Clause, see Section 18.5.1
- **order** Clause, see Section 18.4
- **private** Clause, see Section 7.3.2
- **reduction** Clause, see Section 8.10
- **safelen** Clause, see Section 18.5.2
- **scan** Directive, see Section 8.17
- **simdlen** Clause, see Section 18.5.3

18.5.1 nontemporal Clause

Name: nontemporal	**Properties:** *default*

Arguments

Name	Type	Properties
list	list of variable list item type	*default*

Modifiers

Name	Modifies	Type	Properties
directive-name-modifier	*all arguments*	Keyword: *directive-name* (a directive name)	unique

Directives

simd

Semantics

The **nontemporal** clause specifies that accesses to the storage locations to which the list items refer have low temporal locality across the 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 18.5

18.5.2 `safelen` Clause

Name: safelen	**Properties:** unique

Arguments

Name	Type	Properties
length	expression of integer type	positive, constant

Modifiers

Name	Modifies	Type	Properties
directive-name-modifier	*all arguments*	Keyword: *directive-name* (a directive name)	unique

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 18.5

18.5.3 `simdlen` Clause

Name: simdlen	**Properties:** unique

Arguments

Name	Type	Properties
length	expression of integer type	positive, constant

Modifiers

Name	Modifies	Type	Properties
scaled-modifier	*length*	Complex, name: **scaled** Arguments: ***type*** type name (*default*) ***divisor*** expression of integer type (constant, positive, optional)	*default*
directive-name-modifier	*all arguments*	Keyword: *directive-name* (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 *scaled-modifier*, *type* must be an arithmetic type.

C / C++

Fortran

- In a *scaled-modifier*, *type* must be a numeric intrinsic type.

Fortran

Cross References

- **declare_simd** Directive, see Section 15.7
- **simd** Construct, see Section 18.5

18.6 masked Construct

Name: masked **Category:** executable	**Association:** block **Properties:** thread-limiting, thread-selecting

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 18.6.1
- **masked** Callback, see Section 41.3.3
- OMPT **scope_endpoint** Type, see Section 40.27

18.6.1 `filter` Clause

Name: `filter`	**Properties:** unique

Arguments

Name	Type	Properties
thread_num	expression of integer type	*default*

Modifiers

Name	Modifies	Type	Properties
directive-name-modifier	*all arguments*	Keyword: *directive-name* (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 18.6

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.

Fortran

- A variable must not be private within a **teams** or **parallel** region if it has either **LOCAL_INIT** or **SHARED** locality in a **DO CONCURRENT** loop that is associated with a work-distribution construct, where the **teams** or **parallel** region is a binding region of the corresponding work-distribution region.

Fortran

19.1 `single` Construct

Name: `single` **Category:** executable	**Association:** block **Properties:** cancellation-point, implicit-barrier, implicit-barrier-flushes, implicit-flush, partitioned, task-synchronizing, team-executed, thread-limiting, thread-selecting, work-distribution, worksharing

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.

Tool Callbacks

A thread dispatches a registered **`work`** callback with **`ompt_scope_begin`** as its *endpoint* argument for each occurrence of a *single-begin* event in that thread. Similarly, a thread dispatches a registered **`work`** callback with **`ompt_scope_end`** as its *endpoint* argument for each occurrence of a *single-end* event in that thread. For each of these callbacks, the *work_type* argument is **`ompt_work_single_executor`** if the thread executes the structured block associated with the **`single`** region; otherwise, the *work_type* argument is **`ompt_work_single_other`**.

Cross References

- **allocate** Clause, see Section 13.6
- **copyprivate** Clause, see Section 10.2
- **firstprivate** Clause, see Section 7.3.3
- **nowait** Clause, see Section 23.6
- **private** Clause, see Section 7.3.2
- OMPT **scope_endpoint** Type, see Section 40.27
- **work** Callback, see Section 41.4.1
- OMPT **work** Type, see Section 40.41

19.2 scope Construct

Name: **scope** **Category:** executable	**Association:** block **Properties:** cancellation-point, implicit-barrier, implicit-barrier-flushes, implicit-flush, task-synchronizing, team-executed, thread-limiting, threads-orderable, work-distribution, worksharing

Clauses

allocate, **firstprivate**, **nowait**, **private**, **reduction**

Binding

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

Semantics

The **scope** construct specifies that all threads in a team execute the associated structured block and any additionally specified OpenMP operations. An implicit barrier occurs at the end of a **scope** region if the **nowait** clause does not specify otherwise.

Execution Model Events

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

Tool Callbacks

A thread dispatches a registered **work** callback with **ompt_scope_begin** as its *endpoint* argument and **ompt_work_scope** as its *work_type* argument for each occurrence of a *scope-begin* event in that thread. Similarly, a thread dispatches a registered **work** callback with **ompt_scope_end** as its *endpoint* argument and **ompt_work_scope** as its *work_type* argument for each occurrence of a *scope-end* event in that thread. The callbacks occur in the context of the implicit task.

Cross References

- **allocate** Clause, see Section 13.6
- **firstprivate** Clause, see Section 7.3.3
- **nowait** Clause, see Section 23.6
- **private** Clause, see Section 7.3.2
- **reduction** Clause, see Section 8.10
- OMPT **scope_endpoint** Type, see Section 40.27
- **work** Callback, see Section 41.4.1
- OMPT **work** Type, see Section 40.41

19.3 sections Construct

Name: **sections** **Category:** executable	**Association:** block **Properties:** cancellable, cancellation-point, implicit-barrier, implicit-barrier-flushes, implicit-flush, partitioned, sequentially-orderable, task-synchronizing, team-executed, thread-limiting, work-distribution, worksharing

Separating directives

section

Clauses

allocate, **firstprivate**, **lastprivate**, **nowait**, **private**, **reduction**

Binding

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

Semantics

The **sections** construct is a non-iterative worksharing construct that contains a structured block that consists of a set of structured block sequences that are to be distributed among and executed by the threads in a team. Each structured block sequence is executed by one of the threads in the team in the context of its implicit task. An implicit barrier occurs at the end of a **sections** region if the **nowait** clause does not specify otherwise.

Each structured block sequence in the **sections** construct is preceded by a **section** subsidiary directive except possibly the first sequence, for which a preceding **section** subsidiary directive is optional. The method of scheduling the structured block sequences among the threads in the team is implementation defined.

Execution Model Events

The *sections-begin* event occurs after an implicit task encounters a **sections** construct but before the task executes any structured block sequences of the **sections** region. The *sections-end* event occurs after an implicit task finishes execution of a **sections** region but before it resumes execution of the enclosing context.

Tool Callbacks

A thread dispatches a registered **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 7.3.3
- **lastprivate** Clause, see Section 7.3.4
- **nowait** Clause, see Section 23.6
- **private** Clause, see Section 7.3.2
- **reduction** Clause, see Section 8.10
- OMPT **scope_endpoint** Type, see Section 40.27
- **section** Directive, see Section 19.3.1
- **work** Callback, see Section 41.4.1
- OMPT **work** Type, see Section 40.41

19.3.1 `section` Directive

Name: `section`	Association: separating
Category: subsidiary	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 41.4.2
- **`sections`** Construct, see Section 19.3

Fortran

19.4 `workshare` Construct

Name: `workshare`	Association: block
Category: executable	Properties: cancellation-point, implicit-barrier, implicit-barrier-flushes, implicit-flush, partitioned, task-synchronizing, team-executed, work-distribution, 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 within the associated structured block. This enforcement includes that 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** constructs or statements, evaluation of the mask expression and the masked assignments are each a unit of work.
- For **FORALL** constructs or statements, 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 procedure calls unless either the procedure is pure and elemental or the procedure 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 40.27
- **work** Callback, see Section 41.4.1
- OMPT **work** Type, see Section 40.41

Fortran

Fortran

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 within the associated structured block. This enforcement includes that 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 a procedure that is not pure and elemental, unless it is one of the transformational array intrinsic functions explicitly permitted by the construct.
- 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 40.27
- **target** Construct, see Section 21.8
- **teams** Construct, see Section 18.2
- **work** Callback, see Section 41.4.1
- OMPT **work** Type, see Section 40.41

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, the declaration specifier sequence 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*.

Worksharing-loop constructs have the implicit-barrier property. They also have the cancellation-point property, which occurs as part of the implicit barrier. However, even if the implicit barrier will not occur due to the use of the **`nowait`** clause, the cancellation point still occurs at exit from the construct if the same list item appears in both **`firstprivate`** and **`lastprivate`** clauses.

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 *work_type*	If determined schedule is
`ompt_work_loop`	unknown at runtime
`ompt_work_loop_static`	**`static`**
`ompt_work_loop_dynamic`	**`dynamic`**
`ompt_work_loop_guided`	**`guided`**
`ompt_work_loop_other`	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 41.4.2
- *run-sched-var* ICV, see Table 3.1
- **`nowait`** Clause, see Section 23.6
- **`order`** Clause, see Section 18.4
- **`schedule`** Clause, see Section 19.6.3
- OMPT **`scope_endpoint`** Type, see Section 40.27
- **`work`** Callback, see Section 41.4.1
- OMPT **`work`** Type, see Section 40.41

C / C++

19.6.1 `for` Construct

Name: **`for`** **Category:** executable	**Association:** loop nest **Properties:** cancellable, cancellation-point, context-matching, implicit-barrier, implicit-barrier-flushes, implicit-flush, partitioned, sequentially-orderable, SIMD-partitionable, task-synchronizing, team-executed, work-distribution, worksharing, worksharing-loop

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 7.3.3
- Worksharing-Loop Constructs, see Section 19.6
- **induction** Clause, see Section 8.13
- **lastprivate** Clause, see Section 7.3.4
- **linear** Clause, see Section 8.14
- **nowait** Clause, see Section 23.6
- **order** Clause, see Section 18.4
- **ordered** Clause, see Section 6.4.6
- **private** Clause, see Section 7.3.2
- **reduction** Clause, see Section 8.10
- **scan** Directive, see Section 8.17
- **schedule** Clause, see Section 19.6.3

C / C++

Fortran

19.6.2 do Construct

Name: do **Category:** executable	**Association:** loop nest **Properties:** cancellable, cancellation-point, context-matching, implicit-barrier, implicit-barrier-flushes, implicit-flush, partitioned, sequentially-orderable, SIMD-partitionable, task-synchronizing, team-executed, work-distribution, worksharing, worksharing-loop

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 7.3.3
- Worksharing-Loop Constructs, see Section 19.6
- **induction** Clause, see Section 8.13
- **lastprivate** Clause, see Section 7.3.4
- **linear** Clause, see Section 8.14
- **nowait** Clause, see Section 23.6
- **order** Clause, see Section 18.4
- **ordered** Clause, see Section 6.4.6
- **private** Clause, see Section 7.3.2
- **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: schedule	**Properties:** schedule-specification, unique

Arguments

Name	Type	Properties
kind	Keyword: **auto**, **dynamic**, **guided**, **runtime**, **static**	*default*
chunk_size	expression of integer type	ultimate, optional, positive, region-invariant

Modifiers

Name	Modifies	Type	Properties
ordering-modifier	*kind*	Keyword: **monotonic**, **nonmonotonic**	unique
chunk-modifier	*kind*	Keyword: **simd**	unique
directive-name-modifier	*all arguments*	Keyword: *directive-name* (a directive name)	unique

Directives

do, **for**

Semantics

The **schedule** clause specifies how collapsed iterations of a worksharing-loop construct are divided into chunks, and how these chunks are distributed among threads of the team.

The *chunk_size* expression is evaluated using the original list items of any variables that are made private variables in the worksharing-loop construct. Whether, in what order, or how many times, any side effects of the evaluation of this expression occur is unspecified. The use of a variable in a **schedule** clause expression of a worksharing-loop construct causes an implicit reference to the variable in all enclosing constructs.

If the *kind* argument is **static**, chunks of increasing collapsed iteration numbers are assigned to the threads of the team in a round-robin fashion in the order of the thread number. Each chunk contains *chunk_size* collapsed iterations, except for the chunk that contains the sequentially last iteration, which contains at most *chunk_size* 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 **simd** *chunk-modifier* is specified and the canonical loop nest is associated with a SIMD construct then, for T threads and n collapsed iterations, *chunk_size* is considered to be either $\lceil n/T \rceil$ or $\lfloor n/T \rfloor$.

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 contains at most *chunk_size* 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 contain at least k collapsed iterations (except for the chunk that contains the sequentially last iteration, which contains at most 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. 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 $0 <= r < p$. 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** **Category:** executable	**Association:** loop nest **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 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.
- Each 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 41.4.2
- **dist_schedule** Clause, see Section 19.7.1
- **firstprivate** Clause, see Section 7.3.3
- Consistent Loop Schedules, see Section 6.4.4
- **induction** Clause, see Section 8.13
- **lastprivate** Clause, see Section 7.3.4
- **order** Clause, see Section 18.4
- **private** Clause, see Section 7.3.2
- OMPT **scope_endpoint** Type, see Section 40.27
- **teams** Construct, see Section 18.2

- **work** Callback, see Section 41.4.1
- OMPT **work** Type, see Section 40.41

19.7.1 dist_schedule Clause

Name: dist_schedule	**Properties:** schedule-specification, unique

Arguments

Name	Type	Properties
kind	Keyword: **static**	*default*
chunk_size	expression of integer type	ultimate, optional, positive, region-invariant

Modifiers

Name	Modifies	Type	Properties
directive-name-modifier	*all arguments*	Keyword: *directive-name* (a directive name)	unique

Directives

distribute

Semantics

The **dist_schedule** clause specifies how collapsed iterations of a **distribute** construct are divided into chunks, and how these chunks are distributed among the teams of the league. If *chunk_size* is not specified, the collapsed iteration space is divided into chunks that are approximately equal in size, and at most one chunk is distributed to each initial team of the league. If the *chunk_size* argument is specified, collapsed iterations are divided into chunks of *chunk_size* iterations. The *chunk_size* expression is evaluated using the original list items of any variables that become private variables in the **distribute** construct. Whether, in what order, or how many times, any side effects of the evaluation of this expression occur is unspecified. The use of a variable in a **dist_schedule** clause expression of a **distribute** construct causes an implicit reference to the variable in all enclosing constructs. These chunks are assigned to the initial teams of the league in a round-robin fashion in the order of their team number.

Restrictions

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

- The value of the *chunk_size* expression must be the same for all teams in the league.
- The **dist_schedule** clause cannot be specified if any of the collapse-affected loops is a non-rectangular loop.

Cross References

- **distribute** Construct, see Section 19.7

19.8 `loop` Construct

Name: `loop` **Category:** executable	**Association:** loop nest **Properties:** order-concurrent-nestable, partitioned, sequentially-orderable, simdizable, team-executed, teams-nestable, work-distribution, worksharing

Clauses

`bind`, **`collapse`**, **`lastprivate`**, **`order`**, **`private`**, **`reduction`**

Binding

The **`bind`** clause determines the binding region, which determines the binding thread set.

Semantics

A **`loop`** construct specifies that the collapsed iterations execute in the context of the binding thread set, in an order specified by the **`order`** clause. If the **`order`** clause is not specified, the behavior is as if the **`order`** clause is present and specifies the **`concurrent`** *ordering*. The collapsed iterations are executed as if by the binding thread set, once per instance of the **`loop`** region that is encountered by the binding thread set.

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:

- Each 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 7.3.4
- **order** Clause, see Section 18.4
- **private** Clause, see Section 7.3.2
- **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
binding	Keyword: **parallel**, **teams**, **thread**	*default*

Modifiers

Name	Modifies	Type	Properties
directive-name-modifier	*all arguments*	Keyword: *directive-name* (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 threads that are executing that region. If the binding region is not defined, then the binding thread set is the encountering thread.

Restrictions

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

- If **teams** is specified as *binding* then the corresponding **loop** region must be a strictly nested region of a **teams** region.
- If **teams** is specified as *binding* and the corresponding **loop** region executes on a non-host device then the behavior of a **reduction** clause without the **partial** *reduction-modifier* that appears on the corresponding **loop** construct is unspecified if the construct is not nested inside a **teams** construct.
- If **parallel** is specified as *binding*, the behavior is unspecified if the corresponding **loop** region is a closely nested region of a **simd** region.

Fortran

- If **teams** is specified as *binding* and the corresponding **loop** region executes on a non-host device then the **REDUCE** locality must not be specified in a **DO CONCURRENT** loop that is affected by the corresponding **loop** directive if the construct is not nested inside a **teams** construct.

Fortran

Cross References

- **loop** Construct, see Section 19.8

20 Tasking Constructs

This chapter defines directives and concepts related to explicit tasks.

20.1 `task` Construct

Name: `task` **Category:** executable	**Association:** block **Properties:** explicit-task-generating, implicit-acquire-flush, implicit-acquire-flush-at-entry, implicit-acquire-flush-at-exit, implicit-flush, implicit-release-flush, implicit-release-flush-at-entry, implicit-release-flush-at-exit, parallelism-generating, task-generating, task-scheduling-point, task-scheduling-point-after-task-generation, task-scheduling-point-at-entry, task-scheduling-point-at-exit, task-synchronizing, thread-limiting

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. 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 undeferrable 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 an explicit-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. The *undeferrable-sync-begin* event occurs in each thread that encounters a explicit-task-generating construct that generates an undeferrable task upon suspending the current task region. The *undeferrable-sync-wait-begin* event occurs when a task is suspended to wait for the completion of an undeferrable task. The *undeferrable-sync-wait-end* event occurs when a task ends the waiting interval and execution of an undeferrable task has completed. The *undeferrable-sync-end* event occurs in each thread that encounters a explicit-task-generating construct that generates an undeferrable task upon resuming the execution of the current task region.

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.

A thread dispatches a registered **sync_region** callback with **ompt_sync_region_undeferrable_task** as its *kind* argument and **ompt_scope_begin** as its *endpoint* argument for each occurrence of an *undeferrable-sync-begin* event. Similarly, a thread dispatches a registered **sync_region** callback with **ompt_sync_region_undeferrable_task** as its *kind* argument and **ompt_scope_end** as its *endpoint* argument for each occurrence of an *undeferrable-sync-end* event. These callbacks occur in the context of the task that encountered the **task** construct that generates the undeferrable task.

A thread dispatches a registered **sync_region_wait** callback with **ompt_sync_region_undeferrable_task** as its *kind* argument and **ompt_scope_begin** as its *endpoint* argument for each occurrence of an *undeferrable-sync-wait-begin* event. Similarly, a thread dispatches a registered **sync_region_wait** callback with **ompt_sync_region_undeferrable_task** as its *kind* argument and **ompt_scope_end** as its *endpoint* argument for each occurrence of an *undeferrable-sync-wait-end* event. These callbacks occur in the context of the task that encountered the **task** construct that generates the undeferrable task.

TABLE 20.1: **task_create** Callback Flags Evaluation

Operation	Evaluates to *true*
(*flags* & **ompt_task_explicit**)	Always in the dispatched callback
(*flags* & **ompt_task_importing**)	If the task is an importing task
(*flags* & **ompt_task_exporting**)	If the task is an exporting task
(*flags* & **ompt_task_undeferred**)	If the task is an undeferred task
(*flags* & **ompt_task_undeferrable**)	If the task is an undeferrable task
(*flags* & **ompt_task_final**)	If the task is a final task
(*flags* & **ompt_task_untied**)	If the task is an untied task
(*flags* & **ompt_task_mergeable**)	If the task is a mergeable task
(*flags* & **ompt_task_merged**)	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 7.4
- **depend** Clause, see Section 23.9.6
- **detach** Clause, see Section 20.11

- **`final`** Clause, see Section 20.7
- **`firstprivate`** Clause, see Section 7.3.3
- Task Scheduling, see Section 20.13
- **`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 30.2.1
- **`priority`** Clause, see Section 20.9
- **`private`** Clause, see Section 7.3.2
- **`replayable`** Clause, see Section 20.6
- **`shared`** Clause, see Section 7.3.1
- **`task_create`** Callback, see Section 41.5.1
- OMPT **`task_flag`** Type, see Section 40.37
- **`threadset`** Clause, see Section 20.8
- **`transparent`** Clause, see Section 23.9.7
- **`untied`** Clause, see Section 20.4

20.2 `taskloop` Construct

Name: **`taskloop`** **Category:** executable	**Association:** loop nest **Properties:** explicit-task-generating, implicit-acquire-flush, implicit-acquire-flush-at-entry, implicit-acquire-flush-at-exit, implicit-flush, implicit-release-flush, implicit-release-flush-at-entry, implicit-release-flush-at-exit, implicit-taskgroup, parallelism-generating, SIMD-partitionable, task-generating, task-scheduling-point, task-scheduling-point-after-task-generation, task-scheduling-point-at-entry, task-scheduling-point-at-exit, task-synchronizing

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.

Since the **taskloop** construct has the implicit-taskgroup property, each **taskloop** region executes as if it is enclosed by a **taskgroup** region if the **nogroup** clause is not present. Thus, the construct implies the taskgroup scheduling constraint.

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*, undeferrable 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 output 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 7.4
- **dispatch** Callback, see Section 41.4.2
- **final** Clause, see Section 20.7
- **firstprivate** Clause, see Section 7.3.3
- 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 7.3.4
- **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 7.3.2
- **reduction** Clause, see Section 8.10
- **replayable** Clause, see Section 20.6
- OMPT **scope_endpoint** Type, see Section 40.27
- **shared** Clause, see Section 7.3.1
- **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 23.9.7
- **untied** Clause, see Section 20.4
- **work** Callback, see Section 41.4.1
- OMPT **work** Type, see Section 40.41

20.2.1 grainsize Clause

Name: grainsize	**Properties:** taskgraph-altering, unique

Arguments

Name	Type	Properties
grain-size	expression of integer type	positive

Modifiers

Name	Modifies	Type	Properties
prescriptiveness	*grain-size*	Keyword: **relaxed**, **strict**	unique
directive-name-modifier	*all arguments*	Keyword: *directive-name* (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 specified as **relaxed**, 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. If *prescriptiveness* is not specified then the effect is as if it is specified as **relaxed**.

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: `num_tasks`	**Properties:** taskgraph-altering, unique

Arguments

Name	Type	Properties
num-tasks	expression of integer type	positive

Modifiers

Name	Modifies	Type	Properties
prescriptiveness	*num-tasks*	Keyword: **`relaxed`**, **`strict`**	unique
directive-name-modifier	*all arguments*	Keyword: *directive-name* (a directive name)	unique

Directives

`taskloop`

Semantics

The **`num_tasks`** clause specifies that the **`taskloop`** construct create as many tasks as the minimum of the *num-tasks* expression and the number of logical iterations. Each task must have at least one logical iteration. If *prescriptiveness* is specified as **`strict`** for a **`taskloop`** region with N logical iterations, the logical iterations are partitioned in a balanced manner and each partition is assigned, in order, to a generated task. The partition size is $\lceil N/\textit{num-tasks} \rceil$ until the number of remaining logical iterations divides the number of remaining tasks evenly, at which point the partition size becomes $\lfloor N/\textit{num-tasks} \rfloor$. If *prescriptiveness* is specified as **`relaxed`** or not specified then the partitioning is implementation defined.

Restrictions

Restrictions to the **`num_tasks`** clause are as follows:

- None of the collapse-affected loops may be non-rectangular loops.

Cross References

- **`taskloop`** Construct, see Section 20.2

20.2.3 `task_iteration` Directive

Name: `task_iteration` **Category:** subsidiary	**Association:** unassociated **Properties:** task-synchronizing

Enclosing directives

`taskloop`

Clauses

`affinity`, **`depend`**, **`if`**

Semantics

The **task_iteration** directive is a subsidiary directive that controls the per-iteration task-execution attributes of the generated tasks of its associated **taskloop** construct, which is the innermost enclosing **taskloop** construct, as described below.

For each task-inherited clause specified on the **task_iteration** directive, the behavior is as if each task generated by the enclosing **taskloop** construct is specified with a corresponding clause that has the same *clause-specification*, but adjusted as follows. These clauses are instantiated for each instance of the loop-iteration variables for which the *if-expression* of the **if** clause evaluates to *true*. If an **if** clause is not specified on the **task_iteration** directive, the behavior is as if the *if-expression* evaluates to *true*.

Restrictions

The restrictions to the **task_iteration** directive are as follows:

- Each **task_iteration** directive must appear in the loop body of one of the **taskloop**-affected loops and must precede all statements and directives (except other **task_iteration** directives) in that loop body.
- If a **task_iteration** directive appears in the loop body of one of the **taskloop**-affected loops, no intervening code may occur between any two collapse-affected loops of the **taskloop**-affected loops.

Cross References

- **affinity** Clause, see Section 20.10
- **depend** Clause, see Section 23.9.6
- **if** Clause, see Section 5.6
- **iterator** Modifier, see Section 5.2.6
- **task** Construct, see Section 20.1
- **taskloop** Construct, see Section 20.2

20.3 taskgraph Construct

Name: **taskgraph** **Category:** executable	**Association:** block **Properties:** implicit-taskgroup, task-scheduling-point, task-scheduling-point-at-entry, task-scheduling-point-at-exit, task-synchronizing

Clauses

graph_id, **graph_reset**, **if**, **nogroup**, **priority**

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 an explicit-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 an explicit-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.

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 each 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 evalutes 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.

Since the **taskgraph** construct has the implicit-taskgroup property, each **taskgraph** region executes as if it is enclosed by a **taskgroup** region if the **nogroup** clause is not present, Thus, the construct implies the taskgroup scheduling constraint.

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.

Task scheduling points occur in the encountering task at entry to and at exit from each **taskgraph** region.

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:

- Explicit-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 undeferrable 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.
- Each 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
- **`priority`** Clause, see Section 20.9
- **`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
graph-id-value	expression of integer type	*default*

Modifiers

Name	Modifies	Type	Properties
directive-name-modifier	*all arguments*	Keyword: *directive-name* (a directive name)	unique

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
graph-reset-expression	expression of OpenMP logical type	optional

Modifiers

Name	Modifies	Type	Properties
directive-name-modifier	*all arguments*	Keyword: *directive-name* (a directive name)	unique

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
can_change_threads	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

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
can_merge	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

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 undeferred 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

Arguments

Name	Type	Properties
replayable-expression	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

target, **target_enter_data**, **target_exit_data**, **target_update**, **task**, **taskloop**, **taskwait**

Semantics

If *replayable-expression* evaluates to *true*, the **replayable** clause specifies that the construct on which it appears is a replayable construct. If *replayable-expression* evaluates to *false*, the **replayable** clause specifies that the construct on which it appears is not a replayable construct. If *replayable-expression* is not specified, the effect is as if *replayable-expression* evaluates to *true*.

Cross References

- **target** Construct, see Section 21.8
- **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
- **taskloop** Construct, see Section 20.2
- **taskwait** Construct, see Section 23.5

20.7 `final` Clause

Name: `final`	**Properties:** unique

Arguments

Name	Type	Properties
finalize	expression of OpenMP logical type	*default*

Modifiers

Name	Modifies	Type	Properties
directive-name-modifier	*all arguments*	Keyword: *directive-name* (a directive name)	unique

Directives

task, **taskloop**

Semantics

The **`final`** clause specifies that tasks generated by the construct on which it appears are final tasks if the *finalize* expression evaluates to *true*. All explicit-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:** unique

Arguments

Name	Type	Properties
set	Keyword: **`omp_pool`**, **`omp_team`**	*default*

Modifiers

Name	Modifies	Type	Properties
directive-name-modifier	*all arguments*	Keyword: *directive-name* (a directive name)	unique

Directives

`task`, **`taskloop`**

Semantics

The **`threadset`** clause specifies the set of threads that may execute tasks that are generated by the construct on which it appears. If the *set* argument is **`omp_team`**, the generated tasks may only be selected for execution by threads of the current team. If the *set* argument is **`omp_pool`**, the generated tasks may be selected for execution by 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
priority-value	expression of integer type	constant, non-negative

Modifiers

Name	Modifies	Type	Properties
directive-name-modifier	*all arguments*	Keyword: *directive-name* (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: `affinity`	**Properties:** task-inherited

Arguments

Name	Type	Properties
locator-list	list of locator list item type	*default*

Modifiers

Name	Modifies	Type	Properties
iterator	*all arguments*	Complex, name: `iterator` Arguments: ***iterator-specifier-list*** list of iterator specifier list item type (*default*)	unique
directive-name-modifier	*all arguments*	Keyword: *directive-name* (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 *iterator-identifier* that is defined in an *iterator* modifier 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: `detach`	**Properties:** data-sharing attribute, innermost-leaf, privatization, unique

Arguments

Name	Type	Properties
event-handle	variable of OpenMP event_handle type	*default*

Modifiers

Name	Modifies	Type	Properties
directive-name-modifier	*all arguments*	Keyword: *directive-name* (a directive name)	unique

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.

Fortran

- *event-handle* must not have the **`POINTER`** attribute.
- If *event-handle* has the **`ALLOCATABLE`** attribute, the allocation status must be allocated when the **`task`** construct is encountered, and the allocation status must not be changed, either explicitly or implicitly, in the **`task`** region.

Fortran

Cross References

- OpenMP **`event_handle`** Type, see Section 26.6.1
- **`target_data`** Construct, see Section 21.7
- **`task`** Construct, see Section 20.1

20.12 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* **& `ompt_task_initial`)** and **(***flags* **& `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 41.5.3
- **`parallel_end`** Callback, see Section 41.3.2
- OMPT **`scope_endpoint`** Type, see Section 40.27
- OMPT **`task_flag`** Type, see Section 40.37
- OMPT **`thread`** Type, see Section 40.39
- **`thread_begin`** Callback, see Section 41.1.3
- **`thread_end`** Callback, see Section 41.1.4

20.13 Task Scheduling

20.13.1 Task Scheduling Points

A task scheduling point is a location in a task at which a thread can perform task switching. Whenever a thread reaches a task scheduling point, it may begin or resume execution of a task from its schedulable task set. If the thread is a structured thread, the schedulable task set is the set of tasks bound to the current team. If the thread is an unassigned thread, the schedulable task set consists of any explicit task in the contention group associated with the current OpenMP thread pool. An idle thread, which must be an unassigned thread, is treated as if it is always at a task scheduling point. For other threads, task scheduling points are implied by constructs and routines that have the task-scheduling-point property.

In particular, every explicit-task-generating construct has the task-scheduling-point property. Each such construct has the task-scheduling-point-at-entry property so a task scheduling point occurs in the encountering task after execution of any implicit acquire-at-entry flushes or implicit release-at-entry flushes that occur in it for the construct. Each such construct also has the task-scheduling-point-at-exit property so a task scheduling point occurs in each generated task at its completion and, thus, after execution of any implicit acquire-at-exit flushes or implicit release-at-exit flushes that occur in it. Each such construct also has the task-scheduling-point-after-task-generation property so a task scheduling point occurs in the encountering task after each task associated with the construct is generated.

Details of other task scheduling points are specified in the semantics of constructs and routines that have the task-scheduling-point property or related sections.

20.13.2 `taskyield` Construct

Name: `taskyield`	**Association:** unassociated
Category: executable	**Properties:** task-scheduling-point, task-scheduling-point-at-exit

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.

20.13.3 Task Scheduling and its Constraints

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, which is necessarily in its schedulable task set;
- 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, which further limit the tasks in the schedulable task set that can be selected for execution, are as follows:

1. The tied-task scheduling constraint;
2. The free-agent-thread scheduling constraint; and
3. The synchronization scheduling constraints.

The tied-task scheduling constraint limits the tied tasks in the schedulable task set that can be selected to ensure that task switching does not lead to deadlock. Specifically, if any suspended tasks are tied to the thread and are not suspended in a barrier region, a new explicit tied task may be selected only if it is a descendent task of all of those suspended tasks. Otherwise, any new explicit tied task may be selected.

The free-agent-thread scheduling constraint limits the tasks in the schedulable task set that can be selected for execution by an unassigned thread. Specifically, an unassigned thread shall not start or resume execution of any task if it would result in the total number of free-agent threads in the OpenMP thread pool exceeding *free-agent-thread-limit-var*.

Each synchronization scheduling constraint enforces the semantics of OpenMP synchronization mechanisms.

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, the order in which a thread executes subtasks of different tasks in its schedulable task set is unspecified.

If the encountering task of a task scheduling point at which it will be suspended is untied then an implicit release flush occurs. That flush synchronizes with an implicit acquire flush that occurs in that task when it is resumed.

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.

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 41.5.2

20.13.4 Synchronization Scheduling Constraints

Synchronization scheduling constraints limit the tasks in the schedulable task set that can be selected to ensure that the semantics of a synchronization construct or other OpenMP synchronization mechanism are enforced. Each construct, clause and routine that is associated with a synchronization scheduling constraint has the task-synchronizing property. The synchronization scheduling constraints are the following:

1. The undeferred-task scheduling constraint;
2. The barrier scheduling constraint;
3. The taskgroup scheduling constraint;
4. The taskwait scheduling constraint;
5. The task-dependence scheduling constraint; and
6. The mutually-exclusive-task scheduling constraint.

The undeferred-task scheduling constraint limits the generating tasks in the schedulable task set that can be selected. Specifically, any task that is the generating task of an undeferred task may be selected only if the undeferred task has completed execution.

The barrier scheduling constraint ensures that each barrier, regardless of whether it is an explicit barrier or an implicit barrier, ensures that the conditions are met for execution to proceed beyond the barrier. Specifically it precludes selection at a task scheduling point in the barrier of each encountering task until all threads in the team have reached the barrier and all explicit tasks generated for execution by the team have executed to completion.

The taskgroup scheduling constraint ensures that all tasks in a taskgroup set execute to completion before the taskgroup region completes. Specifically, at the task scheduling point that occurs at exit from any taskgroup region, the encountering task must not be selected until all tasks in the taskgroup set complete execution.

The taskwait scheduling constraint ensures that all child tasks that the encountering task generated before encountering a full taskwait region complete execution before the full taskwait region completes. Specifically, at the task scheduling point that occurs at entry to a full taskwait region, the encountering thread must not select the generated included task of the full taskwait region until all of those child tasks complete execution. The implicit release-at-exit flush of each of those child tasks synchronizes with the implicit acquire-at-entry flush of the generated task.

The task-dependence scheduling constraint ensures that task dependences are enforced. Specifically, it requires that any task in the schedulable task set that is a dependent task may be selected only after task completion has occurred for all of its antecedent tasks.

The mutually-exclusive-task scheduling constraint limits mutually exclusive tasks in the schedulable task set that can be selected such that at most one can execute at a time. Specifically, a task shall not be selected while another task has been selected but has not yet completed if they are mutually exclusive tasks. Once that other task completes and that first task is selected, the implicit release-at-exit flush synchronizes with the implicit acquire-at-entry flush that occurs in the first task.

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
device-type-description	Keyword: **`any`**, **`host`**, **`nohost`**	*default*

Modifiers

Name	Modifies	Type	Properties
directive-name-modifier	*all arguments*	Keyword: *directive-name* (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 15.8.2
- **`declare_target`** Directive, see Section 15.8.1
- **`groupprivate`** Directive, see Section 9.2
- **`target`** Construct, see Section 21.8

21.2 `device` Clause

Name: `device`	**Properties:** ICV-defaulted, unique

Arguments

Name	Type	Properties
device-description	expression of integer type	*default*

Modifiers

Name	Modifies	Type	Properties
device-modifier	*device-description*	Keyword: **`ancestor`**, **`device_num`**	*default*
directive-name-modifier	*all arguments*	Keyword: *directive-name* (a directive name)	unique

Directives

`dispatch`, **`interop`**, **`target`**, **`target_data`**, **`target_enter_data`**, **`target_exit_data`**, **`target_update`**

Semantics

The **`device`** clause identifies the target device that is associated with a device construct.

If **`device_num`** is specified as the *device-modifier*, the *device-description* specifies the device number of the target device. If *device-modifier* does not appear in the clause, the behavior of the clause is as if *device-modifier* is **`device_num`**. If the *device-description* evaluates to **`omp_invalid_device`**, runtime error termination is performed.

If **`ancestor`** is specified as the *device-modifier*, the *device-description* specifies the number of target nesting levels of the target device. Specifically, if the *device-description* evaluates to 1, the target device is the parent device of the enclosing **`target`** region. If the construct on which the **`device`** clause appears is not encountered in a **`target`** region, the current device is treated as the parent device.

Unless otherwise specified, for directives that accept the **`device`** clause, if no **`device`** clause is present, the behavior is as if the **`device`** clause appears with **`device_num`** as *device-modifier* and with a *device-description* that evaluates to the value of the *default-device-var* ICV.

Restrictions

- The **ancestor** *device-modifier* must not appear on the **device** clause on any directive other than the **target** construct.
- If the **ancestor** *device-modifier* is specified, the *device-description* must evaluate to 1 and a **requires** directive with the **reverse_offload** clause must be specified;
- If the **device_num** *device-modifier* is specified and *target-offload-var* is not **mandatory**, *device-description* must evaluate to a conforming device number.

Cross References

- **dispatch** Construct, see Section 15.6
- *target-offload-var* ICV, see Table 3.1
- **interop** Construct, see Section 22.1
- **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

21.3 thread_limit Clause

Name: **thread_limit**	**Properties:** ICV-modifying, target-consistent, space-configuring, unique

Arguments

Name	Type	Properties
threadlim-list	list of expression of integer type	positive

Modifiers

Name	Modifies	Type	Properties
dims-modifier	*all arguments*	Complex, name: **dims** Arguments: ***dimensionality*** expression of integer type (constant, positive)	unique
directive-name-modifier	*all arguments*	Keyword: *directive-name* (a directive name)	unique
prescriptiveness-modifier	*threadlim-list*	Keyword: **relaxed**, **strict**	*default*

Directives

target, **teams**

Semantics

As described in Section 3.5, 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 spatial list for which the value of each list item is positive.

The **thread_limit** clause specifies the spatial limit of the team space of any team formed within the context of the contention group. The *dims-modifier* specifies the limit on the spatial dimensionality of those team spaces. If the *dims-modifier* is not specified, the limit on their spatial dimensionality is one.

The *threadlim-list* argument is a spatial list that specifies the limits on the number of threads in each dimension. If *prescriptiveness-modifier* is specified as **relaxed**, the number of available threads in each dimension that may participate will be less than or equal to the limit specified in the corresponding list item of *threadlim-list*. If the *prescriptiveness-modifier* is specified as **strict**, the number of available threads in each dimension will be equal to the limit specified in the corresponding list item of *threadlim-list*. When an implementation cannot fulfill this request then runtime error termination is performed. In this error termination scenario, the effect is as if an **error** directive has been encountered on which any provided **message** and **severity** clauses and an **at** clause with **execution** as *action-time* are specified. If the *prescriptiveness-modifier* is not specified, the effect is as if **strict** is specified when the *dims-modifier* is present, and as if **relaxed** is specified when the *dims-modifier* is not present.

If the **thread_limit** clause is not specified, and 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 the **relaxed** *prescriptiveness-modifier*, without the *dims-modifier* and with a *threadlim-list* that evaluates to a spatial list with a single implementation defined positive value that is 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 42.2
- **`device_initialize`** Callback, see Section 42.1
- **`device_load`** Callback, see Section 42.3
- **`device_unload`** Callback, see Section 42.4

21.5 `target_enter_data` Construct

Name: `target_enter_data` **Category:** executable	**Association:** unassociated **Properties:** data-mapping, device, device-affecting, explicit-task-generating, implicit-acquire-flush, implicit-acquire-flush-at-entry, implicit-acquire-flush-at-exit, implicit-flush, implicit-release-flush, implicit-release-flush-at-entry, implicit-release-flush-at-exit, map-entering, parallelism-generating, task-generating, task-scheduling-point, task-scheduling-point-after-task-generation, task-scheduling-point-at-entry, task-scheduling-point-at-exit, task-synchronizing

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, the target task is deferrable. 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 11.5) occur when the target task executes.

When an **if** clause is present and *if-expression* evaluates to *false*, the target device is the encountering device, regardless of any other **device** clause on the directive.

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_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_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 encountering 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 encountering 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 23.9.6
- **device** Clause, see Section 21.2
- **if** Clause, see Section 5.6
- **map** Clause, see Section 11.5
- **nowait** Clause, see Section 23.6
- **priority** Clause, see Section 20.9
- **replayable** Clause, see Section 20.6
- OMPT **scope_endpoint** Type, see Section 40.27
- OMPT **target** Type, see Section 40.34
- **target_emi** Callback, see Section 42.8
- **task** Construct, see Section 20.1
- OMPT **task_flag** Type, see Section 40.37

21.6 `target_exit_data` Construct

Name: `target_exit_data` **Category:** executable	**Association:** unassociated **Properties:** data-mapping, device, device-affecting, explicit-task-generating, implicit-acquire-flush, implicit-acquire-flush-at-entry, implicit-acquire-flush-at-exit, implicit-flush, implicit-release-flush, implicit-release-flush-at-entry, implicit-release-flush-at-exit, map-exiting, parallelism-generating, task-generating, task-scheduling-point-after-task-generation, task-scheduling-point-at-entry, task-scheduling-point-at-exit, task-synchronizing

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, the target task is deferrable. 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 11.5) occur when the target task executes.

When an **if** clause is present and *if-expression* evaluates to *false*, the target device is the encountering device, regardless of any other **device** clause on the directive.

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_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_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 encountering 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 encountering 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 23.9.6
- **device** Clause, see Section 21.2
- **if** Clause, see Section 5.6
- **map** Clause, see Section 11.5
- **nowait** Clause, see Section 23.6
- **priority** Clause, see Section 20.9
- **replayable** Clause, see Section 20.6
- OMPT **scope_endpoint** Type, see Section 40.27

- OMPT **target** Type, see Section 40.34
- **target_emi** Callback, see Section 42.8
- **task** Construct, see Section 20.1
- OMPT **task_flag** Type, see Section 40.37

21.7 target_data Construct

Name: target_data **Category:** executable	**Association:** block **Properties:** data-mapping, device, device-affecting, explicit-task-generating, implicit-acquire-flush, implicit-acquire-flush-at-entry, implicit-acquire-flush-at-exit, implicit-flush, implicit-release-flush, implicit-release-flush-at-entry, implicit-release-flush-at-exit, implicit-taskgroup, map-entering, map-exiting, parallelism-generating, sharing-task, simple-composite, task-generating, task-scheduling-point-after-task-generation, task-scheduling-point-at-entry, task-scheduling-point-at-exit, task-synchronizing

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: required	**Members:** **map**, **use_device_addr**, **use_device_ptr**

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 an explicit-task-generating construct, the **target_data** directive generates three explicit 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 is deferrable 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.

Each 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, implying the taskgroup scheduling constraint. Thus, the construct is task-synchronizing and has the implicit-taskgroup property. 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 7.4
- **depend** Clause, see Section 23.9.6
- **detach** Clause, see Section 20.11
- **device** Clause, see Section 21.2
- **firstprivate** Clause, see Section 7.3.3
- **if** Clause, see Section 5.6
- **in_reduction** Clause, see Section 8.12
- **map** Clause, see Section 11.5
- **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 7.3.2
- **shared** Clause, see Section 7.3.1
- **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 23.9.7
- **`use_device_addr`** Clause, see Section 7.3.8
- **`use_device_ptr`** Clause, see Section 7.3.6

21.8 `target` Construct

Name: `target` **Category:** executable	**Association:** block **Properties:** context-matching, data-mapping, device, device-affecting, exception-aborting, explicit-task-generating, groupprivate-instantiating, implicit-acquire-flush, implicit-acquire-flush-at-entry, implicit-acquire-flush-at-exit, implicit-flush, implicit-release-flush, implicit-release-flush-at-entry, implicit-release-flush-at-exit, implicit-task-generating, map-entering, map-exiting, parallelism-generating, team-generating, thread-limiting, task-generating, task-scheduling-point-after-task-generation, task-scheduling-point-at-entry, task-scheduling-point-at-exit, task-synchronizing

Clauses

`allocate`, **`default`**, **`defaultmap`**, **`depend`**, **`device`**, **`device_type`**, **`dyn_groupprivate`**, **`firstprivate`**, **`has_device_addr`**, **`if`**, **`in_reduction`**, **`is_device_ptr`**, **`map`**, **`message`**, **`nowait`**, **`priority`**, **`private`**, **`replayable`**, **`severity`**, **`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 an explicit task, the target task, that encloses a **target** region to be executed on a target device. If the **nowait** clause is present, the target task is deferrable. If the **nowait** clause is not present, the target task is an included task. The **device** and **device_type** clauses determine the device on which to execute the **target** region. To execute the **target** region, the target task generates an initial task, an undeferrable implicit task, that an initial thread executes on the selected device. If a **depend** clause is present, it is associated with the target task. The effect of any **map** clauses occur on entry to and exit from the generated **target** region, as specified in Section 11.5.

The state of memory between the encountering task and the target task is handled through implicit flushes as with any other explicit-task-generating construct. The initial task generated by the **target** construct implies that it is also an implicit-task-generating construct; the target task is the encountering task and the initial task is the generated task for that purpose. The thread-set of the implicit flushes associated with the generation of the initial task consists of the encountering thread and the initial thread that executes the **target** region.

All clauses are evaluated when the **target** construct is encountered. The data environment of the target task is created according to the data-sharing attribute clauses 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 11.5) 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 target device is the encountering device, 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 20.12. 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* **& 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** 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 encountering 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 encountering 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** region and does not have a corresponding list item in the device data environment prior to a task encountering the construct must also appear as a list item in a **map** clause.
- Each list item in a **map** clause that is specified on a **target** construct must have a base variable or base pointer.
- Each 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.

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.
- 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++

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 pointer association queries, pointer assignment and indirect calls.
- Invoking a type-bound procedure of an object on a device other than the device on which the object was constructed results in unspecified behavior, unless it has the **NON_OVERRIDABLE** attribute or the object is accessible and was constructed on the host device.
- The **SAME_TYPE_AS** intrinsic, **EXTEND_TYPE_OF** intrinsic and the **SELECT TYPE** construct have unspecified behavior except on the device on which the object was constructed.

Fortran

Cross References

- **allocate** Clause, see Section 13.6
- **default** Clause, see Section 7.4
- **defaultmap** Clause, see Section 11.6
- **depend** Clause, see Section 23.9.6
- **device** Clause, see Section 21.2
- **device_type** Clause, see Section 21.1
- **dyn_groupprivate** Clause, see Section 13.9
- **firstprivate** Clause, see Section 7.3.3
- **has_device_addr** Clause, see Section 7.3.7
- **if** Clause, see Section 5.6
- **in_reduction** Clause, see Section 8.12
- **is_device_ptr** Clause, see Section 7.3.5
- **map** Clause, see Section 11.5
- **message** Clause, see Section 16.3
- **nowait** Clause, see Section 23.6
- **priority** Clause, see Section 20.9
- **private** Clause, see Section 7.3.2
- **replayable** Clause, see Section 20.6
- OMPT **scope_endpoint** Type, see Section 40.27
- **severity** Clause, see Section 16.4
- OMPT **target** Type, see Section 40.34
- **target_data** Construct, see Section 21.7
- **target_emi** Callback, see Section 42.8
- **target_submit_emi** Callback, see Section 42.10
- **task** Construct, see Section 20.1
- OMPT **task_flag** Type, see Section 40.37
- **thread_limit** Clause, see Section 21.3
- **uses_allocators** Clause, see Section 13.8

21.9 `target_update` Construct

Name: `target_update` **Category:** executable	**Association:** unassociated **Properties:** device, device-affecting, explicit-task-generating, implicit-acquire-flush, implicit-acquire-flush-at-entry, implicit-acquire-flush-at-exit, implicit-flush, implicit-release-flush, implicit-release-flush-at-entry, implicit-release-flush-at-exit, parallelism-generating, task-generating, task-scheduling-point-after-task-generation, task-scheduling-point-at-entry, task-scheduling-point-at-exit, task-synchronizing

Clauses

`depend`, **`device`**, **`from`**, **`if`**, **`nowait`**, **`priority`**, **`replayable`**, **`to`**

Clause set

Properties: required	**Members:** **`from`**, **`to`**

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, the target task is deferrable. 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 encountering 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 encountering 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 23.9.6
- **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 40.27
- OMPT **target** Type, see Section 40.34
- **target_data_op_emi** Callback, see Section 42.7
- **target_emi** Callback, see Section 42.8
- **task** Construct, see Section 20.1
- OMPT **task_flag** Type, see Section 40.37
- **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 33

22.1 `interop` Construct

Name: `interop`	**Association:** unassociated
Category: executable	**Properties:** device, task-synchronizing

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 *interop-type* modifiers that were specified by the **init** clause that initialized that interoperability object.

If the interop-type set includes *targetsync-type*, 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-type*, 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 happens-before order, the foreign task must not begin execution until the generated task completes execution. No ordering is imposed between the encountering thread and either foreign tasks or OpenMP tasks by the **interop** construct.

Restrictions

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

- A **depend** clause must only appear on the directive if the interop-type set includes *targetsync-type*.
- 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 23.9.6
- **destroy** Clause, see Section 5.8
- **device** Clause, see Section 21.2
- **init** Clause, see Section 5.7
- *interop-type* Modifiers, see Section 5.7.1
- **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:** *default*

Arguments

Name	Type	Properties
interop-var	variable of OpenMP interop type	*default*

Modifiers

Name	Modifies	Type	Properties
directive-name-modifier	*all arguments*	Keyword: *directive-name* (a directive name)	unique

Directives

interop

Semantics

The **use** clause specifies the *interop-var* that is used for the effects of the directive on which the clause appears. However, *interop-var* is not initialized, destroyed or otherwise modified. The interop-type 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
- *interop-type* Modifiers, see Section 5.7.1

22.1.3 prefer_type Modifier

Modifiers

Name	Modifies	Type	Properties
prefer-type	*init-var*	Complex, name: **prefer_type** Arguments: ***prefer-type-specification*** list of preference specification list item type (*default*)	complex, unique

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 **ompx_** 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: `hint`	**Properties:** unique

Arguments

Name	Type	Properties
hint-expr	expression of OpenMP sync_hint type	constant

Modifiers

Name	Modifies	Type	Properties
directive-name-modifier	*all arguments*	Keyword: *directive-name* (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:** implicit-acquire-flush-at-entry, implicit-acquire-flush, implicit-flush, implicit-release-flush-at-exit, implicit-release-flush, mutual-exclusion, thread-limiting, thread-exclusive

Arguments

`critical``(`*name*`)`

Name	Type	Properties
name	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. When a thread executes a **`critical`** region that has a given *name*, the behavior is as if the implicit release flush performed on exit from the region synchronizes with the implicit acquire flush performed on entry to the next **`critical`** region with the same *name* that is performed within the same contention group.

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 40.20
- **mutex_acquire** Callback, see Section 41.7.8
- **mutex_acquired** Callback, see Section 41.7.12
- **mutex_released** Callback, see Section 41.7.13
- OpenMP **sync_hint** Type, see Section 26.9.5

23.3 Barriers

A barrier is a point in the execution of an OpenMP program beyond which no thread in a team that encounters it may continue execution until all threads in the team have reached the barrier and all explicit tasks generated for execution by the team have executed to completion.

All barrier regions, regardless of whether they correspond to explicit barriers or implicit barriers, have the implicit-release-flush-at-entry property and the implicit-acquire-flush-at-exit property. The corresponding implicit release flush that occurs in the implicit task of each encountering thread of the barrier synchronizes with every corresponding implicit acquire flush that occurs in the implicit tasks. In addition, the implicit release-at-exit flush in each explicit task that must execute to completion synchronizes with the implicit acquire-at-exit flush in each encountering task. In order to ensure that all associated explicit tasks can execute to completion, all barriers enforce the barrier scheduling constraint at the task scheduling point that occurs after at its implicit release-at-entry flush but before its implicit acquire-at-exit flush in each encountering task.

23.3.1 `barrier` Construct

Name: **barrier** **Category:** executable	**Association:** unassociated **Properties:** cancellation-point, implicit-acquire-flush-at-exit, implicit-acquire-flush, implicit-flush, implicit-release-flush-at-entry, implicit-release-flush, task-scheduling-point, task-scheduling-point-at-entry, task-synchronizing, team-executed

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 construct has the cancellation-point property since it includes a cancellation point.

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.

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.

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 41.6
- OMPT **cancel_flag** Type, see Section 40.7
- OMPT **scope_endpoint** Type, see Section 40.27
- **sync_region** Callback, see Section 41.7.4
- OMPT **sync_region** Type, see Section 40.33
- **sync_region_wait** Callback, see Section 41.7.5

23.3.2 Implicit Barriers

Implicit barriers are barriers that implicitly occur at the end of various regions as defined in the description of the constructs to which they correspond. All such constructs have the implicit-barrier property. Each construct that has the implicit-barrier property also has the implicit-barrier-flushes property, which indicates that implicit flushes are executed within the corresponding implicit barrier regions. Implicit barriers are task scheduling points. All constructs that have the implicit-barrier property have the cancellation-point property since the implicit barrier includes a cancellation point. Additionally, they are all task-synchronizing since the implicit barrier implies the barrier scheduling constraint.

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.

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**.

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 41.6
- OMPT **cancel_flag** Type, see Section 40.7
- **get_parallel_info** Entry Point, see Section 43.15
- OMPT **scope_endpoint** Type, see Section 40.27
- OMPT **state** Type, see Section 40.31
- **sync_region** Callback, see Section 41.7.4
- OMPT **sync_region** Type, see Section 40.33
- **sync_region_wait** Callback, see Section 41.7.5

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. 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, implicit-acquire-flush-at-exit, implicit-acquire-flush, implicit-flush, task-scheduling-point, task-scheduling-point-at-exit, task-synchronizing

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. All tasks in the taskgroup set must complete execution before the taskgroup region completes.

The encountering task has a task scheduling point that occurs before the implicit acquire-at-exit flush of the construct at which the taskgroup scheduling constraint is enforced. The implicit release-at-exit flush of each task in the taskgroup set synchronizes with the implicit acquire-at-exit flush in the encountering task.

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 20.13
- OMPT **scope_endpoint** Type, see Section 40.27
- **sync_region** Callback, see Section 41.7.4
- OMPT **sync_region** Type, see Section 40.33
- **sync_region_wait** Callback, see Section 41.7.5
- **task_reduction** Clause, see Section 8.11

23.5 taskwait Construct

Name: **taskwait** **Category:** executable	**Association:** unassociated **Properties:** explicit-task-generating, implicit-acquire-flush, implicit-acquire-flush-at-entry, implicit-acquire-flush-at-exit, implicit-flush, implicit-release-flush, implicit-release-flush-at-entry, implicit-release-flush-at-exit, task-generating, task-scheduling-point, task-scheduling-point-after-task-generation, task-scheduling-point-at-entry, task-scheduling-point-at-exit, task-synchronizing

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. The construct generates an explicit task with an empty associated structured block to provide the task scheduling points and implicit flushes that implement that synchronization.

If the **nowait** clause is not present, the generated task is a mergeable task and included task, which implies that the undeferred-task scheduling constraint ensures that the generated task completes execution before the encountering task is resumed. If the **nowait** clause is present, the generated task is deferred.

If no **depend** clause is present on the **taskwait** construct then a full taskwait is specified and the **taskwait** region is a full taskwait region. The encountering task is suspended at the task scheduling point that occurs at the entry to the full taskwait region and the taskwait scheduling constraint ensures that all child tasks that the encountering task generated before encountering the full taskwait region complete execution.

If one or more **depend** clauses are present on the **taskwait** construct, the task-dependence scheduling constraint ensures that the generated task is suspended until its predecessor tasks complete execution. If the **nowait** clause is also present, the task-dependence scheduling constraint also ensures that the generated task must complete execution before any subsequently generated task that depends on it 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 full taskwait region. The *taskwait-wait-end* event occurs when a task ends a waiting interval and resumes execution from a full taskwait region. The *taskwait-end* event occurs in a thread when it encounters a full taskwait after the taskwait synchronization on exit from the full 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_undeferred**) 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 23.9.6
- **nowait** Clause, see Section 23.6
- **replayable** Clause, see Section 20.6
- OMPT **scope_endpoint** Type, see Section 40.27
- **sync_region** Callback, see Section 41.7.4
- OMPT **sync_region** Type, see Section 40.33
- **sync_region_wait** Callback, see Section 41.7.5
- **task** Construct, see Section 20.1
- OMPT **task_flag** Type, see Section 40.37
- **task_schedule** Callback, see Section 41.5.2
- OMPT **task_status** Type, see Section 40.38

23.6 nowait Clause

Name: nowait	**Properties:** dispatch-requirement, end-clause, outermost-leaf, task-synchronizing, unique

Arguments

Name	Type	Properties
do_not_synchronize	expression of OpenMP logical type	optional

Modifiers

Name	Modifies	Type	Properties
directive-name-modifier	*all arguments*	Keyword: *directive-name* (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. 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 is deferrable. If the **`nowait`** clause is not specified on a task-generating construct that accepts the clause then the generated task is an included task.

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 15.6
- **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: nogroup	**Properties:** outermost-leaf, task-synchronizing, unique

Arguments

Name	Type	Properties
do_not_synchronize	expression of OpenMP logical type	optional

Modifiers

Name	Modifies	Type	Properties
directive-name-modifier	*all arguments*	Keyword: *directive-name* (a directive name)	unique

Directives

target_data, **taskgraph**, **taskloop**

Semantics

If *do_not_synchronize* evaluates to *true*, the **nogroup** clause removes any implicit taskgroup that would otherwise affect 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 set

Properties: exclusive	**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
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, **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
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, **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
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`, **`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
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, **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: seq_cst	**Properties:** 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, **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 set

Properties: exclusive	**Members:** **Clauses** **read**, **update**, **write**

Directives

atomic

Semantics

The *atomic* clause group defines a set of clauses that defines the semantics for which a directive enforces atomicity. If a construct accepts the *atomic* clause group and no member of the clause group is specified, the effect is as if the **update** clause is specified.

Cross References

- **atomic** 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: update	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 **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: write	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 **write** clause specifies that the **atomic** construct has atomic write semantics, which write the value of the shared variable atomically. If *use_semantics* evaluates to *false*, the effect is as if the **write** clause is not specified. If *use_semantics* is not specified, the effect is as if *use_semantics* evaluates to *true*.

Cross References

- **atomic** Construct, see Section 23.8.5

23.8.3 *extended-atomic* Clauses

Clause set

Properties: *default*	**Members:** **Clauses** **capture**, **compare**, **fail**, **weak**

Directives

atomic

Semantics

The *extended-atomic* clause group defines a set of clauses that extend the atomicity semantics specified by members of the *atomic* clause group.

Restrictions

Restrictions to the *extended-atomic* clause group are as follows:

- The **compare** clause may not be specified such that *use_semantics* evaluates to *false* if the **weak** clause is specified such that *use_semantics* evaluates to *true*.

Cross References

- **atomic** Construct, see Section 23.8.5
- *atomic* Clauses, see Section 23.8.2

23.8.3.1 `capture` Clause

Name: **capture**	**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 **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, 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 **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
memorder	Keyword: **acquire**, **relaxed**, **seq_cst**	*default*

Modifiers

Name	Modifies	Type	Properties
directive-name-modifier	*all arguments*	Keyword: *directive-name* (a directive name)	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
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 **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:** unique

Arguments

Name	Type	Properties
scope-specifier	Keyword: **all**, **cgroup**, **device**	*default*

Modifiers

Name	Modifies	Type	Properties
directive-name-modifier	*all arguments*	Keyword: *directive-name* (a directive name)	unique

Directives

atomic, **flush**

Semantics

The **memscope** clause determines the binding thread set of the region that corresponds to the construct on which it is specified. If the *scope-specifier* is **device**, the binding thread set consists of all threads on the device. If the *scope-specifier* is **cgroup**, the binding thread set consists of all threads that are executing tasks in the contention group. If the *scope-specifier* is **all**, the binding thread set consists of all threads on all devices. Unless otherwise stated, the thread-set of any flushes that are performed in an **atomic** or **flush** region is the same as the binding thread set of the region, as determined by the **memscope** clause.

Restrictions

The restrictions for the **memscope** clause are as follows:

- The binding thread set defined by the *scope-specifier* of the **memscope** clause on an **atomic** construct must be a subset of the atomic scope of the atomically accessed memory.
- The binding thread set defined by the *scope-specifier* of the **memscope** clause on an **atomic** construct must be a subset of all threads that are executing tasks in the contention group if the size of the atomically accessed storage location is not 8, 16, 32, or 64 bits.

Cross References

- **atomic** Construct, see Section 23.8.5
- **flush** Construct, see Section 23.8.6

23.8.5 atomic Construct

Name: **atomic** **Category:** executable	**Association:** block : atomic **Properties:** implicit-acquire-flush-at-exit, implicit-acquire-flush, implicit-flush, implicit-release-flush-at-entry, implicit-release-flush, implicit-strong-flush-at-entry, implicit-strong-flush-at-exit, implicit-strong-flush, mutual-exclusion, order-concurrent-nestable, simdizable

Clause groups

atomic, *extended-atomic*, *memory-order*

Clauses

hint, **memscope**

Binding

The **memscope** clause determines the binding thread set for an **atomic** region. If the **memscope** clause is not present, the behavior is as if the **memscope** clause appeared on the construct with the **device** *scope-specifier*.

Semantics

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. An implicit 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 implicit strong flush on entry to the atomic operation is also an implicit release flush. If the **read** or **update** clause is specified and the effective memory ordering is **acquire**, **acq_rel**, or **seq_cst** then the implicit strong flush on exit from the atomic operation is also an implicit acquire flush. Therefore, if the effective memory ordering is not **relaxed**, implicit release flushes and/or implicit acquire flushes are implied and permit synchronization between the threads without the use of explicit **flush** directives. Specifically, when a thread executes an **atomic** region that includes an implicit release flush that starts a given release sequence, its implicit release-at-entry flush synchronizes with the implicit acquire-at-exit flush of any atomic operation that reads a value written by the preceding atomic operation.

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 35
- 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 40.20
- **mutex_acquire** Callback, see Section 41.7.8
- **mutex_acquired** Callback, see Section 41.7.12
- **mutex_released** Callback, see Section 41.7.13
- **ordered** Construct, see Section 23.10
- **requires** Directive, see Section 16.5

23.8.6 flush Construct

Name: flush **Category:** executable	**Association:** unassociated **Properties:** *default*

Arguments

flush (*list*)

Name	Type	Properties
list	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 41.7.15
- **memscope** Clause, see Section 23.8.4

23.8.7 Implicit Flushes

This section details all implicit flushes, one or more of which occurs in the region that corresponds to any directive or routine that has the implicit-flush property unless otherwise specified. Any directive or routine that has the implicit-acquire-flush property includes one or more implicit acquire flushes unless otherwsise specified. Any directive or routine that has the implicit-release-flush property includes one or more implicit release flushes unless otherwise specified. Any directive or routine that has the implicit-strong-flush property includes one or more implicit strong flushes unless otherwise specified. The thread-set of any implicit flush is all threads on the current device unless otherwise specified.

Additional properties for constructs or routines specify where implicit flushes occur. For any construct or routine with the implicit-acquire-flush-at-entry property, the implicit-release-flush-at-entry property, or the implicit-strong-flush-at-entry property, the corresponding implicit flush occurs at entry to the corresponding region unless otherwise specified. Alternatively, for any construct or routine with the implicit-acquire-flush-at-exit property, the implicit-release-flush-at-exit property, or the implicit-strong-flush-at-exit property, the corresponding implicit flush occurs at exit from the corresponding region unless otherwise

specified. If a directive or routine has multiple implicit-flush properties that specify the implicit flushes occur at the same point within the same region then the implicit flushes are merged into a single implicit flush with all of the specified implicit-flush properties.

Each task-generating construct includes an implicit release-at-entry flush that occurs in its generating task before the construct generates any tasks. Each task-generating construct also includes an implicit acquire-at-entry flush that occurs in each generated task, regardless of whether it is an explicit tasks or an implicit task. The implicit release flush in the generating task synchronizes with the implicit acquire flush in each generated task to ensure that each generated task begins with an appropriate view of the memory.

Similarly, each task includes an implicit release-at-exit flush that ensures its temporary views is forced to memory upon task completion. If generated tasks are undeferrable then the task-generating construct includes an implicit acquire-at-exit flush; the implicit release-at-exit flush of each generated task synchronizes with that acquire flush.

Synchronization constructs, lock-acquiring routines and lock-releasing routines include implicit flushes to ensure memory reflects their occurence. The synchronizes-with relations of implicit flushes are generally described with the semantics of constructs and routines in which they occur.

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 ordered iterations of a loop.

23.9.1 *task-dependence-type* Modifier

Modifiers

Name	Modifies	Type	Properties
task-dependence-type	*all arguments*	Keyword: **`in`**, **`inout`**, **`inoutset`**, **`mutexinoutset`**, **`out`**	*default*

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 23.9.6.

Cross References

- **`depend`** Clause, see Section 23.9.6
- **`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
update-var	variable of OpenMP depend type	*default*

Modifiers

Name	Modifies	Type	Properties
task-dependence-type	*all arguments*	Keyword: **`in`**, **`inout`**, **`inoutset`**, **`mutexinoutset`**, **`out`**	*default*
directive-name-modifier	*all arguments*	Keyword: *directive-name* (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:

- 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-list-type* Modifiers

Modifier set

Properties: exclusive	**Members:** *depobj-modifier*, *loc-modifier*

Modifiers

Name	Modifies	Type	Properties
depobj-modifier	*all arguments*	Keyword: **`depobj`**	*default*
loc-modifier	*all arguments*	Keyword: **`loc`**	*default*

Clauses

`depend`

Semantics

The *depend-list-type* modifier group defines a set of modifiers that determine the type of list items in a depend list argument in the clause on which the modifier appears. The *loc-modifier* indicates that each depend list item is a locator list item. The *depobj-modifier* indicates that each depend list item is an expression of **depend** OpenMP type.

Restrictions

Restrictions to the *depend-list-type* modifiers are as follows:

- Depend list items used in clauses with the *loc-modifier* must be locator list items.
- Depend list items used in clauses with the *depobj-modifier* 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 must not be used in clauses with the *loc-modifier*.

Cross References

- **depend** Clause, see Section 23.9.6

23.9.6 depend Clause

Name: depend	**Properties:** dispatch-requirement, taskgraph-altering, task-inherited, task-synchronizing

Arguments

Name	Type	Properties
depend-list	list of depend list item type	*default*

Modifier groups

depend-list-type

Modifiers

Name	Modifies	Type	Properties
task-dependence-type	*all arguments*	Keyword: **in**, **inout**, **inoutset**, **mutexinoutset**, **out**	*default*
iterator	*all arguments*	Complex, name: **iterator** Arguments: ***iterator-specifier-list*** list of iterator specifier list item type (*default*)	unique
directive-name-modifier	*all arguments*	Keyword: *directive-name* (a directive name)	unique

Modifier set

Properties: exclusive	**Members:** *depobj-modifier*, *task-dependence-type*

Directives

dispatch, **interop**, **target**, **target_data**, **target_enter_data**, **target_exit_data**, **target_update**, **task**, **task_iteration**, **taskwait**

Semantics

The **depend** clause specifies a task dependence between two dependence-compatible tasks: the antecedent task and the dependent task. The implicit release-at-exit flush of each antecedent task of a given dependent task synchronizes with the implicit acquire-at-entry flush of that dependent task.

Task-dependence 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 list items in the *depend-list* argument and the dependence types associated with those list items. The dependence type is given by *task-dependence-type* if the *loc-modifier* is specified, and it is given by each depend object if the *depobj-modifier* is specified. If a *depend-list-type* modifier is not present, the behavior is as if the *loc-modifier* is specified.

One task, *A*, is a preceding dependence-compatible task of another task, *B*, if:

- *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 undeferrable exporting task that is a preceding dependence-compatible task of *B*; or
- *A* is a child task of an exporting task that precedes *B* in taskwait order.

If the *loc-modifier* is specified, 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 *depobj-modifier* is specified, 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 with the *loc-modifier*.

The list items that appear in the **depend** clause may reference any *iterator-identifier* that is defined in an *iterator* modifier in the same clause.

The list items that appear in a **depend** clause with the *loc-modifier* may include array sections or the **omp_all_memory** reserved locator.

C / C++

The list items that appear in a **depend** clause with the *loc-modifier* may use shape-operators.

C / C++

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 must not be zero-length array sections.
- The *task-dependence-type* modifier must appear in the clause if the *loc-modifier* is specified.
- The **omp_all_memory** reserved locator can only be used in a **depend** clause with an **out** or **inout** *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 41.7.1
- **dispatch** Construct, see Section 15.6
- 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 41.7.2
- **task_iteration** Directive, see Section 20.2.3
- **taskwait** Construct, see Section 23.5

23.9.7 transparent Clause

Name: transparent	**Properties:** unique

Arguments

Name	Type	Properties
impex-type	expression of OpenMP impex type	optional

Modifiers

Name	Modifies	Type	Properties
directive-name-modifier	*all arguments*	Keyword: *directive-name* (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 23.9.6
- **target_data** Construct, see Section 21.7
- **task** Construct, see Section 20.1
- **taskloop** Construct, see Section 20.2

23.9.8 `doacross` Clause

Name: `doacross`	**Properties:** required

Arguments

Name	Type	Properties
iteration-specifier	OpenMP iteration specifier	*default*

Modifiers

Name	Modifies	Type	Properties
dependence-type	*iteration-specifier*	Keyword: **`sink`**, **`source`**	required
directive-name-modifier	*all arguments*	Keyword: *directive-name* (a directive name)	unique

Directives

`ordered`

Semantics

The **`doacross`** clause identifies doacross dependences that imply additional constraints on the scheduling of ordered logical iterations. These constraints establish dependences only between ordered iterations. The *iteration-specifier* specifies an ordered 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 ordered iteration is a source iteration and, thus, satisfies doacross dependences that arise from the current ordered 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 ordered iteration is a sink iteration and, thus, has a doacross dependence, where *iteration-specifier* indicates the ordered iteration that satisfies the dependence. If *iteration-specifier* indicates an ordered iteration that does not occur in the ordered 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 ordered 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.

Restrictions

- The construct that corresponds to the binding region of an **ordered** region must specify an **ordered** clause if one of its constituent constructs is a worksharing-loop construct.
- The construct that corresponds to the binding region of an **ordered** region must not specify a **reduction** clause with the *inscan-modifier*.
- The region of a block-associated **ordered** construct must not have a binding region that corresponds to a construct in which a stand-alone **ordered** construct is closely nested.
- 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).
- A doacross-affected loop must be perfectly nested.
- The construct that corresponds to the binding region of an **ordered** region must not specify a **linear** clause.

C++

- A doacross-affected loop must not be a range-based **for** loop.

C++

Cross References

- OMPT **mutex** Type, see Section 40.20
- **mutex_acquire** Callback, see Section 41.7.8
- **mutex_released** Callback, see Section 41.7.13

23.10.1 Stand-alone `ordered` Construct

Name: **ordered** **Category:** executable	**Association:** unassociated **Properties:** implicit-acquire-flush-at-exit, implicit-acquire-flush, implicit-flush, implicit-release-flush-at-entry, implicit-release-flush, mutual-exclusion

Clauses

doacross

Binding

The binding thread set for a stand-alone **ordered** region is the current team. A stand-alone **ordered** region binds to the innermost enclosing worksharing-loop region.

Semantics

The innermost enclosing worksharing-loop construct of a stand-alone **ordered** construct is associated with *n* ordered-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 an ordered 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 ordered 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 ordered iteration encounters an **ordered** construct with a **doacross** clause for which the **source** *dependence-type* is specified. When threads execute stand-alone **ordered** regions, the implicit release-at-entry flush of a stand-alone **ordered** region from a given source iteration synchronizes with the implicit acquire-at-exit flush performed in all stand-alone **ordered** regions that are waiting for dependences on that source iteration to be satisfied.

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 40.10
- **dependences** Callback, see Section 41.7.1
- **doacross** Clause, see Section 23.9.8
- Worksharing-Loop Constructs, see Section 19.6

23.10.2 Block-associated `ordered` Construct

Name: `ordered` **Category:** executable	**Association:** block **Properties:** implicit-acquire-flush-at-entry, implicit-acquire-flush, implicit-flush, implicit-release-flush-at-exit, implicit-release-flush, 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. When a thread executes an **`ordered`** region, the behavior is as if its implicit release-at-exit flush synchronizes with the implicit acquire-at-entry flush performed in the **`ordered`** region encountered in the next collapsed iteration, if it exists.

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 41.7.12
- **ordered** Clause, see Section 6.4.6
- **simd** Construct, see Section 18.5

23.10.3 *parallelization-level* Clauses

Clause set

Properties: *default*	**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

Arguments

Name	Type	Properties
apply-to-threads	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

`ordered`

Semantics

If *apply_to_threads* evaluates to *true*, the effect is as if the **`threads`** *parallelization-level* clause is specified. If *apply_to_threads* evaluates to *false*, the effect is as if the **`threads`** clause is not specified. If *apply_to_threads* is not specified, the effect is as if *apply_to_threads* evaluates to *true*.

Cross References

- Block-associated **`ordered`** Construct, see Section 23.10.2

23.10.3.2 `simd` Clause

Name: `simd`	**Properties:** innermost-leaf, unique

Arguments

Name	Type	Properties
apply-to-simd	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

`ordered`

Semantics

If *apply_to_simd* evaluates to *true*, the effect is as if the **`simd`** *parallelization-level* clause is specified. If *apply_to_simd* evaluates to *false*, the effect is as if the **`simd`** clause is not specified. If *apply_to_simd* is not specified, the effect is as if *apply_to_simd* evaluates to *true*.

Cross References

- 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 set

Properties: required, exclusive	**Members:** **Clauses** **`do`**, **`for`**, **`parallel`**, **`sections`**, **`taskgroup`**

Modifiers

Name	Modifies	Type	Properties
directive-name-modifier	*all arguments*	Keyword: *directive-name* (a directive name)	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*.

Cross References

- **cancel** Construct, see Section 24.2
- **cancellation_point** Construct, see Section 24.4
- **do** Construct, see Section 19.6.2
- **for** Construct, see Section 19.6.1
- **parallel** Construct, see Section 18.1
- **sections** Construct, see Section 19.3
- **taskgroup** Construct, see Section 23.4

24.2 cancel Construct

Name: cancel **Category:** executable	**Association:** unassociated **Properties:** cancellation-point, implicit-flush

Clause groups

cancel-directive-name

Clauses

if

Binding

The binding thread set of the **cancel** region is the current team. The binding region of the **cancel** region is the innermost enclosing region of the type that corresponds to *cancel-directive-name*.

Semantics

The **cancel** construct activates cancellation of the innermost enclosing region of the type specified by *cancel-directive-name*, which must be the *directive-name* of a cancellable construct. Cancellation of the binding region is activated only if the *cancel-var* ICV is *true*, in which case the **cancel** construct causes the encountering task to continue execution at the end of the binding region if *cancel-directive-name* is not **taskgroup**. If the *cancel-var* ICV is *true* and *cancel-directive-name* is **taskgroup**, the encountering task continues execution at the end of the current task region. If the *cancel-var* ICV is *false*, the **cancel** construct is ignored.

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.

Tool Callbacks

A thread dispatches a registered **cancel** callback for each occurrence of a *cancel* event in the context of the encountering task. **(***flags* **& ompt_cancel_activated)** always evaluates to *true* in the dispatched callback; **(***flags* **& ompt_cancel_parallel)** evaluates to *true* in the dispatched callback if *cancel-directive-name* is **parallel**; **(***flags* **& ompt_cancel_sections)** evaluates to *true* in the dispatched callback if *cancel-directive-name* is **sections**; **(***flags* **& ompt_cancel_loop)** evaluates to *true* in the dispatched callback if *cancel-directive-name* is **for** or **do**; and **(***flags* **& ompt_cancel_taskgroup)** evaluates to *true* in the dispatched callback if *cancel-directive-name* is **taskgroup**.

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-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-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

- **cancel** Callback, see Section 41.6
- OMPT **cancel_flag** Type, see Section 40.7
- **declare_reduction** Directive, see Section 8.15
- *cancel-var* ICV, see Table 3.1
- **if** Clause, see Section 5.6
- **nowait** Clause, see Section 23.6
- **ordered** Clause, see Section 6.4.6
- **parallel** Construct, see Section 18.1
- **reduction** Clause, see Section 8.10
- **task** Construct, see Section 20.1
- **taskgroup** Construct, see Section 23.4
- **taskloop** Construct, see Section 20.2

24.3 Cancellation Points

Threads check for active cancellation only at cancellation points, which implicitly occur in constructs that have the cancellation-point property. Constructs that include implicit barriers include a cancellation point in the barrier. Unless otherwise specified, for such constructs, the cancellation point will not occur if the **nowait** clause specifies that the barrier will not occur.

If the *cancel-var* ICV is *true* when a thread reaches a cancellation points 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; and
- If the encountering task is at any other cancellation point, 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 and, thus, execution of the implicit release-at-exit flush of that task region. 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 and, thus, execution of the implicit release-at-exit flush of that task region. 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. The implicit flushes of discarded tasks are not executed.

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 descendent tasks of the implicit tasks bound to the current team are canceled according to the above **taskgroup** cancellation semantics. If the canceled region is not a **parallel** region, no task cancellation occurs.

Execution Model Events

A *cancellation* event occurs at a cancellation point if cancellation is active. A *discarded-task* event occurs for any discarded tasks.

Tool Callbacks

A thread dispatches a registered **cancel** callback for each occurrence of a *cancellation* event in the context of the encountering task. **(***flags* **& ompt_cancel_detected)** always evaluates to *true* in the dispatched callback; **(***flags* **& ompt_cancel_parallel)** evaluates to *true* in the dispatched callback if the innermost **parallel** region is canceled; **(***flags* **& ompt_cancel_sections)** evaluates to *true* in the dispatched callback if the innermost **sections** region is canceled; **(***flags* **& ompt_cancel_loop)** evaluates to *true* in the dispatched callback if the innermost worksharing-loop region is canceled; and **(***flags* **& ompt_cancel_taskgroup)** evaluates to *true* in the dispatched callback if the innermost **taskgroup** region is canceled.

A thread dispatches a registered **cancel** callback with its *task_data* argument pointing to the **data** object associated with the discarded task and with **ompt_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.

Cross References

- **barrier** Construct, see Section 23.3.1
- **cancel** Callback, see Section 41.6
- **cancel** Construct, see Section 24.2
- OMPT **cancel_flag** Type, see Section 40.7
- **cancellation_point** Construct, see Section 24.4
- OMPT **data** Type, see Section 40.8
- *cancel-var* ICV, see Table 3.1
- **nowait** Clause, see Section 23.6
- **parallel** Construct, see Section 18.1
- **task** Construct, see Section 20.1
- **taskgroup** Construct, see Section 23.4
- **taskloop** Construct, see Section 20.2

24.4 cancellation_point Construct

Name: **cancellation_point** **Category:** executable	**Association:** unassociated **Properties:** cancellation-point, implicit-flush

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 24.3.

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 37.1

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>:
  <distributed-loop-composite-directive-name>
  <simd-partitionable-composite-directive-name>
  <simple-composite-directive-name>

<distributed-loop-composite-directive-name>:
  <distribute-directive-name><separator><parallel-loop-combined-directive-name>

<parallel-loop-combined-directive-name>:
  <parallel-directive-name><separator><worksharing-loop-directive-name>

<simd-partitionable-composite-directive-name>:
  <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 explicit-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 18.5

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-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, 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 7.3.3
- **lastprivate** Clause, see Section 7.3.4
- **linear** Clause, see Section 8.14
- **loop** Construct, see Section 19.8
- **map** Clause, see Section 11.5
- **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 7.3.1
- **simd** Construct, see Section 18.5
- **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 18.5
- **taskloop** Construct, see Section 20.2

Part IV

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 **ompx_** 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. Including the C/C++ header file (**omp.h**), the Fortran module file (**omp_lib**) or the deprecated Fortran include file (**omp_lib.h**) must result in the declaration of each routine as well as a definition of each predefined identifier and 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 through 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++

Fortran

- Routines may not be called from **PURE** or **ELEMENTAL** procedures.
- Routines may not be called in **DO CONCURRENT** constructs.

Fortran

26.1 Predefined Identifiers

Predefined Identifiers

Name	Type	Value	Properties
omp_curr_progress_width	integer	see below	positive
omp_default_device	integer	**< -1**	constant
omp_initial_device	integer	**-1**	constant
omp_invalid_device	integer	**< -1**	constant
omp_unassigned_thread	integer	**< -1**	constant
openmp_version	integer	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.

Fortran

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).

Fortran

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: `intptr` **Properties:** omp	**Base Type:** `c_intptr_t`

Type Definition

C / C++

```
typedef intptr_t omp_intptr_t;
```

C / C++

Fortran

```
integer (kind=omp_c_intptr_t_kind)
```

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: `uintptr` **Properties:** C/C++-only, omp	**Base Type:** `c_uintptr_t`

Type Definition

C / C++

```
typedef uintptr_t omp_uintptr_t;
```

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: `sched` **Properties:** omp	**Base Type:** enumeration

Values

Name	Value	Properties
`omp_sched_static`	`0x1`	omp
`omp_sched_dynamic`	`0x2`	omp
`omp_sched_guided`	`0x3`	omp
`omp_sched_auto`	`0x4`	omp
`omp_sched_monotonic`	`0x80000000u`	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++

```
typedef <implementation-defined-integral> omp_event_handle_t;
```

C / C++

Fortran

```
integer (kind=omp_event_handle_kind)
```

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
`omp_interop_none`	0	*default*

Type Definition

C / C++

```
typedef <implementation-defined-integral> 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: interop_fr **Properties:** omp	**Base Type:** enumeration

Values

Name	Value	Properties
omp_ifr_last	N	omp

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 31.3

26.7.3 OpenMP `interop_property` Type

Name: `interop_property` **Properties:** omp	**Base Type:** enumeration

Values

Name	Value	Properties
`omp_ipr_fr_id`	`-1`	omp
`omp_ipr_fr_name`	`-2`	omp
`omp_ipr_vendor`	`-3`	omp
`omp_ipr_vendor_name`	`-4`	omp
`omp_ipr_device_num`	`-5`	omp
`omp_ipr_platform`	`-6`	omp
`omp_ipr_device`	`-7`	omp
`omp_ipr_device_context`	`-8`	omp
`omp_ipr_targetsync`	`-9`	omp
`omp_ipr_first`	`-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

```
integer (kind=omp_interop_property_kind), &
   parameter :: omp_ipr_fr_id = -1
integer (kind=omp_interop_property_kind), &
   parameter :: omp_ipr_fr_name = -2
integer (kind=omp_interop_property_kind), &
   parameter :: omp_ipr_vendor = -3
integer (kind=omp_interop_property_kind), &
   parameter :: omp_ipr_vendor_name = -4
integer (kind=omp_interop_property_kind), &
   parameter :: omp_ipr_device_num = -5
```

TABLE 26.2: Required Values of the **interop_property** OpenMP Type

Enum Name	Contexts	Name	Property
omp_ipr_fr_id	all	**fr_id**	An **intptr_t** value that represents the foreign runtime environment ID of context
omp_ipr_fr_name	all	**fr_name**	C string value that represents the name of the foreign runtime environment of context
omp_ipr_vendor	all	**vendor**	An **intptr_t** that represents the vendor of context
omp_ipr_vendor_name	all	**vendor_name**	C string value that represents the vendor of context
omp_ipr_device_num	all	**device_num**	The OpenMP device number for the device in the range 0 to **omp_get_num_devices** inclusive
omp_ipr_platform	*target*	**platform**	A foreign platform handle usually spanning multiple devices
omp_ipr_device	*target*	**device**	A foreign device handle
omp_ipr_device_context	*target*	**device_context**	A handle to an instance of a foreign device context
omp_ipr_targetsync	*targetsync*	**targetsync**	A handle to a synchronization object of a foreign execution context

```
integer (kind=omp_interop_property_kind), &
   parameter :: omp_ipr_platform = -6
integer (kind=omp_interop_property_kind), &
   parameter :: omp_ipr_device = -7
integer (kind=omp_interop_property_kind), &
   parameter :: omp_ipr_device_context = -8
integer (kind=omp_interop_property_kind), &
   parameter :: omp_ipr_targetsync = -9
integer (kind=omp_interop_property_kind), &
   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 31.3

26.7.4 OpenMP interop_rc Type

Name: interop_rc **Properties:** omp	**Base Type:** enumeration

Values

Name	Value	Properties
omp_irc_no_value	1	omp
omp_irc_success	0	omp
omp_irc_empty	-1	omp
omp_irc_out_of_range	-2	omp
omp_irc_type_int	-3	omp
omp_irc_type_ptr	-4	omp
omp_irc_type_str	-5	omp
omp_irc_other	-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,
  omp_irc_type_ptr     = -4,
  omp_irc_type_str     = -5,
  omp_irc_other        = -6
} omp_interop_rc_t;
```

C / C++

TABLE 26.3: Required Values for the **interop_rc** OpenMP Type

Enum Name	Description
omp_irc_no_value	Valid but no meaningful value available
omp_irc_success	Successful, value is usable
omp_irc_empty	The provided interoperability object is equal to **omp_interop_none**
omp_irc_out_of_range	Property ID is out of range, see Table 26.2
omp_irc_type_int	Property type is **int**; use **omp_get_interop_int**
omp_irc_type_ptr	Property type is pointer; use **omp_get_interop_ptr**
omp_irc_type_str	Property type is string; use **omp_get_interop_str**
omp_irc_other	Other error; use **omp_get_interop_rc_desc**

Fortran

```
integer (kind=omp_interop_rc_kind), &
   parameter :: omp_irc_no_value = 1
integer (kind=omp_interop_rc_kind), &
   parameter :: omp_irc_success = 0
integer (kind=omp_interop_rc_kind), &
   parameter :: omp_irc_empty = -1
integer (kind=omp_interop_rc_kind), &
   parameter :: omp_irc_out_of_range = -2
integer (kind=omp_interop_rc_kind), &
   parameter :: omp_irc_type_int = -3
integer (kind=omp_interop_rc_kind), &
   parameter :: omp_irc_type_ptr = -4
integer (kind=omp_interop_rc_kind), &
   parameter :: omp_irc_type_str = -5
integer (kind=omp_interop_rc_kind), &
   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 33.2
- **omp_get_interop_ptr** Routine, see Section 33.3

- **omp_get_interop_rc_desc** Routine, see Section 33.7
- **omp_get_interop_str** Routine, see Section 33.4

26.8 OpenMP Memory Management Types

This section describes OpenMP types that support memory management.

26.8.1 OpenMP access Type

Name: access **Properties:** omp	**Base Type:** enumeration

Values

Name	Value	Properties
omp_access_cgroup	0	omp
omp_access_pteam	1	omp

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: `allocator_handle` **Properties:** omp	**Base Type:** enumeration

Values

Name	Value	Properties
`omp_null_allocator`	0	omp
`omp_default_mem_alloc`	1	omp
`omp_large_cap_mem_alloc`	2	omp
`omp_const_mem_alloc`	3	omp
`omp_high_bw_mem_alloc`	4	omp
`omp_low_lat_mem_alloc`	5	omp
`omp_cgroup_mem_alloc`	6	omp
`omp_pteam_mem_alloc`	7	omp
`omp_thread_mem_alloc`	8	omp

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
```

```
integer (kind=omp_allocator_handle_kind), &
   parameter :: omp_cgroup_mem_alloc = 6
integer (kind=omp_allocator_handle_kind), &
   parameter :: omp_pteam_mem_alloc = 7
integer (kind=omp_allocator_handle_kind), &
   parameter :: omp_thread_mem_alloc = 8
```

Fortran

The **allocator_handle** OpenMP type represents an allocator as described in Table 13.3. This OpenMP type must be an implementation defined (for C++ possibly scoped) enum type and its valid constants must include those shown above.

C++

The **omp.h** header file also defines a class template that models the memory allocator concept in the **omp::allocator** namespace for each value of the **allocator_handle** OpenMP type. The names in this class do not include either the **omp_** prefix or the **_alloc** suffix. The class template **user_mem** in the same namespace permits this for any non-predefined allocator, requiring a handle of **allocator_handle** OpenMP type as an argument for initialization of the allocator.

C++

26.8.3 OpenMP `alloctrait` Type

Name: `alloctrait` **Properties:** omp	**Base Type:** structure

Fields

Name	Type	Properties
key	alloctrait_key	omp
value	alloctrait_val	omp

Type Definition

C / C++

```
typedef struct omp_alloctrait_t {
  omp_alloctrait_key_t key;
  omp_alloctrait_val_t value;
} omp_alloctrait_t;
```

C / C++

Fortran

```
! omp_alloctrait might not be provided
! in deprecated include file omp_lib.h
type omp_alloctrait
   integer (kind=omp_alloctrait_key_kind) key
   integer (kind=omp_alloctrait_val_kind) value
end type omp_alloctrait;
```

Fortran

TABLE 26.4: Allowed Key-Values for **alloctrait** OpenMP Type

Trait	Key	Allowed Values
sync_hint	omp_atk_sync_hint	omp_atv_contended, omp_atv_uncontended, omp_atv_serialized, omp_atv_private
alignment	omp_atk_alignment	Positive property integer powers of 2
access	omp_atk_access	omp_atv_all, omp_atv_memspace, omp_atv_device, omp_atv_cgroup, omp_atv_pteam, omp_atv_thread
pool_size	omp_atk_pool_size	Any positive property integer
fallback	omp_atk_fallback	omp_atv_default_mem_fb, omp_atv_null_fb, omp_atv_abort_fb, omp_atv_allocator_fb
fb_data	omp_atk_fb_data	An allocator handle
pinned	omp_atk_pinned	omp_atv_true, omp_atv_false
partition	omp_atk_partition	omp_atv_environment, omp_atv_nearest, omp_atv_blocked, omp_atv_interleaved, omp_atv_partitioner

table continued on next page

table continued from previous page

Trait	Key	Allowed Values
`pin_device`	`omp_atk_pin_device`	Any conforming device number
`preferred_device`	`omp_atk_preferred_device`	Any conforming device number
`target_access`	`omp_atk_target_access`	`omp_atv_single`, `omp_atv_multiple`
`atomic_scope`	`omp_atk_atomic_scope`	`omp_atv_all`, `omp_atv_device`
`part_size`	`omp_atk_part_size`	Any positive property integer value
`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: `alloctrait_key`	**Base Type:** enumeration
Properties: omp	

Values

Name	Value	Properties
omp_atk_sync_hint	1	omp
omp_atk_alignment	2	omp
omp_atk_access	3	omp
omp_atk_pool_size	4	omp
omp_atk_fallback	5	omp
omp_atk_fb_data	6	omp
omp_atk_pinned	7	omp
omp_atk_partition	8	omp
omp_atk_pin_device	9	omp
omp_atk_preferred_device	10	omp
omp_atk_device_access	11	omp
omp_atk_target_access	12	omp
omp_atk_atomic_scope	13	omp
omp_atk_part_size	14	omp
omp_atk_partitioner	15	omp
omp_atk_partitioner_arg	16	omp

Type Definition

C / C++

```
typedef enum omp_alloctrait_key_t {
  omp_atk_sync_hint        = 1,
  omp_atk_alignment        = 2,
  omp_atk_access           = 3,
  omp_atk_pool_size        = 4,
  omp_atk_fallback         = 5,
  omp_atk_fb_data          = 6,
  omp_atk_pinned           = 7,
  omp_atk_partition        = 8,
  omp_atk_pin_device       = 9,
  omp_atk_preferred_device = 10,
  omp_atk_device_access    = 11,
  omp_atk_target_access    = 12,
  omp_atk_atomic_scope     = 13,
  omp_atk_part_size        = 14,
  omp_atk_partitioner      = 15,
  omp_atk_partitioner_arg  = 16
} omp_alloctrait_key_t;
```

C / C++

Fortran

```
integer (kind=omp_alloctrait_key_kind), &
   parameter :: omp_atk_sync_hint = 1
integer (kind=omp_alloctrait_key_kind), &
   parameter :: omp_atk_alignment = 2
integer (kind=omp_alloctrait_key_kind), &
   parameter :: omp_atk_access = 3
integer (kind=omp_alloctrait_key_kind), &
   parameter :: omp_atk_pool_size = 4
integer (kind=omp_alloctrait_key_kind), &
   parameter :: omp_atk_fallback = 5
integer (kind=omp_alloctrait_key_kind), &
   parameter :: omp_atk_fb_data = 6
integer (kind=omp_alloctrait_key_kind), &
   parameter :: omp_atk_pinned = 7
integer (kind=omp_alloctrait_key_kind), &
   parameter :: omp_atk_partition = 8
integer (kind=omp_alloctrait_key_kind), &
   parameter :: omp_atk_pin_device = 9
integer (kind=omp_alloctrait_key_kind), &
   parameter :: omp_atk_preferred_device = 10
integer (kind=omp_alloctrait_key_kind), &
   parameter :: omp_atk_device_access = 11
integer (kind=omp_alloctrait_key_kind), &
   parameter :: omp_atk_target_access = 12
integer (kind=omp_alloctrait_key_kind), &
   parameter :: omp_atk_atomic_scope = 13
integer (kind=omp_alloctrait_key_kind), &
   parameter :: omp_atk_part_size = 14
integer (kind=omp_alloctrait_key_kind), &
   parameter :: omp_atk_partitioner = 15
integer (kind=omp_alloctrait_key_kind), &
   parameter :: omp_atk_partitioner_arg = 16
```

Fortran

The **`alloctrait_key`** OpenMP type represents an allocator trait 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

26.8.5 OpenMP `alloctrait_value` Type

Name: `alloctrait_value` **Properties:** omp	**Base Type:** enumeration

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_alloctrait_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_alloctrait_value_t;
```

C / C++

Fortran

```
integer (kind=omp_alloctrait_value_kind), &
   parameter :: omp_atv_default = -1
integer (kind=omp_alloctrait_value_kind), &
   parameter :: omp_atv_false = 0
integer (kind=omp_alloctrait_value_kind), &
   parameter :: omp_atv_true = 1
integer (kind=omp_alloctrait_value_kind), &
   parameter :: omp_atv_contended = 3
integer (kind=omp_alloctrait_value_kind), &
   parameter :: omp_atv_uncontended = 4
integer (kind=omp_alloctrait_value_kind), &
   parameter :: omp_atv_serialized = 5
integer (kind=omp_alloctrait_value_kind), &
   parameter :: omp_atv_private = 6
integer (kind=omp_alloctrait_value_kind), &
   parameter :: omp_atv_device = 7
integer (kind=omp_alloctrait_value_kind), &
   parameter :: omp_atv_thread = 8
integer (kind=omp_alloctrait_value_kind), &
   parameter :: omp_atv_pteam = 9
integer (kind=omp_alloctrait_value_kind), &
   parameter :: omp_atv_cgroup = 10
integer (kind=omp_alloctrait_value_kind), &
   parameter :: omp_atv_default_mem_fb = 11
integer (kind=omp_alloctrait_value_kind), &
```

```
   parameter :: omp_atv_null_fb = 12
integer (kind=omp_alloctrait_value_kind), &
   parameter :: omp_atv_abort_fb = 13
integer (kind=omp_alloctrait_value_kind), &
   parameter :: omp_atv_allocator_fb = 14
integer (kind=omp_alloctrait_value_kind), &
   parameter :: omp_atv_environment = 15
integer (kind=omp_alloctrait_value_kind), &
   parameter :: omp_atv_nearest = 16
integer (kind=omp_alloctrait_value_kind), &
   parameter :: omp_atv_blocked = 17
integer (kind=omp_alloctrait_value_kind), &
   parameter :: omp_atv_interleaved = 18
integer (kind=omp_alloctrait_value_kind), &
   parameter :: omp_atv_all = 19
integer (kind=omp_alloctrait_value_kind), &
   parameter :: omp_atv_single = 20
integer (kind=omp_alloctrait_value_kind), &
   parameter :: omp_atv_multiple = 21
integer (kind=omp_alloctrait_value_kind), &
   parameter :: omp_atv_memspace = 22
integer (kind=omp_alloctrait_value_kind), &
   parameter :: omp_atv_partitioner = 23
```

Fortran

The **`alloctrait_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 `alloctrait_val` Type

Name: `alloctrait_val` **Properties:** omp	**Base Type:** `intptr`

Type Definition

C / C++

```
typedef omp_intptr_t omp_alloctrait_val_t;
```

C / C++

Fortran

```
integer (kind=c_intptr_t)
```

Fortran

The **alloctrait_val** OpenMP type represents the values that may be assigned to the **value** field of the **alloctrait_val** OpenMP type. Any of the semantic values of the **alloctrait_value** OpenMP type may be used for the **alloctrait_val** OpenMP type; in addition, other numeric values may be used for it as appropriate for the specified **key** of the **alloctrait** OpenMP type.

26.8.7 OpenMP mempartition Type

Name: mempartition **Properties:** named-handle, omp, opaque	Base Type: opaque

Type Definition

C / C++

```
typedef <implementation-defined> omp_mempartition_t;
```

C / C++

Fortran

```
integer (kind=omp_mempartition_kind)
```

Fortran

The **mempartition** OpenMP type is an opaque type that represents memory partitions.

26.8.8 OpenMP mempartitioner Type

Name: mempartitioner **Properties:** named-handle, omp, opaque	Base Type: opaque

Type Definition

C / C++

```
typedef <implementation-defined> omp_mempartitioner_t;
```

C / C++

Fortran

```
integer (kind=omp_mempartitioner_kind)
```

Fortran

The **mempartitioner** OpenMP type is an opaque type that represents memory partitioners.

26.8.9 OpenMP `mempartitioner_lifetime` Type

Name: `mempartitioner_lifetime` **Properties:** omp	**Base Type:** enumeration

Values

Name	Value	Properties
`omp_static_mempartition`	1	omp
`omp_allocator_mempartition`	2	omp
`omp_dynamic_mempartition`	3	omp

Type Definition

C / C++

```
typedef enum omp_mempartitioner_lifetime_t {
  omp_static_mempartition    = 1,
  omp_allocator_mempartition = 2,
  omp_dynamic_mempartition   = 3
} omp_mempartitioner_lifetime_t;
```

C / C++

Fortran

```
integer (kind=omp_mempartitioner_lifetime_kind), &
   parameter :: omp_static_mempartition = 1
integer (kind=omp_mempartitioner_lifetime_kind), &
   parameter :: omp_allocator_mempartition = 2
integer (kind=omp_mempartitioner_lifetime_kind), &
   parameter :: omp_dynamic_mempartition = 3
```

Fortran

The **`mempartitioner_lifetime`** OpenMP type represents the lifetime of a memory partitioner. The valid constants for the **`mempartitioner_lifetime`** OpenMP type must include those shown above.

26.8.10 OpenMP `mempartitioner_compute_proc` Type

Name: `mempartitioner_compute_proc` **Category:** subroutine pointer	**Properties:** iso_c_binding, omp

Arguments

Name	Type	Properties
memspace	memspace_handle	omp
allocation_size	c_size_t	iso_c, value
partitioner_arg	alloctrait_val	omp, value
partition	mempartition	C/C++ pointer, omp

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 34.5.3
- **omp_mempartition_set_part** Routine, see Section 34.5.5

26.8.11 OpenMP mempartitioner_release_proc Type

Name: mempartitioner_release_proc **Category:** subroutine pointer	**Properties:** iso_c_binding, omp

Arguments

Name	Type	Properties
partition	mempartition	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 34.5.1

26.8.12 OpenMP memspace_handle Type

Name: memspace_handle Properties: omp	Base Type: enumeration

Values

Name	Value	Properties
omp_default_mem_space	−1	omp
omp_null_mem_space	0	omp
omp_large_cap_mem_space	1	omp
omp_const_mem_space	2	omp
omp_high_bw_mem_space	3	omp
omp_low_lat_mem_space	4	omp
omp_cgroup_mem_space	5	omp

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_cgroup_mem_space    = 5
} 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
integer (kind=omp_memspace_handle_kind), &
   parameter :: omp_cgroup_mem_space = 5
```

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` **Properties:** named-handle, omp, opaque	**Base Type:** `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` **Properties:** omp	**Base Type:** enumeration

Values

Name	Value	Properties
`omp_not_impex`	0	omp
`omp_import`	1	omp
`omp_export`	2	omp
`omp_impex`	3	omp

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

The **impex** OpenMP type is an enumeration type that is used to specify whether the child tasks of a task may form a task dependence with respect to its dependence-compatible tasks. In particular, it is used to identify whether a task is an importing task and/or an exporting task. The valid constants must include those shown above.

Cross References

- **transparent** Clause, see Section 23.9.7

26.9.3 OpenMP lock Type

Name: lock **Properties:** named-handle, opaque	Base Type: opaque

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 35.1.3
- **omp_init_nest_lock_with_hint** Routine, see Section 35.1.4

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` **Properties:** omp	**Base Type:** enumeration

Values

Name	Value	Properties
omp_proc_bind_false	0	omp
omp_proc_bind_true	1	omp
omp_proc_bind_primary	2	omp
omp_proc_bind_close	3	omp
omp_proc_bind_spread	4	omp

Type Definition

C / C++

```
typedef enum omp_proc_bind_t {
  omp_proc_bind_false   = 0,
  omp_proc_bind_true    = 1,
  omp_proc_bind_primary = 2,
  omp_proc_bind_close   = 3,
  omp_proc_bind_spread  = 4
} omp_proc_bind_t;
```

C / C++

Fortran

```
integer (kind=omp_proc_bind_kind), &
   parameter :: omp_proc_bind_false = 0
integer (kind=omp_proc_bind_kind), &
   parameter :: omp_proc_bind_true = 1
integer (kind=omp_proc_bind_kind), &
   parameter :: omp_proc_bind_primary = 2
integer (kind=omp_proc_bind_kind), &
   parameter :: omp_proc_bind_close = 3
integer (kind=omp_proc_bind_kind), &
   parameter :: omp_proc_bind_spread = 4
```

Fortran

The **`proc_bind`** OpenMP type is used in routines that modify or retrieve the value of the *bind-var* ICV. The valid constants for the **`proc_bind`** type must include those shown above.

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: `pause_resource` **Properties:** omp	**Base Type:** enumeration

Values

Name	Value	Properties
`omp_pause_soft`	1	omp
`omp_pause_hard`	2	omp
`omp_pause_stop_tool`	3	omp

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: `control_tool` **Properties:** omp	**Base Type:** enumeration

Values

Name	Value	Properties
`omp_control_tool_start`	1	omp
`omp_control_tool_pause`	2	omp
`omp_control_tool_flush`	3	omp
`omp_control_tool_end`	4	omp
`omp_control_tool_max`	INT32_MAX	omp

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,
  omp_control_tool_max   = INT32_MAX
} omp_control_tool_t;
```

C / C++

Fortran

```
integer (kind=omp_control_tool_kind), &
   parameter :: omp_control_tool_start = 1
integer (kind=omp_control_tool_kind), &
   parameter :: omp_control_tool_pause = 2
integer (kind=omp_control_tool_kind), &
   parameter :: omp_control_tool_flush = 3
integer (kind=omp_control_tool_kind), &
   parameter :: omp_control_tool_end = 4
integer (kind=omp_control_tool_kind), &
   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.

Including the C/C++ header file **omp.h** or **omp-tools.h** must result in the definition of this OpenMP type. The implementation must ensure that including both header files will not result in an error.

TABLE 26.5: Standard Tool Control Commands

Command	Action
omp_control_tool_start	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.
omp_control_tool_pause	Temporarily turn monitoring off. If monitoring is already off, it is idempotent.
omp_control_tool_flush	Flush any data buffered by a tool. This command may be applied whether monitoring is on or off.
omp_control_tool_end	Turn monitoring off permanently; the tool 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
num_threads	integer	positive

Prototypes

C / C++

```
void omp_set_num_threads(int num_threads);
```

C / C++

Fortran

```
subroutine omp_set_num_threads(num_threads)
  integer num_threads
```

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: omp_get_num_threads Category: function	Properties: *default*

Return Type

Name	Type	Properties
<return type>	integer	*default*

Prototypes

C / C++

```
int omp_get_num_threads(void);
```

C / C++

Fortran

```
integer function omp_get_num_threads()
```

Fortran

Effect

The **omp_get_num_threads** routine returns the team size of the team that is executing the parallel region to which the routine region binds.

27.3 omp_get_thread_num Routine

Name: omp_get_thread_num Category: function	Properties: *default*

Return Type

Name	Type	Properties
<return type>	integer	*default*

Prototypes

C / C++

```
int omp_get_thread_num(void);
```

C / C++

Fortran

```
integer function omp_get_thread_num()
```

Fortran

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: omp_get_max_threads **Category:** function	**Properties:** ICV-retrieving

Return Type

Name	Type	Properties
<*return type*>	integer	*default*

Prototypes

C / C++

```
int omp_get_max_threads(void);
```

C / C++

Fortran

```
integer function omp_get_max_threads()
```

Fortran

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: omp_get_thread_limit **Category:** function	**Properties:** ICV-retrieving

Return Type

Name	Type	Properties
<*return type*>	integer	*default*

Prototypes

C / C++

```
int omp_get_thread_limit(void);
```

C / C++

Fortran

```
integer function omp_get_thread_limit()
```

Fortran

Effect

The **omp_get_thread_limit** routine returns the unidimensional 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: omp_in_parallel **Category:** function	**Properties:** *default*

Return Type

Name	Type	Properties
<return type>	logical	*default*

Prototypes

C / C++

```
int omp_in_parallel(void);
```

C / C++

Fortran

```
logical function omp_in_parallel()
```

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: `omp_set_dynamic` **Category:** subroutine	**Properties:** ICV-modifying

Arguments

Name	Type	Properties
dynamic_threads	logical	*default*

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: `omp_get_dynamic` **Category:** function	**Properties:** ICV-retrieving

Return Type

Name	Type	Properties
<return type>	logical	*default*

Prototypes

C / C++

```
int omp_get_dynamic(void);
```

C / C++

Fortran

```
logical function omp_get_dynamic()
```

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: omp_set_schedule **Category:** subroutine	**Properties:** ICV-modifying

Arguments

Name	Type	Properties
kind	sched	omp
chunk_size	integer	*default*

Prototypes

C / C++

```
void omp_set_schedule(omp_sched_t kind, int chunk_size);
```

C / C++

Fortran

```
subroutine omp_set_schedule(kind, chunk_size)
  integer (kind=omp_sched_kind) kind
  integer chunk_size
```

Fortran

Effect

The effect of this routine is to set the value of the *run-sched-var* ICV of the current task to the values specified in the two arguments. Thus, the routine affects the schedule that is applied when **runtime** is used as the schedule type.

The schedule is set to the schedule type that is specified by the first argument *kind*. For the schedule types **omp_sched_static**, **omp_sched_dynamic**, and **omp_sched_guided**, the *chunk_size* is set to the value of the second argument, or to the default *chunk_size* if the value of the second argument is less than 1; for the schedule type **omp_sched_auto**, the second argument is ignored; for implementation defined schedule types, the values and associated meanings of the second argument are implementation defined.

Cross References

- *run-sched-var* ICV, see Table 3.1
- OpenMP **sched** Type, see Section 26.5.1

27.10 omp_get_schedule Routine

Name: omp_get_schedule **Category:** subroutine	**Properties:** ICV-retrieving

Arguments

Name	Type	Properties
kind	sched	C/C++ pointer, omp
chunk_size	integer	C/C++ pointer

Prototypes

C / C++

```
void omp_get_schedule(omp_sched_t *kind, int *chunk_size);
```

C / C++

Fortran

```
subroutine omp_get_schedule(kind, chunk_size)
  integer (kind=omp_sched_kind) kind
  integer chunk_size
```

Fortran

Effect

The **omp_get_schedule** routine returns the *run-sched-var* ICV in the task to which the routine binds. Thus, the routine returns the schedule that is applied when the **runtime** schedule type is used. The first argument *kind* returns the schedule type to be used. If the returned schedule type is **omp_sched_static**, **omp_sched_dynamic**, or **omp_sched_guided**, the second argument, *chunk_size*, returns the chunk size to be used, or a value less than 1 if the default chunk size is to be used. The value returned by the second argument is implementation defined for any other schedule types.

Cross References

- *run-sched-var* ICV, see Table 3.1
- OpenMP **sched** Type, see Section 26.5.1

27.11 `omp_get_supported_active_levels` Routine

Name: `omp_get_supported_active_levels` **Category:** function	**Properties:** *default*

Return Type

Name	Type	Properties
<return type>	integer	*default*

Prototypes

C / C++

```
int omp_get_supported_active_levels(void);
```

C / C++

Fortran

```
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: `omp_set_max_active_levels` **Category:** subroutine	**Properties:** ICV-modifying

Arguments

Name	Type	Properties
max_levels	integer	non-negative

Prototypes

C / C++

```
void omp_set_max_active_levels(int max_levels);
```

C / C++

Fortran

```
subroutine omp_set_max_active_levels(max_levels)
  integer max_levels
```

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: omp_get_max_active_levels **Category:** function	**Properties:** ICV-retrieving

Return Type

Name	Type	Properties
<return type>	integer	*default*

Prototypes

C / C++

```
int omp_get_max_active_levels(void);
```

C / C++

Fortran

```
integer function omp_get_max_active_levels()
```

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: `omp_get_level` **Category:** function	**Properties:** ICV-retrieving

Return Type

Name	Type	Properties
<return type>	integer	*default*

Prototypes

C / C++

```
int omp_get_level(void);
```

C / C++

Fortran

```
integer function omp_get_level()
```

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: `omp_get_ancestor_thread_num` **Category:** function	**Properties:** *default*

Return Type and Arguments

Name	Type	Properties
<return type>	integer	*default*
level	integer	*default*

Prototypes

C / C++

```
int omp_get_ancestor_thread_num(int level);
```

C / C++

Fortran

```
integer function omp_get_ancestor_thread_num(level)
  integer level
```

Fortran

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: omp_get_team_size **Category:** function	**Properties:** *default*

Return Type and Arguments

Name	Type	Properties
<return type>	integer	*default*
level	integer	*default*

Prototypes

C / C++

```
int omp_get_team_size(int level);
```

C / C++

Fortran

```
integer function omp_get_team_size(level)
  integer level
```

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: omp_get_active_level **Category:** function	**Properties:** ICV-retrieving

Return Type

Name	Type	Properties
<return type>	integer	*default*

Prototypes

C / C++

```
int omp_get_active_level(void);
```

C / C++

Fortran

```
integer function omp_get_active_level()
```

Fortran

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: omp_get_num_teams **Category:** function	**Properties:** ICV-retrieving, teams-nestable

Return Type

Name	Type	Properties
<return type>	integer	*default*

Prototypes

C / C++

```
int omp_get_num_teams(void);
```

C / C++

Fortran

```
integer function omp_get_num_teams()
```

Fortran

Effect

The **omp_get_num_teams** routine returns the value of the *league-size-var* ICV, which is the league size 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: omp_set_num_teams **Category:** subroutine	**Properties:** ICV-modifying

Arguments

Name	Type	Properties
num_teams	integer	non-negative

Prototypes

C / C++

```
void omp_set_num_teams(int num_teams);
```

C / C++

Fortran

```
subroutine omp_set_num_teams(num_teams)
  integer num_teams
```

Fortran

Effect

The effect of the **omp_set_num_teams** routine is to set the value of the *nteams-var* ICV of the host device to the value specified in the *num_teams* argument.

Restrictions

Restrictions to the **omp_set_num_teams** 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

- *nteams-var* ICV, see Table 3.1
- **num_teams** Clause, see Section 18.2.1
- **teams** Construct, see Section 18.2

28.3 omp_get_team_num Routine

Name: omp_get_team_num **Category:** function	**Properties:** ICV-retrieving, teams-nestable

Return Type

Name	Type	Properties
<return type>	integer	*default*

Prototypes

C / C++

```
int omp_get_team_num(void);
```

C / C++

Fortran

```
integer function omp_get_team_num()
```

Fortran

Effect

The **omp_get_team_num** routine returns the value of the *team-num-var* ICV, which is the team number of the current team and is an integer between 0 and one less than the value returned by **omp_get_num_teams**, inclusive. The routine returns 0 if it is called outside of a **teams** region.

Cross References

- *team-num-var* ICV, see Table 3.1
- **omp_get_num_teams** Routine, see Section 28.1
- **teams** Construct, see Section 18.2

28.4 omp_get_max_teams Routine

Name: omp_get_max_teams **Category:** function	**Properties:** ICV-retrieving

Return Type

Name	Type	Properties
<return type>	integer	*default*

Prototypes

C / C++

```
int omp_get_max_teams(void);
```

C / C++

Fortran

```
integer function omp_get_max_teams()
```

Fortran

Effect

The **omp_get_max_teams** routine returns the value of the *nteams-var* ICV of the current device. If positive, this value is also an upper bound on the number of teams that can be created by a **teams** construct without a **num_teams** clause that is encountered after execution returns from this routine.

Cross References

- *nteams-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: omp_get_teams_thread_limit **Category:** function	**Properties:** ICV-retrieving

Return Type

Name	Type	Properties
<return type>	integer	*default*

Prototypes

C / C++

```
int omp_get_teams_thread_limit(void);
```

C / C++

Fortran

```
integer function omp_get_teams_thread_limit()
```

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: omp_set_teams_thread_limit **Category:** subroutine	**Properties:** ICV-modifying

Arguments

Name	Type	Properties
thread_limit	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_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

29 Multidimensional Support Routines

This section describes the multidimensional-information routines, which are routines that have the multidimensional-information-routine property.

29.1 `omp_get_num_threads_dim` Routine

Name: `omp_get_num_threads_dim` **Category:** function	**Properties:** multidimensional-information-routine

Return Type and Arguments

Name	Type	Properties
<return type>	integer	*default*
dim	integer	non-negative

Prototypes

C / C++

```
int omp_get_num_threads_dim(int dim);
```

C / C++

Fortran

```
integer function omp_get_num_threads_dim(dim)
  integer dim
```

Fortran

Effect

The **`omp_get_num_threads_dim`** routine returns the element in the multidimensional value of the *team-size-var* ICV associated with dimension *dim*. This value corresponds to the team size for the dimension *dim* in the team space of the team that is executing the parallel region to which the routine region binds. If *dim* corresponds to an inactive spatial dimension in the team space then the routine returns one.

Cross References

- *team-size-var* ICV, see Table 3.1

29.2 `omp_get_thread_num_dim` Routine

Name: `omp_get_thread_num_dim` **Category:** function	**Properties:** multidimensional-information-routine

Return Type and Arguments

Name	Type	Properties
<return type>	integer	*default*
dim	integer	non-negative

Prototypes

C / C++

```
int omp_get_thread_num_dim(int dim);
```

C / C++

Fortran

```
integer function omp_get_thread_num_dim(dim)
  integer dim
```

Fortran

Effect

The **omp_get_thread_num_dim** routine returns the element in the multidimensional value of the *thread-num-var* ICV associated with dimension *dim*. This value corresponds to the thread number of the calling thread for the dimension *dim* within the team space of the team that is executing the parallel region to which the routine region binds. For an assigned thread, the result is an integer value between zero and one less than the value returned by **omp_get_num_threads_dim** for dimension *dim*, inclusive. If *dim* corresponds to an inactive spatial dimension in the team space then the routine returns zero. For an unassigned thread, the routine returns **omp_unassigned_thread**.

Cross References

- Predefined Identifiers, see Section 26.1
- *thread-num-var* ICV, see Table 3.1
- **omp_get_num_threads_dim** Routine, see Section 29.1

29.3 omp_get_num_teams_dim Routine

Name: omp_get_num_teams_dim **Category:** function	**Properties:** ICV-retrieving, multidimensional-information-routine, teams-nestable

Return Type and Arguments

Name	Type	Properties
<return type>	integer	*default*
dim	integer	non-negative

Prototypes

C / C++

```
int omp_get_num_teams_dim(int dim);
```

C / C++

Fortran

```
integer function omp_get_num_teams_dim(dim)
  integer dim
```

Fortran

Effect

The **omp_get_num_teams_dim** routine returns the element in the multidimensional value of the *league-size-var* ICV associated with dimension *dim*. This value corresponds to the league size for the dimension *dim* in the league space of the current **teams** region. If the call to the routine does not occur in a **teams** region or *dim* corresponds to an inactive spatial dimension in the team space then the routine returns one.

Cross References

- *league-size-var* ICV, see Table 3.1
- **teams** Construct, see Section 18.2

29.4 omp_get_team_num_dim Routine

Name: omp_get_team_num_dim **Category:** function	**Properties:** ICV-retrieving, multidimensional-information-routine, teams-nestable

Return Type and Arguments

Name	Type	Properties
<return type>	integer	*default*
dim	integer	non-negative

Prototypes

C / C++

```
int omp_get_team_num_dim(int dim);
```

C / C++

Fortran

```
integer function omp_get_team_num_dim(dim)
  integer dim
```

Fortran

Effect

The **omp_get_team_num_dim** routine returns the element in the multidimensional value of the *team-num-var* ICV associated with dimension *dim*. This value corresponds to the team number for the dimension *dim* within the league space of the current **teams** region. The result is an integer value between zero and one less than the value returned by **omp_get_num_teams_dim** for dimension *dim*, inclusive. If the call to the routine does not occur in a **teams** region or *dim* corresponds to an inactive spatial dimension in the team space then the routine returns zero.

Cross References

- *team-num-var* ICV, see Table 3.1
- **omp_get_num_teams_dim** Routine, see Section 29.3
- **teams** Construct, see Section 18.2

29.5 omp_get_thread_limit_dim Routine

Name: omp_get_thread_limit_dim **Category:** function	**Properties:** ICV-retrieving, multidimensional-information-routine

Return Type and Arguments

Name	Type	Properties
<return type>	integer	*default*
dim	integer	non-negative

Prototypes

C / C++

```
int omp_get_thread_limit_dim(int dim);
```

C / C++

Fortran

```
integer function omp_get_thread_limit_dim(dim)
  integer dim
```

Fortran

Effect

The **omp_get_thread_limit_dim** routine returns the element of the *thread-limit-var* ICV associated with dimension *dim*. This value corresponds to the maximum number of threads that the team space of any team executing in the current contention group may have in that dimension. If *dim* corresponds to a dimension that is not specified in the *thread-limit-var* ICV then the routine returns one.

Cross References

- *thread-limit-var* ICV, see Table 3.1

29.6 omp_get_supported_active_team_dims Routine

Name: `omp_get_supported_active_team_dims` **Category:** function	**Properties:** multidimensional-information-routine

Return Type and Arguments

Name	Type	Properties
<return type>	integer	*default*
device_num	integer	*default*
native_support	logical	*default*

Prototypes

C / C++

```
int omp_get_supported_active_team_dims(int device_num,
  int native_support);
```

C / C++

Fortran

```
integer function omp_get_supported_active_team_dims(device_num, &
  native_support)
  integer device_num
  logical native_support
```

Fortran

Effect

The **omp_get_supported_active_team_dims** routine returns the number of supported active spatial dimensions in team spaces for the device *device_num*. If the *native_support* argument is specified as *true*, the routine returns the number of natively-supported active spatial dimensions by the device.

Restrictions

Restrictions to the **omp_get_supported_active_team_dims** routine are as follows:

- The *device_num* argument must be a conforming device number.
- When called from within a **target** region, the effect is unspecified.

29.7 omp_get_supported_active_league_dims Routine

Name: `omp_get_supported_active_league_dims` **Category:** function	**Properties:** multidimensional-information-routine

Return Type and Arguments

Name	Type	Properties
<return type>	integer	*default*
device_num	integer	*default*
native_support	logical	*default*

Prototypes

C / C++

```
int omp_get_supported_active_league_dims(int device_num,
  int native_support);
```

C / C++

Fortran

```
integer function omp_get_supported_active_league_dims(&
  device_num, native_support)
  integer device_num
  logical native_support
```

Fortran

Effect

The **omp_get_supported_active_league_dims** routine returns the number of supported active spatial dimensions for league spaces in the device *device_num*. If the *native_support* argument is specified as *true*, the routine returns the number of natively-supported active spatial dimensions by the device.

Restrictions

Restrictions to the **omp_get_supported_active_league_dims** routine are as follows:

- The *device_num* argument must be a conforming device number.
- When called from within a **target** region, the effect is unspecified.

29.8 omp_get_supported_threads_dim Routine

Name: omp_get_supported_threads_dim **Category:** function	**Properties:** multidimensional-information-routine

Return Type and Arguments

Name	Type	Properties
<return type>	integer	*default*
device_num	integer	*default*
dim	integer	non-negative

Prototypes

C / C++

```
int omp_get_supported_threads_dim(int device_num, int dim);
```

C / C++

Fortran

```
integer function omp_get_supported_threads_dim(device_num, dim)
  integer device_num, dim
```

Fortran

Effect

The **omp_get_supported_threads_dim** routine returns the maximum number of threads that the implementation may provide in a team space for dimension *dim* on the device *device_num*. If *dim* corresponds to a dimension that is not included in the supported active spatial dimensions then the routine returns one.

Restrictions

Restrictions to the **omp_get_supported_threads_dim** routine are as follows:

- The *device_num* argument must be a conforming device number.
- When called from within a **target** region, the effect is unspecified.

29.9 omp_get_supported_teams_dim Routine

Name: omp_get_supported_teams_dim **Category:** function	**Properties:** multidimensional-information-routine

Return Type and Arguments

Name	Type	Properties
<return type>	integer	*default*
device_num	integer	*default*
dim	integer	non-negative

Prototypes

C / C++

```
int omp_get_supported_teams_dim(int device_num, int dim);
```

C / C++

Fortran

```
integer function omp_get_supported_teams_dim(device_num, dim)
  integer device_num, dim
```

Fortran

Effect

The **omp_get_supported_teams_dim** routine returns the maximum number of teams that the implementation may provide in a league space for dimension *dim* on the device *device_num*. If *dim* corresponds to a dimension that is not included in the supported active spatial dimensions then the routine returns one.

Restrictions

Restrictions to the **omp_get_supported_teams_dim** routine are as follows:

- The *device_num* argument must be a conforming device number.
- When called from within a **target** region, the effect is unspecified.

30 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.

30.1 Tasking Routines

This section describes routines that pertain to OpenMP explicit tasks.

30.1.1 `omp_get_max_task_priority` Routine

Name: `omp_get_max_task_priority` **Category:** function	**Properties:** all-device-threads-binding, ICV-retrieving

Return Type

Name	Type	Properties
<return type>	integer	*default*

Prototypes

C / C++

```
int omp_get_max_task_priority(void);
```

C / C++

Fortran

```
integer function omp_get_max_task_priority()
```

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

30.1.2 `omp_in_explicit_task` Routine

Name: `omp_in_explicit_task` **Category:** function	**Properties:** ICV-retrieving

Return Type

Name	Type	Properties
<return type>	logical	*default*

Prototypes

C / C++

```
int omp_in_explicit_task(void);
```

C / C++

Fortran

```
logical function omp_in_explicit_task()
```

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

30.1.3 `omp_in_final` Routine

Name: `omp_in_final` **Category:** function	**Properties:** ICV-retrieving

Return Type

Name	Type	Properties
<return type>	logical	*default*

Prototypes

C / C++

```
int omp_in_final(void);
```

C / C++

Fortran

```
logical function omp_in_final()
```

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

30.1.4 `omp_is_free_agent` Routine

Name: `omp_is_free_agent` **Category:** function	**Properties:** ICV-retrieving

Return Type

Name	Type	Properties
<return type>	logical	*default*

Prototypes

C / C++

```
int omp_is_free_agent(void);
```

C / C++

Fortran

```
logical function omp_is_free_agent()
```

Fortran

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

30.1.5 `omp_ancestor_is_free_agent` Routine

Name: `omp_ancestor_is_free_agent` **Category:** function	**Properties:** *default*

Return Type and Arguments

Name	Type	Properties
<return type>	logical	*default*
level	integer	*default*

Prototypes

C / C++

```
int omp_ancestor_is_free_agent(int level);
```

C / C++

Fortran

```
logical function omp_ancestor_is_free_agent(level)
  integer level
```

Fortran

Effect

The **omp_ancestor_is_free_agent** routine returns *true* if the ancestor thread of the encountering thread is a free-agent thread, for a given nested level of the encountering thread; otherwise, it returns *false*. If the requested nesting level is outside the range of 0 and the nesting level of the current task, as returned by the **omp_get_level** routine, the routine returns *false*.

Note – When the **omp_ancestor_is_free_agent** routine is called with a value of *level* =**omp_get_level**, the routine has the same effect as the **omp_is_free_agent** routine.

Cross References

- **omp_get_level** Routine, see Section 27.14
- **omp_is_free_agent** Routine, see Section 30.1.4
- **task** Construct, see Section 20.1
- **threadset** Clause, see Section 20.8

30.2 Event Routine

This section describes routines that support OpenMP event objects.

30.2.1 omp_fulfill_event Routine

Name: omp_fulfill_event **Category:** subroutine	**Properties:** implicit-flush, implicit-release-flush, implicit-release-flush-at-exit

Arguments

Name	Type	Properties
event	event_handle	*default*

Prototypes

C / C++

```
void omp_fulfill_event(omp_event_handle_t event);
```

C / C++

Fortran

```
subroutine omp_fulfill_event(event)
  integer (kind=omp_event_handle_kind) event
```

Fortran

Effect

The effect of this routine is to fulfill the event associated with the *event* argument. The effect of fulfilling the event will depend on how the event object was created. The event object is destroyed and cannot be accessed after calling this routine, and the event handle becomes unassociated with any event object. This routine has no effect if the *event* argument corresponds to a completed task.

If the implicit release-at-exit flush of a detachable task synchronizes with an acquire flush, the implicit release-at-exit flush in the instance of the **omp_fulfill_event** routine that fulfills the *allow-completion* event of that detachable task also synchronizes with that acquire flush.

Execution Model Events

The *task-fulfill* event occurs in a thread that executes an **omp_fulfill_event** region before the event is fulfilled if the OpenMP event object was created by a **detach** clause on a task.

Tool Callbacks

A thread dispatches a registered **task_schedule** callback with NULL as its *next_task_data* argument while the argument *prior_task_data* binds to the detachable task for each occurrence of a *task-fulfill* event. If the *task-fulfill* event occurs before the detachable task finished execution of the associated structured block, the callback has **ompt_task_early_fulfill** as its *prior_task_status* argument; otherwise the callback has **ompt_task_late_fulfill** as its *prior_task_status* argument.

Restrictions

Restrictions to the **omp_fulfill_event** routine are as follows:

- 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 41.5.2
- OMPT **task_status** Type, see Section 40.38

31 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.

31.1 `omp_set_default_device` Routine

Name: `omp_set_default_device` **Category:** subroutine	**Properties:** device-information, ICV-modifying

Arguments

Name	Type	Properties
device_num	integer	*default*

Prototypes

C / C++

```
void omp_set_default_device(int device_num);
```

C / C++

Fortran

```
subroutine omp_set_default_device(device_num)
  integer device_num
```

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

31.2 `omp_get_default_device` Routine

Name: `omp_get_default_device` **Category:** function	**Properties:** device-information, ICV-retrieving

Return Type

Name	Type	Properties
<return type>	integer	*default*

Prototypes

C / C++

```
int omp_get_default_device(void);
```

C / C++

Fortran

```
integer function omp_get_default_device()
```

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

31.3 `omp_get_num_devices` Routine

Name: `omp_get_num_devices` **Category:** function	**Properties:** device-information, ICV-retrieving

Return Type

Name	Type	Properties
<return type>	integer	*default*

Prototypes

C / C++

```
int omp_get_num_devices(void);
```

C / C++

Fortran

```
integer function omp_get_num_devices()
```

Fortran

Effect

The **omp_get_num_devices** routine returns the value of the *num-devices-var* ICV, which is the number of available non-host devices onto which code or data may be offloaded. When called from within a **target** region the effect of this routine is unspecified.

Cross References

- *num-devices-var* ICV, see Table 3.1
- **target** Construct, see Section 21.8

31.4 omp_get_device_num Routine

Name: omp_get_device_num **Category:** function	**Properties:** device-information

Return Type

Name	Type	Properties
<return type>	integer	*default*

Prototypes

C / C++

```
int omp_get_device_num(void);
```

C / C++

Fortran

```
integer function omp_get_device_num()
```

Fortran

Effect

The **omp_get_device_num** routine returns the value of the *device-num-var* ICV, which is the device number of the device on which the encountering thread is executing. When called on the 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

31.5 omp_get_num_procs Routine

Name: omp_get_num_procs **Category:** function	**Properties:** all-device-threads-binding, device-information, ICV-retrieving

Return Type

Name	Type	Properties
<return type>	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

31.6 omp_get_max_progress_width Routine

Name: omp_get_max_progress_width **Category:** function	**Properties:** device-information

Return Type and Arguments

Name	Type	Properties
<return type>	integer	*default*
device_num	integer	*default*

Prototypes

C / C++

```
int omp_get_max_progress_width(int device_num);
```

C / C++

Fortran

```
integer function omp_get_max_progress_width(device_num)
  integer device_num
```

Fortran

Effect

The **omp_get_max_progress_width** routine returns the maximum size, in terms of hardware threads, of progress units on the device specified by *device_num*. When called from within a **target** region the effect of this routine is unspecified.

31.7 `omp_get_device_from_uid` Routine

Name: `omp_get_device_from_uid` **Category:** function	**Properties:** device-information

Return Type and Arguments

Name	Type	Properties
<return type>	integer	*default*
uid	char	pointer, intent(in)

Prototypes

C / C++

```
int omp_get_device_from_uid(const char *uid);
```

C / C++

Fortran

```
integer function omp_get_device_from_uid(uid)
  character(len=*), intent(in) :: uid
```

Fortran

Effect

The **`omp_get_device_from_uid`** routine returns the device number associated with the device specified by the *uid*; if no device with that *uid* is available, the value of **`omp_invalid_device`** 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 31.8

31.8 `omp_get_uid_from_device` Routine

Name: `omp_get_uid_from_device` **Category:** function	**Properties:** device-information

Return Type and Arguments

Name	Type	Properties
<return type>	const char	pointer
device_num	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 31.7

31.9 omp_is_initial_device Routine

Name: omp_is_initial_device **Category:** function	**Properties:** device-information

Return Type

Name	Type	Properties
<return type>	logical	*default*

Prototypes

C / C++

```
int omp_is_initial_device(void);
```

C / C++

Fortran

```
logical function omp_is_initial_device()
```

Fortran

Effect

The **omp_is_initial_device** routine returns *true* if the current task is executing on the host device; otherwise, it returns *false*.

31.10 `omp_get_initial_device` Routine

Name: `omp_get_initial_device` **Category:** function	**Properties:** device-information

Return Type

Name	Type	Properties
<return type>	integer	*default*

Prototypes

C / C++

```
int omp_get_initial_device(void);
```

C / C++

Fortran

```
integer function omp_get_initial_device()
```

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

31.11 `omp_get_device_num_teams` Routine

Name: `omp_get_device_num_teams` **Category:** function	**Properties:** device-information, ICV-retrieving

Return Type and Arguments

Name	Type	Properties
<return type>	integer	*default*
device_num	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 *nteams-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

- *nteams-var* ICV, see Table 3.1
- **num_teams** Clause, see Section 18.2.1
- **teams** Construct, see Section 18.2

31.12 omp_set_device_num_teams Routine

Name: omp_set_device_num_teams **Category:** subroutine	**Properties:** device-information, ICV-modifying

Arguments

Name	Type	Properties
num_teams	integer	non-negative
device_num	integer	*default*

Prototypes

C / C++

```
void omp_set_device_num_teams(int num_teams, int device_num);
```

C / C++

Fortran

```
subroutine omp_set_device_num_teams(num_teams, device_num)
  integer num_teams, device_num
```

Fortran

Effect

The effect of the **omp_set_device_num_teams** routine is to set the value of the *nteams-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

- *nteams-var* ICV, see Table 3.1
- **num_teams** Clause, see Section 18.2.1
- **teams** Construct, see Section 18.2

31.13 omp_get_device_teams_thread_limit Routine

Name: **omp_get_device_teams_thread_limit** **Category:** function	**Properties:** device-information, ICV-retrieving

Return Type and Arguments

Name	Type	Properties
<return type>	integer	*default*
device_num	integer	*default*

Prototypes

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

Effect

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.

Cross References

- *teams-thread-limit-var* ICV, see Table 3.1
- **teams** Construct, see Section 18.2

31.14 omp_set_device_teams_thread_limit Routine

Name: omp_set_device_teams_thread_limit **Category:** subroutine	**Properties:** device-information, ICV-modifying

Arguments

Name	Type	Properties
thread_limit	integer	positive
device_num	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

31.15 omp_get_filtered_devices Routine

Name: omp_get_filtered_devices **Category:** function	**Properties:** device-information

Return Type and Arguments

Name	Type	Properties
<return type>	integer	*default*
filter	char	pointer, intent(in)
ndevs	integer	intent(in), non-negative
devs	integer	intent(out), pointer

Prototypes

C / C++

```
int omp_get_filtered_devices(const char *filter, int ndevs,
  int *devs);
```

C / C++

Fortran

```
integer function omp_get_filtered_devices(filter, ndevs, devs)
  character(len=*), intent(in) :: filter
  integer, intent(in) :: ndevs
  integer, intent(out) :: devs(*)
```

Fortran

Effect

The **omp_get_filtered_devices** routine returns the number of available devices, which may include the host device, that fulfill the trait specification expressed in the *filter* argument. The trait specification is as described for **OMP_AVAILABLE_DEVICES**, except that, in addition, **host** is permitted as *kind-name*. If the number of matching devices is positive, the device numbers of the devices that fulfill the specification are returned in the array *devs*. These device numbers are sorted in ascending order. If the return value is larger than *ndevs*, the list of device numbers is truncated. If *ndevs* is zero, no values are returned in *devs* and it may be NULL.

Restrictions

Restrictions to the **omp_get_filtered_devices** routine are as follows:

- The routine must only be invoked on the host device.
- The *devs* argument must point to an array that contains at least *ndevs* values.

Cross References

- **OMP_AVAILABLE_DEVICES**, see Section 4.3.7

31.16 omp_get_device_distances Routine

Name: omp_get_device_distances **Category:** subroutine	**Properties:** device-information

Arguments

Name	Type	Properties
ndevs	integer	intent(in), positive
devs	integer	intent(in), pointer
distances	integer	intent(out), pointer

Prototypes

C / C++

```
void omp_get_device_distances(int ndevs, const int *devs,
  int *distances);
```

C / C++

Fortran

```
subroutine omp_get_device_distances(ndevs, devs, distances)
  integer, intent(in) :: ndevs, devs(*)
  integer, intent(out) :: distances(*)
```

Fortran

Effect

The **omp_get_device_distances** routine returns *ndevs* non-negative values in the array *distances* that may be used to guide device selection for device constructs. The value of each returned array element is an implementation defined relative distance between the place on which the routine is executing and the device with the device number specified by the corresponding array element of the *devs* argument. If a value of the *devs* argument specifies the host device, the corresponding distance is 0.

Restrictions

Restrictions to the **omp_get_device_distances** routine are as follows:

- The routine must only be invoked on the host device.
- The *devs* and *distances* arguments must point to arrays that contain at least *ndevs* values.
- The *ndevs* values of the array to which *devs* points must be conforming device numbers, which may include the host device but not **omp_invalid_device**.

32 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* **& ompt_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.

Cross References

- **target** Construct, see Section 21.8
- **task** Construct, see Section 20.1
- OMPT **task_flag** Type, see Section 40.37

32.1 Asynchronous Device Memory Routines

Some device memory routines have the asynchronous-device routine property. The target task that is generated by the call to an asynchronous device routines is deferrable. Task dependences are expressed with zero or more OpenMP depend objects. The dependences are specified by passing the number of depend objects followed by an array of the objects. The generated target task is not a dependent task if the program passes in a count of zero for *depobj_count*. The *depobj_list* argument is ignored if the value of *depobj_count* is zero.

Execution Model Events

Events associated with task dependences that result from *depobj_list* are the same as for a **depend** clause with the **depobj** *task-dependence-type* defined in Section 23.9.6.

Tool Callbacks

Callbacks associated with events for task dependences are the same as for the **depend** clause defined in Section 23.9.6.

Cross References

- **depend** Clause, see Section 23.9.6
- **depobj** Construct, see Section 23.9.3

32.2 Device Memory Information Routines

This section describes routines that have the device-memory-information routine property. These device-memory-information routines provide information about device pointers, which can be determined without directly accessing the target device; thus, they do not create a target task.

32.2.1 omp_target_is_present Routine

Name: omp_target_is_present **Category:** function	**Properties:** device-memory-information-routine, device-memory-routine, iso_c_binding

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

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

Effect

The **omp_target_is_present** routine returns a non-zero value if *device_num* refers to the host device or if *ptr* refers to storage that has corresponding storage in the device data environment of device *device_num*. Otherwise, the routine returns zero. If *ptr* is NULL. the routine returns zero. Thus, the **omp_target_is_present** routine tests whether a host pointer refers to storage that is mapped to a given device.

Restrictions

Restrictions to the **omp_target_is_present** routine are as follows:

- The value of *ptr* must be a valid host pointer or NULL.

32.2.2 omp_target_is_accessible Routine

Name: omp_target_is_accessible **Category:** function	**Properties:** device-memory-information-routine, device-memory-routine, iso_c_binding

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

Prototypes

C / C++

```
int omp_target_is_accessible(const void *ptr, size_t size,
  int device_num);
```

C / C++

Fortran

```
integer (kind=c_int) function omp_target_is_accessible(ptr, &
  size, device_num) bind(c)
  use, intrinsic :: iso_c_binding, only : c_int, c_ptr, &
    c_size_t
  type (c_ptr), value, intent(in) :: ptr
  integer (kind=c_size_t), value :: size
  integer (kind=c_int), value :: device_num
```

Fortran

Effect

The **omp_target_is_accessible** routine returns a non-zero value if the storage of *size* bytes that corresponds to the address range starting at the address given by *ptr* is accessible from device *device_num*. Otherwise, it returns zero. If *ptr* is NULL or the implementation cannot guarantee accessibility, the routine returns zero. The value of *ptr* is interpreted as an address in the address space of the specified device.

32.2.3 omp_get_mapped_ptr Routine

Name: omp_get_mapped_ptr **Category:** function	**Properties:** device-memory-information-routine, device-memory-routine, iso_c_binding

Return Type and Arguments

Name	Type	Properties
<return type>	c_ptr	*default*
ptr	c_ptr	intent(in), iso_c, value
device_num	c_int	iso_c, value

Prototypes

C / C++

```
void  *omp_get_mapped_ptr(const void *ptr, int device_num);
```

C / C++

Fortran

```
type (c_ptr) function omp_get_mapped_ptr(ptr, device_num) &
  bind(c)
  use, intrinsic :: iso_c_binding, only : c_ptr, c_int
  type (c_ptr), value, intent(in) :: ptr
  integer (kind=c_int), value :: device_num
```

Fortran

Effect

The **omp_get_mapped_ptr** routine returns the associated device pointer for host pointer *ptr* on device *device_num*. A call to this routine for a pointer that is not NULL and does not have an associated pointer on the given device will return NULL. The routine returns NULL if unsuccessful. Otherwise it returns the device pointer, which is *ptr* if *device_num* specifies the host device.

Cross References

- **omp_get_initial_device** Routine, see Section 31.10

32.3 omp_target_alloc Routine

Name: omp_target_alloc **Category:** function	**Properties:** device-memory-routine, generating-task-binding, iso_c_binding

Return Type and Arguments

Name	Type	Properties
<return type>	c_ptr	*default*
size	c_size_t	iso_c, value
device_num	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 7.3.5).

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 7.3.5
- **omp_target_free** Routine, see Section 32.4
- OMPT **scope_endpoint** Type, see Section 40.27
- **target_data_op_emi** Callback, see Section 42.7

32.4 omp_target_free Routine

Name: omp_target_free **Category:** subroutine	**Properties:** device-memory-routine, generating-task-binding, iso_c_binding

Arguments

Name	Type	Properties
device_ptr	c_ptr	iso_c, value
device_num	c_int	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.

Restrictions

Restrictions to the **omp_target_free** routine are as follows:

- The value of *device_ptr* must be NULL or have been returned by **omp_target_alloc**.

Cross References

- **omp_target_alloc** Routine, see Section 32.3
- OMPT **scope_endpoint** Type, see Section 40.27
- **target_data_op_emi** Callback, see Section 42.7

32.5 omp_target_associate_ptr Routine

Name: omp_target_associate_ptr **Category:** function	**Properties:** device-memory-routine, generating-task-binding, iso_c_binding

Return Type and Arguments

Name	Type	Properties
<return type>	c_int	*default*
host_ptr	c_ptr	intent(in), iso_c, value
device_ptr	c_ptr	intent(in), iso_c, value
size	c_size_t	iso_c, value
device_offset	c_size_t	iso_c, value
device_num	c_int	iso_c, value

Prototypes

C / C++

```
int omp_target_associate_ptr(const void *host_ptr,
  const void *device_ptr, size_t size, size_t device_offset,
  int device_num);
```

C / C++

Fortran

```
integer (kind=c_int) function omp_target_associate_ptr(host_ptr, &
  device_ptr, size, device_offset, device_num) bind(c)
  use, intrinsic :: iso_c_binding, only : c_int, c_ptr, &
    c_size_t
  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 32.3
- **omp_target_disassociate_ptr** Routine, see Section 32.6
- **omp_target_is_present** Routine, see Section 32.2.1
- OMPT **scope_endpoint** Type, see Section 40.27
- **target_data_op_emi** Callback, see Section 42.7

32.6 omp_target_disassociate_ptr Routine

Name: omp_target_disassociate_ptr **Category:** function	**Properties:** device-memory-routine, generating-task-binding, iso_c_binding

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

Prototypes

C / C++

```
int omp_target_disassociate_ptr(const void *ptr, int device_num);
```

C / C++

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
```

Fortran

Effect

The **omp_target_disassociate_ptr** removes the association of the host pointer *ptr* with a device pointer in the device data environment of device *device_num*. A call to this routine on a pointer that is not NULL and does not have associated data on the given device results in unspecified behavior. The reference count of the mapping is reduced to zero, regardless of its current value. The routine returns zero if successful. Otherwise it returns a non-zero value.

Execution Model Events

The *target-data-disassociate* event occurs before a thread initiates a device pointer disassociation on a target device.

Tool Callbacks

A thread dispatches a registered **target_data_op_emi** callback with **ompt_scope_beginend** as its *endpoint* argument for each occurrence of a *target-data-disassociate* event in that thread.

Cross References

- OMPT **scope_endpoint** Type, see Section 40.27
- **target_data_op_emi** Callback, see Section 42.7

32.7 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 40.27
- **`target_data_op_emi`** Callback, see Section 42.7

32.7.1 `omp_target_memcpy` Routine

Name: `omp_target_memcpy` **Category:** function	**Properties:** device-memory-routine, flat-memory-copying, generating-task-binding, iso_c_binding, memory-copying, task-scheduling-point, task-scheduling-point-at-entry, task-scheduling-point-at-exit

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
length	c_size_t	iso_c, value
dst_offset	c_size_t	iso_c, value
src_offset	c_size_t	iso_c, value
dst_device_num	c_int	iso_c, value
src_device_num	c_int	iso_c, value

Prototypes

C / C++

```
int omp_target_memcpy(void *dst, const void *src, size_t length,
  size_t dst_offset, size_t src_offset, int dst_device_num,
  int src_device_num);
```

C / C++

Fortran

```
integer (kind=c_int) function omp_target_memcpy(dst, src, &
  length, dst_offset, src_offset, dst_device_num, &
  src_device_num) bind(c)
  use, intrinsic :: iso_c_binding, only : c_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
  integer (kind=c_int), value :: dst_device_num, src_device_num
```

Fortran

Effect

As a flat-memory-copying routine, the effect of the **omp_target_memcpy** routine is as described in Section 32.7. This effect includes the associated tool events and callbacks defined in that section.

Cross References

- Memory Copying Routines, see Section 32.7

32.7.2 `omp_target_memcpy_rect` Routine

Name: `omp_target_memcpy_rect` **Category:** function	**Properties:** device-memory-routine, generating-task-binding, iso_c_binding, memory-copying, rectangular-memory-copying, task-scheduling-point, task-scheduling-point-at-entry, task-scheduling-point-at-exit

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

Prototypes

C / C++

```
int omp_target_memcpy_rect(void *dst, const void *src,
  size_t element_size, int num_dims, const size_t *volume,
  const size_t *dst_offsets, const size_t *src_offsets,
  const size_t *dst_dimensions, const size_t *src_dimensions,
  int dst_device_num, int src_device_num);
```

C / C++

Fortran

```
integer (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(*)
```

Fortran

Effect

As a rectangular-memory-copying routine, the effect of the **omp_target_memcpy_rect** routine is as described in Section 32.7. This effect includes the associated tool events and callbacks defined in that section.

Cross References

- Memory Copying Routines, see Section 32.7

32.7.3 omp_target_memcpy_async Routine

Name: omp_target_memcpy_async **Category:** function	**Properties:** asynchronous-device-routine, device-memory-routine, flat-memory-copying, generating-task-binding, iso_c_binding, memory-copying, overloaded, task-scheduling-point, task-scheduling-point-at-entry, task-scheduling-point-at-exit

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
length	c_size_t	iso_c, value
dst_offset	c_size_t	iso_c, value
src_offset	c_size_t	iso_c, value
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

```
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++

```
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 = 0,
  omp_depend_t *depobj_list = NULL);
```

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
  integer (kind=c_int), value :: dst_device_num, src_device_num, &
    depobj_count
  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 32.7. 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 32.1.

Cross References

- Asynchronous Device Memory Routines, see Section 32.1
- Memory Copying Routines, see Section 32.7

32.7.4 omp_target_memcpy_rect_async Routine

Name: omp_target_memcpy_rect_async **Category:** function	**Properties:** asynchronous-device-routine, device-memory-routine, generating-task-binding, iso_c_binding, memory-copying, rectangular-memory-copying, overloaded, task-scheduling-point, task-scheduling-point-at-entry, task-scheduling-point-at-exit

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

```
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++

```
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 = 0,
  omp_depend_t *depobj_list = NULL);
```

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 32.7. 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 32.1.

Cross References

- Asynchronous Device Memory Routines, see Section 32.1
- Memory Copying Routines, see Section 32.7

32.8 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 40.27
- **target_data_op_emi** Callback, see Section 42.7

32.8.1 omp_target_memset Routine

Name: omp_target_memset **Category:** function	**Properties:** device-memory-routine, generating-task-binding, iso_c_binding, memory-setting, task-scheduling-point, task-scheduling-point-at-entry, task-scheduling-point-at-exit

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

Prototypes

C / C++

```
void  *omp_target_memset(void *ptr, int val, size_t count,
  int device_num);
```

C / C++

Fortran

```
type (c_ptr) function omp_target_memset(ptr, val, count, &
  device_num) 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
  integer (kind=c_size_t), value :: count
```

Fortran

Effect

As a memory-setting routine, the effect of the **omp_target_memset** routine is as described in Section 32.8. This effect includes the tool events and callbacks defined in that section.

Cross References

- Memory Setting Routines, see Section 32.8

32.8.2 omp_target_memset_async Routine

Name: omp_target_memset_async **Category:** function	**Properties:** asynchronous-device-routine, device-memory-routine, generating-task-binding, iso_c_binding, memory-setting, overloaded, task-scheduling-point, task-scheduling-point-at-entry, task-scheduling-point-at-exit

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

```
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++

```
void  *omp_target_memset_async(void *ptr, int val, size_t count,
  int device_num, int depobj_count = 0,
  omp_depend_t *depobj_list = NULL);
```

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 32.8. 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 32.1.

Cross References

- Asynchronous Device Memory Routines, see Section 32.1
- Memory Setting Routines, see Section 32.8

33 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

33.1 omp_get_num_interop_properties Routine

Name: omp_get_num_interop_properties Category: function	Properties: interoperability-routine

Return Type and Arguments

Name	Type	Properties
<return type>	integer	*default*
interop	interop	intent(in)

Prototypes

C / C++

```
int omp_get_num_interop_properties(const omp_interop_t interop);
```

C / C++

Fortran

```
integer function omp_get_num_interop_properties(interop)
  integer (kind=omp_interop_kind), intent(in) :: interop
```

Fortran

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

33.2 omp_get_interop_int Routine

Name: omp_get_interop_int Category: function	Properties: interoperability-property-retrieving, interoperability-routine, overloaded

Return Type and Arguments

Name	Type	Properties
<return type>	intptr	*default*
interop	interop	omp, opaque, intent(in)
property_id	interop_property	omp
ret_code	interop_rc	omp, intent(out), optional

Prototypes

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++

```
omp_intptr_t omp_get_interop_int(const omp_interop_t interop,
  omp_interop_property_t property_id,
  omp_interop_rc_t *ret_code = NULL);
```

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 33.1

33.3 omp_get_interop_ptr Routine

Name: omp_get_interop_ptr **Category:** function	**Properties:** interoperability-property-retrieving, interoperability-routine, overloaded

Return Type and Arguments

Name	Type	Properties
<return type>	c_ptr	*default*
interop	interop	omp, opaque, intent(in)
property_id	interop_property	omp
ret_code	interop_rc	omp, intent(out), optional

Prototypes

C

```
void  *omp_get_interop_ptr(const omp_interop_t interop,
  omp_interop_property_t property_id, omp_interop_rc_t *ret_code);
```

C

C++

```
void  *omp_get_interop_ptr(const omp_interop_t interop,
  omp_interop_property_t property_id,
  omp_interop_rc_t *ret_code = NULL);
```

C++

Fortran

```
type (c_ptr) function omp_get_interop_ptr(interop, property_id, &
  ret_code)
  use, intrinsic :: iso_c_binding, only : c_ptr
  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_ptr** routine is an interoperability-property-retrieving routine that retrieves an interoperability property of pointer 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 33.1

33.4 `omp_get_interop_str` Routine

Name: `omp_get_interop_str` **Category:** function	**Properties:** interoperability-property-retrieving, interoperability-routine, overloaded

Return Type and Arguments

Name	Type	Properties
<return type>	const char	pointer
interop	interop	omp, opaque, intent(in)
property_id	interop_property	omp
ret_code	interop_rc	omp, intent(out), optional

Prototypes

C

```
const char  *omp_get_interop_str(const omp_interop_t interop,
  omp_interop_property_t property_id, omp_interop_rc_t *ret_code);
```

C

C++

```
const char  *omp_get_interop_str(const omp_interop_t interop,
  omp_interop_property_t property_id,
  omp_interop_rc_t *ret_code = NULL);
```

C++

Fortran

```
character(:) function omp_get_interop_str(interop, property_id, &
  ret_code)
  pointer :: omp_get_interop_str
  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_str`** routine is an interoperability-property-retrieving routine that retrieves an interoperability string property type as a string, 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 33.1

33.5 omp_get_interop_name Routine

Name: omp_get_interop_name **Category:** function	**Properties:** interoperability-routine

Return Type and Arguments

Name	Type	Properties
<return type>	const char	pointer
interop	interop	omp, opaque, intent(in)
property_id	interop_property	omp

Prototypes

C / C++

```
const char  *omp_get_interop_name(const omp_interop_t interop,
  omp_interop_property_t property_id);
```

C / C++

Fortran

```
character(:) function omp_get_interop_name(interop, property_id)
  pointer :: omp_get_interop_name
  integer (kind=omp_interop_kind), intent(in) :: interop
  integer (kind=omp_interop_property_kind) property_id
```

Fortran

Effect

The **omp_get_interop_name** routine returns, as a string, the name of the interoperability property identified by *property_id*. Property names for non-implementation defined interoperability properties are listed in Table 26.2. If the *property_id* is less than **omp_ipr_first** or greater than or equal to **omp_get_num_interop_properties**(*interop*), NULL is returned.

Cross References

- OpenMP **interop** Type, see Section 26.7.1
- OpenMP **interop_property** Type, see Section 26.7.3
- **omp_get_num_interop_properties** Routine, see Section 33.1

33.6 `omp_get_interop_type_desc` Routine

Name: `omp_get_interop_type_desc` **Category:** function	**Properties:** interoperability-routine

Return Type and Arguments

Name	Type	Properties
<return type>	const char	pointer
interop	interop	omp, opaque, intent(in)
property_id	interop_property	omp

Prototypes

C / C++

```
const char  *omp_get_interop_type_desc(
  const omp_interop_t interop, omp_interop_property_t property_id);
```

C / C++

Fortran

```
character(:) function omp_get_interop_type_desc(interop, &
  property_id)
  pointer :: omp_get_interop_type_desc
  integer (kind=omp_interop_kind), intent(in) :: interop
  integer (kind=omp_interop_property_kind) property_id
```

Fortran

Effect

The **`omp_get_interop_type_desc`** routine returns a string that describes the type of the interoperability property identified by *property_id* in human-readable form. The description may contain a valid type declaration, possibly followed by a description or name of the type. If *interop* has the value **`omp_interop_none`**, or if the *property_id* is less than **`omp_ipr_first`** or greater than or equal to **`omp_get_num_interop_properties`**`(`*interop*`)`, NULL is returned.

Cross References

- OpenMP **`interop`** Type, see Section 26.7.1
- OpenMP **`interop_property`** Type, see Section 26.7.3
- **`omp_get_num_interop_properties`** Routine, see Section 33.1

33.7 `omp_get_interop_rc_desc` Routine

Name: `omp_get_interop_rc_desc` **Category:** function	**Properties:** interoperability-routine

Return Type and Arguments

Name	Type	Properties
<return type>	const char	pointer
interop	interop	omp, opaque, intent(in)
ret_code	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 33.1

34 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

34.1 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.

Restrictions

The restrictions to memory-space-retrieving routines are as follows:

- These routines must only be invoked on the host device.
- The *memspace* argument must be one of the predefined memory spaces.

- For any memory-space-retrieving routine that has a *devs* argument, the argument must point to an array that contains at least *ndevs* values.
- For any memory-space-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 Spaces, see Section 13.1
- **`requires`** Directive, see Section 16.5
- **`target`** Construct, see Section 21.8

34.1.1 `omp_get_devices_memspace` Routine

Name: `omp_get_devices_memspace` **Category:** function	**Properties:** all-device-threads-binding, memory-management-routine, memory-space-retrieving

Return Type and Arguments

Name	Type	Properties
<return type>	memspace_handle	*default*
ndevs	integer	intent(in), positive
devs	integer	intent(in), pointer
memspace	memspace_handle	intent(in), omp

Prototypes

C / C++

```
omp_memspace_handle_t omp_get_devices_memspace(int ndevs,
  const int *devs, omp_memspace_handle_t memspace);
```

C / C++

Fortran

```
integer (kind=omp_memspace_handle_kind) function &
  omp_get_devices_memspace(ndevs, devs, memspace)
  integer, intent(in) :: ndevs, devs(*)
  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 34.1
- OpenMP **memspace_handle** Type, see Section 26.8.12

34.1.2 omp_get_device_memspace Routine

Name: omp_get_device_memspace **Category:** function	**Properties:** all-device-threads-binding, memory-management-routine, memory-space-retrieving

Return Type and Arguments

Name	Type	Properties
<return type>	memspace_handle	*default*
dev	integer	intent(in)
memspace	memspace_handle	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 34.1
- OpenMP **memspace_handle** Type, see Section 26.8.12

34.1.3 omp_get_devices_and_host_memspace Routine

Name: omp_get_devices_and_host_memspace **Category:** function	**Properties:** all-device-threads-binding, memory-management-routine, memory-space-retrieving

Return Type and Arguments

Name	Type	Properties
<return type>	memspace_handle	*default*
ndevs	integer	intent(in), positive
devs	integer	intent(in), pointer
memspace	memspace_handle	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 34.1
- OpenMP **memspace_handle** Type, see Section 26.8.12

34.1.4 omp_get_device_and_host_memspace Routine

Name: **omp_get_device_and_host_memspace** **Category:** function	**Properties:** all-device-threads-binding, memory-management-routine, memory-space-retrieving

Return Type and Arguments

Name	Type	Properties
<return type>	memspace_handle	*default*
dev	integer	intent(in)
memspace	memspace_handle	intent(in), omp

Prototypes

C / C++

```
omp_memspace_handle_t omp_get_device_and_host_memspace(int dev,
  omp_memspace_handle_t memspace);
```

C / C++

Fortran

```
integer (kind=omp_memspace_handle_kind) function &
  omp_get_device_and_host_memspace(dev, memspace)
  integer, intent(in) :: dev
  integer (kind=omp_memspace_handle_kind), intent(in) :: memspace
```

Fortran

Effect

The **omp_get_device_and_host_memspace** routine is a memory-space-retrieving routine. The devices selected by the routine are the host device and the device specified in the *dev* argument.

Cross References

- Memory Space Retrieving Routines, see Section 34.1
- OpenMP **memspace_handle** Type, see Section 26.8.12

34.1.5 omp_get_devices_all_memspace Routine

Name: omp_get_devices_all_memspace **Category:** function	**Properties:** all-device-threads-binding, memory-management-routine, memory-space-retrieving

Return Type and Arguments

Name	Type	Properties
<return type>	memspace_handle	*default*
memspace	memspace_handle	intent(in), omp

Prototypes

C / C++

```
omp_memspace_handle_t omp_get_devices_all_memspace(
  omp_memspace_handle_t memspace);
```

C / C++

Fortran

```
integer (kind=omp_memspace_handle_kind) function &
  omp_get_devices_all_memspace(memspace)
  integer (kind=omp_memspace_handle_kind), intent(in) :: memspace
```

Fortran

Effect

The **omp_get_devices_all_memspace** routine is a memory-space-retrieving routine. The devices selected by the routine are all available devices.

Cross References

- Memory Space Retrieving Routines, see Section 34.1
- OpenMP **memspace_handle** Type, see Section 26.8.12

34.2 omp_get_memspace_num_resources Routine

Name: omp_get_memspace_num_resources **Category:** function	**Properties:** all-device-threads-binding, memory-management-routine

Return Type and Arguments

Name	Type	Properties
<return type>	integer	*default*
memspace	memspace_handle	intent(in), omp

Prototypes

C / C++

```
int omp_get_memspace_num_resources(
  omp_memspace_handle_t memspace);
```

C / C++

Fortran

```
integer function omp_get_memspace_num_resources(memspace)
  integer (kind=omp_memspace_handle_kind), intent(in) :: memspace
```

Fortran

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

34.3 `omp_get_memspace_pagesize` Routine

Name: `omp_get_memspace_pagesize` **Category:** function	**Properties:** all-device-threads-binding, iso_c_binding, memory-management-routine

Return Type and Arguments

Name	Type	Properties
<return type>	c_size_t	*default*
memspace	memspace_handle	intent(in), omp

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.

Restrictions

The restrictions to the **`omp_get_memspace_pagesize`** 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

34.4 `omp_get_submemspace` Routine

Name: `omp_get_submemspace` **Category:** function	**Properties:** all-device-threads-binding, memory-management-routine

Return Type and Arguments

Name	Type	Properties
<return type>	memspace_handle	*default*
memspace	memspace_handle	intent(in), omp
num_resources	integer	intent(in), non-negative
resources	integer	intent(in), pointer

Prototypes

C / C++

```
omp_memspace_handle_t omp_get_submemspace(
  omp_memspace_handle_t memspace, int num_resources,
  const int *resources);
```

C / C++

Fortran

```
integer (kind=omp_memspace_handle_kind) function &
  omp_get_submemspace(memspace, num_resources, resources)
  integer (kind=omp_memspace_handle_kind), intent(in) :: memspace
  integer, intent(in) :: num_resources, resources(*)
```

Fortran

Effect

The **omp_get_submemspace** routine is a memory-management routine that returns a new memory space that contains a subset of the resources of the original memory space. The new memory space represents only the resources of the memory space represented by the *memspace* handle that are specified by the *resources* argument. If *num_resources* is zero or a memory space cannot be created for the requested resources, the special value **omp_null_mem_space** is returned.

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

34.5 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.

34.5.1 `omp_init_mempartitioner` Routine

Name: `omp_init_mempartitioner` **Category:** subroutine	**Properties:** all-device-threads-binding, memory-management-routine, memory-partitioning

Arguments

Name	Type	Properties
partitioner	mempartitioner	C/C++ pointer, omp, intent(out)
lifetime	mempartitioner_lifetime	omp, intent(in)
compute_proc	mempartitioner_compute_proc	omp, procedure
release_proc	mempartitioner_release_proc	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

34.5.2 omp_destroy_mempartitioner Routine

Name: omp_destroy_mempartitioner **Category:** subroutine	**Properties:** all-device-threads-binding, memory-management-routine, memory-partitioning

Arguments

Name	Type	Properties
partitioner	mempartitioner	C/C++ pointer, omp, intent(in)

Prototypes

C / C++

```
void omp_destroy_mempartitioner(
  const omp_mempartitioner_t *partitioner);
```

C / C++

Fortran

```
subroutine omp_destroy_mempartitioner(partitioner)
  integer (kind=omp_mempartitioner_kind), intent(in) :: &
    partitioner
```

Fortran

Effect

The effect of the **omp_destroy_mempartitioner** routine is to uninitialize 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

34.5.3 omp_init_mempartition Routine

Name: omp_init_mempartition **Category:** subroutine	**Properties:** all-device-threads-binding, iso_c_binding, memory-management-routine, memory-partitioning

Arguments

Name	Type	Properties
partition	mempartition	C/C++ pointer, omp, intent(out)
nparts	c_size_t	intent(in), iso_c, intent(in)
user_data	c_ptr	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* memory 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

34.5.4 omp_destroy_mempartition Routine

Name: omp_destroy_mempartition **Category:** subroutine	**Properties:** all-device-threads-binding, memory-management-routine, memory-partitioning

Arguments

Name	Type	Properties
partition	mempartition	C/C++ pointer, omp, intent(in)

Prototypes

C / C++

```
void omp_destroy_mempartition(
  const omp_mempartition_t *partition);
```

C / C++

Fortran

```
subroutine omp_destroy_mempartition(partition)
  integer (kind=omp_mempartition_kind), intent(in) :: partition
```

Fortran

Effect

The effect of the **omp_destroy_mempartition** routine is to uninitialize 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

34.5.5 omp_mempartition_set_part Routine

Name: omp_mempartition_set_part **Category:** function	**Properties:** all-device-threads-binding, iso_c_binding, memory-management-routine, memory-partitioning

Return Type and Arguments

Name	Type	Properties
<return type>	integer	*default*
partition	mempartition	C/C++ pointer, omp, intent(out)
part	c_size_t	intent(in), iso_c
resource	integer	intent(in), iso_c
size	c_size_t	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 memory 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

34.5.6 omp_mempartition_get_user_data Routine

Name: **omp_mempartition_get_user_data** **Category:** function	**Properties:** all-device-threads-binding, iso_c_binding, memory-management-routine, memory-partitioning

Return Type and Arguments

Name	Type	Properties
<return type>	c_ptr	*default*
partition	mempartition	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.

Restrictions

The restrictions to the **omp_mempartition_get_user_data** 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

34.6 omp_init_allocator Routine

Name: omp_init_allocator **Category:** function	**Properties:** all-device-threads-binding, memory-management-routine

Return Type and Arguments

Name	Type	Properties
<return type>	allocator_handle	*default*
memspace	memspace_handle	intent(in), omp
ntraits	integer	intent(in)
traits	alloctrait	intent(in), pointer, omp

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

Effect

The **omp_init_allocator** routine creates a new allocator that is associated with the *memspace* memory space and returns a handle to it. All allocations through the created allocator will behave according to the allocator traits specified in the *traits* argument. The number of traits in the *traits* argument is specified by the *ntraits* argument. If the special **omp_atv_default** value 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

34.7 omp_destroy_allocator Routine

Name: omp_destroy_allocator **Category:** subroutine	**Properties:** all-device-threads-binding, memory-management-routine

Arguments

Name	Type	Properties
allocator	allocator_handle	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

34.8 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

34.8.1 `omp_get_devices_allocator` Routine

Name: `omp_get_devices_allocator` **Category:** function	**Properties:** all-device-threads-binding, memory-management-routine, memory-allocator-retrieving

Return Type and Arguments

Name	Type	Properties
<return type>	allocator_handle	*default*
ndevs	integer	intent(in), positive
devs	integer	intent(in), pointer
memspace	memspace_handle	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 34.8
- OpenMP **`memspace_handle`** Type, see Section 26.8.12

34.8.2 `omp_get_device_allocator` Routine

Name: `omp_get_device_allocator` **Category:** function	**Properties:** all-device-threads-binding, memory-management-routine, memory-allocator-retrieving

Return Type and Arguments

Name	Type	Properties
<return type>	allocator_handle	*default*
dev	integer	intent(in)
memspace	memspace_handle	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 34.8
- OpenMP **`memspace_handle`** Type, see Section 26.8.12

34.8.3 `omp_get_devices_and_host_allocator` Routine

Name: `omp_get_devices_and_host_allocator` **Category:** function	**Properties:** all-device-threads-binding, memory-management-routine, memory-allocator-retrieving

Return Type and Arguments

Name	Type	Properties
<return type>	allocator_handle	*default*
ndevs	integer	intent(in), positive
devs	integer	intent(in), pointer
memspace	memspace_handle	intent(in), omp

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

Effect

The **omp_get_devices_and_host_allocator** routine is a memory-allocator-retrieving routine. The devices selected by the routine are the host device and those specified in the *devs* argument.

Cross References

- OpenMP **allocator_handle** Type, see Section 26.8.2
- Memory Allocator Retrieving Routines, see Section 34.8
- OpenMP **memspace_handle** Type, see Section 26.8.12

34.8.4 omp_get_device_and_host_allocator Routine

Name: **omp_get_device_and_host_allocator** **Category:** function	**Properties:** all-device-threads-binding, memory-management-routine, memory-allocator-retrieving

Return Type and Arguments

Name	Type	Properties
<return type>	allocator_handle	*default*
dev	integer	intent(in)
memspace	memspace_handle	intent(in), omp

Prototypes

C / C++

```
omp_allocator_handle_t omp_get_device_and_host_allocator(int dev,
  omp_memspace_handle_t memspace);
```

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
```

Fortran

Effect

The **omp_get_device_and_host_allocator** routine is a memory-allocator-retrieving routine. The devices selected by the routine are the host device and the device specified in the *dev* argument.

Cross References

- OpenMP **allocator_handle** Type, see Section 26.8.2
- Memory Allocator Retrieving Routines, see Section 34.8
- OpenMP **memspace_handle** Type, see Section 26.8.12

34.8.5 omp_get_devices_all_allocator Routine

Name: omp_get_devices_all_allocator **Category:** function	**Properties:** all-device-threads-binding, memory-management-routine, memory-allocator-retrieving

Return Type and Arguments

Name	Type	Properties
<return type>	allocator_handle	*default*
memspace	memspace_handle	intent(in), omp

Prototypes

C / C++

```
omp_allocator_handle_t omp_get_devices_all_allocator(
  omp_memspace_handle_t memspace);
```

C / C++

Fortran

```
integer (kind=omp_allocator_handle_kind) function &
  omp_get_devices_all_allocator(memspace)
  integer (kind=omp_memspace_handle_kind), intent(in) :: memspace
```

Fortran

Effect

The **omp_get_devices_all_allocator** routine is a memory-allocator-retrieving routine. The devices selected by the routine are all available devices.

Cross References

- OpenMP **allocator_handle** Type, see Section 26.8.2
- Memory Space Retrieving Routines, see Section 34.1
- OpenMP **memspace_handle** Type, see Section 26.8.12

34.9 omp_set_default_allocator Routine

Name: omp_set_default_allocator **Category:** subroutine	**Properties:** binding-implicit-task-binding, memory-management-routine

Arguments

Name	Type	Properties
allocator	allocator_handle	omp, intent(in)

Prototypes

C / C++

```
void omp_set_default_allocator(omp_allocator_handle_t allocator);
```

C / C++

Fortran

```
subroutine omp_set_default_allocator(allocator)
  integer (kind=omp_allocator_handle_kind), intent(in) :: &
    allocator
```

Fortran

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

34.10 `omp_get_default_allocator` Routine

Name: `omp_get_default_allocator` **Category:** function	**Properties:** binding-implicit-task-binding, memory-management-routine

Return Type

Name	Type	Properties
<return type>	allocator_handle	*default*

Prototypes

C / C++

```
omp_allocator_handle_t omp_get_default_allocator(void);
```

C / C++

Fortran

```
integer (kind=omp_allocator_handle_kind) function &
  omp_get_default_allocator()
```

Fortran

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

34.11 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 34.7
- **requires** Directive, see Section 16.5
- **target** Construct, see Section 21.8

34.11.1 omp_alloc Routine

Name: omp_alloc **Category:** function	**Properties:** iso_c_binding, memory-allocating-routine, memory-management-routine, overloaded, raw-memory-allocating-routine

Return Type and Arguments

Name	Type	Properties
<return type>	c_ptr	*default*
size	c_size_t	iso_c, value
allocator	allocator_handle	value, omp

Prototypes

C

```
void  *omp_alloc(size_t size, omp_allocator_handle_t allocator);
```

C

C++

```
void  *omp_alloc(size_t size,
  omp_allocator_handle_t allocator = omp_null_allocator);
```

C++

Fortran

```
type (c_ptr) function omp_alloc(size, allocator) bind(c)
  use, intrinsic :: iso_c_binding, only : c_ptr, c_size_t
  integer (kind=c_size_t), value :: size
  integer (kind=omp_allocator_handle_kind), value :: allocator
```

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 34.11

34.11.2 omp_aligned_alloc Routine

Name: omp_aligned_alloc **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
<return type>	c_ptr	*default*
alignment	c_size_t	iso_c, value
size	c_size_t	iso_c, value
allocator	allocator_handle	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 34.11

34.11.3 omp_calloc Routine

Name: omp_calloc **Category:** function	**Properties:** iso_c_binding, memory-allocating-routine, memory-management-routine, overloaded, zeroed-memory-allocating-routine

Return Type and Arguments

Name	Type	Properties
<return type>	c_ptr	*default*
nmemb	c_size_t	iso_c, value
size	c_size_t	iso_c, value
allocator	allocator_handle	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 34.11

34.11.4 `omp_aligned_calloc` Routine

Name: `omp_aligned_calloc` **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
<return type>	c_ptr	*default*
alignment	c_size_t	iso_c, value
nmemb	c_size_t	iso_c, value
size	c_size_t	iso_c, value
allocator	allocator_handle	value, omp

Prototypes

C

```
void  *omp_aligned_calloc(size_t alignment, size_t nmemb,
  size_t size, omp_allocator_handle_t allocator);
```

C

C++

```
void  *omp_aligned_calloc(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_calloc(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_calloc`** 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 34.11

34.11.5 `omp_realloc` Routine

Name: `omp_realloc` **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>	c_ptr	*default*
ptr	c_ptr	iso_c, value
size	c_size_t	iso_c, value
allocator	allocator_handle	value, omp
free_allocator	allocator_handle	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 34.11

34.12 `omp_free` Routine

Name: `omp_free` **Category:** subroutine	**Properties:** iso_c_binding, memory-management-routine, overloaded

Arguments

Name	Type	Properties
ptr	c_ptr	iso_c, value
allocator	allocator_handle	value, omp

Prototypes

C

```
void omp_free(void *ptr, omp_allocator_handle_t allocator);
```

C

C++

```
void omp_free(void *ptr,
  omp_allocator_handle_t allocator = omp_null_allocator);
```

C++

Fortran

```
subroutine omp_free(ptr, allocator) bind(c)
  use, intrinsic :: iso_c_binding, only : c_ptr
  type (c_ptr), value :: ptr
  integer (kind=omp_allocator_handle_kind), value :: allocator
```

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 **`omp_free`** routine has the overloaded property since it is an overloaded routine 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 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 34.11
- Memory Allocators, see Section 13.2
- **omp_destroy_allocator** Routine, see Section 34.7

34.13 Grouppprivate Memory Routines

This section describes the grouppprivate-information routines, which are routines that have the grouppprivate-information-routine property. Some of these routines are dynamic-grouppprivate-information routines, which are routines that have the dynamic-grouppprivate-information-routine property. The dynamic-grouppprivate-information routines include routines to retrieve a pointer to and the size in bytes of the dynamic grouppprivate block.

All dynamic-grouppprivate-information routines take a *access_group* argument, which specifies the access group type of the dynamic grouppprivate 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 grouppprivate block associated with that access group.

If the value of the *access_group* argument is **omp_access_cgroup**, the dynamic-grouppprivate-information routine returns information about the dynamic grouppprivate block associated with the contention group of the calling task.

C++

The C++ versions of the grouppprivate-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++ versions of the **omp_get_dyn_gprivate_ptr** and **omp_get_dyn_gprivate_nofb_ptr** routines, the *offset* argument may also be omitted, in which case the effect is as if zero is specified.

C++

Fortran

For the Fortran versions of the grouppprivate-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 versions of the **omp_get_dyn_gprivate_ptr** and **omp_get_dyn_gprivate_nofb_ptr** routines, the *offset* argument is also optional and may be omitted. If the *offset* argument is omitted, the effect is as if zero is specified as the argument.

Fortran

Restrictions

Restrictions to the dynamic-groupprivate-information routines are as follows:

- A task may call a dynamic-groupprivate-information routine with an access group type as the *access_group* argument only if it belongs to an access group of that access group type.

34.13.1 `omp_get_dyn_gprivate_ptr` Routine

Name: `omp_get_dyn_gprivate_ptr` **Category:** function	**Properties:** dynamic-groupprivate-information-routine, groupprivate-information-routine, memory-management-routine, overloaded

Return Type and Arguments

Name	Type	Properties
<return type>	c_ptr	*default*
offset	c_size_t	iso_c, value, optional
access_group	access	value, omp, optional

Prototypes

C

```
void  *omp_get_dyn_gprivate_ptr(size_t offset,
  omp_access_t access_group);
```

C

C++

```
void  *omp_get_dyn_gprivate_ptr(size_t offset = 0,
  omp_access_t access_group = omp_access_cgroup);
```

C++

Fortran

```
type (c_ptr) function omp_get_dyn_gprivate_ptr(offset, &
  access_group)
  use, intrinsic :: iso_c_binding, only : c_ptr, c_size_t
  integer (kind=c_size_t), value, optional :: offset
  integer (kind=omp_access_kind), value, optional :: access_group
```

Fortran

Effect

The **`omp_get_dyn_gprivate_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. When no dynamic groupprivate block exists, the routine returns NULL.

Restrictions

Restrictions to the **omp_get_dyn_gprivate_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.

34.13.2 omp_get_dyn_gprivate_nofb_ptr Routine

Name: omp_get_dyn_gprivate_nofb_ptr **Category:** function	**Properties:** dynamic-groupprivate-information-routine, groupprivate-information-routine, memory-management-routine, overloaded

Return Type and Arguments

Name	Type	Properties
<return type>	c_ptr	*default*
offset	c_size_t	iso_c, value, optional
access_group	access	value, omp, optional

Prototypes

C

```
void  *omp_get_dyn_gprivate_nofb_ptr(size_t offset,
  omp_access_t access_group);
```

C

C++

```
void  *omp_get_dyn_gprivate_nofb_ptr(size_t offset = 0,
  omp_access_t access_group = omp_access_cgroup);
```

C++

Fortran

```
type (c_ptr) function omp_get_dyn_gprivate_nofb_ptr(offset, &
  access_group)
  use, intrinsic :: iso_c_binding, only : c_ptr, c_size_t
  integer (kind=c_size_t), value, optional :: offset
  integer (kind=omp_access_kind), value, optional :: access_group
```

Fortran

Effect

The **omp_get_dyn_gprivate_nofb_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, if this belongs to the groupprivate memory space, and otherwise returns NULL. When no dynamic groupprivate block exists, the routine returns NULL.

Restrictions

Restrictions to the **omp_get_dyn_gprivate_nofb_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 positive and the dynamic groupprivate block does not belong to the groupprivate memory space, 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.

34.13.3 omp_get_dyn_gprivate_size Routine

Name: omp_get_dyn_gprivate_size **Category:** function	**Properties:** dynamic-groupprivate-information-routine, groupprivate-information-routine, memory-management-routine, overloaded

Return Type and Arguments

Name	Type	Properties
<return type>	c_size_t	*default*
access_group	access	value, omp, optional

Prototypes

C

```
size_t omp_get_dyn_gprivate_size(omp_access_t access_group);
```

C

C++

```
size_t omp_get_dyn_gprivate_size(
  omp_access_t access_group = omp_access_cgroup);
```

C++

Fortran

```
integer (kind=c_size_t) function omp_get_dyn_gprivate_size(&
  access_group)
  use, intrinsic :: iso_c_binding, only : c_size_t
  integer (kind=omp_access_kind), value, optional :: access_group
```

Fortran

Effect

The **omp_get_dyn_gprivate_size** routine returns the size of the dynamic groupprivate block, if it exists, and otherwise returns zero. If the dynamic groupprivate block exists, this value is the result of the evaluation of the *size* argument that was specified in the corresponding **dyn_groupprivate** clause that instantiated the dynamic groupprivate block.

34.13.4 omp_get_dyn_gprivate_memspace Routine

Name: omp_get_dyn_gprivate_memspace **Category:** function	**Properties:** dynamic-groupprivate-information-routine, groupprivate-information-routine, memory-management-routine, overloaded

Return Type and Arguments

Name	Type	Properties
<return type>	memspace_handle	*default*
access_group	access	value, omp, optional

Prototypes

C

```
omp_memspace_handle_t omp_get_dyn_gprivate_memspace(
  omp_access_t access_group);
```

C

C++

```
omp_memspace_handle_t omp_get_dyn_gprivate_memspace(
  omp_access_t access_group = omp_access_cgroup);
```

C++

Fortran

```
integer (kind=omp_memspace_handle_kind) function &
  omp_get_dyn_gprivate_memspace(access_group)
  integer (kind=omp_access_kind), value, optional :: access_group
```

Fortran

Effect

The **omp_get_dyn_gprivate_memspace** routine returns the handle of the memory space that was used to instantiate the dynamic groupprivate block, if it exists, and otherwise returns **omp_null_mem_space**.

34.13.5 omp_get_gprivate_limit Routine

Name: omp_get_gprivate_limit **Category:** function	**Properties:** groupprivate-information-routine, memory-management-routine, overloaded

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

Prototypes

C

```
size_t omp_get_gprivate_limit(int device_num,
  omp_access_t access_group);
```

C

C++

```
size_t omp_get_gprivate_limit(int device_num,
  omp_access_t access_group = omp_access_cgroup);
```

C++

Fortran

```
integer (kind=c_size_t) function omp_get_gprivate_limit(&
  device_num, access_group)
  use, intrinsic :: iso_c_binding, only : c_size_t, c_int
  integer (kind=c_int), value :: device_num
  integer (kind=omp_access_kind), value, optional :: access_group
```

Fortran

Effect

The **omp_get_gprivate_limit** routine returns the groupprivate absolute limit in bytes of the groupprivate memory space that groupprivate variables and dynamic groupprivate blocks with the access group type *access_group* may use in any region that corresponds to a groupprivate-instantiating construct that targets the device *device_num*. The device *device_num* may be any host device or non-host device.

Restrictions

Restrictions to device memory routines are as follows:

- The *device_num* argument must be a conforming device number.
- When called from within a **target** region, the effect is unspecified.

Cross References

- **dyn_groupprivate** Clause, see Section 13.9
- **target** Construct, see Section 21.8
- **teams** Construct, see Section 18.2

35 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.

35.1 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.

35.1.1 `omp_init_lock` Routine

Name: `omp_init_lock` **Category:** subroutine	**Properties:** all-contention-group-tasks-binding, lock-initializing, simple-lock

Arguments

Name	Type	Properties
svar	lock	C/C++ pointer, omp

Prototypes

C / C++

```
void omp_init_lock(omp_lock_t *svar);
```

C / C++

Fortran

```
subroutine omp_init_lock(svar)
  integer (kind=omp_lock_kind) svar
```

Fortran

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 41.7.9
- OMPT **mutex** Type, see Section 40.20

35.1.2 omp_init_nest_lock Routine

Name: omp_init_nest_lock **Category:** subroutine	**Properties:** all-contention-group-tasks-binding, lock-initializing, nestable-lock

Arguments

Name	Type	Properties
nvar	nest_lock	C/C++ pointer, omp

Prototypes

C / C++

```
void omp_init_nest_lock(omp_nest_lock_t *nvar);
```

C / C++

Fortran

```
subroutine omp_init_nest_lock(nvar)
  integer (kind=omp_nest_lock_kind) nvar
```

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 41.7.9
- OMPT **mutex** Type, see Section 40.20
- OpenMP **nest_lock** Type, see Section 26.9.4

35.1.3 `omp_init_lock_with_hint` Routine

Name: `omp_init_lock_with_hint` **Category:** subroutine	**Properties:** all-contention-group-tasks-binding, lock-initializing, simple-lock

Arguments

Name	Type	Properties
svar	lock	C/C++ pointer, omp
hint	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 41.7.9
- OMPT **`mutex`** Type, see Section 40.20
- OpenMP **`sync_hint`** Type, see Section 26.9.5

35.1.4 `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

Arguments

Name	Type	Properties
nvar	nest_lock	C/C++ pointer, omp
hint	sync_hint	omp

Prototypes

C / C++

```
void omp_init_nest_lock_with_hint(omp_nest_lock_t *nvar,
  omp_sync_hint_t hint);
```

C / C++

Fortran

```
subroutine omp_init_nest_lock_with_hint(nvar, hint)
  integer (kind=omp_nest_lock_kind) nvar
  integer (kind=omp_sync_hint_kind) hint
```

Fortran

Effect

The **`omp_init_nest_lock_with_hint`** routine is a lock-initializing routine.

Execution Model Events

The *nest-lock-init-with-hint* event occurs in a thread that executes an **`omp_init_nest_lock`** region after initialization of the lock, but before it finishes the region.

Tool Callbacks

A thread dispatches a registered **`lock_init`** callback with the same value for its *hint* argument as the *hint* argument of the call to **`omp_init_nest_lock_with_hint`** and **`ompt_mutex_nest_lock`** as the *kind* argument for each occurrence of a *nest-lock-init-with-hint* event in that thread This callback occurs in the task that encounters the routine.

Cross References

- **`lock_init`** Callback, see Section 41.7.9
- OMPT **`mutex`** Type, see Section 40.20
- OpenMP **`nest_lock`** Type, see Section 26.9.4
- OpenMP **`sync_hint`** Type, see Section 26.9.5

35.2 Lock Destroying Routines

Lock-destroying routines are routines with the lock-destroying property. These routines deactivate the lock by setting it to the uninitialized state.

Restrictions

Restrictions to lock-destroying routines are as follows:

- A lock-destroying routine must not access a lock that is not in the unlocked state.

35.2.1 `omp_destroy_lock` Routine

Name: `omp_destroy_lock` **Category:** subroutine	**Properties:** all-contention-group-tasks-binding, lock-destroying, simple-lock

Arguments

Name	Type	Properties
svar	lock	C/C++ pointer, omp

Prototypes

C / C++

```
void omp_destroy_lock(omp_lock_t *svar);
```

C / C++

Fortran

```
subroutine omp_destroy_lock(svar)
  integer (kind=omp_lock_kind) svar
```

Fortran

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 41.7.11
- OMPT **`mutex`** Type, see Section 40.20

35.2.2 `omp_destroy_nest_lock` Routine

Name: `omp_destroy_nest_lock` **Category:** subroutine	**Properties:** all-contention-group-tasks-binding, lock-destroying, nestable-lock

Arguments

Name	Type	Properties
nvar	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 41.7.11
- OMPT **`mutex`** Type, see Section 40.20
- OpenMP **`nest_lock`** Type, see Section 26.9.4

35.3 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 that the lock guards. 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.

35.3.1 `omp_set_lock` Routine

Name: `omp_set_lock` **Category:** subroutine	**Properties:** all-contention-group-tasks-binding, implicit-acquire-flush-at-exit, implicit-acquire-flush, implicit-flush, lock-acquiring, simple-lock

Arguments

Name	Type	Properties
svar	lock	C/C++ pointer, omp

Prototypes

C / C++

```
void omp_set_lock(omp_lock_t *svar);
```

C / C++

Fortran

```
subroutine omp_set_lock(svar)
  integer (kind=omp_lock_kind) svar
```

Fortran

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 40.20
- **mutex_acquire** Callback, see Section 41.7.8
- **mutex_acquired** Callback, see Section 41.7.12

35.3.2 omp_set_nest_lock Routine

Name: omp_set_nest_lock **Category:** subroutine	**Properties:** all-contention-group-tasks-binding, implicit-acquire-flush-at-exit, implicit-acquire-flush, implicit-flush, lock-acquiring, nestable-lock

Arguments

Name	Type	Properties
nvar	nest_lock	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.

Tool Callbacks

A thread dispatches a registered **mutex_acquire** callback for each occurrence of a *nest-lock-acquire* event in that thread. A thread dispatches a registered **mutex_acquired** callback for each occurrence of a *nest-lock-acquired* event in that thread. A thread dispatches a registered **nest_lock** callback with **ompt_scope_begin** as its *endpoint* argument for each occurrence of a *nest-lock-owned* event in that thread. These callbacks occur in the task that encounters the **omp_set_nest_lock** routine and their *kind* argument is **ompt_mutex_nest_lock**.

Cross References

- OMPT **mutex** Type, see Section 40.20
- **mutex_acquire** Callback, see Section 41.7.8
- **mutex_acquired** Callback, see Section 41.7.12
- **nest_lock** Callback, see Section 41.7.14
- OpenMP **nest_lock** Type, see Section 26.9.4
- OMPT **scope_endpoint** Type, see Section 40.27

35.4 Lock Releasing Routines

Lock-releasing routines are routines with the lock-releasing property. These routines provide a means of unsetting locks. If the effect of a lock-releasing routine changes the lock state to the unlocked state and one or more task regions were effectively suspended because the lock was unavailable, the effect is that one task is chosen and given ownership of the lock.

Restrictions

Restrictions to lock-releasing routines are as follows:

- A lock-releasing routine must not access a lock that is not in the locked state.
- A lock-releasing routine must not access a lock that is owned by a task other than the encountering task.

35.4.1 `omp_unset_lock` Routine

Name: `omp_unset_lock` **Category:** subroutine	**Properties:** all-contention-group-tasks-binding, implicit-flush, implicit-release-flush-at-exit, implicit-release-flush, lock-releasing, simple-lock

Arguments

Name	Type	Properties
svar	lock	C/C++ pointer, omp

Prototypes

C / C++

```
void omp_unset_lock(omp_lock_t *svar);
```

C / C++

Fortran

```
subroutine omp_unset_lock(svar)
  integer (kind=omp_lock_kind) svar
```

Fortran

Effect

The **`omp_unset_lock`** routine changes the lock state to the unlocked state. When a thread executes an executes an **`omp_unset_lock`** region, the behavior is as if an implicit release-at-exit flush is performed during the region that synchronizes with the implicit acquire-at-entry flush that is performed during execution of the next **`omp_set_lock`** or **`omp_test_lock`** region that changes the lock state to the locked state.

Execution Model Events

The *lock-release* event occurs in a thread that executes an **`omp_unset_lock`** region after it releases the associated lock but before it finishes the region.

Tool Callbacks

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.

Cross References

- OpenMP **`lock`** Type, see Section 26.9.3
- OMPT **`mutex`** Type, see Section 40.20
- **`mutex_released`** Callback, see Section 41.7.13
- **`omp_set_lock`** Routine, see Section 35.3.1
- **`omp_test_lock`** Routine, see Section 35.5.1

35.4.2 `omp_unset_nest_lock` Routine

Name: `omp_unset_nest_lock` **Category:** subroutine	**Properties:** all-contention-group-tasks-binding, implicit-flush, implicit-release-flush-at-exit, implicit-release-flush, lock-releasing, nestable-lock

Arguments

Name	Type	Properties
nvar	nest_lock	C/C++ pointer, omp

Prototypes

C / C++

```
void omp_unset_nest_lock(omp_nest_lock_t *nvar);
```

C / C++

Fortran

```
subroutine omp_unset_nest_lock(nvar)
  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. When a thread executes an **`omp_unset_nest_lock`** region that changes the lock state to the unlocked state, the behavior is as if an implicit release-at-exit flush is performed during the region that synchronizes with the implicit acquire-at-entry flush that is performed during execution of the next **`omp_set_nest_lock`** or **`omp_test_nest_lock`** region that changes the lock state to the locked 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 40.20
- **mutex_released** Callback, see Section 41.7.13
- **nest_lock** Callback, see Section 41.7.14
- OpenMP **nest_lock** Type, see Section 26.9.4
- **omp_set_nest_lock** Routine, see Section 35.3.2
- **omp_test_nest_lock** Routine, see Section 35.5.2
- OMPT **scope_endpoint** Type, see Section 40.27

35.5 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.

35.5.1 omp_test_lock Routine

Name: omp_test_lock **Category:** function	**Properties:** all-contention-group-tasks-binding, implicit-acquire-flush-at-exit, implicit-acquire-flush, implicit-flush, lock-testing, simple-lock

Return Type and Arguments

Name	Type	Properties
<return type>	logical	*default*
svar	lock	C/C++ pointer, omp

Prototypes

C / C++

```
int omp_test_lock(omp_lock_t *svar);
```

C / C++

Fortran

```
logical function omp_test_lock(svar)
  integer (kind=omp_lock_kind) svar
```

Fortran

Effect

The **omp_test_lock** routine returns *true* if it successfully acquires the lock; otherwise, it returns *false*.

Execution Model Events

The *lock-test* event occurs in a thread that executes an **omp_test_lock** region before the associated lock is tested. The *lock-test-acquired* event occurs in a thread that executes an **omp_test_lock** region before it finishes the region if the associated lock was acquired.

Tool Callbacks

A thread dispatches a registered **mutex_acquire** callback for each occurrence of a *lock-test* event in that thread. A thread dispatches a registered **mutex_acquired** callback for each occurrence of a *lock-test-acquired* event in that thread. These callbacks occur in the encountering task and their *kind* argument is **ompt_mutex_test_lock**.

Restrictions

Restrictions to **omp_test_lock** routines are as follows:

- An **omp_test_lock** routine must not access a lock that is already owned by the encountering task.

Cross References

- OpenMP **lock** Type, see Section 26.9.3
- OMPT **mutex** Type, see Section 40.20
- **mutex_acquire** Callback, see Section 41.7.8
- **mutex_acquired** Callback, see Section 41.7.12

35.5.2 omp_test_nest_lock Routine

Name: omp_test_nest_lock **Category:** function	**Properties:** all-contention-group-tasks-binding, implicit-acquire-flush-at-exit, implicit-acquire-flush, implicit-flush, lock-testing, nestable-lock

Return Type and Arguments

Name	Type	Properties
<return type>	integer	*default*
nvar	nest_lock	C/C++ pointer, omp

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

Effect

The **omp_test_nest_lock** routine returns the new nesting count if it successfully sets the lock; otherwise, it returns zero.

Execution Model Events

The *nest-lock-test* event occurs in a thread that executes an **omp_test_nest_lock** region before the associated lock is tested. The *nest-lock-test-acquired* event occurs in a thread that executes an **omp_test_nest_lock** region before it finishes the region if the associated lock was acquired and the thread did not already own the lock. The *nest-lock-owned* event occurs in a thread that executes an **omp_test_nest_lock** region before it finishes the region after the nesting count is incremented if the thread already owned the lock.

Tool Callbacks

A thread dispatches a registered **mutex_acquire** callback for each occurrence of a *nest-lock-test* event in that thread. A thread dispatches a registered **mutex_acquired** callback for each occurrence of a *nest-lock-test-acquired* event in that thread. A thread dispatches a registered **nest_lock** callback with **ompt_scope_begin** as its *endpoint* argument for each occurrence of a *nest-lock-owned* event in that thread. These callbacks occur in the encountering task and their *kind* argument is **ompt_mutex_test_nest_lock**.

Cross References

- OMPT **mutex** Type, see Section 40.20
- **mutex_acquire** Callback, see Section 41.7.8
- **mutex_acquired** Callback, see Section 41.7.12
- **nest_lock** Callback, see Section 41.7.14
- OpenMP **nest_lock** Type, see Section 26.9.4
- OMPT **scope_endpoint** Type, see Section 40.27

36 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.

36.1 `omp_get_proc_bind` Routine

Name: `omp_get_proc_bind` **Category:** function	**Properties:** ICV-retrieving

Return Type

Name	Type	Properties
<return type>	proc_bind	*default*

Prototypes

C / C++

```
omp_proc_bind_t omp_get_proc_bind(void);
```

C / C++

Fortran

```
integer (kind=omp_proc_bind_kind) function omp_get_proc_bind()
```

Fortran

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

36.2 omp_get_num_places Routine

Name: omp_get_num_places **Category:** function	**Properties:** all-device-threads-binding

Return Type

Name	Type	Properties
<return type>	integer	*default*

Prototypes

C / C++

```
int omp_get_num_places(void);
```

C / C++

Fortran

```
integer function omp_get_num_places()
```

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-group-var* ICV in the execution environment of the initial task.

Cross References

- *place-group-var* ICV, see Table 3.1

36.3 omp_get_place_num_procs Routine

Name: omp_get_place_num_procs **Category:** function	**Properties:** all-device-threads-binding, ICV-retrieving

Return Type and Arguments

Name	Type	Properties
<return type>	integer	*default*
place_num	integer	*default*

Prototypes

C / C++

```
int omp_get_place_num_procs(int place_num);
```

C / C++

Fortran

```
integer function omp_get_place_num_procs(place_num)
  integer place_num
```

Fortran

Effect

The **omp_get_place_num_procs** routine returns the number of processors associated with the place numbered *place_num* as per the *place-group-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-group-var* ICV, see Table 3.1
- **omp_get_num_places** Routine, see Section 36.2

36.4 omp_get_place_proc_ids Routine

Name: omp_get_place_proc_ids **Category:** subroutine	**Properties:** all-device-threads-binding, ICV-retrieving

Arguments

Name	Type	Properties
place_num	integer	*default*
ids	integer	pointer

Prototypes

C / C++

```
void omp_get_place_proc_ids(int place_num, int *ids);
```

C / C++

Fortran

```
subroutine omp_get_place_proc_ids(place_num, ids)
  integer place_num, ids(*)
```

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-group-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 36.2
- **omp_get_place_num_procs** Routine, see Section 36.3

36.5 `omp_get_place_num` Routine

Name: `omp_get_place_num` **Category:** function	**Properties:** *default*

Return Type

Name	Type	Properties
<return type>	integer	*default*

Prototypes

C / C++

```
int omp_get_place_num(void);
```

C / C++

Fortran

```
integer function omp_get_place_num()
```

Fortran

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 36.2

36.6 `omp_get_partition_num_places` Routine

Name: `omp_get_partition_num_places` **Category:** function	**Properties:** ICV-retrieving

Return Type

Name	Type	Properties
<return type>	integer	*default*

Prototypes

C / C++

```
int omp_get_partition_num_places(void);
```

C / C++

Fortran

```
integer function omp_get_partition_num_places()
```

Fortran

Effect

The **omp_get_partition_num_places** routine returns the number of places in the *place-group-var* ICV.

Cross References

- *place-group-var* ICV, see Table 3.1

36.7 omp_get_partition_place_nums Routine

Name: omp_get_partition_place_nums **Category:** subroutine	**Properties:** ICV-retrieving

Arguments

Name	Type	Properties
place_nums	integer	pointer

Prototypes

C / C++

```
void omp_get_partition_place_nums(int *place_nums);
```

C / C++

Fortran

```
subroutine omp_get_partition_place_nums(place_nums)
  integer place_nums(*)
```

Fortran

Effect

The **omp_get_partition_place_nums** routine returns the list of place numbers that correspond to the places in the *place-group-var* ICV of the innermost implicit task. The array must be sufficiently large to contain **omp_get_partition_num_places** integers; otherwise, the behavior is unspecified.

Cross References

- *place-group-var* ICV, see Table 3.1
- **omp_get_partition_num_places** Routine, see Section 36.6

36.8 omp_set_affinity_format Routine

Name: omp_set_affinity_format **Category:** subroutine	**Properties:** ICV-modifying

Arguments

Name	Type	Properties
format	char	pointer, intent(in)

Prototypes

C / C++

```
void omp_set_affinity_format(const char *format);
```

C / C++

Fortran

```
subroutine omp_set_affinity_format(format)
  character(len=*), intent(in) :: format
```

Fortran

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

36.9 omp_get_affinity_format Routine

Name: omp_get_affinity_format **Category:** function	**Properties:** ICV-retrieving

Return Type and Arguments

Name	Type	Properties
<return type>	size_t	*default*
buffer	char	pointer, intent(out)
size	size_t	C/C++-only

Prototypes

C / C++

```
size_t omp_get_affinity_format(char *buffer, size_t size);
```

C / C++

Fortran

```
integer function omp_get_affinity_format(buffer)
  character(len=*), intent(out) :: buffer
```

Fortran

Effect

C / C++

The **omp_get_affinity_format** routine returns the number of characters in the *affinity-format-var* ICV on the current device, excluding the terminating null byte (**'\0'**) and, if *size* is non-zero, writes the value of the *affinity-format-var* ICV on the current device to *buffer* followed by a null byte. If the return value is larger or equal to *size*, the affinity format specification is truncated, with the terminating null byte stored to ***buffer* [*size*-1]**. If *size* is zero, nothing is stored and *buffer* may be NULL.

C / C++

Fortran

The **omp_get_affinity_format** routine returns the number of characters that are required to hold the *affinity-format-var* ICV on the current device and writes the value of the *affinity-format-var* ICV on the current device to *buffer*. If the return value is larger than **len(*buffer*)**, the affinity format specification is truncated.

Fortran

If the *buffer* argument does not conform to the specified format then the result is implementation defined.

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

36.10 omp_display_affinity Routine

Name: omp_display_affinity **Category:** subroutine	**Properties:** *default*

Arguments

Name	Type	Properties
format	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

36.11 `omp_capture_affinity` Routine

Name: `omp_capture_affinity` **Category:** function	**Properties:** *default*

Return Type and Arguments

Name	Type	Properties
<return type>	size_t	*default*
buffer	char	pointer, intent(out)
size	size_t	C/C++-only
format	char	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

37 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.

37.1 `omp_get_cancellation` Routine

Name: `omp_get_cancellation` **Category:** function	**Properties:** ICV-retrieving

Return Type

Name	Type	Properties
<return type>	logical	*default*

Prototypes

C / C++

```
int omp_get_cancellation(void);
```

C / C++

Fortran

```
logical function omp_get_cancellation()
```

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

37.2 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.

37.2.1 `omp_pause_resource` Routine

Name: `omp_pause_resource` **Category:** function	**Properties:** all-tasks-binding, resource-relinquishing

Return Type and Arguments

Name	Type	Properties
<return type>	integer	*default*
kind	pause_resource	*default*
device_num	integer	*default*

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

37.2.2 omp_pause_resource_all Routine

Name: omp_pause_resource_all **Category:** function	**Properties:** all-tasks-binding, resource-relinquishing

Return Type and Arguments

Name	Type	Properties
<return type>	integer	*default*
kind	pause_resource	*default*

Prototypes

C / C++

```
int omp_pause_resource_all(omp_pause_resource_t kind);
```

C / C++

Fortran

```
integer function omp_pause_resource_all(kind)
  integer (kind=omp_pause_resource_kind) kind
```

Fortran

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 37.2.1
- OpenMP **pause_resource** Type, see Section 26.11.1

37.3 Timing Routines

This section describes routines that support a portable wall clock timer.

37.3.1 omp_get_wtime Routine

Name: omp_get_wtime **Category:** function	**Properties:** *default*

Return Type

Name	Type	Properties
<return type>	double	*default*

Prototypes

C / C++

```
double omp_get_wtime(void);
```

C / C++

Fortran

```
double precision function omp_get_wtime()
```

Fortran

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.

37.3.2 omp_get_wtick Routine

Name: omp_get_wtick **Category:** function	**Properties:** *default*

Return Type

Name	Type	Properties
<return type>	double	*default*

Prototypes

C / C++

```
double omp_get_wtick(void);
```

C / C++

Fortran

```
double precision function omp_get_wtick()
```

Fortran

Effect

The **omp_get_wtick** routine returns the precision of the timer used by **omp_get_wtime** as a value equal to the number of seconds between successive clock ticks. The return value of the **omp_get_wtick** routine is not guaranteed to be consistent across any set of threads.

Cross References

- **omp_get_wtime** Routine, see Section 37.3.1

37.4 omp_display_env Routine

Name: omp_display_env **Category:** subroutine	**Properties:** *default*

Arguments

Name	Type	Properties
verbose	logical	intent(in)

Prototypes

C / C++

```
void omp_display_env(int verbose);
```

C / C++

Fortran

```
subroutine omp_display_env(verbose)
  logical, intent(in) :: verbose
```

Fortran

Effect

Each time that the **omp_display_env** routine is invoked, the runtime system prints the OpenMP version number and the initial values of the ICVs associated with the environment variables described in Chapter 4. The displayed values are the values of the ICVs after they have been modified according to the environment variable settings and before the execution of any construct 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

38 Tool Support Routines

This chapter describes the OpenMP API routines that support the use of OpenMP tool interfaces.

38.1 `omp_control_tool` Routine

Name: `omp_control_tool` **Category:** function	**Properties:** *default*

Return Type and Arguments

Name	Type	Properties
<return type>	control_tool_result	*default*
command	control_tool	omp
modifier	integer	*default*
arg	void	C/C++-only, C/C++ pointer

Prototypes

C / C++

```
omp_control_tool_result_t omp_control_tool(
  omp_control_tool_t command, int modifier, void *arg);
```

C / C++

Fortran

```
integer (kind=omp_control_tool_result_kind) function &
  omp_control_tool(command, modifier)
  integer (kind=omp_control_tool_kind) command
  integer modifier
```

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 41.8
- OpenMP **control_tool** Type, see Section 26.12.1
- OpenMP **control_tool_result** Type, see Section 26.12.2
- OMPT Overview, see Chapter 39

Part V

OMPT

39 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.

39.1 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++

39.2 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.

39.2.1 ompt_start_tool Procedure

Name: ompt_start_tool **Category:** function	**Properties:** C-only, OMPT

Return Type and Arguments

Name	Type	Properties
<return type>	start_tool_result	pointer, OMPT
omp_version	integer	unsigned
runtime_version	char	intent(in), pointer

Prototypes

C

```
ompt_start_tool_result_t  *ompt_start_tool(
  unsigned int omp_version, const char *runtime_version);
```

C

Semantics

For a tool to use the OMPT interface that an OpenMP implementation provides, the tool must define a globally-visible implementation of the **ompt_start_tool** procedure. The tool indicates that it will use the OMPT interface that an OpenMP implementation provides by returning a non-null pointer to a **start_tool_result** OMPT type structure from the **ompt_start_tool** implementation that it provides. The **start_tool_result** structure contains pointers to **initialize** and **finalize** callbacks as well as a tool data word that an OpenMP implementation must pass by reference to these callbacks. A tool may return NULL from **ompt_start_tool** to indicate that it will not use the OMPT interface in a particular execution.

A tool may use the *omp_version* argument to determine if it is compatible with the OMPT interface that the OpenMP implementation provides. The *omp_version* argument is the value of the **_OPENMP** version macro associated with the OpenMP implementation. This value identifies the version that an implementation supports, which specifies the version of the OMPT interface that it supports. The *runtime_version* argument is a version string that unambiguously identifies the OpenMP implementation. The version string at least identifies the OpenMP vendor using the key **vendor**, the OpenMP API version as **omp-api-version=***YYYYMM*, the OMPT API version as **ompt-api-version=***YYYYMM*, the OpenMP library release version as **omp-runtime-version=***major*[**.***minor*[**.***patch*[*suffix*]]], and the date of the latest source code change to the OpenMP runtime as **omp-runtime-source-date=***YYYYMMDD*. The version string may contain additional implementation defined key-value pairs.

If a tool returns a non-null pointer to a **start_tool_result** OMPT type structure, an OpenMP implementation will call the OMPT-tool initializer specified by the **initialize** field in this structure before beginning execution of any construct or completing execution of any routine; the OpenMP implementation will call the OMPT-tool finalizer specified by the **finalize** field in this structure when the OpenMP implementation shuts down.

Restrictions

Restrictions to **ompt_start_tool** procedures are as follows:

- The *runtime_version* argument must be an immutable string that is defined for the lifetime of a program execution.

Cross References

- **`finalize`** Callback, see Section 41.1.2
- **`initialize`** Callback, see Section 41.1.1
- OMPT **`start_tool_result`** Type, see Section 40.30

39.2.2 Determining Whether to Initialize a First-Party Tool

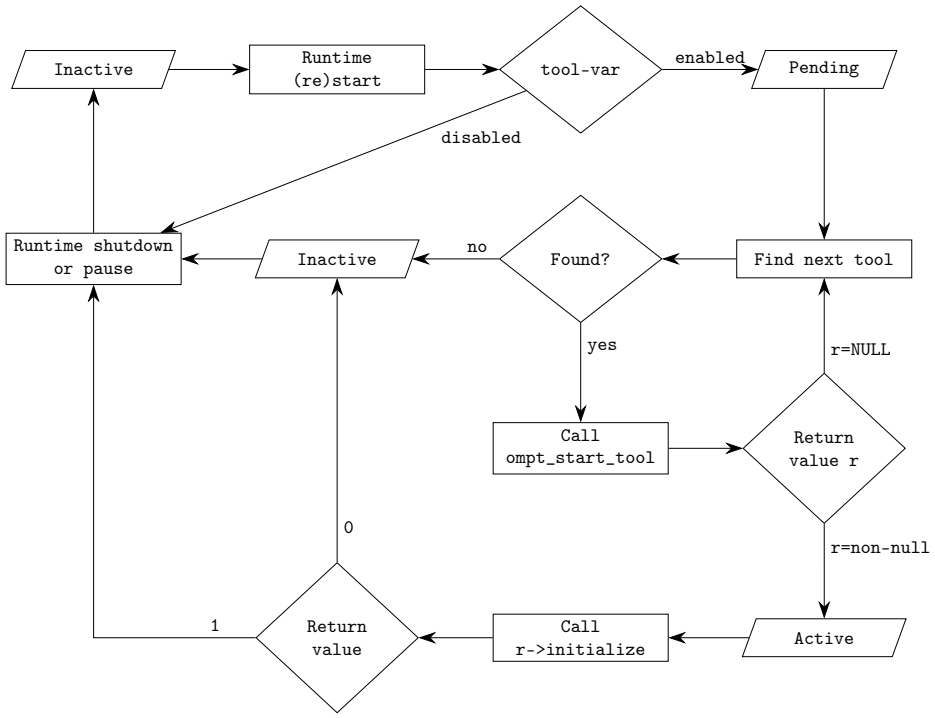

FIGURE 39.1: First-Party Tool Activation Flow Chart

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 39.2.1
- OMPT **start_tool_result** Type, see Section 40.30

39.2.3 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 39.2.3.1. 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 39.2.4.

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 43.3
- **enumerate_states** Entry Point, see Section 43.2
- Binding Entry Points, see Section 39.2.3.1
- **initialize** Callback, see Section 41.1.1
- **ompt_start_tool** Procedure, see Section 39.2.1
- **set_callback** Entry Point, see Section 43.4
- OMPT **start_tool_result** Type, see Section 40.30

39.2.3.1 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 39.1 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 39.1 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 39.1: 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 43.3
- **enumerate_states** Entry Point, see Section 43.2
- **finalize_tool** Entry Point, see Section 43.21
- **function_lookup** Entry Point, see Section 43.1
- **get_callback** Entry Point, see Section 43.6
- **get_num_devices** Entry Point, see Section 43.19
- **get_num_places** Entry Point, see Section 43.9
- **get_num_procs** Entry Point, see Section 43.8
- **get_parallel_info** Entry Point, see Section 43.15
- **get_partition_place_nums** Entry Point, see Section 43.12
- **get_place_num** Entry Point, see Section 43.11
- **get_place_proc_ids** Entry Point, see Section 43.10

- **get_proc_id** Entry Point, see Section 43.13
- **get_state** Entry Point, see Section 43.14
- **get_target_info** Entry Point, see Section 43.18
- **get_task_info** Entry Point, see Section 43.16
- **get_task_memory** Entry Point, see Section 43.17
- **get_thread_data** Entry Point, see Section 43.7
- **get_unique_id** Entry Point, see Section 43.20
- **initialize** Callback, see Section 41.1.1
- **set_callback** Entry Point, see Section 43.4

39.2.4 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 39.2, **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 39.2, 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 barrier, **taskwait**, and **taskgroup** regions. Some events, for example, those that correspond to **sync_region_wait**, use both techniques.

TABLE 39.2: Callbacks for which **set_callback** Must Return **ompt_set_always**

Callback Name
thread_begin
thread_end
parallel_begin
parallel_end
task_create
task_schedule
implicit_task
target_data_op_emi
target_emi
target_submit_emi
control_tool
device_initialize
device_finalize
device_load
device_unload
error

Cross References

- **get_callback** Entry Point, see Section 43.6
- **initialize** Callback, see Section 41.1.1
- OMPT **scope_endpoint** Type, see Section 40.27
- **set_callback** Entry Point, see Section 43.4
- OMPT **set_result** Type, see Section 40.28

39.2.5 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.

TABLE 39.3: 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

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 39.3 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 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 44.12
- **buffer_complete** Callback, see Section 42.6
- **buffer_request** Callback, see Section 42.5
- **device_finalize** Callback, see Section 42.2
- **device_initialize** Callback, see Section 42.1
- **device_load** Callback, see Section 42.3
- **device_unload** Callback, see Section 42.4
- **flush_trace** Entry Point, see Section 44.10
- **function_lookup** Entry Point, see Section 43.1
- **get_buffer_limits** Entry Point, see Section 44.7
- **get_device_num_procs** Entry Point, see Section 44.1
- **get_device_time** Entry Point, see Section 44.3
- **get_record_abstract** Entry Point, see Section 44.16
- **get_record_native** Entry Point, see Section 44.15
- **get_record_ompt** Entry Point, see Section 44.14
- **get_record_type** Entry Point, see Section 44.13
- **pause_trace** Entry Point, see Section 44.9
- OMPT **record_abstract** Type, see Section 40.24

- OMPT **set_result** Type, see Section 40.28
- **set_trace_native** Entry Point, see Section 44.6
- **set_trace_ompt** Entry Point, see Section 44.5
- **start_trace** Entry Point, see Section 44.8
- **stop_trace** Entry Point, see Section 44.11
- **translate_time** Entry Point, see Section 44.4

39.3 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 41.1.2
- **ompt_start_tool** Procedure, see Section 39.2.1
- OMPT **start_tool_result** Type, see Section 40.30

40 OMPT Data Types

This chapter specifies OMPT types. Including the **omp-tools.h** C/C++ header file must result in the definition of these OMPT types.

C / C++

40.1 OMPT Predefined Identifiers

Predefined Identifiers

Name	Value	Properties
ompt_addr_none	**~0**	*default*
ompt_mutex_impl_none	**0**	*default*

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++

40.2 OMPT any_record_ompt Type

Name: any_record_ompt **Properties:** C/C++-only, OMPT	**Base Type:** union

Fields

Name	Type	Properties
thread_begin	thread_begin	C/C++-only
parallel_begin	parallel_begin	C/C++-only
parallel_end	parallel_end	C/C++-only
work	work	C/C++-only
dispatch	dispatch	C/C++-only
task_create	task_create	C/C++-only
dependences	dependences	C/C++-only
task_dependence	task_dependence	C/C++-only
task_schedule	task_schedule	C/C++-only
implicit_task	implicit_task	C/C++-only
masked	masked	C/C++-only
sync_region	sync_region	C/C++-only
mutex_acquire	mutex_acquire	C/C++-only
mutex	mutex	C/C++-only
nest_lock	nest_lock	C/C++-only
flush	flush	C/C++-only
cancel	cancel	C/C++-only
target_emi	target_emi	C/C++-only
target_data_op_emi	target_data_op_emi	C/C++-only
target_map_emi	target_map_emi	C/C++-only
target_submit_emi	target_submit_emi	C/C++-only
control_tool	control_tool	C/C++-only
error	error	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;
  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++

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 40.26

40.3 OMPT buffer Type

Name: buffer **Properties:** C/C++-only, OMPT, opaque	Base Type: void

Type Definition

C / C++

```
typedef void ompt_buffer_t;
```

C / C++

Semantics

The **buffer** OMPT type represents a handle for a device buffer.

40.4 OMPT buffer_cursor Type

Name: buffer_cursor **Properties:** C/C++-only, OMPT, opaque	Base Type: c_uint64_t

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.

40.5 OMPT `callback` Type

Name: `callback` **Category:** subroutine 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.

40.6 OMPT `callbacks` Type

Name: `callbacks`	**Base Type:** enumeration
Properties: C/C++-only, OMPT	

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

C / C++

```
typedef enum ompt_callbacks_t {
  ompt_callback_thread_begin       = 1,
  ompt_callback_thread_end         = 2,
  ompt_callback_parallel_begin     = 3,
  ompt_callback_parallel_end       = 4,
  ompt_callback_task_create        = 5,
  ompt_callback_task_schedule      = 6,
  ompt_callback_implicit_task      = 7,
  ompt_callback_control_tool       = 11,
  ompt_callback_device_initialize  = 12,
  ompt_callback_device_finalize    = 13,
  ompt_callback_device_load        = 14,
  ompt_callback_device_unload      = 15,
  ompt_callback_sync_region_wait   = 16,
  ompt_callback_mutex_released     = 17,
  ompt_callback_dependences        = 18,
  ompt_callback_task_dependence    = 19,
  ompt_callback_work               = 20,
  ompt_callback_masked             = 21,
  ompt_callback_sync_region        = 23,
  ompt_callback_lock_init          = 24,
  ompt_callback_lock_destroy       = 25,
  ompt_callback_mutex_acquire      = 26,
  ompt_callback_mutex_acquired     = 27,
  ompt_callback_nest_lock          = 28,
  ompt_callback_flush              = 29,
  ompt_callback_cancel             = 30,
  ompt_callback_reduction          = 31,
  ompt_callback_dispatch           = 32,
  ompt_callback_target_emi         = 33,
  ompt_callback_target_data_op_emi = 34,
  ompt_callback_target_submit_emi  = 35,
  ompt_callback_target_map_emi     = 36,
  ompt_callback_error              = 37
} ompt_callbacks_t;
```

C / C++

Additional information

The following instances and associated values of the **callbacks** OMPT type are also defined: **ompt_callback_target**, with value 8; **ompt_callback_target_data_op**, with value 9; **ompt_callback_target_submit**, with value 10; and **ompt_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.

40.7 OMPT `cancel_flag` Type

Name: `cancel_flag` **Properties:** C/C++-only, OMPT	**Base Type:** enumeration

Values

Name	Value	Properties
`ompt_cancel_parallel`	`0x01`	C/C++-only, OMPT
`ompt_cancel_sections`	`0x02`	C/C++-only, OMPT
`ompt_cancel_loop`	`0x04`	C/C++-only, OMPT
`ompt_cancel_taskgroup`	`0x08`	C/C++-only, OMPT
`ompt_cancel_activated`	`0x10`	C/C++-only, OMPT
`ompt_cancel_detected`	`0x20`	C/C++-only, OMPT
`ompt_cancel_discarded_task`	`0x40`	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.

40.8 OMPT `data` Type

Name: `data` **Properties:** C/C++-only, OMPT	**Base Type:** union

Fields

Name	Type	Properties
value	c_uint64_t	*default*
ptr	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**.

40.9 OMPT dependence Type

Name: dependence **Properties:** C/C++-only, OMPT	**Base Type:** structure

Fields

Name	Type	Properties
variable	data	C/C++-only
dependence_type	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 40.8
- OMPT **dependence_type** Type, see Section 40.10

40.10 OMPT dependence_type Type

Name: dependence_type **Properties:** C/C++-only, OMPT	**Base Type:** enumeration

Values

Name	Value	Properties
ompt_dependence_type_in	1	C/C++-only, OMPT
ompt_dependence_type_out	2	C/C++-only, OMPT
ompt_dependence_type_inout	3	C/C++-only, OMPT
ompt_dependence_type_mutexinoutset	4	C/C++-only, OMPT
ompt_dependence_type_source	5	C/C++-only, OMPT
ompt_dependence_type_sink	6	C/C++-only, OMPT
ompt_dependence_type_inoutset	7	C/C++-only, OMPT
ompt_dependence_type_out_all_memory	34	C/C++-only, OMPT
ompt_dependence_type_inout_all_memory	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++

Semantics

The **dependence_type** OMPT type defines task dependence type values. The **ompt_dependence_type_in**, **ompt_dependence_type_out**, **ompt_dependence_type_inout**, **ompt_dependence_type_mutexinoutset**, **ompt_dependence_type_inoutset**, **ompt_dependence_type_out_all_memory**, and **ompt_dependence_type_inout_all_memory** values represent the task dependence type present in a **depend** clause while the **ompt_dependence_type_source** and **ompt_dependence_type_sink** values represent the *dependence-type* present in a **doacross** clause. The **ompt_dependence_type_out_all_memory** and **ompt_dependence_type_inout_all_memory** represent task dependences for which the **omp_all_memory** reserved locator is specified.

40.11 OMPT device Type

Name: device **Properties:** C/C++-only, OMPT, opaque	Base Type: void

Type Definition

C / C++

```
typedef void ompt_device_t;
```

C / C++

Semantics

The **device** OMPT type represents a device.

40.12 OMPT device_time Type

Name: device_time **Properties:** C/C++-only, OMPT, opaque	Base Type: c_uint64_t

Predefined Identifiers

Name	Value	Properties
ompt_time_none	0	C/C++-only, OMPT

Type Definition

C / C++

```
typedef uint64_t ompt_device_time_t;
```

C / C++

Semantics

The **device_time** OMPT type represents raw device time values; **ompt_time_none** represents an unknown or unspecified time.

40.13 OMPT `dispatch` Type

Name: `dispatch`	**Base Type:** enumeration
Properties: C/C++-only, OMPT, overlapping-type-name	

Values

Name	Value	Properties
`ompt_dispatch_iteration`	1	C/C++-only, OMPT
`ompt_dispatch_section`	2	C/C++-only, OMPT
`ompt_dispatch_ws_loop_chunk`	3	C/C++-only, OMPT
`ompt_dispatch_taskloop_chunk`	4	C/C++-only, OMPT
`ompt_dispatch_distribute_chunk`	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.

40.14 OMPT `dispatch_chunk` Type

Name: `dispatch_chunk`	**Base Type:** structure
Properties: C/C++-only, OMPT	

Fields

Name	Type	Properties
start	c_uint64_t	*default*
iterations	c_uint64_t	*default*

Type Definition

C / C++

```
typedef struct ompt_dispatch_chunk_t {
  uint64_t start;
  uint64_t iterations;
} ompt_dispatch_chunk_t;
```

C / C++

Semantics

The **dispatch_chunk** OMPT type represents chunk information for a dispatched chunk. The **start** field specifies the first logical iteration of the chunk and the **iterations** field specifies the number of logical iterations in the chunk. Whether the chunk of a **taskloop** region is contiguous is implementation defined.

40.15 OMPT frame Type

Name: **frame** **Properties:** C/C++-only, OMPT	**Base Type:** structure

Fields

Name	Type	Properties
exit_frame	data	C/C++-only, OMPT
enter_frame	data	C/C++-only, OMPT
exit_frame_flags	integer	*default*
enter_frame_flags	integer	*default*

Type Definition

C / C++

```
typedef struct ompt_frame_t {
  ompt_data_t exit_frame;
  ompt_data_t enter_frame;
  int exit_frame_flags;
  int enter_frame_flags;
} ompt_frame_t;
```

C / C++

Semantics

The **frame** OMPT type describes procedure frame information for a task. Each **frame** object is associated with the task to which the procedure frames belong. Every task that is not a merged task with one or more frames on the stack of a native thread, whether an initial task, an implicit task, an explicit task, or a target task, has an associated **frame** object.

The **exit_frame** field contains information to identify the first procedure frame executing the task region. The **exit_frame** for the **frame** object associated with the initial task that is not nested inside any OpenMP construct is **ompt_data_none**. The **enter_frame** field contains information to identify the latest still active procedure frame executing the task region before entering the OpenMP runtime implementation or before executing a different task. If a task with frames on the stack is not executing implementation code in the OpenMP runtime, the value of **enter_frame** for its associated **frame** object is **ompt_data_none**.

For the frame indicated by **exit_frame** (**enter_frame**), the **exit_frame_flags** (**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 40.8
- OMPT **`frame_flag`** Type, see Section 40.16
- **`get_task_info`** Entry Point, see Section 43.16

40.16 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,
  ompt_frame_stackaddress = 0x30
} 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.

40.17 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 40.24

40.18 OMPT id Type

Name: id **Properties:** C/C++-only, OMPT	**Base Type: c_uint64_t**

Predefined Identifiers

Name	Value	Properties
ompt_id_none	0	C/C++-only, OMPT

Type Definition

C / C++

```
typedef uint64_t ompt_id_t;
```

C / C++

Semantics

The **id** OMPT type is used to provide various identifiers to tools; **ompt_id_none** is used when the specific ID is unknown or unavailable. When tracing asynchronous activity on devices, identifiers enable tools to correlate device regions and operations that the host device initiates with associated activities on a target device. In addition, OMPT provides identifiers to refer to parallel regions and tasks that execute on a device.

Restrictions

Restrictions to the **id** OMPT type are as follows:

- Identifiers created on each device must be unique from the time an OpenMP implementation is initialized until it is shut down. Identifiers for each device region and target data operation instance that the host device initiates must be unique over time on the host device. Identifiers for instances of parallel regions and task regions that execute on a device must be unique over time within that device.

40.19 OMPT `interface_fn` Type

Name: `interface_fn` **Category:** subroutine pointer	**Properties:** C/C++-only, OMPT

Type Signature

C / C++

```
typedef void (*ompt_interface_fn_t) (void);
```

C / C++

Semantics

The **interface_fn** OMPT type serves as a generic function pointer that the **function_lookup** entry point returns to provide access to a tool to entry points by name.

40.20 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++

```
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;
```

C / C++

Semantics

The **`mutex`** OMPT type defines the valid mutex values.

40.21 OMPT `native_mon_flag` Type

Name: `native_mon_flag`	**Base Type:** enumeration
Properties: C/C++-only, OMPT	

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

Type Definition

C / C++

```
typedef enum ompt_native_mon_flag_t {
  ompt_native_data_motion_explicit = 0x01,
  ompt_native_data_motion_implicit = 0x02,
  ompt_native_kernel_invocation    = 0x04,
  ompt_native_kernel_execution     = 0x08,
  ompt_native_driver               = 0x10,
  ompt_native_runtime              = 0x20,
  ompt_native_overhead             = 0x40,
  ompt_native_idleness             = 0x80
} ompt_native_mon_flag_t;
```

C / C++

Semantics

The **native_mon_flag** OMPT type defines the valid native monitoring flag values.

40.22 OMPT parallel_flag Type

Name: parallel_flag **Properties:** C/C++-only, OMPT	**Base Type:** enumeration

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

Type Definition

C / C++

```
typedef enum ompt_parallel_flag_t {
  ompt_parallel_invoker_program = 0x00000001,
  ompt_parallel_invoker_runtime = 0x00000002,
  ompt_parallel_league          = 0x40000000,
  ompt_parallel_team            = 0x80000000
} ompt_parallel_flag_t;
```

C / C++

Semantics

The **parallel_flag** OMPT type defines valid invoker values, which indicate how the code that implements the associated structured block of the region is invoked or encountered. The **ompt_parallel_invoker_program** value indicates that the encountering thread for a **parallel** or **teams** region executes code to implement its associated structured block as if directly invoked or encountered in application code. The **ompt_parallel_invoker_runtime** value indicates that the encountering thread for a **parallel** or **teams** region invokes the code that implements its associated structured block from the runtime. The **ompt_parallel_league** value indicates that the callback is invoked due to the creation of a league of teams by a **teams** construct. The **ompt_parallel_team** value indicates that the callback is invoked due to the creation of a team of threads by a **parallel** construct.

40.23 OMPT record Type

Name: record **Properties:** C/C++-only, OMPT	**Base Type:** enumeration

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

Type Definition

C / C++

```
typedef enum ompt_record_t {
  ompt_record_ompt    = 1,
  ompt_record_native  = 2,
  ompt_record_invalid = 3
} ompt_record_t;
```

C / C++

Semantics

The **record** OMPT type indicates the integer codes that identify OMPT trace record formats.

40.24 OMPT record_abstract Type

Name: record_abstract **Properties:** C/C++-only, OMPT	**Base Type:** structure

Fields

Name	Type	Properties
rclass	record_native	C/C++-only, OMPT
type	char	common-field, intent(in), pointer
start_time	device_time	C/C++-only, OMPT
end_time	device_time	C/C++-only, OMPT
hwid	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 **ompt_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 40.12
- OMPT **hwid** Type, see Section 40.17
- OMPT **record_native** Type, see Section 40.25

40.25 OMPT `record_native` Type

Name: `record_native` **Properties:** C/C++-only, OMPT	Base Type: enumeration

Values

Name	Value	Properties
`ompt_record_native_info`	`1`	C/C++-only, OMPT
`ompt_record_native_event`	`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.

40.26 OMPT `record_ompt` Type

Name: `record_ompt` **Properties:** C/C++-only, OMPT	Base Type: structure

Fields

Name	Type	Properties
type	callbacks	C/C++-only, common-field, OMPT
time	device_time	C/C++-only, OMPT
thread_id	id	C/C++-only, OMPT
target_id	id	C/C++-only, OMPT
record	any_record_ompt	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 40.2
- OMPT **callbacks** Type, see Section 40.6
- OMPT **device_time** Type, see Section 40.12
- OMPT **id** Type, see Section 40.18

40.27 OMPT scope_endpoint Type

Name: scope_endpoint **Properties:** C/C++-only, OMPT	**Base Type:** enumeration

Values

Name	Value	Properties
ompt_scope_begin	1	C/C++-only, OMPT
ompt_scope_end	2	C/C++-only, OMPT
ompt_scope_beginend	3	C/C++-only, OMPT

Type Definition

C / C++

```
typedef enum ompt_scope_endpoint_t {
  ompt_scope_begin    = 1,
  ompt_scope_end      = 2,
  ompt_scope_beginend = 3
} ompt_scope_endpoint_t;
```

C / C++

Summary

The **scope_endpoint** OMPT type defines valid region endpoint values. If the OpenMP implementation cannot perform callback dispatch at the end of the region then **ompt_scope_beginend** may be used in place of **ompt_scope_begin** to simultaneously represent both region endpoints.

40.28 OMPT set_result Type

Name: set_result **Properties:** C/C++-only, OMPT	**Base Type:** enumeration

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

Type Definition

C / C++

```
typedef enum ompt_set_result_t {
  ompt_set_error             = 0,
  ompt_set_never             = 1,
  ompt_set_impossible        = 2,
  ompt_set_sometimes         = 3,
  ompt_set_sometimes_paired  = 4,
  ompt_set_always            = 5
} ompt_set_result_t;
```

C / C++

Summary

The **set_result** OMPT type corresponds to values that the **set_callback**, **set_trace_ompt** and **set_trace_native** entry points return. Its values indicate several possible outcomes. The **ompt_set_error** value indicates that the associated call failed. Otherwise, the value indicates when an event may occur and, when appropriate, callback dispatch leads to the invocation of the callback. The **ompt_set_never** value indicates that the event will never occur or that the callback will never be invoked at runtime. The **ompt_set_impossible** value indicates that the event may occur but that tracing of it is not possible. The **ompt_set_sometimes** value indicates that the event may occur and, for an implementation defined subset of associated event occurrences, will be traced or the callback will be invoked at runtime. The **ompt_set_sometimes_paired** value indicates the same result as **ompt_set_sometimes** and, in addition, that a callback with an *endpoint* value of **ompt_scope_begin** will be invoked if and only if the same callback with an *endpoint* value of **ompt_scope_end** will also be invoked sometime in the future. The **ompt_set_always** value indicates that, whenever an associated event occurs, it will be traced or the callback will be invoked.

Cross References

- OMPT **scope_endpoint** Type, see Section 40.27
- **set_callback** Entry Point, see Section 43.4
- **set_trace_native** Entry Point, see Section 44.6
- **set_trace_ompt** Entry Point, see Section 44.5

40.29 OMPT severity Type

Name: severity **Properties:** C/C++-only, OMPT	**Base Type:** enumeration

Values

Name	Value	Properties
ompt_warning	1	C/C++-only, OMPT
ompt_fatal	2	C/C++-only, OMPT

Type Definition

C / C++

```
typedef enum ompt_severity_t {
  ompt_warning = 1,
  ompt_fatal   = 2
} ompt_severity_t;
```

C / C++

Semantics

The **severity** OMPT type defines severity values.

40.30 OMPT start_tool_result Type

Name: start_tool_result **Properties:** C/C++-only, OMPT	**Base Type:** structure

Fields

Name	Type	Properties
initialize	initialize	C/C++-only, OMPT
finalize	finalize	C/C++-only, OMPT
tool_data	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 40.8
- **finalize** Callback, see Section 41.1.2
- **initialize** Callback, see Section 41.1.1
- **ompt_start_tool** Procedure, see Section 39.2.1

40.31 OMPT `state` Type

Name: `state` **Properties:** C/C++-only, OMPT	**Base Type:** enumeration

Values

Name	Value	Properties
`ompt_state_work_serial`	`0x000`	C/C++-only, OMPT
`ompt_state_work_parallel`	`0x001`	C/C++-only, OMPT
`ompt_state_work_reduction`	`0x002`	C/C++-only, OMPT
`ompt_state_work_free_agent`	`0x003`	C/C++-only, OMPT
`ompt_state_work_induction`	`0x004`	C/C++-only, OMPT
`ompt_state_wait_barrier_implicit_parallel`	`0x011`	C/C++-only, OMPT
`ompt_state_wait_barrier_implicit_workshare`	`0x012`	C/C++-only, OMPT
`ompt_state_wait_barrier_explicit`	`0x014`	C/C++-only, OMPT
`ompt_state_wait_barrier_implementation`	`0x015`	C/C++-only, OMPT
`ompt_state_wait_barrier_teams`	`0x016`	C/C++-only, OMPT
`ompt_state_wait_taskwait`	`0x020`	C/C++-only, OMPT
`ompt_state_wait_taskgroup`	`0x021`	C/C++-only, OMPT
`ompt_state_wait_mutex`	`0x040`	C/C++-only, OMPT
`ompt_state_wait_lock`	`0x041`	C/C++-only, OMPT
`ompt_state_wait_critical`	`0x042`	C/C++-only, OMPT
`ompt_state_wait_atomic`	`0x043`	C/C++-only, OMPT
`ompt_state_wait_ordered`	`0x044`	C/C++-only, OMPT
`ompt_state_wait_target`	`0x080`	C/C++-only, OMPT
`ompt_state_wait_target_map`	`0x081`	C/C++-only, OMPT
`ompt_state_wait_target_update`	`0x082`	C/C++-only, OMPT
`ompt_state_idle`	`0x100`	C/C++-only, OMPT
`ompt_state_overhead`	`0x101`	C/C++-only, OMPT
`ompt_state_undefined`	`0x102`	C/C++-only, OMPT

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,
  ompt_state_wait_barrier_teams               = 0x016,
```

```
  ompt_state_wait_taskwait                = 0x020,
  ompt_state_wait_taskgroup               = 0x021,
  ompt_state_wait_mutex                   = 0x040,
  ompt_state_wait_lock                    = 0x041,
  ompt_state_wait_critical                = 0x042,
  ompt_state_wait_atomic                  = 0x043,
  ompt_state_wait_ordered                 = 0x044,
  ompt_state_wait_target                  = 0x080,
  ompt_state_wait_target_map              = 0x081,
  ompt_state_wait_target_update           = 0x082,
  ompt_state_idle                         = 0x100,
  ompt_state_overhead                     = 0x101,
  ompt_state_undefined                    = 0x102
} ompt_state_t;
```

C / C++

Semantics

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 that the thread is waiting at a barrier at the end of a **teams** region. The value **ompt_state_wait_taskwait** indicates that the thread is waiting at a **taskwait** construct.

The **ompt_state_wait_taskgroup** value indicates that the thread is waiting at the end of a **taskgroup** construct. The **ompt_state_wait_mutex** value indicates that the thread is waiting for a mutex of an unspecified type. The **ompt_state_wait_lock** value indicates that the thread is waiting for a lock or nestable lock. The **ompt_state_wait_critical** value indicates that the thread is waiting to enter a **critical** region. The **ompt_state_wait_atomic** value indicates that the thread is waiting to enter an **atomic** region. The **ompt_state_wait_ordered** value indicates that the thread is waiting to enter an **ordered** region. The **ompt_state_wait_target** value indicates that the thread is waiting for a **target** region to complete. The **ompt_state_wait_target_map** value indicates that the thread is waiting for a mapping operation to complete. An implementation may report **ompt_state_wait_target** for **target_data** constructs. The **ompt_state_wait_target_update** value indicates that the thread is waiting for a **target_update** operation to complete. An implementation may report **ompt_state_wait_target** for **target_update** constructs. The **ompt_state_idle** value indicates that the native thread is an idle thread, that is, it is an unassigned thread that is not a free-agent thread. The **ompt_state_overhead** value indicates that the thread is in the overhead state at any point while executing within the OpenMP runtime, except while waiting at a synchronization point. The **ompt_state_undefined** value indicates that the native thread is not created by the OpenMP implementation.

40.32 OMPT `subvolume` Type

Name: `subvolume` **Properties:** C/C++-only, OMPT	**Base Type:** structure

Fields

Name	Type	Properties
base	c_ptr	C/C++-only, intent(in), value
size	c_uint64_t	value
num_dims	c_uint64_t	value, positive
volume	c_uint64_t	C/C++-only, intent(in), pointer
offsets	c_uint64_t	C/C++-only, intent(in), pointer
dimensions	c_uint64_t	C/C++-only, intent(in), pointer

Type Definition

C / C++

```
typedef struct ompt_subvolume_t {
  const void *base;
  uint64_t size;
  uint64_t num_dims;
  const uint64_t *volume;
  const uint64_t *offsets;
  const uint64_t *dimensions;
} 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 32.7

40.33 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
`ompt_sync_region_undeferrable_task`	11	C/C++-only, OMPT

Type Definition

C / C++

```
typedef enum ompt_sync_region_t {
  ompt_sync_region_barrier_explicit           = 3,
  ompt_sync_region_barrier_implementation     = 4,
  ompt_sync_region_taskwait                   = 5,
  ompt_sync_region_taskgroup                  = 6,
  ompt_sync_region_reduction                  = 7,
  ompt_sync_region_barrier_implicit_workshare = 8,
  ompt_sync_region_barrier_implicit_parallel  = 9,
  ompt_sync_region_barrier_teams              = 10,
  ompt_sync_region_undeferrable_task          = 11
} ompt_sync_region_t;
```

C / C++

Semantics

The **`sync_region`** OMPT type defines the valid synchronization region values.

40.34 OMPT `target` Type

Name: `target`	**Base Type:** enumeration
Properties: C/C++-only, OMPT	

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++

```
typedef enum ompt_target_t {
  ompt_target                   = 1,
  ompt_target_enter_data        = 2,
  ompt_target_exit_data         = 3,
  ompt_target_update            = 4,
  ompt_target_nowait            = 9,
  ompt_target_enter_data_nowait = 10,
  ompt_target_exit_data_nowait  = 11,
  ompt_target_update_nowait     = 12
} ompt_target_t;
```

C / C++

Semantics

The **`target`** OMPT type defines valid values to identify device constructs.

40.35 OMPT `target_data_op` Type

Name: `target_data_op`	**Base Type:** enumeration
Properties: C/C++-only, OMPT	

Values

Name	Value	Properties
`ompt_target_data_alloc`	1	C/C++-only, OMPT
`ompt_target_data_delete`	4	C/C++-only, OMPT
`ompt_target_data_associate`	5	C/C++-only, OMPT
`ompt_target_data_disassociate`	6	C/C++-only, OMPT
`ompt_target_data_transfer`	7	C/C++-only, OMPT
`ompt_target_data_memset`	8	C/C++-only, OMPT
`ompt_target_data_transfer_rect`	9	C/C++-only, OMPT
`ompt_target_data_alloc_async`	17	C/C++-only, OMPT
`ompt_target_data_delete_async`	20	C/C++-only, OMPT
`ompt_target_data_transfer_async`	23	C/C++-only, OMPT
`ompt_target_data_memset_async`	24	C/C++-only, OMPT
`ompt_target_data_transfer_rect_async`	25	C/C++-only, OMPT
`ompt_target_data_implementation_specific`	64	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_implementation_specific = 64
} 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 the **target_data_op_emi** callbacks. The five least significant bits of the *optype* argument can indicate 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.

The **(***optype* **&** **ompt_target_data_implementation_specific)** evaluates to *true*, if and only if the OpenMP implementation inserts the target data operation.

40.36 OMPT `target_map_flag` Type

Name: `target_map_flag`	Base Type: enumeration
Properties: C/C++-only, OMPT	

Values

Name	Value	Properties
`ompt_target_map_flag_to`	`0x01`	C/C++-only, OMPT
`ompt_target_map_flag_from`	`0x02`	C/C++-only, OMPT
`ompt_target_map_flag_alloc`	`0x04`	C/C++-only, OMPT
`ompt_target_map_flag_release`	`0x08`	C/C++-only, OMPT
`ompt_target_map_flag_delete`	`0x10`	C/C++-only, OMPT
`ompt_target_map_flag_implicit`	`0x20`	C/C++-only, OMPT
`ompt_target_map_flag_always`	`0x40`	C/C++-only, OMPT
`ompt_target_map_flag_present`	`0x80`	C/C++-only, OMPT
`ompt_target_map_flag_close`	`0x100`	C/C++-only, OMPT
`ompt_target_map_flag_shared`	`0x200`	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.

40.37 OMPT `task_flag` Type

Name: `task_flag` **Properties:** C/C++-only, OMPT	**Base Type:** enumeration

Values

Name	Value	Properties
`ompt_task_initial`	`0x00000001`	C/C++-only, OMPT
`ompt_task_implicit`	`0x00000002`	C/C++-only, OMPT
`ompt_task_explicit`	`0x00000004`	C/C++-only, OMPT
`ompt_task_target`	`0x00000008`	C/C++-only, OMPT
`ompt_task_taskwait`	`0x00000010`	C/C++-only, OMPT
`ompt_task_undeferrable`	`0x01000000`	C/C++-only, OMPT
`ompt_task_importing`	`0x02000000`	C/C++-only, OMPT
`ompt_task_exporting`	`0x04000000`	C/C++-only, OMPT
`ompt_task_undeferred`	`0x08000000`	C/C++-only, OMPT
`ompt_task_untied`	`0x10000000`	C/C++-only, OMPT
`ompt_task_final`	`0x20000000`	C/C++-only, OMPT
`ompt_task_mergeable`	`0x40000000`	C/C++-only, OMPT
`ompt_task_merged`	`0x80000000`	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_undeferrable = 0x01000000,
  ompt_task_importing    = 0x02000000,
  ompt_task_exporting    = 0x04000000,
  ompt_task_undeferred   = 0x08000000,
  ompt_task_untied       = 0x10000000,
  ompt_task_final        = 0x20000000,
  ompt_task_mergeable    = 0x40000000,
  ompt_task_merged       = 0x80000000
} ompt_task_flag_t;
```

C / C++

Semantics

The **`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.

40.38 OMPT `task_status` Type

Name: `task_status` **Properties:** C/C++-only, OMPT	**Base Type:** enumeration

Values

Name	Value	Properties
`ompt_task_complete`	1	C/C++-only, OMPT
`ompt_task_yield`	2	C/C++-only, OMPT
`ompt_task_cancel`	3	C/C++-only, OMPT
`ompt_task_detach`	4	C/C++-only, OMPT
`ompt_task_early_fulfill`	5	C/C++-only, OMPT
`ompt_task_late_fulfill`	6	C/C++-only, OMPT
`ompt_task_switch`	7	C/C++-only, OMPT
`ompt_taskwait_complete`	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 **ompt_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.

40.39 OMPT thread Type

Name: thread **Properties:** C/C++-only, OMPT	**Base Type:** enumeration

Values

Name	Value	Properties
ompt_thread_initial	1	C/C++-only, OMPT
ompt_thread_worker	2	C/C++-only, OMPT
ompt_thread_other	3	C/C++-only, OMPT
ompt_thread_unknown	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**.

40.40 OMPT wait_id Type

Name: wait_id **Properties:** C/C++-only, OMPT	**Base Type: c_uint64_t**

Predefined Identifiers

Name	Value	Properties
ompt_wait_id_none	0	C/C++-only, OMPT

Type Definition

C / C++

```
typedef uint64_t ompt_wait_id_t;
```

C / C++

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.

40.41 OMPT `work` Type

Name: `work`	Base Type: enumeration
Properties: C/C++-only, OMPT, overlapping-type-name	

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.

41 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. Including the C/C++ header file **omp-tools.h** must result in the definition of 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 **ompt_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.

41.1 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.

41.1.1 `initialize` Callback

Name: `initialize` **Category:** function	**Properties:** C/C++-only, OMPT

Return Type and Arguments

Name	Type	Properties
<return type>	integer	*default*
lookup	function_lookup	OMPT
initial_device_num	integer	*default*
tool_data	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 39.2.3.

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 40.8
- **omp_get_initial_device** Routine, see Section 31.10
- **ompt_start_tool** Procedure, see Section 39.2.1
- OMPT **start_tool_result** Type, see Section 40.30

41.1.2 finalize Callback

Name: **finalize** **Category:** subroutine	**Properties:** C/C++-only, OMPT

Arguments

Name	Type	Properties
tool_data	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 40.8
- **ompt_start_tool** Procedure, see Section 39.2.1
- OMPT **start_tool_result** Type, see Section 40.30

41.1.3 thread_begin Callback

Name: **thread_begin** **Category:** subroutine	**Properties:** C/C++-only, OMPT

Arguments

Name	Type	Properties
thread_type	thread	OMPT
thread_data	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);
```

C / C++

Trace Record

C / C++

```
typedef struct ompt_record_thread_begin_t {
  ompt_thread_t thread_type;
} ompt_record_thread_begin_t;
```

C / C++

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 40.8
- OMPT **thread** Type, see Section 40.39

41.1.4 thread_end Callback

Name: thread_end **Category:** subroutine	**Properties:** C/C++-only, OMPT

Arguments

Name	Type	Properties
thread_data	data	OMPT, pointer

Type Signature

C / C++

```
typedef void (*ompt_callback_thread_end_t) (
  ompt_data_t *thread_data);
```

C / C++

Semantics

A tool provides a **thread_end** callback, which has the **thread_end** OMPT type, that the OpenMP implementation dispatches when native threads are destroyed. The binding of the *thread_data* argument is the thread that will be destroyed.

Cross References

- OMPT **data** Type, see Section 40.8

41.2 error Callback

Name: **error** **Category:** subroutine	**Properties:** C/C++-only, OMPT

Arguments

Name	Type	Properties
severity	severity	OMPT
message	char	intent(in), pointer
length	size_t	*default*
codeptr_ra	void	intent(in), pointer

Type Signature

C / C++

```
typedef void (*ompt_callback_error_t) (ompt_severity_t severity,
  const char *message, size_t length, const void *codeptr_ra);
```

C / C++

Trace Record

C / C++

```
typedef struct ompt_record_error_t {
  ompt_severity_t severity;
  const char *message;
  size_t length;
  const void *codeptr_ra;
} ompt_record_error_t;
```

C / C++

Semantics

A tool provides an **error** callback, which has the **error** OMPT type, that the OpenMP implementation dispatches when an **error** directive is encountered for which the *action-time* argument of the **at** clause is specified as **execution**. The *severity* argument passes the specified severity level. The *message* argument passes the C string from the **message** clause. The *length* argument provides the length of the C string.

Cross References

- **error** Directive, see Section 16.1
- OMPT **severity** Type, see Section 40.29

41.3 Parallelism Generation Callback Signatures

This section describes callbacks that are related to constructs for generating and controlling parallelism.

41.3.1 `parallel_begin` Callback

Name: `parallel_begin` **Category:** subroutine	**Properties:** C/C++-only, OMPT, selectable

Arguments

Name	Type	Properties
encountering_task_data	data	OMPT, pointer
encountering_task_frame	frame	intent(in), OMPT, pointer, untraced-argument
parallel_data	data	OMPT, pointer
requested_parallelism	integer	unsigned
flags	integer	selector
codeptr_ra	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++

Semantics

A tool provides a **`parallel_begin`** callback, which has the **`parallel_begin`** OMPT type, that the OpenMP implementation dispatches when a **`parallel`** or **`teams`** region starts. The *requested_parallelism* argument indicates the number of threads or teams that the user requested.

The *flags* argument indicates whether the code for the region is inlined into the application or invoked by the runtime and also whether the region is a **parallel** or **teams** region. Valid values for *flags* are a disjunction of elements in the **parallel_flag** OMPT type.

Cross References

- OMPT **data** Type, see Section 40.8
- OMPT **frame** Type, see Section 40.15
- OMPT **id** Type, see Section 40.18
- **parallel** Construct, see Section 18.1
- OMPT **parallel_flag** Type, see Section 40.22
- **teams** Construct, see Section 18.2

41.3.2 parallel_end Callback

Name: **parallel_end** **Category:** subroutine	**Properties:** C/C++-only, OMPT, selectable

Arguments

Name	Type	Properties
parallel_data	data	OMPT, pointer
encountering_task_data	data	OMPT, pointer
flags	integer	selector
codeptr_ra	void	intent(in), pointer

Type Signature

C / C++

```
typedef void (*ompt_callback_parallel_end_t) (
  ompt_data_t *parallel_data, ompt_data_t *encountering_task_data,
  int flags, const void *codeptr_ra);
```

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 40.8
- OMPT **id** Type, see Section 40.18
- **parallel** Construct, see Section 18.1
- OMPT **parallel_flag** Type, see Section 40.22
- **teams** Construct, see Section 18.2

41.3.3 masked Callback

Name: masked **Category:** subroutine	**Properties:** C/C++-only, OMPT

Arguments

Name	Type	Properties
endpoint	scope_endpoint	OMPT
parallel_data	data	OMPT, pointer
task_data	data	OMPT, pointer
codeptr_ra	void	intent(in), pointer

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 40.8
- **masked** Construct, see Section 18.6
- OMPT **id** Type, see Section 40.18
- OMPT **scope_endpoint** Type, see Section 40.27

41.4 Work Distribution Callback Signatures

This section describes callbacks that are related to work-distribution constructs.

41.4.1 work Callback

Name: **work** **Category:** subroutine	**Properties:** C/C++-only, OMPT, overlapping-type-name, selectable

Arguments

Name	Type	Properties
work_type	work	OMPT, overlapping-type-name, selector
endpoint	scope_endpoint	OMPT
parallel_data	data	OMPT, pointer
task_data	data	OMPT, pointer
count	c_uint64_t	*default*
codeptr_ra	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 40.8
- Work-Distribution Constructs, see Chapter 19
- OMPT **id** Type, see Section 40.18
- OMPT **scope_endpoint** Type, see Section 40.27
- **taskloop** Construct, see Section 20.2
- OMPT **work** Type, see Section 40.41

41.4.2 dispatch Callback

Name: dispatch **Category:** subroutine	**Properties:** C/C++-only, OMPT, overlapping-type-name

Arguments

Name	Type	Properties
parallel_data	data	OMPT, pointer
task_data	data	OMPT, pointer
kind	dispatch	OMPT, overlapping-type-name
instance	data	OMPT

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 40.8
- OMPT **dispatch** Type, see Section 40.13
- OMPT **dispatch_chunk** Type, see Section 40.14
- Worksharing-Loop Constructs, see Section 19.6
- OMPT **id** Type, see Section 40.18
- **sections** Construct, see Section 19.3
- **taskloop** Construct, see Section 20.2

41.5 Tasking Callback Signatures

This section describes callbacks that are related to tasks.

41.5.1 `task_create` Callback

Name: `task_create` **Category:** subroutine	**Properties:** C/C++-only, OMPT

Arguments

Name	Type	Properties
encountering_task_data	data	OMPT, pointer
encountering_task_frame	frame	intent(in), OMPT, pointer, untraced-argument
new_task_data	data	OMPT, pointer
flags	integer	*default*
has_dependences	integer	*default*
codeptr_ra	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 40.8
- OMPT **frame** Type, see Section 40.15
- Initial Task, see Section 20.12
- OMPT **id** Type, see Section 40.18
- **task** Construct, see Section 20.1
- OMPT **task_flag** Type, see Section 40.37

41.5.2 task_schedule Callback

Name: task_schedule **Category:** subroutine	**Properties:** C/C++-only, OMPT

Arguments

Name	Type	Properties
prior_task_data	data	OMPT, pointer
prior_task_status	task_status	OMPT
next_task_data	data	OMPT, pointer

Type Signature

C / C++

```
typedef void (*ompt_callback_task_schedule_t) (
  ompt_data_t *prior_task_data,
  ompt_task_status_t prior_task_status,
  ompt_data_t *next_task_data);
```

C / C++

Trace Record

C / C++

```
typedef struct ompt_record_task_schedule_t {
  ompt_id_t prior_task_id;
  ompt_task_status_t prior_task_status;
  ompt_id_t next_task_id;
} ompt_record_task_schedule_t;
```

C / C++

Semantics

A tool provides a **task_schedule** callback, which has the **task_schedule** OMPT type, that the OpenMP implementation dispatches when task scheduling decisions are made. The binding of the *prior_task_data* argument is the task that arrived at the task scheduling point. This argument can be NULL if no task was active when the next task is scheduled. The *prior_task_status*

argument indicates the status of that prior task. The binding of the *next_task_data* argument is the task that is resumed at the task scheduling point. This argument is NULL if the callback is dispatched for a *task-fulfill* event or if the callback signals completion of a **`taskwait`** construct. This argument can be NULL if no task was active when the prior task was scheduled.

Cross References

- OMPT **`data`** Type, see Section 40.8
- Task Scheduling, see Section 20.13
- OMPT **`id`** Type, see Section 40.18
- OMPT **`task_status`** Type, see Section 40.38

41.5.3 `implicit_task` Callback

Name: `implicit_task` **Category:** subroutine	**Properties:** C/C++-only, OMPT, selectable

Arguments

Name	Type	Properties
endpoint	scope_endpoint	OMPT
parallel_data	data	OMPT, pointer
task_data	data	OMPT, pointer
actual_parallelism	integer	unsigned
index	integer	unsigned
flags	integer	selector

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);
```

C / C++

Trace Record

C / C++

```
typedef struct ompt_record_implicit_task_t {
  ompt_scope_endpoint_t endpoint;
  ompt_id_t parallel_id;
  ompt_id_t task_id;
  unsigned int actual_parallelism;
  unsigned int index;
  int flags;
} ompt_record_implicit_task_t;
```

C / C++

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 40.8
- OMPT **id** Type, see Section 40.18
- **parallel** Construct, see Section 18.1
- OMPT **scope_endpoint** Type, see Section 40.27
- OMPT **task_flag** Type, see Section 40.37
- **teams** Construct, see Section 18.2

41.6 cancel Callback

Name: **cancel** **Category:** subroutine	**Properties:** C/C++-only, OMPT

Arguments

Name	Type	Properties
task_data	data	OMPT, pointer
flags	integer	*default*
codeptr_ra	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 40.7
- OMPT **data** Type, see Section 40.8
- OMPT **id** Type, see Section 40.18

41.7 Synchronization Callback Signatures

This section describes callbacks that are related to synchronization constructs and clauses.

41.7.1 dependences Callback

Name: dependences **Category:** subroutine	**Properties:** C/C++-only, OMPT

Arguments

Name	Type	Properties
task_data	data	OMPT, pointer
deps	dependence	intent(in), pointer
ndeps	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 40.8
- **depend** Clause, see Section 23.9.6
- OMPT **dependence** Type, see Section 40.9
- OMPT **id** Type, see Section 40.18
- Stand-alone **ordered** Construct, see Section 23.10.1

41.7.2 task_dependence Callback

Name: task_dependence **Category:** subroutine	**Properties:** C/C++-only, OMPT

Arguments

Name	Type	Properties
src_task_data	data	OMPT, pointer
sink_task_data	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 40.8
- **depend** Clause, see Section 23.9.6
- OMPT **id** Type, see Section 40.18

41.7.3 OMPT sync_region Type

Name: sync_region **Category:** subroutine pointer	**Properties:** C/C++-only, OMPT, overlapping-type-name

Arguments

Name	Type	Properties
kind	sync_region	OMPT, selector
endpoint	scope_endpoint	OMPT
parallel_data	data	OMPT, pointer
task_data	data	OMPT, pointer
codeptr_ra	void	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++

Semantics

Callbacks that have the **sync_region** OMPT type are synchronizing-region callbacks, which each have the synchronizing-region property. A tool provides these callbacks to mark the beginning and end of regions that have synchronizing semantics. The *kind* argument, which has the **sync_region** OMPT type, indicates the kind of synchronization.

Cross References

- OMPT **data** Type, see Section 40.8
- OMPT **id** Type, see Section 40.18
- OMPT **scope_endpoint** Type, see Section 40.27
- OMPT **sync_region** Type, see Section 40.33

41.7.4 sync_region Callback

Name: **sync_region** **Category:** subroutine	**Properties:** C/C++-only, common-type-callback, synchronizing-region, OMPT, selectable

Type Signature

sync_region

Semantics

A tool provides a **sync_region** callback, which has the **sync_region** OMPT type, that the OpenMP implementation dispatches when barrier regions, **taskwait** regions, **taskgroup** regions, and synchronization caused by an undeferrable task begin and end. For the *implicit-barrier-end* event at the end of a parallel region, *parallel_data* argument is NULL.

Cross References

- **barrier** Construct, see Section 23.3.1
- Implicit Barriers, see Section 23.3.2
- OMPT **sync_region** Type, see Section 41.7.3
- **taskgroup** Construct, see Section 23.4
- **taskwait** Construct, see Section 23.5

41.7.5 sync_region_wait Callback

Name: **sync_region_wait** **Category:** subroutine	**Properties:** C/C++-only, common-type-callback, synchronizing-region, OMPT, selectable

Type Signature

sync_region

Semantics

A tool provides a **sync_region_wait** callback, which has the **sync_region** OMPT type, that the OpenMP implementation dispatches when waiting begins and ends for barrier regions, **taskwait** regions, **taskgroup** regions, and synchronization caused by an undeferrable task. For the *implicit-barrier-wait-begin* and *implicit-barrier-wait-end* events at the end of a parallel region, whether *parallel_data* is NULL or is the current parallel region is implementation defined.

Cross References

- **barrier** Construct, see Section 23.3.1
- Implicit Barriers, see Section 23.3.2
- OMPT **sync_region** Type, see Section 41.7.3
- **taskgroup** Construct, see Section 23.4
- **taskwait** Construct, see Section 23.5

41.7.6 `reduction` Callback

Name: `reduction`	**Properties:** C/C++-only, common-type-callback, synchronizing-region, OMPT
Category: subroutine	

Type Signature

`sync_region`

Semantics

A tool provides a **`reduction`** callback, which is a synchronizing-region callback, that the OpenMP implementation dispatches when it performs reductions.

Cross References

- Properties Common to All Reduction Clauses, see Section 8.6
- OMPT **`sync_region`** Type, see Section 41.7.3

41.7.7 OMPT `mutex_acquire` Type

Name: `mutex_acquire`	**Properties:** C/C++-only, OMPT
Category: subroutine pointer	

Arguments

Name	Type	Properties
kind	mutex	OMPT, overlapping-type-name, selector
hint	integer	unsigned
impl	integer	unsigned
wait_id	wait_id	OMPT
codeptr_ra	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);
```

C / C++

Trace Record

C / C++

```
typedef struct ompt_record_mutex_acquire_t {
  ompt_mutex_t kind;
  unsigned int hint;
  unsigned int impl;
  ompt_wait_id_t wait_id;
  const void *codeptr_ra;
} ompt_record_mutex_acquire_t;
```

C / C++

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 40.20
- OMPT **wait_id** Type, see Section 40.40

41.7.8 mutex_acquire Callback

Name: **mutex_acquire** **Category:** subroutine	**Properties:** C/C++-only, common-type-callback, mutex-acquiring, OMPT, selectable

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 41.7.7

41.7.9 `lock_init` Callback

Name: `lock_init`	**Properties:** C/C++-only, common-type-callback, mutex-acquiring, OMPT, selectable
Category: subroutine	

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 41.7.7

41.7.10 OMPT `mutex` Type

Name: `mutex`	**Properties:** C/C++-only, OMPT, overlapping-type-name
Category: subroutine pointer	

Arguments

Name	Type	Properties
kind	mutex	OMPT, overlapping-type-name, selector
wait_id	wait_id	OMPT
codeptr_ra	void	intent(in), pointer

Type Signature

C / C++

```
typedef void (*ompt_callback_mutex_t) (ompt_mutex_t kind,
  ompt_wait_id_t wait_id, const void *codeptr_ra);
```

C / C++

Trace Record

C / C++

```
typedef struct ompt_record_mutex_t {
  ompt_mutex_t kind;
  ompt_wait_id_t wait_id;
  const void *codeptr_ra;
} ompt_record_mutex_t;
```

C / C++

Semantics

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.

Cross References

- Lock Acquiring Routines, see Section 35.3
- Lock Destroying Routines, see Section 35.2
- Lock Releasing Routines, see Section 35.4
- Lock Testing Routines, see Section 35.5
- OMPT **mutex** Type, see Section 40.20
- OMPT **wait_id** Type, see Section 40.40

41.7.11 lock_destroy Callback

Name: lock_destroy **Category:** subroutine	**Properties:** C/C++-only, common-type-callback, mutex-execution, OMPT, selectable

Type Signature

```
mutex
```

Semantics

A tool provides a **lock_destroy** callback, which has the **mutex** OMPT type, that the OpenMP implementation dispatches when it executes a lock-destroying routine.

Cross References

- Lock Destroying Routines, see Section 35.2
- OMPT **mutex** Type, see Section 41.7.10

41.7.12 mutex_acquired Callback

Name: mutex_acquired **Category:** subroutine	**Properties:** C/C++-only, common-type-callback, mutex-execution, OMPT, selectable

Type Signature

```
mutex
```

Semantics

A tool provides a **mutex_acquired** callback, which has the **mutex** OMPT type, that the OpenMP implementation dispatches when the structured block associated with a mutual-exclusion construct begins execution or when a region guarded by a lock-acquiring routine or lock-testing routine begins execution.

Cross References

- Lock Acquiring Routines, see Section 35.3
- Lock Testing Routines, see Section 35.5
- OMPT **mutex** Type, see Section 41.7.10

41.7.13 mutex_released Callback

Name: **mutex_released** **Category:** subroutine	**Properties:** C/C++-only, common-type-callback, mutex-execution, OMPT, selectable

Type Signature

mutex

Semantics

A tool provides a **mutex_released** callback, which has the **mutex** OMPT type, that the OpenMP implementation dispatches when the structured block associated with a mutual-exclusion construct completes execution or, similarly, when a region that a lock-releasing routine guards completes execution.

Cross References

- Lock Releasing Routines, see Section 35.4
- OMPT **mutex** Type, see Section 41.7.10

41.7.14 nest_lock Callback

Name: **nest_lock** **Category:** subroutine	**Properties:** C/C++-only, OMPT

Arguments

Name	Type	Properties
endpoint	scope_endpoint	OMPT
wait_id	wait_id	OMPT
codeptr_ra	void	intent(in), pointer

Type Signature

C / C++

```
typedef void (*ompt_callback_nest_lock_t) (
  ompt_scope_endpoint_t endpoint, ompt_wait_id_t wait_id,
  const void *codeptr_ra);
```

C / C++

Trace Record

C / C++

```
typedef struct ompt_record_nest_lock_t {
  ompt_scope_endpoint_t endpoint;
  ompt_wait_id_t wait_id;
  const void *codeptr_ra;
} ompt_record_nest_lock_t;
```

C / C++

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 40.27
- OMPT **wait_id** Type, see Section 40.40

41.7.15 flush Callback

Name: flush **Category:** subroutine	**Properties:** C/C++-only, OMPT

Arguments

Name	Type	Properties
thread_data	data	OMPT, pointer, untraced-argument
codeptr_ra	void	intent(in), pointer

Type Signature

C / C++

```
typedef void (*ompt_callback_flush_t) (ompt_data_t *thread_data,
  const void *codeptr_ra);
```

C / C++

Trace Record

C / C++

```
typedef struct ompt_record_flush_t {
  const void *codeptr_ra;
} ompt_record_flush_t;
```

C / C++

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 40.8
- **flush** Construct, see Section 23.8.6

41.8 control_tool Callback

Name: control_tool **Category:** function	**Properties:** C/C++-only, OMPT

Return Type and Arguments

Name	Type	Properties
<return type>	integer	*default*
command	c_uint64_t	*default*
modifier	c_uint64_t	*default*
arg	c_ptr	iso_c, value, untraced-argument
codeptr_ra	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 38.1

42 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. Including the C/C++ header file **`omp-tools.h`** must result in the definition of the types that this chapter defines.

42.1 `device_initialize` Callback

Name: `device_initialize` **Category:** subroutine	**Properties:** C/C++-only, device-tracing, OMPT

Arguments

Name	Type	Properties
device_num	integer	*default*
type	char	intent(in), pointer
device	device	OMPT, opaque, pointer
lookup	function_lookup	OMPT
documentation	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 39.2.5.

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 40.11
- **`function_lookup`** Entry Point, see Section 43.1

42.2 `device_finalize` Callback

Name: **`device_finalize`** **Category:** subroutine	**Properties:** C/C++-only, device-tracing, OMPT

Arguments

Name	Type	Properties
device_num	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 42.6

42.3 device_load Callback

Name: **device_load** **Category:** subroutine	**Properties:** C/C++-only, device-tracing, OMPT

Arguments

Name	Type	Properties
device_num	integer	*default*
filename	char	intent(in), pointer
offset_in_file	c_int64_t	iso_c, value
vma_in_file	c_ptr	iso_c, value
bytes	c_size_t	iso_c, value
host_addr	c_ptr	iso_c, value
device_addr	c_ptr	iso_c, value
module_id	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.

42.4 `device_unload` Callback

Name: `device_unload` **Category:** subroutine	**Properties:** C/C++-only, device-tracing, OMPT

Arguments

Name	Type	Properties
device_num	integer	*default*
module_id	c_uint64_t	*default*

Type Signature

C / C++

```
typedef void (*ompt_callback_device_unload_t) (int device_num,
  uint64_t module_id);
```

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.

42.5 `buffer_request` Callback

Name: `buffer_request` **Category:** subroutine	**Properties:** C/C++-only, device-tracing, OMPT

Arguments

Name	Type	Properties
device_num	integer	*default*
buffer	buffer	pointer-to-pointer
bytes	size_t	pointer

Type Signature

C / C++

```
typedef void (*ompt_callback_buffer_request_t) (int device_num,
  ompt_buffer_t **buffer, size_t *bytes);
```

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 40.3
- **get_buffer_limits** Entry Point, see Section 44.7

42.6 buffer_complete Callback

Name: **buffer_complete** **Category:** subroutine	**Properties:** C/C++-only, device-tracing, OMPT

Arguments

Name	Type	Properties
device_num	integer	*default*
buffer	buffer	pointer
bytes	size_t	*default*
begin	buffer_cursor	OMPT, opaque
buffer_owned	integer	*default*

Type Signature

C / C++

```
typedef void (*ompt_callback_buffer_complete_t) (int device_num,
  ompt_buffer_t *buffer, size_t bytes, ompt_buffer_cursor_t begin,
  int buffer_owned);
```

C / C++

Semantics

A 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 in the buffer. The *buffer_owned* argument is 1 if the data to which *buffer* points can be deleted by the callback and 0 otherwise. If multiple devices accumulate events into a single buffer, this callback may be invoked with a pointer to one or more trace records in a shared buffer with *buffer_owned* equal to zero.

Typically, a tool will iterate through the trace records in the buffer and process them. The OpenMP implementation makes these callbacks on a native thread that is not an OpenMP thread so these **buffer_complete** callbacks are not required to be async signal safe.

Restrictions

Restrictions on **buffer_complete** callbacks are as follows:

- The callback must not delete the buffer if *buffer_owned* is zero.

Cross References

- OMPT **buffer** Type, see Section 40.3
- OMPT **buffer_cursor** Type, see Section 40.4

42.7 target_data_op_emi Callback

Name: target_data_op_emi **Category:** subroutine	**Properties:** C/C++-only, device-tracing, OMPT, selectable

Arguments

Name	Type	Properties
endpoint	scope_endpoint	OMPT, untraced-argument
target_task_data	data	OMPT, pointer, untraced-argument
target_data	data	OMPT, pointer, untraced-argument
host_op_id	id	OMPT, pointer
optype	target_data_op	OMPT, selector
dev1_addr	c_ptr	iso_c, value
dev1_device_num	integer	*default*
dev2_addr	c_ptr	iso_c, value
dev2_device_num	integer	*default*
bytes	size_t	*default*
codeptr_ra	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 optype, 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 optype;
  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 42.1 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 42.1. The *dev2_addr* argument indicates the data address on the device given by Table 42.1. 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 42.1. 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.

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.

TABLE 42.1: Association of dev1 and dev2 arguments for target data operations

Data op	dev1	dev2
allocate	host/none	device
transfer	*from* device	*to* device
delete	host/none	device
associate	host	device
disassociate	host	device
memset	none	device

Cross References

- OMPT **data** Type, see Section 40.8
- OMPT **device_time** Type, see Section 40.12
- OMPT **id** Type, see Section 40.18
- **map** Clause, see Section 11.5
- OMPT **scope_endpoint** Type, see Section 40.27
- OMPT **target_data_op** Type, see Section 40.35

42.8 target_emi Callback

Name: target_emi **Category:** subroutine	**Properties:** C/C++-only, device-tracing, OMPT, selectable

Arguments

Name	Type	Properties
kind	target	OMPT, selector
endpoint	scope_endpoint	OMPT
device_num	integer	*default*
task_data	data	OMPT, pointer
target_task_data	data	OMPT, pointer, untraced-argument
target_data	data	OMPT, pointer
codeptr_ra	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 40.8
- OMPT **id** Type, see Section 40.18
- OMPT **scope_endpoint** Type, see Section 40.27
- **target** Construct, see Section 21.8
- OMPT **target** Type, see Section 40.34
- **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

42.9 target_map_emi Callback

Name: target_map_emi **Category:** subroutine	**Properties:** C/C++-only, device-tracing, OMPT

Arguments

Name	Type	Properties
target_data	data	OMPT, pointer
nitems	integer	unsigned
host_addr	void	pointer-to-pointer
device_addr	void	pointer-to-pointer
bytes	size_t	pointer
mapping_flags	integer	unsigned, pointer
codeptr_ra	void	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 40.8
- OMPT **id** Type, see Section 40.18
- **map** Clause, see Section 11.5
- **target_data_op_emi** Callback, see Section 42.7
- OMPT **target_map_flag** Type, see Section 40.36

42.10 target_submit_emi Callback

Name: target_submit_emi **Category:** subroutine	**Properties:** C/C++-only, device-tracing, OMPT

Arguments

Name	Type	Properties
endpoint	scope_endpoint	OMPT, untraced-argument
target_data	data	OMPT, pointer, untraced-argument
host_op_id	id	OMPT, pointer
requested_num_teams	integer	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 40.8
- OMPT **device_time** Type, see Section 40.12
- OMPT **id** Type, see Section 40.18
- OMPT **scope_endpoint** Type, see Section 40.27
- **target** Construct, see Section 21.8

43 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.

43.1 function_lookup Entry Point

Name: function_lookup **Category:** function	**Properties:** C/C++-only, OMPT

Return Type and Arguments

Name	Type	Properties
<return type>	interface_fn	*default*
interface_function_name	char	intent(in), pointer

Type Signature

C / C++

```
typedef ompt_interface_fn_t (*ompt_function_lookup_t) (
  const char *interface_function_name);
```

C / C++

Semantics

The **function_lookup** entry point, which has the **function_lookup** OMPT type, enables tools to look up pointers to OMPT entry points by name. When an OpenMP implementation invokes the **initialize** callback to configure the OMPT callback interface, it provides an entry point that provides pointers to other entry points that implement routines that are part of the OMPT callback interface. Alternatively, when it invokes a **device_initialize** callback to configure the OMPT tracing interface for a device, it provides an entry point that provides pointers to entry points that implement tracing control routines appropriate for that device.

For these entry points, the *interface_function_name* argument is a C string that represents the name of the entry point to look up. If the name is unknown to the implementation, the entry point returns NULL. In a compliant implementation, the entry point that is provided by the **initialize** callback returns a valid function pointer for any entry point name listed in Table 39.1. Similarly, in a compliant implementation, the entry point that is provided by the **device_initialize** callback returns non-NULL function pointers for any entry point name listed in Table 39.3, except for **set_trace_ompt** and **get_record_ompt**, as described in Section 39.2.5.

Cross References

- **device_initialize** Callback, see Section 42.1
- Binding Entry Points, see Section 39.2.3.1
- Tracing Activity on Target Devices, see Section 39.2.5
- **initialize** Callback, see Section 41.1.1
- OMPT **interface_fn** Type, see Section 40.19

43.2 enumerate_states Entry Point

Name: **enumerate_states** **Category:** function	**Properties:** C/C++-only, OMPT

Return Type and Arguments

Name	Type	Properties
<return type>	integer	*default*
current_state	integer	*default*
next_state	integer	pointer
next_state_name	const char	intent(out), pointer-to-pointer

Type Signature

C / C++

```
typedef int (*ompt_enumerate_states_t) (int current_state,
  int *next_state, const char **next_state_name);
```

C / C++

Semantics

An OpenMP implementation may support only a subset of the thread states that the **`state`** OMPT type defines. An OpenMP implementation may also support implementation defined states. The **`enumerate_states`** entry point, which has the **`enumerate_states`** OMPT type, is the entry point that enables a tool to enumerate the supported thread states.

When a supported thread state is passed as *current_state*, the entry point assigns the next thread state in the enumeration to the variable passed by reference in *next_state* and assigns the name associated with that state to the character pointer passed by reference in *next_state_name*; the returned string is immutable and defined for the lifetime of program execution. Whenever one or more states are left in the enumeration, the **`enumerate_states`** entry point returns 1. When the last state in the enumeration is passed as *current_state*, **`enumerate_states`** returns 0, which indicates that the enumeration is complete.

To begin enumerating the supported states, a tool should pass **`ompt_state_undefined`** as *current_state*. Subsequent invocations of **`enumerate_states`** should pass the value assigned to the variable that was passed by reference in *next_state* to the previous call. The **`ompt_state_undefined`** value is returned to indicate an invalid thread state.

Cross References

- OMPT **`state`** Type, see Section 40.31

43.3 `enumerate_mutex_impls` Entry Point

Name: `enumerate_mutex_impls` **Category:** function	**Properties:** C/C++-only, OMPT

Return Type and Arguments

Name	Type	Properties
<return type>	integer	*default*
current_impl	integer	*default*
next_impl	integer	pointer
next_impl_name	const char	intent(out), pointer-to-pointer

Type Signature

C / C++

```
typedef int (*ompt_enumerate_mutex_impls_t) (int current_impl,
  int *next_impl, const char **next_impl_name);
```

C / C++

Semantics

Mutual exclusion for locks, **critical** regions, and **atomic** regions may be implemented in several ways. The **enumerate_mutex_impls** entry point, which has the **enumerate_mutex_impls** OMPT type, enables a tool to enumerate the supported mutual exclusion implementations.

When a supported mutex implementation is passed as *current_impl*, the entry point assigns the next mutex implementation in the enumeration to the variable passed by reference in *next_impl* and assigns the name associated with that mutex implementation to the character pointer passed by reference in *next_impl_name*; the returned string is immutable and defined for the lifetime of program execution. Whenever one or more mutex implementations are left in the enumeration, the **enumerate_mutex_impls** entry point returns 1. When the last mutex implementation in the enumeration is passed as *current_impl*, the entry point returns 0, which indicates that the enumeration is complete.

To begin enumerating the supported mutex implementations, a tool should pass **ompt_mutex_impl_none** as *current_impl*. Subsequent invocations of **enumerate_mutex_impls** should pass the value assigned to the variable that was passed by reference in *next_impl* to the previous call. The value **ompt_mutex_impl_none** is returned to indicate an invalid mutex implementation.

43.4 set_callback Entry Point

Name: set_callback **Category:** function	**Properties:** C/C++-only, OMPT

Return Type and Arguments

Name	Type	Properties
<return type>	set_result	*default*
event	callbacks	OMPT
callback	callback	OMPT

Type Signature

C / C++

```
typedef ompt_set_result_t (*ompt_set_callback_t) (
  ompt_callbacks_t event, ompt_callback_t callback);
```

C / C++

Semantics

OpenMP implementations can use callbacks to indicate the occurrence of events during the execution of an OpenMP program. The **`set_callback`** entry point, which has the **`set_callback`** OMPT type, enables a tool to register the callback indicated by the *callback* argument for the event indicated by the *event* argument on the current device. The return value of **`set_callback`** indicates the outcome of registering the callback and may be any value in the **`set_result`** OMPT type except **`ompt_set_impossible`**. If *callback* is NULL then callbacks associated with *event* are disabled whether set by this entry point or the **`set_selective_callback`** entry point. If callbacks are successfully disabled then **`ompt_set_always`** is returned.

Restrictions

Restrictions on the **`set_callback`** entry point are as follows:

- The type signature for *callback* must match the type signature appropriate for the event.

Cross References

- OMPT **`callback`** Type, see Section 40.5
- OMPT **`callbacks`** Type, see Section 40.6
- Monitoring Activity on the Host with OMPT, see Section 39.2.4
- OMPT **`set_result`** Type, see Section 40.28

43.5 `set_selective_callback` Entry Point

Name: **`set_selective_callback`** **Category:** function	**Properties:** C/C++-only, OMPT

Return Type and Arguments

Name	Type	Properties
<return type>	set_result	*default*
selection	c_uint64_t	*default*
event	callbacks	OMPT
callback	callback	OMPT

Type Signature

C / C++

```
typedef ompt_set_result_t (*ompt_set_selective_callback_t) (
  uint64_t selection, ompt_callbacks_t event,
  ompt_callback_t callback);
```

C / C++

Semantics

The **set_selective_callback** entry point, which has the **set_selective_callback** OMPT type, enables a tool to register a selectable callback indicated by the *callback* argument for the event indicated by the *event* argument on the current device for a selective occurence indicated by the *selection* argument. Callback dispatch will only happen if it would pass a value to the selector argument of the callback that was also passed to the *selection* argument of this function during registration. If *callback* is NULL then the callback dispatch associated with the *event* and *selection* arguments is disabled, whether set by this entry point or the **set_callback** entry point. The return value of **set_selective_callback** indicates the outcome of registering the callback and may be any value in the **set_result** OMPT type except **ompt_set_impossible**.

Restrictions

Restrictions on the **set_selective_callback** entry point are as follows:

- The type signature for *callback* must match the type signature appropriate for the event.
- The *callback* argument must be the same for all calls that use the same event argument.

Cross References

- OMPT **callback** Type, see Section 40.5
- OMPT **callbacks** Type, see Section 40.6
- Monitoring Activity on the Host with OMPT, see Section 39.2.4
- OMPT **set_result** Type, see Section 40.28

43.6 get_callback Entry Point

Name: **get_callback** **Category:** function	**Properties:** C/C++-only, OMPT

Return Type and Arguments

Name	Type	Properties
<return type>	integer	*default*
event	callbacks	OMPT
callback	callback	OMPT, pointer

Type Signature

C / C++

```
typedef int (*ompt_get_callback_t) (ompt_callbacks_t event,
  ompt_callback_t *callback);
```

C / C++

Semantics

The **get_callback** entry point, which has the **get_callback** OMPT type, enables a tool to retrieve a pointer to a registered callback (if any) that an OpenMP implementation invokes when a host event occurs. If the callback that is registered for the event that is specified by the *event* argument is not NULL, the pointer to the callback is assigned to the variable passed by reference in *callback* and **get_callback** returns 1; otherwise, it returns 0. If **get_callback** returns 0, the value of the variable passed by reference as *callback* is undefined.

Restrictions

Restrictions on the **get_callback** entry point are as follows:

- The *callback* argument must not be NULL and must point to valid storage.

Cross References

- OMPT **callback** Type, see Section 40.5
- OMPT **callbacks** Type, see Section 40.6
- **set_callback** Entry Point, see Section 43.4

43.7 get_thread_data Entry Point

Name: get_thread_data **Category:** function	**Properties:** async-signal-safe, C/C++-only, OMPT

Return Type

Name	Type	Properties
<return type>	data	pointer

Type Signature

C / C++

```
typedef ompt_data_t  *(*ompt_get_thread_data_t) (void);
```

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 40.8

43.8 `get_num_procs` Entry Point

Name: `get_num_procs` **Category:** function	**Properties:** all-device-threads-binding, async-signal-safe, C/C++-only, OMPT

Return Type

Name	Type	Properties
<return type>	integer	*default*

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.

43.9 `get_num_places` Entry Point

Name: `get_num_places` **Category:** function	**Properties:** all-device-threads-binding, async-signal-safe, C/C++-only, OMPT

Return Type

Name	Type	Properties
<return type>	integer	*default*

Type Signature

C / C++

```
typedef int (*ompt_get_num_places_t) (void);
```

C / C++

Semantics

The **`get_num_places`** entry point, which has the **`get_num_places`** OMPT type, enables a tool to retrieve the number of places in the place list. This value is equal to the number of places in the *place-group-var* ICV in the execution environment of the initial task. The binding thread set of this entry point is all threads on the host device.

Cross References

- **`OMP_PLACES`**, see Section 4.1.6
- *place-group-var* ICV, see Table 3.1

43.10 `get_place_proc_ids` Entry Point

Name: `get_place_proc_ids` **Category:** function	**Properties:** all-device-threads-binding, C/C++-only, OMPT

Return Type and Arguments

Name	Type	Properties
<return type>	integer	*default*
place_num	integer	*default*
ids_size	integer	*default*
ids	integer	pointer

Type Signature

C / C++

```
typedef int (*ompt_get_place_proc_ids_t) (int place_num,
  int ids_size, int *ids);
```

C / C++

Semantics

The **`get_place_proc_ids`** entry point, which has the **`get_place_proc_ids`** OMPT type, enables a tool to retrieve the numerical identifiers of each processor that is associated with the place specified by the *place_num* argument. The *ids* argument is an array in which the entry point can return a vector of processor identifiers in the specified place; these identifiers are non-negative, and their meaning is implementation defined. The *ids_size* argument indicates the size of the result 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.

43.11 `get_place_num` Entry Point

Name: `get_place_num` **Category:** function	**Properties:** async-signal-safe, C/C++-only, OMPT

Return Type

Name	Type	Properties
<return type>	integer	*default*

Type Signature

C / C++

```
typedef int (*ompt_get_place_num_t) (void);
```

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.

43.12 `get_partition_place_nums` Entry Point

Name: `get_partition_place_nums` **Category:** function	**Properties:** async-signal-safe, C/C++-only, OMPT

Return Type and Arguments

Name	Type	Properties
<return type>	integer	*default*
place_nums_size	integer	*default*
place_nums	integer	pointer

Type Signature

C / C++

```
typedef int (*ompt_get_partition_place_nums_t) (
  int place_nums_size, int *place_nums);
```

C / C++

Semantics

The **`get_partition_place_nums`** entry point, which has the **`get_partition_place_nums`** OMPT type, enables a tool to retrieve a list of place numbers that correspond to the places in the *place-group-var* ICV of the innermost implicit task. The *place_nums* argument is an array in which the entry point can return a vector of place identifiers. The *place_nums_size* argument indicates the size of that array.

If the *place_nums* array of size *place_nums_size* is large enough to contain all identifiers then they are returned in *place_nums* and their order in the array is implementation defined. Otherwise, if the *place_nums* array is too small, the values in *place_nums* when the entry point returns are undefined. The entry point always returns the number of places in the *place-group-var* ICV of the innermost implicit task.

Cross References

- **`OMP_PLACES`**, see Section 4.1.6
- *place-group-var* ICV, see Table 3.1

43.13 `get_proc_id` Entry Point

Name: `get_proc_id` **Category:** function	**Properties:** async-signal-safe, C/C++-only, OMPT

Return Type

Name	Type	Properties
<return type>	integer	*default*

Type Signature

C / C++

```
typedef int (*ompt_get_proc_id_t) (void);
```

C / C++

The **`get_proc_id`** entry point, which has the **`get_proc_id`** OMPT type, enables a tool to retrieve the numerical identifier of the processor of the encountering thread. A defined numerical identifier is non-negative, and its meaning is implementation defined. A negative number indicates a failure to retrieve the numerical identifier.

43.14 `get_state` Entry Point

Name: `get_state` **Category:** function	**Properties:** async-signal-safe, C/C++-only, OMPT

Return Type and Arguments

Name	Type	Properties
<return type>	integer	*default*
wait_id	wait_id	OMPT, pointer

Type Signature

C / C++

```
typedef int (*ompt_get_state_t) (ompt_wait_id_t *wait_id);
```

C / C++

Semantics

Each thread has an associated state and a wait identifier. If the thread state indicates that the thread is waiting for mutual exclusion then its wait identifier contains a handle that indicates the data object upon which the thread is waiting. The **`get_state`** entry point, which has the **`get_state`** OMPT type, enables a tool to retrieve the state and the wait identifier of the encountering thread. The returned value may be any one of the states predefined by the **`state`** OMPT type or a value that represents an implementation defined state. The tool may obtain a string representation for each state with the **`enumerate_states`** entry point. If the returned state indicates that the thread is waiting for a lock, nestable lock, **`critical`** region, **`atomic`** region, or **`ordered`** region and the wait identifier passed as the *wait_id* argument is not NULL then the value of the wait identifier is assigned to that argument, which is a pointer to a handle. If the returned state is not one of the specified wait states then the value of that handle is undefined after the call.

Restrictions

Restrictions on the **get_state** entry point are as follows:

- The *wait_id* argument must be a reference to a variable of the **wait_id** OMPT type or NULL.

Cross References

- **enumerate_states** Entry Point, see Section 43.2
- OMPT **state** Type, see Section 40.31
- OMPT **wait_id** Type, see Section 40.40

43.15 get_parallel_info Entry Point

Name: get_parallel_info **Category:** function	**Properties:** async-signal-safe, C/C++-only, OMPT

Return Type and Arguments

Name	Type	Properties
<return type>	integer	*default*
ancestor_level	integer	*default*
parallel_data	data	OMPT, pointer-to-pointer
team_size	integer	pointer

Type Signature

C / C++

```
typedef int (*ompt_get_parallel_info_t) (int ancestor_level,
  ompt_data_t **parallel_data, int *team_size);
```

C / C++

Semantics

During execution, an OpenMP program may employ nested parallel regions. The **get_parallel_info** entry point, which has the **get_parallel_info** OMPT type, enables a tool to retrieve information about the current parallel region and any enclosing parallel regions for the current execution context.

The *ancestor_level* argument specifies the parallel region of interest by its ancestor level. Ancestor level 0 refers to the innermost parallel region; information about enclosing parallel regions may be obtained using larger values for *ancestor_level*. Information about a parallel region may not be available if the ancestor level is 0; otherwise it must be available if a parallel region exists at the

specified ancestor level. The entry point returns 2 if a parallel region exists at the specified ancestor level and the information is available, 1 if a parallel region exists at the specified ancestor level but the information is currently unavailable, and 0 otherwise. The *parallel_data* argument returns the parallel data if the argument is not NULL. The *team_size* argument returns the team size if the argument is not NULL. If no parallel 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_parallel_info`** returns.

A tool may use the pointer to the data object of a parallel region that it obtains from this entry point to inspect or to modify the value of the data object. When a parallel region is created, its data object will be initialized with the value **`ompt_data_none`**. Between a *parallel-begin* event and an *implicit-task-begin* event, a call to **`get_parallel_info`** with an *ancestor_level* value of 0 may return information about the outer team or the new team. If a thread is in the **`ompt_state_wait_barrier_implicit_parallel`** state then a call to **`get_parallel_info`** may return a pointer to a copy of the specified parallel region's *parallel_data* rather than a pointer to the data word for the region itself. This convention enables the primary thread for a parallel region to free storage for the region immediately after the region ends, yet avoid having some other thread in the team that is executing the region potentially reference the *parallel_data* object for the region after it has been freed.

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 40.8
- OMPT **`state`** Type, see Section 40.31

43.16 `get_task_info` Entry Point

Name: **`get_task_info`** **Category:** function	**Properties:** async-signal-safe, C/C++-only, OMPT

Return Type and Arguments

Name	Type	Properties
<return type>	integer	*default*
ancestor_level	integer	*default*
flags	integer	pointer
task_data	data	OMPT, pointer-to-pointer
task_frame	frame	OMPT, pointer-to-pointer
parallel_data	data	OMPT, pointer-to-pointer
thread_num	integer	pointer

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.

Restrictions

Restrictions on the **get_task_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 40.8
- OMPT **frame** Type, see Section 40.15
- OMPT **task_flag** Type, see Section 40.37

43.17 get_task_memory Entry Point

Name: **get_task_memory** **Category:** function	**Properties:** async-signal-safe, C/C++-only, OMPT

Return Type and Arguments

Name	Type	Properties
<return type>	integer	*default*
addr	void	pointer-to-pointer
size	size_t	pointer
block	integer	*default*

Type Signature

C / C++

```
typedef int (*ompt_get_task_memory_t) (void **addr, size_t *size,
  int block);
```

C / C++

Semantics

During execution, a thread may be executing a task. The OpenMP implementation must preserve the data environment from the generation of the task for its execution. The **`get_task_memory`** entry point, which has the **`get_task_memory`** OMPT type, enables a tool to retrieve information about memory ranges that store the data environment for the encountering task. Multiple memory ranges may be used to store these data. The *addr* argument is a pointer to a void pointer return value to provide the start address of a memory range. The *size* argument is a pointer to a size type return value to provide the size of the memory range. The *block* argument, which is an integer value to specify the memory block of interest, supports iteration over the memory ranges. The **`get_task_memory`** entry point returns one if more memory ranges are available, and zero otherwise. If no memory is used for a task, *size* is set to zero. In this case, the value to which *addr* points is undefined.

43.18 `get_target_info` Entry Point

Name: `get_target_info` **Category:** function	**Properties:** async-signal-safe, C/C++-only, OMPT

Return Type and Arguments

Name	Type	Properties
<return type>	integer	*default*
device_num	c_uint64_t	pointer
target_id	id	OMPT, pointer
host_op_id	id	OMPT, pointer-to-pointer

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 40.18

43.19 get_num_devices Entry Point

Name: get_num_devices **Category:** function	**Properties:** async-signal-safe, C/C++-only, OMPT

Return Type

Name	Type	Properties
<return type>	integer	*default*

Type Signature

C / C++

```
typedef int (*ompt_get_num_devices_t) (void);
```

C / C++

Semantics

The **get_num_devices** entry point, which has the **get_num_devices** OMPT type, is the entry point that enables a tool to retrieve the number of devices available to an OpenMP program.

43.20 get_unique_id Entry Point

Name: get_unique_id **Category:** function	**Properties:** async-signal-safe, C/C++-only, OMPT

Return Type

Name	Type	Properties
<return type>	c_uint64_t	*default*

Type Signature

C / C++

```
typedef uint64_t (*ompt_get_unique_id_t) (void);
```

C / C++

Semantics

The **`get_unique_id`** entry point, which has the **`get_unique_id`** OMPT type, enables a tool to retrieve a number that is unique for the duration of an OpenMP program. Successive invocations may not result in consecutive or even increasing numbers.

43.21 `finalize_tool` Entry Point

Name: `finalize_tool` **Category:** subroutine	**Properties:** C/C++-only, OMPT

Type Signature

C / C++

```
typedef void (*ompt_finalize_tool_t) (void);
```

C / C++

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.

44 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. An implementation might provide different sets of these entry points for each device through the **`function_lookup`** entry point provided in the **`device_initialize`** callback.

Several OMPT entry points have a *device* argument. This argument is a pointer to an OpenMP object that represents the target device. The **`device_initialize`** callback provides the pointer to this device object that identifies the device being addressed.

44.1 `get_device_num_procs` Entry Point

Name: `get_device_num_procs` **Category:** function	**Properties:** C/C++-only, OMPT

Return Type and Arguments

Name	Type	Properties
<return type>	integer	*default*
device	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 40.11

44.2 `get_uid_from_device` Entry Point

Name: `get_uid_from_device` **Category:** function	**Properties:** C/C++-only, OMPT

Return Type and Arguments

Name	Type	Properties
<return type>	const char	pointer
device	device	OMPT, pointer

Type Signature

C / C++

```
typedef const char  *(*ompt_get_uid_from_device_t) (
  ompt_device_t *device);
```

C / C++

Semantics

The **get_uid_from_device** entry point returns the implementation defined unique identifier string that identifies the device specified by *device*.

Cross References

- OMPT **device** Type, see Section 40.11
- **omp_get_uid_from_device** Routine, see Section 31.8

44.3 get_device_time Entry Point

Name: get_device_time **Category:** function	**Properties:** C/C++-only, OMPT

Return Type and Arguments

Name	Type	Properties
<return type>	device_time	*default*
device	device	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 40.11
- OMPT **device_time** Type, see Section 40.12

44.4 `translate_time` Entry Point

Name: `translate_time` **Category:** function	**Properties:** C/C++-only, OMPT

Return Type and Arguments

Name	Type	Properties
<return type>	double	*default*
device	device	OMPT, pointer
time	device_time	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 40.11
- OMPT **`device_time`** Type, see Section 40.12
- **`omp_get_wtime`** Routine, see Section 37.3.1

44.5 `set_trace_ompt` Entry Point

Name: `set_trace_ompt` **Category:** function	**Properties:** C/C++-only, OMPT

Return Type and Arguments

Name	Type	Properties
<return type>	set_result	*default*
device	device	OMPT, pointer
enable	integer	OMPT, unsigned
etype	integer	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 40.6
- OMPT **device** Type, see Section 40.11
- Tracing Activity on Target Devices, see Section 39.2.5
- OMPT **set_result** Type, see Section 40.28

44.6 set_trace_native Entry Point

Name: **set_trace_native** **Category:** function	**Properties:** C/C++-only, OMPT

Return Type and Arguments

Name	Type	Properties
<return type>	set_result	*default*
device	device	OMPT, pointer
enable	integer	*default*
flags	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 40.11
- Tracing Activity on Target Devices, see Section 39.2.5
- OMPT **native_mon_flag** Type, see Section 40.21
- OMPT **set_result** Type, see Section 40.28

44.7 get_buffer_limits Entry Point

Name: **get_buffer_limits** **Category:** subroutine	**Properties:** C/C++-only, OMPT

Arguments

Name	Type	Properties
device	device	OMPT, pointer
max_concurrent_allocs	integer	pointer
recommended_bytes	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 42.6
- **buffer_request** Callback, see Section 42.5
- OMPT **device** Type, see Section 40.11
- **start_trace** Entry Point, see Section 44.8
- **stop_trace** Entry Point, see Section 44.11

44.8 start_trace Entry Point

Name: start_trace **Category:** function	**Properties:** C/C++-only, OMPT

Return Type and Arguments

Name	Type	Properties
<return type>	integer	*default*
device	device	OMPT, pointer
request	buffer_request	OMPT, procedure
complete	buffer_complete	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 42.6
- **buffer_request** Callback, see Section 42.5
- OMPT **device** Type, see Section 40.11

44.9 pause_trace Entry Point

Name: pause_trace **Category:** function	**Properties:** C/C++-only, OMPT

Return Type and Arguments

Name	Type	Properties
<return type>	integer	*default*
device	device	OMPT, pointer
begin_pause	integer	*default*

Type Signature

C / C++

```
typedef int (*ompt_pause_trace_t) (ompt_device_t *device,
  int begin_pause);
```

C / C++

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 40.11

44.10 flush_trace Entry Point

Name: flush_trace **Category:** function	**Properties:** C/C++-only, OMPT

Return Type and Arguments

Name	Type	Properties
<return type>	integer	*default*
device	device	OMPT, pointer

Type Signature

C / C++

```
typedef int (*ompt_flush_trace_t) (ompt_device_t *device);
```

C / C++

Semantics

The **flush_trace** entry point, which has the **flush_trace** OMPT type, enables a tool to cause the OpenMP implementation to issue a sequence of zero or more **buffer_complete** callbacks to deliver all trace records that have been collected prior to the flush for the specified device. An invocation of **flush_trace** returns one if the entry point succeeds and zero otherwise.

Cross References

- OMPT **device** Type, see Section 40.11

44.11 stop_trace Entry Point

Name: stop_trace **Category:** function	**Properties:** C/C++-only, OMPT

Return Type and Arguments

Name	Type	Properties
<return type>	integer	*default*
device	device	OMPT, pointer

Type Signature

C / C++

```
typedef int (*ompt_stop_trace_t) (ompt_device_t *device);
```

C / C++

Semantics

The **stop_trace** entry point, which has the **stop_trace** OMPT type, provides a superset of the functionality of the **flush_trace** entry point. Specifically, the **stop_trace** entry point stops tracing for the specified device in addition to flushing pending trace records. An invocation of **stop_trace** returns one if the entry point succeeds and zero otherwise.

Cross References

- OMPT **device** Type, see Section 40.11
- **flush_trace** Entry Point, see Section 44.10

44.12 `advance_buffer_cursor` Entry Point

Name: `advance_buffer_cursor` **Category:** function	**Properties:** C/C++-only, OMPT

Return Type and Arguments

Name	Type	Properties
<return type>	integer	*default*
device	device	OMPT, pointer
buffer	buffer	OMPT, pointer
size	size_t	*default*
current	buffer_cursor	OMPT, opaque
next	buffer_cursor	OMPT, opaque, pointer

Type Signature

C / C++

```
typedef int (*ompt_advance_buffer_cursor_t) (
  ompt_device_t *device, ompt_buffer_t *buffer, size_t size,
  ompt_buffer_cursor_t current, ompt_buffer_cursor_t *next);
```

C / C++

Semantics

The **`advance_buffer_cursor`** entry point, which has the **`advance_buffer_cursor`** OMPT type, enables a tool to advance the trace buffer pointer for the buffer that the *buffer* argument indicates to the next trace record. The *size* argument indicates the size of *buffer* in bytes. The *current* argument is an OpenMP object that indicates the current position, while the *next* argument returns an OpenMP object with the next value. An invocation of **`advance_buffer_cursor`** returns *true* if the advance is successful and the next position in the buffer is valid. Otherwise, it returns *false*.

Cross References

- OMPT **`buffer`** Type, see Section 40.3
- OMPT **`buffer_cursor`** Type, see Section 40.4
- OMPT **`device`** Type, see Section 40.11

44.13 `get_record_type` Entry Point

Name: `get_record_type` **Category:** function	**Properties:** C/C++-only, OMPT

Return Type and Arguments

Name	Type	Properties
<return type>	record	*default*
buffer	buffer	OMPT, pointer
current	buffer_cursor	OMPT

Type Signature

C / C++

```
typedef ompt_record_t (*ompt_get_record_type_t) (
  ompt_buffer_t *buffer, ompt_buffer_cursor_t current);
```

C / C++

Semantics

Trace records for a device may be in one of two forms: native trace format, which may be device-specific, or standard trace format, in which each trace record corresponds to an OpenMP *event* and most fields in the trace record structure are the arguments that would be passed to the callback for the event. For the buffer specified by the *buffer* argument, the **`get_record_type`** entry point, which has the **`get_record_type`** OMPT type, enables a tool to inspect the type of a trace record at the position that the *current* argument specifies and to determine whether the trace record is an OMPT trace record, a native trace record, or is an invalid record, which is returned if the cursor is out of bounds.

Cross References

- OMPT **`buffer`** Type, see Section 40.3
- OMPT **`buffer_cursor`** Type, see Section 40.4
- OMPT **`record`** Type, see Section 40.23

44.14 `get_record_ompt` Entry Point

Name: `get_record_ompt` **Category:** function	**Properties:** C/C++-only, OMPT

Return Type and Arguments

Name	Type	Properties
<return type>	record_ompt	pointer
buffer	buffer	OMPT, pointer
current	buffer_cursor	OMPT, opaque

Type Signature

C / C++

```
typedef ompt_record_ompt_t  *(*ompt_get_record_ompt_t) (
  ompt_buffer_t *buffer, ompt_buffer_cursor_t current);
```

C / C++

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 40.2
- OMPT **buffer** Type, see Section 40.3
- OMPT **buffer_cursor** Type, see Section 40.4
- OMPT **device** Type, see Section 40.11
- OMPT **record_ompt** Type, see Section 40.26

44.15 get_record_native Entry Point

Name: get_record_native **Category:** function	**Properties:** C/C++-only, OMPT

Return Type and Arguments

Name	Type	Properties
<return type>	c_ptr	*default*
buffer	buffer	OMPT, pointer
current	buffer_cursor	OMPT, opaque
host_op_id	id	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++

Semantics

The **get_record_native** entry point, which has the **get_record_native** OMPT type, enables a tool to obtain a pointer to a native 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 a native event depends upon its type. The *current* argument is an OpenMP object that indicates the position from which to extract the trace record. If the entry point returns a non-null pointer result, it will also set the object to which the *host_op_id* argument points to a host-side identifier for the operation that is associated with the trace record on the target device and was created when the operation was initiated by the host device. 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 **buffer** Type, see Section 40.3
- OMPT **buffer_cursor** Type, see Section 40.4
- OMPT **id** Type, see Section 40.18

44.16 get_record_abstract Entry Point

Name: get_record_abstract **Category:** function	**Properties:** C/C++-only, OMPT

Return Type and Arguments

Name	Type	Properties
<return type>	record_abstract	pointer
native_record	void	pointer

Type Signature

C / C++

```
typedef ompt_record_abstract_t  *
  (*ompt_get_record_abstract_t) (void *native_record);
```

C / C++

Semantics

An OpenMP implementation may execute on a device that logs trace records in a native trace format that a tool cannot interpret directly. The **get_record_abstract** entry point, which has the **get_record_abstract** OMPT type, enables a tool to translate a native trace record to which the *native_record* argument points into a standard form.

Cross References

- OMPT **record_abstract** Type, see Section 40.24

Part VI

OMPD

45 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 45.3.2 and Section 45.3.3. 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 49. 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.

45.1 OMPD Interfaces Definitions

C / C++

A compliant implementation must supply a set of definitions for the OMPD 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 OMPD 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 OMPD third-party tool interface routines are external symbols with **C** linkage.

C / C++

45.2 Thread and Signal Safety

The OMPD library does not need to be reentrant. The third-party tool must ensure that only one native thread enters the OMPD library at a time. The OMPD library must not install signal handlers or otherwise interfere with the signal configuration of the third-party tool.

45.3 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.

45.3.1 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 OMPD if the *debug-var* ICV is set to **enabled**.

Cross References

- *debug-var* ICV, see Table 3.1

45.3.2 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.

45.3.3 ompd_dll_locations_valid Breakpoint

Format

C

```
void ompd_dll_locations_valid(void);
```

C

Semantics

Since **ompd_dll_locations** may not be a static variable, it may require runtime initialization. The OpenMP runtime notifies third-party tools that **ompd_dll_locations** is valid by having execution pass through a location that the symbol **ompd_dll_locations_valid** identifies. If **ompd_dll_locations** is NULL, a third-party tool can place a breakpoint at **ompd_dll_locations_valid** to be notified that **ompd_dll_locations** is initialized. In practice, the symbol **ompd_dll_locations_valid** may not be a function; instead, it may be a labeled machine instruction through which execution passes once the vector is valid.

46 OMPD Data Types

This chapter defines OMPD types, which support interactions with the OMPD library and provide information about the device architecture.

46.1 OMPD `addr` Type

Name: `addr` **Properties:** C/C++-only, OMPD	Base Type: `c_uint64_t`

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.

46.2 OMPD `address` Type

Name: `address` **Properties:** C/C++-only, OMPD	Base Type: structure

Fields

Name	Type	Properties
segment	seg	C/C++-only, OMPD
address	addr	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 46.1
- OMPD **seg** Type, see Section 46.10

46.3 OMPD address_space_context Type

Name: address_space_context Properties: C/C++-only, handle, OMPD	Base Type: aspace_cont

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.

46.4 OMPD callbacks Type

Name: callbacks Properties: C/C++-only, OMPD	Base Type: structure

Fields

Name	Type	Properties
alloc_memory	memory_alloc	C-only, OMPD
free_memory	memory_free	C-only, OMPD
print_string	print_string	C-only, OMPD
sizeof_type	sizeof	C-only, OMPD
symbol_addr_lookup	symbol_addr	C-only, OMPD
read_memory	memory_read	C-only, OMPD
write_memory	memory_write	C-only, OMPD
read_string	memory_read	C-only, OMPD
device_to_host	device_host	C-only, OMPD
host_to_device	device_host	C-only, OMPD
get_thread_context_for_thread_id	get_thread_context_for_thread_id	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 47.1.1
- **device_to_host** Callback, see Section 47.4.2
- **free_memory** Callback, see Section 47.1.2
- **get_thread_context_for_thread_id** Callback, see Section 47.3.1
- **host_to_device** Callback, see Section 47.4.3
- **ompd_initialize** Routine, see Section 48.1.1
- **print_string** Callback, see Section 47.5
- **read_memory** Callback, see Section 47.2.2.1
- **read_string** Callback, see Section 47.2.2.2
- **sizeof_type** Callback, see Section 47.3.2
- **symbol_addr_lookup** Callback, see Section 47.2.1
- **write_memory** Callback, see Section 47.2.3

46.5 OMPD device Type

Name: device **Properties:** C/C++-only, OMPD	Base Type: c_uint64_t

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/.

46.6 OMPD `device_type_sizes` Type

Name: `device_type_sizes` **Properties:** C/C++-only, OMPD	**Base Type:** structure

Fields

Name	Type	Properties
sizeof_char	c_uint8_t	C/C++-only, OMPD
sizeof_short	c_uint8_t	C/C++-only, OMPD
sizeof_int	c_uint8_t	C/C++-only, OMPD
sizeof_long	c_uint8_t	C/C++-only, OMPD
sizeof_long_long	c_uint8_t	C/C++-only, OMPD
sizeof_pointer	c_uint8_t	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 47.3.2

46.7 OMPD `frame_info` Type

Name: `frame_info` **Properties:** C/C++-only, OMPD	**Base Type:** structure

Fields

Name	Type	Properties
frame_address	address	C/C++-only, OMPD
frame_flag	word	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 46.2
- OMPT **frame_flag** Type, see Section 40.16
- OMPD **word** Type, see Section 46.17

46.8 OMPD icv_id Type

Name: icv_id **Properties:** C/C++-only, OMPD	**Base Type: c_uint64_t**

Predefined Identifiers

Name	Value	Properties
ompd_icv_undefined	0	C/C++-only, OMPD

Type Definition

C / C++

```
typedef uint64_t ompd_icv_id_t;
```

C / C++

Semantics

The **icv_id** OMPD type identifies ICVs.

46.9 OMPD `rc` Type

Name: `rc` **Properties:** C/C++-only, OMPD	**Base Type:** enumeration

Values

Name	Value	Properties
`ompd_rc_ok`	0	C-only, OMPD
`ompd_rc_unavailable`	1	C-only, OMPD
`ompd_rc_stale_handle`	2	C-only, OMPD
`ompd_rc_bad_input`	3	C-only, OMPD
`ompd_rc_error`	4	C-only, OMPD
`ompd_rc_unsupported`	5	C-only, OMPD
`ompd_rc_needs_state_tracking`	6	C-only, OMPD
`ompd_rc_incompatible`	7	C-only, OMPD
`ompd_rc_device_read_error`	8	C-only, OMPD
`ompd_rc_device_write_error`	9	C-only, OMPD
`ompd_rc_nomem`	10	C-only, OMPD
`ompd_rc_incomplete`	11	C-only, OMPD
`ompd_rc_callback_error`	12	C-only, OMPD
`ompd_rc_incompatible_handle`	13	C-only, OMPD
`ompd_rc_uninitialized`	14	C-only, OMPD

Type Definition

C / C++

```
typedef enum ompd_rc_t {
  ompd_rc_ok                    = 0,
  ompd_rc_unavailable           = 1,
  ompd_rc_stale_handle          = 2,
  ompd_rc_bad_input             = 3,
  ompd_rc_error                 = 4,
  ompd_rc_unsupported           = 5,
  ompd_rc_needs_state_tracking  = 6,
  ompd_rc_incompatible          = 7,
  ompd_rc_device_read_error     = 8,
  ompd_rc_device_write_error    = 9,
  ompd_rc_nomem                 = 10,
  ompd_rc_incomplete            = 11,
  ompd_rc_callback_error        = 12,
  ompd_rc_incompatible_handle   = 13,
  ompd_rc_uninitialized         = 14
} 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;
- **ompd_rc_incompatible_handle**: The specified handle was incompatible with the routine or callback; and
- **ompd_rc_uninitialized**: The interface is not initialized when the function is called.

46.10 OMPD seg Type

Name: seg **Properties:** C/C++-only, OMPD	**Base Type: c_uint64_t**

Predefined Identifiers

Name	Value	Properties
ompd_segment_none	0	C/C++-only, OMPD

Type Definition

C / C++

```
typedef uint64_t ompd_seg_t;
```

C / C++

Semantics

The **seg** OMPD type represents a segment value as an unsigned integer.

46.11 OMPD scope Type

Name: scope **Properties:** C/C++-only, OMPD	**Base Type:** enumeration

Values

Name	Value	Properties
ompd_scope_global	1	C-only, OMPD
ompd_scope_address_space	2	C-only, OMPD
ompd_scope_thread	3	C-only, OMPD
ompd_scope_parallel	4	C-only, OMPD
ompd_scope_implicit_task	5	C-only, OMPD
ompd_scope_task	6	C-only, OMPD
ompd_scope_teams	7	C-only, OMPD
ompd_scope_target	8	C-only, OMPD

Type Definition

C / C++

```
typedef enum ompd_scope_t {
  ompd_scope_global        = 1,
  ompd_scope_address_space = 2,
  ompd_scope_thread        = 3,
  ompd_scope_parallel      = 4,
  ompd_scope_implicit_task = 5,
  ompd_scope_task          = 6,
  ompd_scope_teams         = 7,
  ompd_scope_target        = 8
} ompd_scope_t;
```

C / C++

Semantics

The **scope** OMPD type is used for scope handles to identify OpenMP scopes, including those related to parallel regions and tasks. When used in an OMPD routine or OMPD callback procedure, the **scope** OMPD type and the OMPD handle must match according to Table 46.1.

TABLE 46.1: Mapping of Scope Type and OMPD Handles

Scope types	Handles
`ompd_scope_global`	Address space handle for the host device
`ompd_scope_address_space`	Any address space handle
`ompd_scope_thread`	Any native thread handle
`ompd_scope_parallel`	Any parallel handle
`ompd_scope_implicit_task`	Task handle for an implicit task
`ompd_scope_teams`	Parallel handle for an implicit parallel region generated from a `teams` construct
`ompd_scope_target`	Parallel handle for an implicit parallel region generated from a `target` construct
`ompd_scope_task`	Any task handle

46.12 OMPD `size` Type

Name: `size` **Properties:** C/C++-only, OMPD	**Base Type:** `c_uint64_t`

Type Definition

C / C++

```
typedef uint64_t ompd_size_t;
```

C / C++

The `size` OMPD type specifies the number of bytes in opaque data objects that are passed across the OMPD API.

46.13 OMPD `team_generator` Type

Name: `team_generator` **Properties:** C/C++-only, OMPD	**Base Type:** enumeration

Values

Name	Value	Properties
`ompd_generator_program`	`0`	C-only, OMPD
`ompd_generator_parallel`	`1`	C-only, OMPD
`ompd_generator_teams`	`2`	C-only, OMPD
`ompd_generator_target`	`3`	C-only, OMPD

Type Definition

C / C++

```
typedef enum ompd_team_generator_t {
  ompd_generator_program  = 0,
  ompd_generator_parallel = 1,
  ompd_generator_teams    = 2,
  ompd_generator_target   = 3
} ompd_team_generator_t;
```

C / C++

Semantics

The **team_generator** OMPD type represents the value of the *team-generator-var* ICV. The **ompd_generator_program** value indicates that the team is the initial team created at the start of the OpenMP program. The **ompd_generator_parallel**, **ompd_generator_teams**, and **ompd_generator_target** values indicate that the team was created by an encountered **parallel** construct, **teams** construct, or **target** construct, respectively.

46.14 OMPD thread_context Type

Name: thread_context **Properties:** C/C++-only, handle, OMPD	Base Type: thread_cont

Type Definition

C / C++

```
typedef struct _ompd_thread_cont ompd_thread_context_t;
```

C / C++

Semantics

A third-party tool uses the **thread_context** OMPD type, which represents handles that are opaque to the OMPD library and that uniquely identify a native thread of the OpenMP process that it is monitoring.

46.15 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/.

46.16 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 40.40

46.17 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.

46.18 OMPD Handle Types

The OMPD library defines handles, which have OMPD 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 OMPD library. The structures do not need to be defined; instead, the OMPD library must cast incoming (pointers to) handles to the appropriate internal, private types.

Each OMPD routine or OMPD 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 OMPD 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 49) or it releases the handle, the handle may no longer be referenced.

46.18.1 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.

46.18.2 OMPD `parallel_handle` Type

Name: `parallel_handle` **Properties:** C/C++-only, handle, OMPD	Base Type: `parallel_handle`

Type Definition

C / C++

```
typedef struct _ompd_parallel_handle ompd_parallel_handle_t;
```

C / C++

Semantics

The **`parallel_handle`** OMPD type is used for parallel handles that represent parallel regions.

46.18.3 OMPD `task_handle` Type

Name: `task_handle`	Base Type: `task_handle`
Properties: C/C++-only, handle, OMPD	

Type Definition

C / C++

```
typedef struct _ompd_task_handle ompd_task_handle_t;
```

C / C++

Semantics

The **`task_handle`** OMPD type is used for handles that represent tasks.

46.18.4 OMPD `thread_handle` Type

Name: `thread_handle`	Base Type: `thread_handle`
Properties: C/C++-only, handle, OMPD	

Type Definition

C / C++

```
typedef struct _ompd_thread_handle ompd_thread_handle_t;
```

C / C++

Semantics

The **`thread_handle`** OMPD type is used for handles that represent threads.

47 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`**.

47.1 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 47.1.1
- **`free_memory`** Callback, see Section 47.1.2

47.1.1 `alloc_memory` Callback

Name: `alloc_memory` **Category:** function	**Properties:** C-only, OMPD

Return Type and Arguments

Name	Type	Properties
<return type>	rc	*default*
nbytes	size	*default*
ptr	void	pointer-to-pointer

Type Signature

C

```
typedef ompd_rc_t (*ompd_callback_memory_alloc_fn_t) (
  ompd_size_t nbytes, void **ptr);
```

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 46.9
- OMPD **`size`** Type, see Section 46.12

47.1.2 `free_memory` Callback

Name: `free_memory` **Category:** function	**Properties:** C-only, OMPD

Return Type and Arguments

Name	Type	Properties
<return type>	rc	*default*
ptr	void	pointer

Type Signature

C

```
typedef ompd_rc_t (*ompd_callback_memory_free_fn_t) (void *ptr);
```

C

Semantics

A tool provides a **free_memory** callback, which has the **memory_free** OMPD type, that the OMPD library may call to deallocate memory that was obtained from a prior call to the **alloc_memory** callback. The *ptr* argument is the address of the block to be deallocated.

Cross References

- **alloc_memory** Callback, see Section 47.1.1
- OMPD **rc** Type, see Section 46.9

47.2 Accessing Program or Runtime Memory

The OMPD library cannot directly read from or write to memory of the OpenMP program. Instead the OMPD library must use callbacks into the third-party tool that perform the operation.

47.2.1 symbol_addr_lookup Callback

Name: symbol_addr_lookup **Category:** function	**Properties:** C-only, OMPD

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

C

```
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);
```

C

Semantics

A tool provides a **symbol_addr_lookup** callback, which has the **symbol_addr** OMPD type, that the OMPD 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 OMPD 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 OMPD 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 46.2
- OMPD **address_space_context** Type, see Section 46.3
- OMPD **rc** Type, see Section 46.9
- OMPD **thread_context** Type, see Section 46.14

47.2.2 OMPD memory_read Type

Name: memory_read **Category:** function pointer	**Properties:** C-only, OMPD

Return Type and Arguments

Name	Type	Properties
<return type>	rc	*default*
address_space_context	address_space_context	pointer
thread_context	thread_context	pointer
addr	address	intent(in), pointer
nbytes	size	*default*
buffer	void	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 OMPD 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 46.2
- OMPD **address_space_context** Type, see Section 46.3
- **device_to_host** Callback, see Section 47.4.2
- OMPD **rc** Type, see Section 46.9
- OMPD **size** Type, see Section 46.12
- OMPD **thread_context** Type, see Section 46.14

47.2.2.1 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 46.2
- OMPD **address_space_context** Type, see Section 46.3
- OMPD **memory_read** Type, see Section 47.2.2

47.2.2.2 `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 46.9
- OMPD **`size`** Type, see Section 46.12

47.2.3 `write_memory` Callback

Name: `write_memory` **Category:** function	**Properties:** C-only, OMPD

Return Type and Arguments

Name	Type	Properties
<return type>	rc	*default*
address_space_context	address_space_context	pointer
thread_context	thread_context	pointer
addr	address	intent(in), pointer
nbytes	size	*default*
buffer	void	pointer

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

Semantics

A tool provides a **write_memory** callback, which has the **memory_write** OMPD type, that the OMPD library may call to have the tool write a block of data to a location within an address space from a provided buffer. The address to which the data are to be written in the OpenMP program that *address_space_context* specifies is given by *addr*. The *nbytes* argument is the number of bytes to be transferred. The *thread_context* argument for global memory accesses should be NULL. If it is a non-null value, then *thread_context* identifies the native thread context for the memory access for the purpose of accessing thread local storage.

The data to be written are passed through *buffer*, which is allocated and owned by the OMPD library. The contents of the buffer are unstructured, raw bytes. The OMPD library must use the **host_to_device** callback to perform any transformations such as byte-swapping that may be necessary to render the data into a form that is compatible with the OpenMP runtime.

In addition to the general return codes for OMPD callbacks, **write_memory** callbacks may also return the following return codes:

- **ompd_rc_error** if unallocated memory is reached while writing *nbytes*.

Cross References

- OMPD **address** Type, see Section 46.2
- OMPD **address_space_context** Type, see Section 46.3
- **host_to_device** Callback, see Section 47.4.3
- OMPD **rc** Type, see Section 46.9
- OMPD **size** Type, see Section 46.12
- OMPD **thread_context** Type, see Section 46.14

47.3 Context Management and Navigation

Summary

A tool provides callbacks to manage and to navigate tool context relationships.

47.3.1 get_thread_context_for_thread_id Callback

Name: **get_thread_context_for_thread_id** **Category:** function	**Properties:** C-only, OMPD

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

C

```
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);
```

C

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

- OMPD **address_space_context** Type, see Section 46.3
- OMPD **rc** Type, see Section 46.9
- OMPD **size** Type, see Section 46.12
- OMPD **thread_context** Type, see Section 46.14
- OMPD **thread_id** Type, see Section 46.15

47.3.2 sizeof_type Callback

Name: sizeof_type **Category:** function	**Properties:** C-only, OMPD

Return Type and Arguments

Name	Type	Properties
<return type>	rc	*default*
address_space_context	address_space_context	pointer
sizes	device_type_sizes	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** OMPD type, that the OMPD 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

- OMPD **address_space_context** Type, see Section 46.3
- OMPD **device_type_sizes** Type, see Section 46.6
- OMPD **rc** Type, see Section 46.9

47.4 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.

47.4.1 OMPD `device_host` Type

Name: `device_host` **Category:** function pointer	**Properties:** C-only, OMPD

Return Type and Arguments

Name	Type	Properties
<return type>	rc	*default*
address_space_context	address_space_context	pointer
input	void	intent(in), pointer
unit_size	size	*default*
count	size	*default*
output	void	pointer

Type Signature

C

```
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);
```

C

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.

Cross References

- OMPD **address_space_context** Type, see Section 46.3
- **device_to_host** Callback, see Section 47.4.2
- **host_to_device** Callback, see Section 47.4.3
- OMPD **rc** Type, see Section 46.9
- OMPD **size** Type, see Section 46.12

47.4.2 device_to_host Callback

Name: device_to_host **Category:** function	**Properties:** C-only, common-type-callback, device-translating, OMPD

Type Signature

```
device_host
```

Semantics

The **device_to_host** is the device-translating callback that translates data that is read from the OpenMP program.

Cross References

- OMPD **device_host** Type, see Section 47.4.1

47.4.3 host_to_device Callback

Name: host_to_device **Category:** function	**Properties:** C-only, common-type-callback, device-translating, OMPD

Type Signature

```
device_host
```

Semantics

The **host_to_device** is the device-translating callback that translates data that is to be written to the OpenMP program.

Cross References

- OMPD **device_host** Type, see Section 47.4.1

47.5 print_string Callback

Name: print_string **Category:** function	**Properties:** C-only, OMPD

Return Type and Arguments

Name	Type	Properties
<return type>	rc	*default*
string	char	intent(in), pointer
category	integer	*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 46.9

48 OMPD Routines

This chapter defines the OMPD routines, which are routines that have the OMPD property and, thus, are provided by the OMPD library to be used by third-party tools. Some OMPD 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.

48.1 OMPD Library Initialization and Finalization

The OMPD library must be initialized exactly once after it is loaded, and finalized exactly once before it is unloaded. Per OpenMP process or core file initialization and finalization are also required. Once loaded, the tool can determine the version of the OMPD API that the library supports by calling **`ompd_get_api_version`**. 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 OMPD API that the library supports. If the tool does not support the version that **`ompd_get_api_version`** returns, it may attempt to call **`ompd_initialize`** with a different version.

Cross References

- **`ompd_get_api_version`** Routine, see Section 48.1.2
- **`ompd_initialize`** Routine, see Section 48.1.1

48.1.1 `ompd_initialize` Routine

Name: `ompd_initialize` **Category:** function	**Properties:** C-only, OMPD

Return Type and Arguments

Name	Type	Properties
<return type>	rc	*default*
api_version	word	*default*
callbacks	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 46.4
- **ompd_get_api_version** Routine, see Section 48.1.2
- OMPD **rc** Type, see Section 46.9
- OMPD **word** Type, see Section 46.17

48.1.2 ompd_get_api_version Routine

Name: ompd_get_api_version **Category:** function	**Properties:** C-only, OMPD

Return Type and Arguments

Name	Type	Properties
<return type>	rc	*default*
api_version	word	pointer

Prototypes

C

```
ompd_rc_t ompd_get_api_version(ompd_word_t *api_version);
```

C

Semantics

The tool may call the **ompd_get_api_version** routine to obtain the latest OMPD API version number of the OMPD library. The OMPD API version number is equal to the value of the **_OPENMP** macro defined in the associated OpenMP implementation, if the C preprocessor

is supported. If the associated OpenMP implementation compiles Fortran codes without the use of a C preprocessor, the OMPD API version number is equal to the value of the **openmp_version** predefined identifier. The latest version number is returned into the location to which the *version* argument points.

Cross References

- **ompd_initialize** Routine, see Section 48.1.1
- OMPD **rc** Type, see Section 46.9
- OMPD **word** Type, see Section 46.17

48.1.3 ompd_get_version_string Routine

Name: ompd_get_version_string **Category:** function	**Properties:** C-only, OMPD

Return Type and Arguments

Name	Type	Properties
<return type>	rc	*default*
string	const char	intent(out), pointer-to-pointer

Prototypes

C

```
ompd_rc_t ompd_get_version_string(const char **string);
```

C

Semantics

The **ompd_get_version_string** routine returns a pointer to a version string that unambiguously identifies the OMPD library. This version string at least identifies the OpenMP vendor using the key **vendor**, the API version as **ompd-api-version=***YYYYMM*, an OMPD library release version as **ompd-library-version=***major[.minor[.patch[suffix]]]*, the OpenMP library release version as **omp-runtime-version=***major[.minor[.patch[suffix]]]*, the date of the latest source code change to the OpenMP runtime as **omp-runtime-source-date=***YYYYMMDD* and the date of the latest source code change to the OMPD library as **ompd-library-source-date=***YYYYMMDD*. The version string may contain additional implementation defined key-value pairs.

A pointer to a descriptive version string is placed into the location to which the *string* output argument points. The OMPD library owns the string that the OMPD library returns; the tool must not modify or release this string. The string remains valid for as long as the library is loaded. The **ompd_get_version_string** routine may be called before **ompd_initialize**. Accordingly, the OMPD library must not use heap or stack memory for the string.

The signatures of **ompd_get_api_version** and **ompd_get_version_string** are guaranteed not to change in future versions of OMPD. In contrast, the type definitions and prototypes in the rest of OMPD do not carry the same guarantee. Therefore a tool that uses OMPD should check the version of the loaded OMPD library before it calls any other OMPD routine.

Cross References

- OMPD **address_space_handle** Type, see Section 46.18.1
- **ompd_get_api_version** Routine, see Section 48.1.2
- OMPD **rc** Type, see Section 46.9

48.1.4 ompd_finalize Routine

Name: ompd_finalize **Category:** function	**Properties:** C-only, OMPD

Return Type

Name	Type	Properties
<return type>	rc	*default*

Prototypes

C

```
ompd_rc_t ompd_finalize(void);
```

C

Semantics

When the tool is finished with the OMPD library, it should call **ompd_finalize** before it unloads the library. The call to the **ompd_finalize** routine must be the last OMPD call that the tool makes before it unloads the library. This routine allows the OMPD library to free any resources that it may be holding. The OMPD library may implement a *finalizer* section, which executes as the library is unloaded and therefore after the call to **ompd_finalize**. During finalization, the OMPD library may use the callbacks that the tool provided earlier during the call to **ompd_initialize**. In addition to the return codes permitted for all OMPD routines, this routine returns **ompd_rc_uninitialized** if the OMPD library is not initialized.

Cross References

- OMPD **rc** Type, see Section 46.9

48.2 Process Initialization and Finalization

48.2.1 ompd_process_initialize Routine

Name: ompd_process_initialize **Category:** function	**Properties:** C-only, OMPD

Return Type and Arguments

Name	Type	Properties
<return type>	rc	*default*
context	address_space_context	opaque, pointer
host_handle	address_space_handle	opaque, pointer-to-pointer

Prototypes

C

```
ompd_rc_t ompd_process_initialize(
  ompd_address_space_context_t *context,
  ompd_address_space_handle_t **host_handle);
```

C

Semantics

A tool calls **ompd_process_initialize** to obtain an address space handle for the host device when it initializes a session on an OpenMP process or core file. On return from **ompd_process_initialize**, the tool owns the address space handle, which it must release with **ompd_rel_address_space_handle**. The initialization function must be called before any OMPD operations are performed on the OpenMP process or core file. This routine allows the OMPD library to confirm that it can handle the OpenMP process or core file that *context* identifies.

The *context* argument is an opaque handle that the tool provides to address an address space from the host device. On return, the *host_handle* argument provides an opaque handle to the tool for this address space, which the tool must release when it is no longer needed.

In addition to the return codes permitted for all OMPD routines, this routine returns **ompd_rc_incompatible** if the OMPD library is incompatible with the runtime library loaded in the process.

Cross References

- OMPD **address_space_context** Type, see Section 46.3
- OMPD **address_space_handle** Type, see Section 46.18.1
- **ompd_rel_address_space_handle** Routine, see Section 48.8.1
- OMPD **rc** Type, see Section 46.9

48.2.2 ompd_device_initialize Routine

Name: ompd_device_initialize **Category:** function	**Properties:** C-only, OMPD

Return Type and Arguments

Name	Type	Properties
<return type>	rc	*default*
host_handle	address_space_handle	opaque, pointer
device_context	address_space_context	opaque, pointer
kind	device	*default*
sizeof_id	size	pointer
id	void	pointer
device_handle	address_space_handle	opaque, pointer-to-pointer

Prototypes

C

```
ompd_rc_t ompd_device_initialize(
  ompd_address_space_handle_t *host_handle,
  ompd_address_space_context_t *device_context,
  ompd_device_t kind, ompd_size_t *sizeof_id, void *id,
  ompd_address_space_handle_t **device_handle);
```

C

Semantics

A tool calls **ompd_device_initialize** to obtain an address space handle for a non-host device that has at least one active **target** region. On return from **ompd_device_initialize**, the tool owns the address space handle. 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.

Cross References

- OMPD **address_space_context** Type, see Section 46.3
- OMPD **address_space_handle** Type, see Section 46.18.1
- OMPD **device** Type, see Section 46.5
- OMPD **rc** Type, see Section 46.9
- OMPD **size** Type, see Section 46.12

48.2.3 ompd_get_device_thread_id_kinds Routine

Name: **ompd_get_device_thread_id_kinds** **Category:** function	**Properties:** C-only, OMPD

Return Type and Arguments

Name	Type	Properties
<return type>	rc	*default*
device_handle	address_space_handle	opaque, pointer
kinds	thread_id	pointer-to-pointer
thread_id_sizes	size	pointer-to-pointer
count	integer	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 46.18.1
- **ompd_device_initialize** Routine, see Section 48.2.2
- **ompd_process_initialize** Routine, see Section 48.2.1
- OMPD **rc** Type, see Section 46.9
- OMPD **size** Type, see Section 46.12
- OMPD **thread_id** Type, see Section 46.15

48.3 Address Space Information

48.3.1 ompd_get_omp_version Routine

Name: ompd_get_omp_version **Category:** function	**Properties:** C-only, OMPD

Return Type and Arguments

Name	Type	Properties
<return type>	rc	*default*
address_space	address_space_handle	opaque, pointer
omp_version	word	pointer

Prototypes

C

```
ompd_rc_t ompd_get_omp_version(
  ompd_address_space_handle_t *address_space,
  ompd_word_t *omp_version);
```

C

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 46.18.1
- OMPD **rc** Type, see Section 46.9
- OMPD **word** Type, see Section 46.17

48.3.2 ompd_get_omp_version_string Routine

Name: **ompd_get_omp_version_string** **Category:** function	**Properties:** C-only, OMPD

Return Type and Arguments

Name	Type	Properties
<return type>	rc	*default*
address_space	address_space_handle	opaque, pointer
string	const char	intent(out), pointer-to-pointer

Prototypes

C

```
ompd_rc_t ompd_get_omp_version_string(
  ompd_address_space_handle_t *address_space, const char **string);
```

C

Semantics

The **ompd_get_omp_version_string** routine returns a descriptive string for the OpenMP API version that is associated with an address space. The *address_space* argument is an opaque handle that the tool provides to reference the address space of a process or device. A pointer to a descriptive version string is placed into the location to which the *string* output argument points. After returning from the routine, the tool owns the string. The OMPD library must use the memory allocation callback that the tool provides to allocate the string storage. The tool is responsible for releasing the memory.

Cross References

- OMPD Handle Types, see Section 46.18
- OMPD **rc** Type, see Section 46.9

48.4 Thread Handle Routines

48.4.1 ompd_get_thread_in_parallel Routine

Name: ompd_get_thread_in_parallel **Category:** function	**Properties:** C-only, OMPD

Return Type and Arguments

Name	Type	Properties
<return type>	rc	*default*
parallel_handle	parallel_handle	opaque, pointer
thread_num	integer	*default*
thread_handle	thread_handle	opaque, pointer-to-pointer

Prototypes

C

```
ompd_rc_t ompd_get_thread_in_parallel(
  ompd_parallel_handle_t *parallel_handle, int thread_num,
  ompd_thread_handle_t **thread_handle);
```

C

Semantics

The **ompd_get_thread_in_parallel** routine enables a tool to obtain handles for OpenMP threads that are associated with a parallel region. A successful invocation of **ompd_get_thread_in_parallel** returns a pointer to a native thread handle in the location to which *thread_handle* points. This routine yields meaningful results only if all OpenMP threads in the team that is executing the parallel region are stopped.

The *parallel_handle* argument is an opaque handle for a parallel region and selects the parallel region on which to operate. The *thread_num* argument represents the thread number and selects the thread, the handle for which is to be returned. On return, the *thread_handle* argument is a handle for the selected thread.

This routine returns **ompd_rc_bad_input** if the *thread_num* argument is greater than or equal to the *team-size-var* ICV or negative, in which case the value returned in *thread_handle* is invalid.

Cross References

- **ompd_get_icv_from_scope** Routine, see Section 48.11.2
- OMPD **parallel_handle** Type, see Section 46.18.2
- OMPD **rc** Type, see Section 46.9
- OMPD **thread_handle** Type, see Section 46.18.4

48.4.2 ompd_get_thread_handle Routine

Name: ompd_get_thread_handle **Category:** function	**Properties:** C-only, OMPD

Return Type and Arguments

Name	Type	Properties
<return type>	rc	*default*
handle	address_space_handle	pointer
kind	thread_id	*default*
sizeof_thread_id	size	*default*
thread_id	void	intent(in), pointer
thread_handle	thread_handle	pointer-to-pointer

Prototypes

C

```
ompd_rc_t ompd_get_thread_handle(
  ompd_address_space_handle_t *handle, ompd_thread_id_t kind,
  ompd_size_t sizeof_thread_id, const void *thread_id,
  ompd_thread_handle_t **thread_handle);
```

C

Semantics

The **ompd_get_thread_handle** routine maps a native thread to a native thread handle. Further, the routine determines if the native thread identifier to which *thread_id* points represents an OpenMP thread. If so, the routine returns **ompd_rc_ok** and the location to which *thread_handle* points is set to the native thread handle for the native thread to which the OpenMP thread is mapped.

The *handle* argument is a handle that the tool provides to reference an address space. The *kind*, *sizeof_thread_id*, and *thread_id* arguments represent a native thread identifier. On return, the *thread_handle* argument provides a handle to the native thread within the provided address space.

The native thread identifier to which *thread_id* points must be valid for the duration of the call to the routine. If the OMPD library must retain the native thread identifier, it must copy it.

This routine returns **ompd_rc_bad_input** if a different value in *sizeof_thread_id* is expected for a thread kind of *kind*. In addition to the return codes permitted for all OMPD routines, this routine returns **ompd_rc_unsupported** if the *kind* of thread is not supported and it returns **ompd_rc_unavailable** if the native thread is not an OpenMP thread.

Cross References

- OMPD **address_space_handle** Type, see Section 46.18.1
- OMPD **rc** Type, see Section 46.9
- OMPD **size** Type, see Section 46.12
- OMPD **thread_handle** Type, see Section 46.18.4
- OMPD **thread_id** Type, see Section 46.15

48.4.3 ompd_get_thread_id Routine

Name: ompd_get_thread_id **Category:** function	**Properties:** C-only, OMPD

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

Prototypes

C

```
ompd_rc_t ompd_get_thread_id(ompd_thread_handle_t *thread_handle,
  ompd_thread_id_t kind, ompd_size_t sizeof_thread_id,
  void *thread_id);
```

C

Semantics

The **ompd_get_thread_id** routine maps a native thread handle to a native thread identifier. This routine yields meaningful results only if the referenced OpenMP thread is stopped. The *thread_handle* argument is a native thread handle. The *kind* argument represents the native thread identifier. The *sizeof_thread_id* argument represents the size of the native thread identifier. On return, the *thread_id* argument is a buffer that represents a native thread identifier.

This routine returns **ompd_rc_bad_input** if a different value in *sizeof_thread_id* is expected for a native thread kind of *kind*. In addition to the return codes permitted for all OMPD routines, this routine returns **ompd_rc_unsupported** if the *kind* of native thread is not supported.

Cross References

- OMPD **rc** Type, see Section 46.9
- OMPD **size** Type, see Section 46.12
- OMPD **thread_handle** Type, see Section 46.18.4
- OMPD **thread_id** Type, see Section 46.15

48.4.4 ompd_get_device_from_thread Routine

Name: ompd_get_device_from_thread **Category:** function	**Properties:** C-only, OMPD

Return Type and Arguments

Name	Type	Properties
<return type>	rc	*default*
thread_handle	thread_handle	pointer
device	address_space_handle	pointer-to-pointer

Prototypes

C

```
ompd_rc_t ompd_get_device_from_thread(
  ompd_thread_handle_t *thread_handle,
  ompd_address_space_handle_t **device);
```

C

Semantics

The **ompd_get_device_from_thread** routine obtains a pointer to the address space handle for a device on which an OpenMP thread is executing. The returned pointer will be the same as the address space handle pointer that was previously returned by a call to **ompd_process_initialize** (for a host device) or a call to **ompd_device_initialize** (for a non-host device). This routine yields meaningful results only if the referenced OpenMP thread is stopped.

The *thread_handle* argument is a pointer to a native thread handle that represents a native thread to which an OpenMP thread is mapped. On return, the *device* argument is the address of a pointer to an address space handle.

Cross References

- OMPD **address_space_handle** Type, see Section 46.18.1
- OMPD **rc** Type, see Section 46.9
- OMPD **thread_handle** Type, see Section 46.18.4

48.5 Parallel Region Handle Routines

48.5.1 `ompd_get_curr_parallel_handle` Routine

Name: `ompd_get_curr_parallel_handle` **Category:** function	**Properties:** C-only, OMPD

Return Type and Arguments

Name	Type	Properties
<return type>	rc	*default*
thread_handle	thread_handle	opaque, pointer
parallel_handle	parallel_handle	opaque, pointer-to-pointer

Prototypes

C

```
ompd_rc_t ompd_get_curr_parallel_handle(
  ompd_thread_handle_t *thread_handle,
  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 48.8.2
- OMPD **`parallel_handle`** Type, see Section 46.18.2
- OMPD **`rc`** Type, see Section 46.9
- OMPD **`thread_handle`** Type, see Section 46.18.4

48.5.2 `ompd_get_enclosing_parallel_handle` Routine

Name: `ompd_get_enclosing_parallel_handle` **Category:** function	**Properties:** C-only, OMPD

Return Type and Arguments

Name	Type	Properties
<return type>	rc	*default*
parallel_handle	parallel_handle	opaque, pointer
enclosing_parallel_handle	parallel_handle	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

Semantics

The **ompd_get_enclosing_parallel_handle** routine enables a tool to obtain a pointer to the parallel handle for the parallel region that encloses the parallel region that *parallel_handle* specifies. This routine yields meaningful results only if at least one thread in the team that is executing the parallel region is stopped. A pointer to the parallel handle for the enclosing region is returned in the location to which *enclosing_parallel_handle* points. After a call to this routine, the tool owns the handle; the tool must release the handle with **ompd_rel_parallel_handle** when it is no longer required. The *parallel_handle* argument is an opaque handle for a parallel region that selects the parallel region on which to operate.

In addition to the return codes permitted for all OMPD routines, this routine returns **ompd_rc_unavailable** if no enclosing parallel region exists.

Cross References

- **ompd_rel_parallel_handle** Routine, see Section 48.8.2
- OMPD **parallel_handle** Type, see Section 46.18.2
- OMPD **rc** Type, see Section 46.9

48.5.3 ompd_get_task_parallel_handle Routine

Name: ompd_get_task_parallel_handle **Category:** function	**Properties:** C-only, OMPD

Return Type and Arguments

Name	Type	Properties
<return type>	rc	*default*
task_handle	task_handle	pointer
task_parallel_handle	parallel_handle	pointer-to-pointer

Prototypes

C

```
ompd_rc_t ompd_get_task_parallel_handle(
  ompd_task_handle_t *task_handle,
  ompd_parallel_handle_t **task_parallel_handle);
```

C

Semantics

The **ompd_get_task_parallel_handle** routine enables a tool to obtain a pointer to the parallel handle for the parallel region that encloses the task region that *task_handle* specifies. This routine yields meaningful results only if at least one thread in the team that is executing the parallel region is stopped. A pointer to the parallel handle is returned in the location to which *task_parallel_handle* points. The tool owns that parallel handle, which it must release with **ompd_rel_parallel_handle**.

Cross References

- **ompd_rel_parallel_handle** Routine, see Section 48.8.2
- OMPD **parallel_handle** Type, see Section 46.18.2
- OMPD **rc** Type, see Section 46.9
- OMPD **task_handle** Type, see Section 46.18.3

48.6 Task Handle Routines

48.6.1 ompd_get_curr_task_handle Routine

Name: ompd_get_curr_task_handle **Category:** function	**Properties:** C-only, OMPD

Return Type and Arguments

Name	Type	Properties
<return type>	rc	*default*
thread_handle	thread_handle	opaque, pointer
task_handle	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 48.8.3
- OMPD **rc** Type, see Section 46.9
- OMPD **task_handle** Type, see Section 46.18.3
- OMPD **thread_handle** Type, see Section 46.18.4

48.6.2 ompd_get_generating_task_handle Routine

Name: **ompd_get_generating_task_handle** **Category:** function	**Properties:** C-only, OMPD

Return Type and Arguments

Name	Type	Properties
<return type>	rc	*default*
task_handle	task_handle	pointer
generating_task_handle	task_handle	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 48.8.3
- OMPD **rc** Type, see Section 46.9
- OMPD **task_handle** Type, see Section 46.18.3

48.6.3 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

```
ompd_rc_t ompd_get_scheduling_task_handle(
  ompd_task_handle_t *task_handle,
  ompd_task_handle_t **scheduling_task_handle);
```

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 selected. 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 48.8.3
- OMPD **rc** Type, see Section 46.9
- OMPD **task_handle** Type, see Section 46.18.3

48.6.4 ompd_get_task_in_parallel Routine

Name: **ompd_get_task_in_parallel** **Category:** function	**Properties:** C-only, OMPD

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

Prototypes

C

```
ompd_rc_t ompd_get_task_in_parallel(
  ompd_parallel_handle_t *parallel_handle, int thread_num,
  ompd_task_handle_t **task_handle);
```

C

Semantics

The **ompd_get_task_in_parallel** routine obtains handles for the implicit tasks that are associated with a parallel region. A successful invocation of **ompd_get_task_in_parallel** returns a pointer to a task handle in the location to which *task_handle* points. This routine yields meaningful results only if all OpenMP threads in the parallel region are stopped. The *parallel_handle* argument is an opaque handle that selects the parallel region on which to operate. The *thread_num* argument selects the implicit task of the team to be returned. The *thread_num* argument is equal to the *thread-num-var* ICV value of the selected implicit task. This routine returns **ompd_rc_bad_input** if the *thread_num* argument is greater than or equal to the *team-size-var* ICV or negative.

Cross References

- **ompd_get_icv_from_scope** Routine, see Section 48.11.2
- OMPD **parallel_handle** Type, see Section 46.18.2
- OMPD **rc** Type, see Section 46.9
- OMPD **task_handle** Type, see Section 46.18.3

48.6.5 ompd_get_task_function Routine

Name: ompd_get_task_function **Category:** function	**Properties:** C-only, OMPD

Return Type and Arguments

Name	Type	Properties
<return type>	rc	*default*
task_handle	task_handle	opaque, pointer
entry_point	address	pointer

Prototypes

C

```
ompd_rc_t ompd_get_task_function(ompd_task_handle_t *task_handle,
  ompd_address_t *entry_point);
```

C

Semantics

The **ompd_get_task_function** routine returns the entry point of the code that corresponds to the body of code that the task executes. This routine returns meaningful results only if the thread that is executing the task that *task_handle* specifies is stopped while executing the task. That argument is an opaque handle that selects the task on which to operate. On return, the *entry_point* argument is set to an address that describes the beginning of application code that executes the task region.

Cross References

- OMPD **address** Type, see Section 46.2
- OMPD **rc** Type, see Section 46.9
- OMPD **task_handle** Type, see Section 46.18.3

48.6.6 ompd_get_task_frame Routine

Name: ompd_get_task_frame **Category:** function	**Properties:** C-only, OMPD

Return Type and Arguments

Name	Type	Properties
<return type>	rc	*default*
task_handle	task_handle	pointer
exit_frame	frame_info	pointer
enter_frame	frame_info	pointer

Prototypes

C

```
ompd_rc_t ompd_get_task_frame(ompd_task_handle_t *task_handle,
  ompd_frame_info_t *exit_frame, ompd_frame_info_t *enter_frame);
```

C

Semantics

The **ompd_get_task_frame** routine extracts the frame pointers of a task. An OpenMP implementation maintains an object of **frame** OMPT type for every implicit task and explicit task. The **ompd_get_task_frame** routine extracts the **enter_frame** and **exit_frame** fields of the **frame** object of the task that *task_handle* identifies. 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 *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 46.2
- OMPT **frame** Type, see Section 40.15
- OMPD **frame_info** Type, see Section 46.7
- OMPD **rc** Type, see Section 46.9
- OMPD **task_handle** Type, see Section 46.18.3

48.7 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.

48.7.1 ompd_parallel_handle_compare Routine

Name: **ompd_parallel_handle_compare** **Category:** function	**Properties:** C-only, handle-comparing, OMPD

Return Type and Arguments

Name	Type	Properties
<return type>	rc	*default*
parallel_handle_1	parallel_handle	opaque, pointer
parallel_handle_2	parallel_handle	opaque, pointer
cmp_value	integer	pointer

Prototypes

C

```
ompd_rc_t ompd_parallel_handle_compare(
  ompd_parallel_handle_t *parallel_handle_1,
  ompd_parallel_handle_t *parallel_handle_2, int *cmp_value);
```

C

Semantics

The **ompd_parallel_handle_compare** routine compares two parallel handles. The *parallel_handle_1* and *parallel_handle_2* arguments are parallel handles that correspond to parallel regions.

Cross References

- OMPD **parallel_handle** Type, see Section 46.18.2
- OMPD **rc** Type, see Section 46.9

48.7.2 ompd_task_handle_compare Routine

Name: ompd_task_handle_compare **Category:** function	**Properties:** C-only, handle-comparing, OMPD

Return Type and Arguments

Name	Type	Properties
<return type>	rc	*default*
task_handle_1	task_handle	opaque, pointer
task_handle_2	task_handle	opaque, pointer
cmp_value	integer	pointer

Prototypes

C

```
ompd_rc_t ompd_task_handle_compare(
  ompd_task_handle_t *task_handle_1,
  ompd_task_handle_t *task_handle_2, int *cmp_value);
```

C

Semantics

The **ompd_task_handle_compare** routine compares two task handles. The *task_handle_1* and *task_handle_2* arguments are task handles that correspond to tasks.

Cross References

- OMPD **rc** Type, see Section 46.9
- OMPD **task_handle** Type, see Section 46.18.3

48.7.3 ompd_thread_handle_compare Routine

Name: ompd_thread_handle_compare **Category:** function	**Properties:** C-only, handle-comparing, OMPD

Return Type and Arguments

Name	Type	Properties
<return type>	rc	*default*
thread_handle_1	thread_handle	opaque, pointer
thread_handle_2	thread_handle	opaque, pointer
cmp_value	integer	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 46.9
- OMPD **thread_handle** Type, see Section 46.18.4

48.8 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.

48.8.1 ompd_rel_address_space_handle Routine

Name: ompd_rel_address_space_handle **Category:** function	**Properties:** C-only, handle-releasing, OMPD

Return Type and Arguments

Name	Type	Properties
<return type>	rc	*default*
handle	address_space_handle	opaque, pointer

Prototypes

C

```
ompd_rc_t ompd_rel_address_space_handle(
  ompd_address_space_handle_t *handle);
```

C

Semantics

A tool calls **ompd_rel_address_space_handle** to release an address space handle.

Cross References

- OMPD **address_space_handle** Type, see Section 46.18.1
- OMPD **rc** Type, see Section 46.9

48.8.2 ompd_rel_parallel_handle Routine

Name: ompd_rel_parallel_handle **Category:** function	**Properties:** C-only, handle-releasing, OMPD

Return Type and Arguments

Name	Type	Properties
<return type>	rc	*default*
parallel_handle	parallel_handle	opaque, pointer

Prototypes

C

```
ompd_rc_t ompd_rel_parallel_handle(
  ompd_parallel_handle_t *parallel_handle);
```

C

Semantics

A tool calls **ompd_rel_parallel_handle** to release a parallel handle.

Cross References

- OMPD **parallel_handle** Type, see Section 46.18.2
- OMPD **rc** Type, see Section 46.9

48.8.3 ompd_rel_task_handle Routine

Name: ompd_rel_task_handle **Category:** function	**Properties:** C-only, handle-releasing, OMPD

Return Type and Arguments

Name	Type	Properties
<return type>	rc	*default*
task_handle	task_handle	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 46.9
- OMPD **task_handle** Type, see Section 46.18.3

48.8.4 ompd_rel_thread_handle Routine

Name: ompd_rel_thread_handle **Category:** function	**Properties:** C-only, handle-releasing, OMPD

Return Type and Arguments

Name	Type	Properties
<return type>	rc	*default*
thread_handle	thread_handle	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 46.9
- OMPD **thread_handle** Type, see Section 46.18.4

48.9 Querying Thread States

48.9.1 `ompd_enumerate_states` Routine

Name: `ompd_enumerate_states` **Category:** function	**Properties:** C-only, OMPD

Return Type and Arguments

Name	Type	Properties
<return type>	rc	*default*
address_space_handle	address_space_handle	opaque, pointer
current_state	word	*default*
next_state	word	pointer
next_state_name	const char	intent(out), pointer-to-pointer
more_enums	word	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. On return, the location to which the *more_enums* argument points has the value 1 whenever one or more states are left in the enumeration. On return, the location to which the *more_enums* argument points has the value 0 when *current_state* is the last state in the enumeration.

The *address_space_handle* argument identifies the address space. The *current_state* argument must be a thread state that the OpenMP implementation supports. To begin enumerating the supported states, a tool should pass **`ompt_state_undefined`** as the value of *current_state*. Subsequent calls to **`ompd_enumerate_states`** by the tool should pass the value that the routine returned in the *next_state* argument. This routine returns **`ompd_rc_bad_input`** if an unknown value is provided in *current_state*.

Cross References

- OMPD **address_space_handle** Type, see Section 46.18.1
- OMPD **rc** Type, see Section 46.9
- OMPT **state** Type, see Section 40.31
- OMPD **word** Type, see Section 46.17

48.9.2 ompd_get_state Routine

Name: ompd_get_state **Category:** function	**Properties:** C-only, OMPD

Return Type and Arguments

Name	Type	Properties
<return type>	rc	*default*
thread_handle	thread_handle	opaque, pointer
state	word	pointer
wait_id	wait_id	pointer

Prototypes

C

```
ompd_rc_t ompd_get_state(ompd_thread_handle_t *thread_handle,
  ompd_word_t *state, ompd_wait_id_t *wait_id);
```

C

Semantics

The **ompd_get_state** routine returns the state of an OpenMP thread. This routine yields meaningful results only if the referenced thread is stopped. The *thread_handle* argument identifies the thread. The *state* argument represents the state of that thread as represented by a value that **ompd_enumerate_states** returns. On return, if the *wait_id* argument is a non-null value then it points to a handle that corresponds to the *wait_id* wait identifier of the thread. If the thread state is not one of the specified wait states, the value to which *wait_id* points is undefined.

Cross References

- **ompd_enumerate_states** Routine, see Section 48.9.1
- OMPD **rc** Type, see Section 46.9
- OMPD **thread_handle** Type, see Section 46.18.4
- OMPD **wait_id** Type, see Section 46.16
- OMPD **word** Type, see Section 46.17

48.10 Display Control Variables

48.10.1 `ompd_get_display_control_vars` Routine

Name: `ompd_get_display_control_vars` **Category:** function	**Properties:** C-only, OMPD

Return Type and Arguments

Name	Type	Properties
<return type>	rc	*default*
address_space_handle	address_space_handle	opaque, pointer
control_vars	const char * const *	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 OMPD library must satisfy the termination constraints; it may use static or dynamic memory for the vector and/or the strings and is unconstrained in how it arranges them in memory. If it uses dynamic memory then the OMPD library must use the **`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 46.18.1
- **`ompd_initialize`** Routine, see Section 48.1.1
- **`ompd_rel_display_control_vars`** Routine, see Section 48.10.2
- OMPD **`rc`** Type, see Section 46.9

48.10.2 `ompd_rel_display_control_vars` Routine

Name: `ompd_rel_display_control_vars` **Category:** function	**Properties:** C-only, OMPD

Return Type and Arguments

Name	Type	Properties
<return type>	rc	*default*
control_vars	const char * const *	pointer

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 48.10.1
- OMPD **`rc`** Type, see Section 46.9

48.11 Accessing Scope-Specific Information

48.11.1 `ompd_enumerate_icvs` Routine

Name: `ompd_enumerate_icvs` **Category:** function	**Properties:** C-only, OMPD

Return Type and Arguments

Name	Type	Properties
<return type>	rc	*default*
handle	address_space_handle	opaque, pointer
current	icv_id	*default*
next_id	icv_id	pointer
next_icv_name	const char	intent(out), pointer-to-pointer
next_scope	scope	pointer
more	integer	pointer

Prototypes

C

```
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);
```

C

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 Table 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 46.18.1
- OMPD **icv_id** Type, see Section 46.8
- OMPD **rc** Type, see Section 46.9
- OMPD **scope** Type, see Section 46.11

48.11.2 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	*default*
handle	void	opaque, pointer
scope	scope	*default*
icv_id	icv_id	*default*
icv_value	word	pointer

Prototypes

C

```
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);
```

C

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 OMPD 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 46.8 nor the scope for *icv_id* as identified by **ompd_enumerate_icvs**.

Cross References

- OMPD Handle Types, see Section 46.18
- OMPD **icv_id** Type, see Section 46.8
- **ompd_enumerate_icvs** Routine, see Section 48.11.1
- OMPD **rc** Type, see Section 46.9
- OMPD **scope** Type, see Section 46.11
- OMPD **word** Type, see Section 46.17

48.11.3 ompd_get_icv_string_from_scope Routine

Name: ompd_get_icv_string_from_scope **Category:** function	**Properties:** C-only, OMPD

Return Type and Arguments

Name	Type	Properties
<return type>	rc	*default*
handle	void	opaque, pointer
scope	scope	*default*
icv_id	icv_id	*default*
icv_string	const char	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 46.8 or if it does not match the scope for *icv_id* as identified by **ompd_enumerate_icvs**.

Cross References

- OMPD Handle Types, see Section 46.18
- OMPD **icv_id** Type, see Section 46.8
- **ompd_enumerate_icvs** Routine, see Section 48.11.1
- OMPD **rc** Type, see Section 46.9
- OMPD **scope** Type, see Section 46.11

48.11.4 ompd_get_tool_data Routine

Name: ompd_get_tool_data **Category:** function	**Properties:** C-only, OMPD

Return Type and Arguments

Name	Type	Properties
<return type>	rc	*default*
handle	void	opaque, pointer
scope	scope	*default*
value	word	pointer
ptr	address	pointer

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 46.2
- OMPT **data** Type, see Section 40.8
- OMPD Handle Types, see Section 46.18
- OMPD **rc** Type, see Section 46.9
- OMPD **scope** Type, see Section 46.11
- OMPD **word** Type, see Section 46.17

49 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.

OMPD 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.

49.1 `ompd_bp_thread_begin` Breakpoint

Format

C

```
void ompd_bp_thread_begin(void);
```

C

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.

49.2 `ompd_bp_thread_end` Breakpoint

Format

C

```
void ompd_bp_thread_end(void);
```

C

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 48.8.4

49.3 ompd_bp_device_begin Breakpoint

Format

C

```
void ompd_bp_device_begin(void);
```

C

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

49.4 ompd_bp_device_end Breakpoint

Format

C

```
void ompd_bp_device_end(void);
```

C

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 48.8.1

49.5 `ompd_bp_parallel_begin` Breakpoint

Format

C

```
void ompd_bp_parallel_begin(void);
```

C

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 48.5.1
- **`ompd_get_curr_task_handle`** Routine, see Section 48.6.1
- **`parallel`** Construct, see Section 18.1

49.6 `ompd_bp_parallel_end` Breakpoint

Format

C

```
void ompd_bp_parallel_end(void);
```

C

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 48.5.1
- **ompd_get_curr_task_handle** Routine, see Section 48.6.1
- **ompd_rel_parallel_handle** Routine, see Section 48.8.2
- **parallel** Construct, see Section 18.1

49.7 ompd_bp_teams_begin Breakpoint

Format

C

```
void ompd_bp_teams_begin(void);
```

C

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 48.5.1
- **ompd_get_curr_task_handle** Routine, see Section 48.6.1
- **teams** Construct, see Section 18.2

49.8 ompd_bp_teams_end Breakpoint

Format

C

```
void ompd_bp_teams_end(void);
```

C

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 48.5.1
- **ompd_get_curr_task_handle** Routine, see Section 48.6.1
- **ompd_rel_parallel_handle** Routine, see Section 48.8.2
- **teams** Construct, see Section 18.2

49.9 ompd_bp_task_begin Breakpoint

Format

C

```
void ompd_bp_task_begin(void);
```

C

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 selected to execute.

Cross References

- **ompd_get_curr_task_handle** Routine, see Section 48.6.1

49.10 ompd_bp_task_end Breakpoint

Format

C

```
void ompd_bp_task_end(void);
```

C

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 48.6.1
- **ompd_rel_task_handle** Routine, see Section 48.8.3

49.11 ompd_bp_target_begin Breakpoint

Format

C

```
void ompd_bp_target_begin(void);
```

C

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 48.5.1
- **ompd_get_curr_task_handle** Routine, see Section 48.6.1
- **target** Construct, see Section 21.8

49.12 ompd_bp_target_end Breakpoint

Format

C

```
void ompd_bp_target_end(void);
```

C

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 48.5.1
- **ompd_get_curr_task_handle** Routine, see Section 48.6.1
- **ompd_rel_parallel_handle** Routine, see Section 48.8.2
- **target** Construct, see Section 21.8

Part VII

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-group-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 **`ompx`** 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 **`ompx`** 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 **ompx** 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 7.3.5).

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 7.3.7 and Section 7.3.8).

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 11.3).

Fortran

Chapter 10:

- None.

Chapter 11:

Fortran

- **`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 11.3).

Fortran

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_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 15.5).

- **`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 15.7).

- **Declare target directives**: Whether the same version is generated for different devices, or whether a version that is called in a **`target`** region differs from the version that is called outside a **`target`** region, is implementation defined (see Section 15.8).

Chapter 16:

- **`requires` directive**: Support for any feature specified by a *requirement* clause on a **`requires`** directive is implementation defined (see Section 16.5).

- **`message` clause**: Support for a non-constant expression in the *msg-string* argument within a **`target`** region is implementation defined (see Section 16.3).

Chapter 17:

- **`shift` clause**: If the **`shift`** clause is specified without a *shift-list* argument then the *shift-list* values are implementation defined (see Section 17.2).

- **`flatten` construct**: When applied to a non-rectangular loops, the number of inserted empty iterations is implementation defined (see Section 17.4).

- **`flatten` construct**: The particular integer type used to compute the iteration count for the flattened loop is implementation defined (see Section 17.4).

- **`stripe` construct**: If a resulting offsetting loop and a 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 17.9).

- **`tile` construct**: If a resulting grid loop and a 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 17.10).

- **`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 17.10).

- **`unroll` construct**: If no clauses are specified, if and how the loop is unrolled is implementation defined. If the **`partial`** clause is specified without an *unroll-factor* argument then the unroll factor is a positive integer that is implementation defined (see Section 17.11).

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**: The place to which a free-agent thread is bound is implementation defined, except where otherwise specified. 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 team-worker threads and free-agent threads 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 *nteams-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 18.5).
- **`simdlen` 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. Unless the **`num_tasks`** clause is specified with *prescriptiveness* set to **`strict`**, the partitioning of the logical iterations is implementation defined (see Section 20.2).

C++

- **`taskloop` construct**: For firstprivate variables of class type, the number of invocations of copy constructors to perform the initialization is implementation defined (see Section 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).

C++

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 30:

- None.

Chapter 31:

- The **`omp_get_device_distances`** routine returns *ndevs* non-negative values in the array *distances* that are implementation defined relative distances between the place on which the routine is executing and the devices with the device numbers specified by the corresponding array elements of the *devs* argument.

Chapter 32:

- **Rectangular-memory-copying routine**: The maximum number of dimensions supported is implementation defined, but must be at least three (see Section 32.7).

Chapter 33:

- None.

Chapter 34:

- None.

Chapter 35:

- **Lock routines**: If a lock contains a synchronization hint, the effect of the hint is implementation defined (see Chapter 35).

Chapter 36:

- **`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 36.4).

- **`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 36.8).
- **`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 36.9).
- **`omp_display_affinity` routine**: If the *format* argument does not conform to the specified format then the result is implementation defined (see Section 36.10).
- **`omp_capture_affinity` routine**: If the *format* argument does not conform to the specified format then the result is implementation defined (see Section 36.11).

Chapter 37:

- **`omp_display_env` routine**: Whether ICVs with the same value are combined or displayed in multiple lines is implementation defined (see Section 37.4).

Chapter 38:

- None.

Chapter 39:

- **Tool callbacks**: If a tool attempts to register a callback not listed in Table 39.2, whether the registered callback may never, sometimes or always invoke this callback for the associated events is implementation defined (see Section 39.2.4).
- **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 39.2.5).
- **`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 39.2.5).

Chapter 40:

- **`dispatch_chunk` OMPT type**: Whether the chunk of a **`taskloop`** region is contiguous is implementation defined (see Section 40.14).
- **`record_abstract` OMPT type**: The meaning of a **`hwid`** value for a device is implementation defined (see Section 40.24).
- **`state` OMPT type**: The set of OMPT thread states supported is implementation defined (see Section 40.31).

Chapter 41:

- **`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 41.7.5).

Chapter 42:

- **`target_data_op_emi` callbacks**: Whether *dev1_addr* or *dev2_addr* points to an intermediate buffer in some operations is implementation defined (see Section 42.7).

Chapter 43:

- **`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 43.10).
- **`get_partition_place_nums` entry point**: The order of the identifiers returned in the *place_nums* array is implementation defined (see Section 43.12).
- **`get_proc_id` entry point**: The meaning of the numerical identifier returned is implementation defined (see Section 43.13).

Chapter 44:

- None.

Chapter 45:

- None.

Chapter 46:

- None.

Chapter 47:

- **`print_string` callback**: The value of the *category* argument is implementation defined (see Section 47.5).

Chapter 48:

- **handle-comparing routines**: For all types of handles, the means by which two handles are ordered is implementation defined (see Section 48.7).

Chapter 49:

- 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 40.6).
- 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 40.35).
- The **target_data_op**, **target**, **target_map** and **target_submit** callbacks and the associated trace record OMPT type names were deprecated (see Section 42.7, Section 42.8, Section 42.9 and Section 42.10).

B.2 Version 6.0 to 6.1 Differences

- Spatial ICVs (see Section 3.4) were defined and the *dims-modifier* (see Section 18.3) and related routines (see Chapter 29) were added to support multidimensional execution spaces.
- The abstract name specification for the **OMP_PLACES** environment variable was extended to support explicit selection of resources within the place group given by the abstract name (see Section 4.1.6).
- To support thread affinity policies with respect to hierarchical hardware abstractions, syntax for the **OMP_PLACES** environment variable was extended to specify the place list as a hierarchy of place groups, and the behavior of the **spread** affinity policy was updated to apply to this hierarchy (see Section 4.1.6 and Section 18.1.3).
- To allow places and place groups to be reserved for structured threads, syntax for the **OMP_PLACES** environment variable was extended with the structured operator, and the behavior of thread affinity policies was updated to respect these reservations when determining place assignment (see Section 4.1.6 and Section 18.1.3).
- The **OMP_PROC_BIND** environment variable was extended to accept either a **reset** or **noreset** modifier to specify whether the affinity of the encountering thread should be restored at the end of a **parallel** region (see Section 4.1.7).

Fortran

- The *optional-modifier* was added to the *variable-selector* modifier group for the **default** and **defaultmap** clauses to support selection of variables with the **OPTIONAL** attribute (see Section 5.9).

Fortran

- The syntax for atomic structured blocks was extended to permit implicit and explicit type conversion operations on the right-hand side of the assignment (see Section 6.3.3).
- 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 17.5).
- The **implicit** *default-behavior* was added to the **default** and **defaultmap** clauses to specify an implicitly determined default behavior for selected variables and is preferred over the **default** *default-behavior* that is available for the **defaultmap** clause (see Section 7.4 and Section 11.6).
- The *per-thread-modifier* was added to the **reduction** clause to support reduction computations that compute a result for each encountering thread of the construct without combining them across those threads (see Section 8.10).

Fortran

- The **map** clause was extended to support polymorphic variables (see Section 11.5).

Fortran

- 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 11.5.5).
- 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 34.13).
- The **vendor** trait selector was made available for the **device** selector set and **target_device** selector set by adding the corresponding trait to the device trait set and target device trait set (see Section 15.1).
- Support of the **begin declare_variant** directive was extended to Fortran (see Section 15.5).
- The **shift** clause was added to the **tile** construct to adjust the alignment of the grid (see Section 17.2).
- The *sizes-selector* modifier was added to the **sizes** clause to specify the number of tiles (see Section 17.3).
- The **flatten** directive was added to coalesce multiple loops of a loop nest into a single loop (see Section 17.4).
- To support the use of scalable SIMD instructions, the *scaled-modifier* was added to the **simdlen** clause (see Section 18.5.3).
- The restriction that the first chunk must have at least *new_chunk_size* collapsed iterations when the *simd* modifier is used in the **schedule** clause was removed (see Section 19.6.3).
- When the **schedule** clause does not have the *chunk_size* specified, an implementation may use the *chunk_size* of any chunks into which it subdivided the collapsed iteration space for the purpose of computing new_chunk_size (see Section 19.6.3).

Fortran

- The invocation of type-bound procedures inside a **target** region was clarified (see Section 21.8).

Fortran

- The rules for dependence-compatible tasks were updated so that the **taskwait** construct can imply task dependences between the child tasks of an exporting task and a subsequent sibling task of the exporting task (see Section 23.9.6).
- The *depobj-modifier* and *loc-modifier* were added to the **depend** clause, permitting specification of depend objects that are not locator list items (see Section 23.9.6).
- The **omp_default_device** predefined identifier was added and is a conforming device number (see Section 26.1).

C / C++

- The class template **user_mem** has been added to the **omp::allocator** namespace that models the memory allocator concept and permits any non-predefined allocator (see Section 26.8.2).

C / C++

- The **omp_get_filtered_devices** and **omp_get_device_distances** routines were added to support thread-to-device affinity (see Section 31.15 and Section 31.16).
- The wording for the **omp_target_is_accessible** routine was updated to clarify when it returns zero (see Section 32.2.2).
- The **set_selective_callback** entry point was added to the OMPT general entry points so that a first-party tool can register callbacks selectively for specific events (see Section 43.5).
- The **get_uid_from_device** entry point was added to the OMPT device tracing interface so that a first-party tool can ask for the unique identifier of a device (see Section 44.2).

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 3.5) 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).
- The environment variable **OMP_PLACES** was extended to support an increment between consecutive places when creating a place list from an abstract name (see Section 4.1.6).
- The environment variable **OMP_AVAILABLE_DEVICES** was added and the environment variable **OMP_DEFAULT_DEVICE** was extended to support device selection by traits (see Section 4.3.7 and Section 4.3.8).
- The *uid* trait was added to the permissible traits in the environment variables **OMP_AVAILABLE_DEVICES** and **OMP_DEFAULT_DEVICE** and to the target device trait set (see Section 4.3.7, Section 4.3.8 and Section 15.2).

- The environment variable **OMP_THREADS_RESERVE** was added to reserve a number of structured threads and free-agent threads (see Section 4.3.10).

C++

- The **decl** attribute was added to improve the attribute syntax for declarative directives (see Section 5.1).

C++

C

- The OpenMP directive syntax was extended to include C attribute specifiers (see Section 5.1).

C

Fortran

- Support for directives with the pure property in **DO CONCURRENT** constructs has been added (see Section 5.1).

Fortran

- 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).
- The **adjust_args** clause was extended to support positional specification of arguments (see Section 5.2.1 and Section 15.5.2)

Fortran

- 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).

Fortran

C / C++

- 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).

C / C++

- 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 15.5.3).
- 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 7.4).
- 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 7.3.6 and Section 7.3.8).
- The *fallback* modifier was added to the **use_device_ptr** clause and the *need-device-ptr* modifier was extended to accept a *fallback* argument to control the behavior when device pointer lookup or conversion fails (see Section 7.3.6 and Section 15.5.2.1).
- The *saved* modifier, the **replayable** clause, and the **taskgraph** construct were added to support the recording and efficient replay execution of a sequence of explicit-task-generating constructs (see Section 7.5, 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 11.5.1).
- 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 11.5.4 and Section 11.5).
- 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 11.5).
- 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 11.5 and Section 11.6).
- The **map** clause was extended to permit mapping of assumed-size arrays (see Section 11.5).
- The **delete** keyword on the **map** clause was reformulated to be the *delete-modifier* (see Section 11.5).

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 11.3).

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 15.5.2).
- The **dispatch** construct was extended with the **interop** clause to support appending arguments specific to a call site (see Section 15.6 and Section 15.6.1).

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 15.8.1).

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 output 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 17.11 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 17.3).
- The **fuse** construct was added to fuse two or more loops in a canonical loop sequence (see Section 17.5).
- The **interchange** construct was added to permute the order of loops in a loop nest (see Section 17.6).
- The **reverse** construct was added to reverse the iteration order of a loop (see Section 17.7).
- The **split** loop-transforming construct was added to apply index-set splitting to canonical loop nests (see Section 17.8).

- The **stripe** loop-transforming construct was added to apply striping to canonical loop nests (see Section 17.9).
- The **tile** construct was extended to allow grid loops and tile loops to be affected by the same construct (see Section 17.10).
- 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 addded (see Section 31.6).
- 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 18.4).
- 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 18.4).
- 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 18.4).

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 23.9.6).
- The **threadset** clause was added to explicit-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 23.9.7).
- 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 30.2.1).
- 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 30.1.4 and Section 30.1.5).
- 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 31.7 and Section 31.8).
- 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 *nteams-var* and *teams-thread-limit-var* ICVs for specific devices (see Section 31.11, Section 31.12, Section 31.13, and Section 31.14).
- 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 32.8.1 and Section 32.8.2).

Fortran

- Fortran versions of the runtime routines to operate on interoperability objects were added (see Chapter 33).

Fortran

- New routines were added to obtain memory spaces and memory allocators to allocate remote and shared memory (see Chapter 34).
- The **omp_get_memspace_num_resources** routine was added to support querying the number of available resources of a memory space (see Section 34.2).
- The **omp_get_memspace_pagesize** routine was added to obtain the page size supported by a given memory space (see Section 34.3).
- The **omp_get_submemspace** routine was added to obtain a memory space with a subset of the original storage resources (see Section 34.4).
- 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 34.5).
- 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 39.2.4, Section 42.7, Section 42.8, Section 42.9 and Section 42.10).
- 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 40.35). 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 44.7).

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 30.1.2).

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 11.5).

- 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 11.5 and Section 11.7).
- 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 11.5).
- 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 11.3 and Section 15.8).

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 15.6).

Fortran

C / C++

- To improve overall syntax consistency and to reduce redundancy, the delimited form of the **declare_target** directive was deprecated (see Section 15.8.2).

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 18.4).
- 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 23.9.8).
- 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 23.9.6).

- Support for **private** and **firstprivate** as an argument to the **default** clause in C and C++ was added (see Section 7.4).
- 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 7.3.7 and Section 21.8).
- Support was added so that iterators may be defined and used in **map** clauses (see Section 11.5) 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 11.6).
- 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 15.5).
- The **declare_variant** directive was extended with new clauses (**adjust_args** and **append_args**) that support adjustment of the interface between the original function and its function variants (see Section 15.5.4).
- 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 15.6).
- Support was added for indirect calls to the device version of a procedure in **target** regions (see Section 15.8).
- 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 18.6).
- 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 33).
- 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 23.9.6).
- 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 32.2.2).
- 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 32.2.3).
- To support asynchronous device memory management, **omp_target_memcpy_async** and **omp_target_memcpy_rect_async** routines were added (see Section 32.7.3 and Section 32.7.4).
- The **omp_calloc**, **omp_realloc**, **omp_aligned_alloc** and **omp_aligned_calloc** routines were added (see Chapter 34).
- The **omp_display_env** routine was added to provide information about ICVs and settings of environment variables (see Section 37.4).

- 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 40.27).
- The **ompt_state_wait_barrier_implementation** and **ompt_state_wait_barrier_teams** values were added to the **state** OMPT type (see Section 40.31).
- 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 40.33).
- Values for asynchronous data transfers were added to the **target_data_op** OMPT type (see Section 40.35).
- The **error** callback was added (see Section 41.2).
- The **target_data_op_emi**, **target_emi**, **target_map_emi**, and **target_submit_emi** callbacks were added to support external monitoring interfaces (see Section 42.7, Section 42.8, Section 42.9 and Section 42.10).

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 36.8), **omp_get_affinity_format** (see Section 36.9), **omp_display_affinity** (see Section 36.10), and **omp_capture_affinity** (see Section 36.11) 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 discontiguous storage, and the **target_update** construct (see Section 21.9) and the **depend** clause (see Section 23.9.6) 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 7.3.4).
- 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 7.3.6, Section 7.3.8 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 11.5).
- 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 11.5).

- All uses of the **map** clause (see Section 11.5), as well as the **to** and **from** clauses on the **target_update** construct (see Section 21.9) and the **depend** clause on explicit-task-generating constructs (see Section 23.9.6) were extended to allow any lvalue expression as a list item for C/C++.
- The **defaultmap** clause (see Section 11.6) 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 11.7).
- Predefined memory spaces, predefined memory allocators and allocator traits and directives, clauses and routines (see Chapter 13 and Chapter 34) to use them were added to support different kinds of memories.
- Metadirectives (see Section 15.4) and declare variant directives (see Section 15.5) 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 15.8).
- To reduce programmer effort, implicit declare target directives for some procedures were added (see Section 15.8 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 18.4 and Section 19.8).
- The collapse of affected loops that are imperfectly nested loops was defined for **simd** constructs (see Section 18.5), 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 18.5) 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 30.2.1) 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 discontiguous storage.
- The **taskwait** construct was extended to accept the **depend** clause (see Section 23.5 and Section 23.9.6).
- 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 23.9.6).
- The **depend** clause (see Section 23.9.6) 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 31.4) 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 37.2.1 and Section 37.2.2).

- Support for a first-party tool interface (see Chapter 39) was added.
- Support for a third-party tool interface (see Chapter 45) 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 to support stand-alone **ordered** directives (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 7.3.2, Section 7.3.3, Section 11.6 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 7.3.5, Section 7.3.6, 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 11.5).
- To support unstructured data mapping for devices, the **map** clause (see Section 11.5) 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 15.8).
- 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 18.5).
- 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 18.5, 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 23.9.6).
- 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 stand-alone **ordered** directives (see Section 23.9.6).
- 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 30.1.1).
- Device memory routines were added to allow explicit memory allocations, deallocations and transfers and memory associations (see Chapter 32).
- 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 35.1.3 and Section 35.1.4).
- Query routines for thread affinity were added (see Section 36.2 to Section 36.7).
- 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 36.1) 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 24.4), and the **omp_get_cancellation** routine (see Section 37.1) 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 31) 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 18.5).
- Implementation defined task scheduling points for untied tasks were removed (see Section 20.13).
- 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 23.9.6) 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 7.3.3).
- Data environment restrictions were changed to allow Fortran pointers in **firstprivate** (see Section 7.3.3) and **lastprivate** (see Section 7.3.4) 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 20.13.2).
- 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 30.1.3).
- 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 7.3.2, Section 7.3.3, 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 7.3.2, Section 7.3.3, Section 7.3.4, Section 8.10, Section 10.1 and Section 10.2).
- Support for **firstprivate** was added to the **default** clause in Fortran (see Section 7.4).
- 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 7.3.2).
- 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).
- The **omp_get_team_size** routine was added to return the size of the team to which the ancestor thread of the current thread belongs (see Section 27.16).
- The **omp_get_active_level** routine was added to return the number of active parallel regions that enclose the task that contains the call (see Section 27.17).
- Lock ownership was defined in terms of tasks instead of threads (see Chapter 35).

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 18.4).
- 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 18.4).
- The only OpenMP constructs that can be encountered during execution of a **simd** region are SIMDizable constructs (see Section 18.5).
- 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 an explicit-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 24.4).

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*

```
taskloop-composite-dir-name:
    taskloop simd-dir-name

worksharing-loop-composite-dir-name:
    for simd-dir-name
    do simd-dir-name

parallel-combined-dir-name:
    parallel partitioned-worksharing-dir-name
    parallel simd-dir-name
    parallel target-task-generating-dir-name
    parallel task-dir-name
    parallel taskloop-dir-name
    parallel thread-selecting-dir-name

target-combined-dir-name:
    target loop-dir-name
    target parallel-dir-name
    target simd-dir-name
    target task-dir-name
    target taskloop-dir-name
    target teams-dir-name

target_data-combined-dir-name:
    target_data loop-dir-name
    target_data parallel-dir-name
    target_data simd-dir-name

task-combined-dir-name:
    task loop-dir-name
    task parallel-dir-name
    task simd-dir-name

teams-combined-dir-name:
    teams parallel-dir-name
    teams partitioned-nonworksharing-workdist-dir-name
    teams simd-dir-name
    teams target-task-generating-dir-name
    teams task-dir-name
    teams taskloop-dir-name

masked-combined-dir-name:
    masked loop-dir-name
```

```
    masked parallel-dir-name
    masked simd-dir-name
    masked target-task-generating-dir-name
    masked task-dir-name
    masked taskloop-dir-name

single-combined-dir-name:
    single loop-dir-name
    single parallel-dir-name
    single simd-dir-name
    single target-task-generating-dir-name
    single task-dir-name
    single taskloop-dir-name

parallel-worksharing-loop-dir-name:
    parallel worksharing-loop-dir-name

simd-dir-name:
    simd

partitioned-worksharing-dir-name:
    loop-dir-name
    single-dir-name
    worksharing-loop-dir-name
    sections
    workshare

target-task-generating-dir-name:
    target_data-dir-name
    target-dir-name
    target_enter_data
    target_exit_data
    target_update

task-dir-name:
    task-combined-dir-name
    task

taskloop-dir-name:
    taskloop-composite-dir-name
    taskloop

thread-selecting-dir-name:
```

```
    masked-dir-name
    single-dir-name

loop-dir-name:
    loop

parallel-dir-name:
    parallel-combined-dir-name
    parallel

teams-dir-name:
    teams-combined-dir-name
    teams

partitioned-nonworksharing-workdist-dir-name:
    distribute-dir-name
    loop-dir-name
    workdistribute

worksharing-loop-dir-name:
    worksharing-loop-composite-dir-name
    for
    do

single-dir-name:
    single-combined-dir-name
    single

target_data-dir-name:
    target_data-combined-dir-name
    target_data

target-dir-name:
    target-combined-dir-name
    target

masked-dir-name:
    masked-combined-dir-name
    masked

distribute-dir-name:
    distribute-composite-dir-name
    distribute
```

Index

Symbols
_OPENMP macro, 155, 165, 190

A
absent, 399
access type, 589
acq_rel, 531
acquire, 532
acquire flush, 12
adjust-op, 367
adjust_args, 366
affinity, 427
affinity, 486
aggregate maps , 320
align, 340
aligned, 351
alloc_memory, 898
allocate, 341, 343
allocator, 341
allocator_handle type, 590
allocators, 346
alloctrait type, 591
alloctrait_key type, 593
alloctrait_val type, 598
alloctrait_value type, 596
append_args, 369
apply Clause, 407
array mapping, 320
array section mapping, 320
array sections, 185
array shaping, 184
assume, 404
assumed maps, 320
assumed-size array mapping, 320
assumed-type variable mapping, 320
assumes, 403
assumption clauses, 398
assumption directives, 398
asynchronous device memory routines, 658
at, 389
atomic, 541
atomic, 535
atomic construct, 957
atomic_default_mem_order, 393
attach-modifier, 318
attribute clauses, 240
attributes, data-mapping, 305
attributes, data-sharing, 233
auto, 459

B
barrier, 521
barrier, implicit, 523
base language format, 205
begin declare_target, 385
begin declare_variant, 372
begin metadirective, 363
begin assumes, 404
bind, 465
branch, 379
buffer_complete, 838
buffer_request, 837

C
callback_device_host_fn, 907
device_finalize, 835
callbacks, 884
cancel, 564, 820
cancel-directive-name, 563
cancellation constructs, 563
    cancel, 564
    cancellation_point, 568

cancellation points, 566
cancellation_point, 568
canonical loop nest form, 218
canonical loop sequence form, 225
capture, 537
capture, atomic, 541
clause format, 176
clauses
absent, 399
acq_rel, 531
acquire, 532
adjust_args, 366
affinity, 486
align, 340
aligned, 351
allocate, 343
allocator, 341
append_args, 369
apply Clause, 407
assumption, 398
at, 389
atomic, 535
atomic_default_mem_order, 393
attribute data-sharing, 240
bind, 465
branch, 379
cancel-directive-name, 563
capture, 537
collapse, 228
collector, 288
combiner, 284
compare, 538
contains, 400
copyin, 301
copyprivate, 302
counts, 415
data copying, 301
data motion, 330
data-sharing, 240
default, 253
defaultmap, 326
depend, 551
depth, 230
destroy, 201
detach, 487
device, 494
device_safesync, 397
device_type, 493
dist_schedule, 463
doacross, 556
dyn_groupprivate, 348
dynamic_allocators, 393
enter, 308
exclusive, 291
extended-atomic, 537
fail, 538
filter, 443
final, 483
firstprivate, 243
from, 332
full, 419
grainsize, 474
graph_id, 480
graph_reset, 480
has_device_addr, 251
hint, 518
holds, 400
if Clause, 198
in_reduction, 275
inbranch, 380
inclusive, 291
indirect, 386
induction, 276
inductor, 287
init, 199
init_complete, 292
initializer, 284
interop, 375
is_device_ptr, 249
lastprivate, 246
linear, 278
link, 310
local, 300
looprange, 231
map, 310

match, 366
memory-order, 531
memscope, 540
mergeable, 482
message, 389
no_openmp, 401
no_openmp_constructs, 401
no_openmp_routines, 402
no_parallelism, 403
nocontext, 377
nogroup, 530
nontemporal, 439
notinbranch, 380
novariants, 376
nowait, 528
num_tasks, 475
num_teams, 434
num_threads, 425
order, 437
ordered, 229
otherwise, 362
parallelization-level, 561
partial, 420
permutation, 413
priority, 485
private, 242
proc_bind, 429
read, 535
reduction, 269
relaxed, 533
release, 533
replayable, 482
requirement, 392
reverse_offload, 394
safelen, 440
safesync, 430
schedule, 458
self_maps, 397
seq_cst, 534
severity, 390
shared, 241
simd, 562
simdlen, 440
shift, 408
sizes, 409
task_reduction, 274
thread_limit, 495
threads, 562
threadset, 484
to, 331
transparent, 555
unified_address, 395
unified_shared_memory, 396
uniform, 351
untied, 481
update, 536, 550
use, 515
use_device_addr, 252
use_device_ptr, 250
uses_allocators, 347
weak, 539
when, 361
write, 536
collapse, 228
combined and composite directive names, 570
compare, 538
compare, atomic, 541
compilation sentinels, 190, 191
compliance, 16
composition of constructs, 570
compound construct semantics, 576
compound directive names, 570
conditional compilation, 190
consistent loop schedules, 228
construct syntax, 166
constructs
allocators, 346
atomic, 541
barrier, 521
cancel, 564
cancellation constructs, 563
cancellation_point, 568
compound constructs, 576
critical, 519
depobj, 549

device constructs, 493
dispatch, 374
distribute, 461
do, 457
flatten, 410
flush, 545
for, 456
fuse, 411
interchange, 412
interop, 514
loop, 464
masked, 442
ordered, 557, 558, 560
parallel, 421
reverse, 414
scope, 446
sections, 447
simd, 438
single, 445
split, 414
stripe, 416
target, 505
target_data, 502
target_enter_data, 497
target_exit_data, 500
target_update, 511
task, 467
task_iteration, 475
taskgraph, 476
taskgroup, 524
tasking constructs, 467
taskloop, 470
taskwait, 526
taskyield, 489
teams, 431
tile, 417
unroll, 418
work-distribution, 444
workdistribute, 452
workshare, 449
worksharing, 444
worksharing-loop construct, 454
contains, 400
control_tool, 832
control_tool type, 611
control_tool_result type, 612
controlling OpenMP thread affinity, 427
copyin, 301
copyprivate, 302
counts, 415
critical, 519

D

data copying clauses, 301
data mapping, 305
data motion clauses, 330
data sharing, 233
data-localizing control, 293
data-sharing attribute clauses, 240
data-sharing attribute rules, 233
Declare Target, 381
declare variant, 364
declare_induction, 285
declare_mapper, 327
declare_reduction, 281
declare_simd, 377
declare_target, 382
declare_variant, 370
default, 253
defaultmap, 326
depend, 551
depend object, 549
depend type, 604
depend-list-type, 550
dependences, 548
dependences, 821
depobj, 549
deprecated features, 962
depth, 230
destroy, 201
detach, 487
device, 494
device constructs
device constructs, 493
target, 505
target_update, 511
target_enter_data, 497

target_exit_data, 500
device data environments overview, 9
device directives, 493
device information routines, 645
device memory information routines, 658
device memory routines, 657
device_initialize, 834
device_load, 836
device_safesync, 397
device_to_host, 908
device_type, 493
device_unload, 837
dims-modifier, 436
directive format, 168
directive syntax, 166
directive-name-modifier, 192
directives, 990
allocate, 341
assume, 404
assumes, 403
assumptions, 398
begin assumes, 404
begin declare_target, 385
begin declare_variant, 372
begin metadirective, 363
Declare Target, 381
declare variant, 364
declare_induction, 285
declare_mapper, 327
declare_reduction, 281
declare_simd, 377
declare_target, 382
declare_variant, 370
error, 388
groupprivate, 297
memory management directives, 335
metadirective, 360, 363
nothing, 405
requires, 391
scan Directive, 288
section, 449
threadprivate, 293
variant directives, 354
dispatch, 374, 815
dist_schedule, 463
distribute, 461
do, 457
doacross, 556
dyn_groupprivate, 348
dynamic, 459
dynamic thread adjustment, 956
dynamic_allocators, 393

E
enter, 308
environment variables, 144
OMP_AFFINITY_FORMAT, 155
OMP_ALLOCATOR, 162
OMP_AVAILABLE_DEVICES, 157
OMP_CANCELLATION, 157
OMP_DEBUG, 165
OMP_DEFAULT_DEVICE, 159
OMP_DISPLAY_AFFINITY, 155
OMP_DISPLAY_ENV, 165
OMP_DYNAMIC, 146
OMP_MAX_ACTIVE_LEVELS, 147
OMP_MAX_TASK_PRIORITY, 161
OMP_NUM_TEAMS, 152
OMP_NUM_THREADS, 146
OMP_PLACES, 148
OMP_PROC_BIND, 150
OMP_SCHEDULE, 153
OMP_STACKSIZE, 153
OMP_TARGET_OFFLOAD, 159
OMP_TEAMS_THREAD_LIMIT, 152
OMP_THREAD_LIMIT, 147
OMP_THREADS_RESERVE, 160
OMP_TOOL, 163
OMP_TOOL_LIBRARIES, 163
OMP_TOOL_VERBOSE_INIT, 164
OMP_WAIT_POLICY, 154
event, 642
event callback registration, 763
event routines, 642
event_handle type, 583
exclusive, 291
execution control, 748

execution model, 2
extended-atomic, 537

F
fail, 538
features history, 962
filter, 443
final, 483
firstprivate, 243
fixed source form conditional compilation sentinels, 191
fixed source form directives, 175
flatten, 410
flush, 545, 831
flush operation, 10
flush synchronization, 12
flush-set, 10
for, 456
frames, 780
free source form conditional compilation sentinel, 190
free source form directives, 175
free_memory, 898
from, 332
full, 419
fuse, 411

G
general OpenMP types, 580
get_thread_context_for_thread_id, 904
glossary, 20
grainsize, 474
graph_id, 480
graph_reset, 480
groupprivate, 297
groupprivate routines, 717
guided, 459

H
happens before, 12
has_device_addr, 251
header files, 578
hint, 518
history of features, 962
holds, 400
host_to_device, 908

I
ICVs (internal control variables), 131
if Clause, 198
impex type, 604
implementation, 950
implicit barrier, 523
implicit data-mapping, 325
implicit data-mapping attribute rules, 305
implicit flushes, 547
implicit_task, 819
in_reduction, 275
inbranch, 380
include files, 578
inclusive, 291
indirect, 386
induction, 276
inductor, 287
informational and utility directives, 388
init, 199
init_complete, 292
interchange, 412
internal control variables, 950
internal control variables (ICVs), 131
interop, 375
interop type, 583
interop-type, 201
interop_rc type, 584, 585, 587
interoperability, 514
Interoperability routines, 677
intptr type, 580
intrinsic identifiers, 191
introduction, 2
is_device_ptr, 249
iterators, 187

L
lastprivate, 246
linear, 278
link, 310
list item privatization, 237

local, 300
lock routines, 723
lock type, 605
lock_destroy, 829
lock_init, 828
loop, 464
loop concepts, 218
loop iteration spaces, 226
loop iteration vectors, 226
loop-transforming constructs, 406
looprange, 231

M
map, 310
map type decay, 316
map-type, 315
map-type-modifying, 316
mappable storage blocks, 323
mapper, 307
mapper identifiers, 307
masked, 442, 813
match, 366
memory allocator retrieving routines, 702
memory allocators, 336
memory copying routines, 666
memory management, 335
memory management directives
    memory management directives, 335
memory management routines, 685
memory model, 7
memory setting routines, 673
memory space retrieving routines, 685, 709
memory spaces, 335
memory-order, 531
memory_read, 901
mempartition type, 599
mempartitioner routines, 693
mempartitioner type, 599
mempartitioner_compute_proc type, 600
mempartitioner_lifetime type, 600
mempartitioner_release_proc type, 602
memscope, 540
memspace_handle type, 603
mergeable, 482
message, 389
metadirective, 360
metadirective, 363
modifier
    attach-modifier*attach-modifier*, 318
    dims-modifier*dims-modifier*, 436
    directive-name-modifier*directive-name-modifier*, 192
    map-type*map-type*, 315
    reduction-identifier*reduction-identifier*, 268
    ref-modifier*ref-modifier*, 317
    task-dependence-type*task-dependence-type*, 548
modifiers
    adjust-op, 367
    depend-list-type, 550
    interop-type, 201
    map-type-modifying, 316
    reduction-type, 272
    variable-selector, 202
modifying and retrieving ICV values, 137
modifying ICVs, 134
multidimensional support routines, 631
mutex_acquire, 826, 827
mutex_acquired, 828, 829
mutex_released, 830

N
nest_lock, 830
nest_lock type, 606
nesting, 988
no_openmp, 401
no_openmp_constructs, 401
no_openmp_routines, 402
no_parallelism, 403
nocontext, 377
nogroup, 530
nontemporal, 439
normative references, 16
nothing, 405
notinbranch, 380

novariants, 376
nowait, 528
num_tasks, 475
num_teams, 434
num_threads, 425

O
OMP_AFFINITY_FORMAT, 155
omp_aligned_alloc, 712
omp_caligned_alloc, 714
omp_alloc, 711
OMP_ALLOCATOR, 162
omp_ancestor_is_free_agent, 641
OMP_AVAILABLE_DEVICES, 157
omp_calloc, 713
OMP_CANCELLATION, 157
omp_capture_affinity, 746
OMP_DEBUG, 165
OMP_DEFAULT_DEVICE, 159
omp_destroy_allocator, 701
omp_destroy_lock, 728
omp_destroy_mempartition, 696
omp_destroy_mempartitioner, 694
omp_destroy_nest_lock, 729
OMP_DISPLAY_AFFINITY, 155
omp_display_affinity, 745
OMP_DISPLAY_ENV, 165
omp_display_env, 752
OMP_DYNAMIC, 146
omp_free, 716
omp_fulfill_event, 642
omp_get_active_level, 625
omp_get_affinity_format, 743
omp_get_ancestor_thread_num, 623
omp_get_cancellation, 748
omp_get_default_allocator, 708
omp_get_default_device, 646
omp_get_device_allocator, 704
omp_get_device_and_host_allocator, 705
omp_get_device_and_host_memspace, 688
omp_get_device_distances, 656
omp_get_device_from_uid, 649
omp_get_device_memspace, 687
omp_get_device_num, 647
omp_get_device_num_teams, 651
omp_get_device_teams_thread_limit, 653
omp_get_devices_all_allocator, 706
omp_get_devices_all_memspace, 689
omp_get_devices_allocator, 703
omp_get_devices_and_host_allocator, 704
omp_get_devices_and_host_memspace, 687
omp_get_devices_memspace, 686
omp_get_dyn_gprivate_memspace, 721
omp_get_dyn_gprivate_nofb_ptr, 719
omp_get_dyn_gprivate_ptr, 718
omp_get_dyn_gprivate_size, 720
omp_get_dynamic, 618
omp_get_filtered_devices, 655
omp_get_gprivate_limit, 721
omp_get_initial_device, 651
omp_get_interop_int, 678
omp_get_interop_name, 682
omp_get_interop_ptr, 679
omp_get_interop_rc_desc, 683
omp_get_interop_str, 681
omp_get_interop_type_desc, 683
omp_get_level, 623
omp_get_mapped_ptr, 660
omp_get_max_active_levels, 622
omp_get_max_progress_width, 648
omp_get_max_task_priority, 639
omp_get_max_teams, 628
omp_get_max_threads, 616
omp_get_memspaces_num_resources, 690
omp_get_num_devices, 646
omp_get_num_interop_properties, 678

omp_get_num_places, 739
omp_get_num_procs, 647
omp_get_num_teams, 626
omp_get_num_teams_dim, 632
omp_get_num_threads, 615
omp_get_num_threads_dim, 631
omp_get_partition_num_places, 741
omp_get_partition_place_nums, 742
omp_get_place_num, 741
omp_get_place_num_procs, 739
omp_get_place_proc_ids, 740
omp_get_proc_bind, 738
omp_get_schedule, 620
omp_get_submemspace, 691
omp_get_supported_active_league_dims, 635
omp_get_supported_active_levels, 621
omp_get_supported_active_team_dims, 635
omp_supported_teams_dim, 637
omp_get_supported_threads_dim, 636
omp_get_team_num, 627
omp_get_team_num_dim, 633
omp_get_team_size, 624
omp_get_teams_thread_limit, 629
omp_get_thread_limit, 616
omp_get_thread_limit_dim, 634
omp_get_thread_num, 615
omp_get_thread_num_dim, 631
omp_get_uid_from_device, 649
omp_get_wtick, 751
omp_get_wtime, 751
omp_in_explicit_task, 640
omp_in_final, 640
omp_in_parallel, 617
omp_init_allocator, 699
omp_init_lock, 724, 726
omp_init_mempartition, 695
omp_init_mempartitioner, 693
omp_init_nest_lock, 725, 727
omp_is_free_agent, 641
omp_is_initial_device, 650
OMP_MAX_ACTIVE_LEVELS, 147
OMP_MAX_TASK_PRIORITY, 161
omp_mempartion_get_user_data, 698
omp_mempartion_set_part, 697
omp_memspace_get_pagesize, 691
OMP_NUM_TEAMS, 152
OMP_NUM_THREADS, 146
omp_pause_resource, 749
omp_pause_resource_all, 750
OMP_PLACES, 148
omp_pool, 484
OMP_PROC_BIND, 150
omp_realloc, 715
OMP_SCHEDULE, 153
omp_set_affinity_format, 742
omp_set_default_allocator, 707
omp_set_default_device, 645
omp_set_device_num_teams, 652
omp_set_device_teams_thread_limit, 654
omp_set_dynamic, 618
omp_set_lock, 730
omp_set_max_active_levels, 621
omp_set_nest_lock, 731
omp_set_num_teams, 626
omp_set_num_threads, 614
omp_set_schedule, 619
omp_set_teams_thread_limit, 629
OMP_STACKSIZE, 153
omp_target_alloc, 661
omp_target_associate_ptr, 663
omp_target_disassociate_ptr, 665
omp_target_free, 662
omp_target_is_accessible, 659
omp_target_is_present, 658
omp_target_memcpy, 667
omp_target_memcpy_async, 670
omp_target_memcpy_rect, 669
omp_target_memcpy_rect_async,

671
omp_target_memset, 674
omp_target_memset_async, 675
OMP_TARGET_OFFLOAD, 159
omp_team, 484
OMP_TEAMS_THREAD_LIMIT, 152
omp_test_lock, 735
omp_test_nest_lock, 736
OMP_THREAD_LIMIT, 147
OMP_THREADS_RESERVE, 160
OMP_TOOL, 163
OMP_TOOL_LIBRARIES, 163
OMP_TOOL_VERBOSE_INIT, 164
omp_unset_lock, 733
omp_unset_nest_lock, 734
OMP_WAIT_POLICY, 154
ompd_bp_device_begin, 943
ompd_bp_device_end, 943
ompd_bp_parallel_begin, 944
ompd_bp_parallel_end, 944
ompd_bp_target_begin, 947
ompd_bp_target_end, 947
ompd_bp_task_begin, 946
ompd_bp_task_end, 946
ompd_bp_teams_begin, 945
ompd_bp_teams_end, 945
ompd_bp_thread_begin, 942
ompd_bp_thread_end, 942
ompd_dll_locations_valid, 882
ompd_dll_locations, 881
OMPT predefined identifiers, 769
ompt_callback_error_t, 810
OpenMP affinity support types, 608
OpenMP allocator structured blocks, 209
OpenMP atomic structured blocks, 210
OpenMP compliance, 16
OpenMP context-specific structured blocks, 208
OpenMP function dispatch structured blocks, 209
OpenMP interoparability support types, 583
OpenMP list arguments, 181
OpenMP operations, 184
OpenMP parallel region support types, 581
OpenMP resource relinquishing types, 609
OpenMP stylized expressions, 207
OpenMP synchronization types, 604
OpenMP tasking support types, 583
OpenMP tool types, 611
OpenMP types, 205
order, 437
ordered, 229, 557, 558, 560
otherwise, 362

P

parallel, 421
parallel region support routines, 614
parallel_begin, 811
parallel_end, 812
parallelism generating constructs, 421
parallelization-level, 561
partial, 420
pause_resource type, 609
permutation, 413
predefined identifiers, 579
prefer_type, 516
print_string, 908
priority, 485
private, 242
proc_bind, 429
proc_bind type, 608

R

rc, 889
read, 535
read, atomic, 541
read_memory, 902
read_string, 903
collector, 288
combiner, 284
initializer, 284
reduction, 269, 826
reduction and induction data control, 256
task-dependence-type, 268
reduction-type, 272
ref-modifier, 317
relaxed, 533

release, 533
release flush, 12
replayable, 482
requirement, 392
requires, 391
reserved locators, 184
resource relinquishing routines, 749
reverse, 414
reverse_offload, 394
routine argument properties, 580
routine bindings, 579
runtime, 460
runtime library definitions, 578

S
safelen, 440
safesync, 430
saved, 255
scan Directive, 288
sched type, 581
schedule, 458
scheduling, 490
scope, 446
section, 449
sections, 447
self_maps, 397
seq_cst, 534
severity, 390
shared, 241
shift, 408
simd, 438, 562
simd data control, 351
simdlen, 440
single, 445
sizeof_type, 906
sizes, 409
split, 414
stand-alone directives, 174
static, 459
storage block mapping, 323
strip, 416
strong flush, 10
structure mapping, 320
structured blocks, 208
symbol_addr_lookup, 899
sync_region, 823, 824
sync_region_wait, 825
synchronization constructs, 518
synchronization constructs and clauses, 518
synchronization hint type, 606

T
target, 505, 842
target asynchronous device memory routines, 658
target grouppprivate routines, 717
target memory copying routines, 666
target memory information routines, 658
target memory routines, 657
target memory setting routines, 673
target_data, 502
target_data_op, 839
target_data_op_emi, 839
target_emi, 842
target_map, 844
target_map_emi, 844
target_submit, 846
target_submit_emi, 846
target_update, 511
task, 467
task scheduling, 490
task-dependence-type, 548
task_create, 817
task_dependence, 822
task_iteration, 475
task_reduction, 274
task_schedule, 818
taskgraph, 476
taskgroup, 524
tasking constructs, 467
tasking routines, 639
tasking support, 639
taskloop, 470
taskwait, 526
taskyield, 489
teams, 431
teams region routines, 626
thread affinity, 427

thread affinity routines, 738
thread_begin, 808
thread_end, 809
thread_limit, 495
threadprivate, 293
threads, 562
threadset, 484
tile, 417
timer, 751
timing routines, 751
to, 331
tool control, 754
tool initialization, 760
tool interfaces definitions, 757, 881
tool support, 754
tools header files, 757, 881
tracing device activity, 764
transparent, 555
types
- **access**, 589
- **allocator_handle**, 590
- **alloctrait**, 591
- **alloctrait_key**, 593
- **impex**, 604
- **alloctrait_val**, 598
- **alloctrait_value**, 596
- **control_tool**, 611
- **control_tool_result**, 612
- **depend**, 604
- **event_handle**, 583
- **interop_rc**, 584, 585, 587
- **interop**, 583
- **intptr**, 580
- **lock**, 605
- **mempartition**, 599
- **mempartitioner**, 599
- **mempartitioner_compute_proc**, 600
- **mempartitioner_lifetime**, 600
- **mempartitioner_release_proc**, 602
- **memspace_handle**, 603
- **nest_lock**, 606
- **pause_resource**, 609
- **proc_bind**, 608
- **sched**, 581
- **sync_hint**, 606
- **uintptr**, 581

U
uintptr type, 581
unified_address, 395
unified_shared_memory, 396
uniform, 351
unroll, 418
untied, 481
update, 536, 550
update, atomic, 541
use, 515
use_device_addr, 252
use_device_ptr, 250
uses_allocators, 347

V
variable-selector, 202
variables, environment, 144
variant directives, 354

W
wait identifier, 804
wall clock timer, 751
error, 388
weak, 539
when, 361
work, 814
work-distribution
- constructs, 444

work-distribution constructs, 444
workdistribute, 452
workshare, 449
worksharing
- constructs, 444

worksharing constructs, 444
worksharing-loop construct, 454
write, 536
write, atomic, 541
write_memory, 903